Crowdfunded Microfinance*

Runtong Lin
The University of Hong Kong
ryanlin@connect.hku.hk

Tse-Chun Lin
The University of Hong Kong
tsechunlin@hku.hk

This draft: February 2025

Abstract

We model crowdfunding as a device that commits stakeholders to their ESG preferences. We test the model predictions in the context of crowdfunded microfinance by constructing a novel dataset of partnerships between the Kiva crowdfunding platform and 112 microfinance institutions (MFIs) worldwide. In each partnership, Kiva crowdfunders extend MFI-intermediated loans to necessity entrepreneurs in developing countries. Kiva lenders have strong non-pecuniary preferences: they collectively prefer funding female entrepreneurs and demand zero interest on their loans. In our triple-difference framework, we show that the MFIs that crowdfund a large share of their loan portfolios have more non-performing loans, lower asset utilization rates, and higher labor costs than low take-up counterparts. Importantly, the gender gap in financial inclusion narrows, primarily due to less inclusive MFIs catching up. Our results suggest a costly financial inclusion driven by entrepreneurs left behind by banks but picked up by the Kiva lenders with ESG preferences.

JEL: D64, G21, G51, L31, N30, O16, O35, O57, R51

Keywords: Crowdfunding, Financial inclusion, Microfinance, Financial technology (FinTech), Household Finance.

*For valuable comments, we thank Vesa Pursiainen, Yang Ji (discussant), Francesco D'Acunto, Andrea Vedolin, Bohui Zhang, the participants at the 2024 Knut Wicksell Conference on Crypto and Fintech, the 2024 Research in Behaviroal Finance Conference, the 2024 HKU Enterprise Sustainability Conference, the 2024 New Zealand Finance Meeting, and seminar participants at the University of Hong Kong. All errors are ours.

1 Introduction

The recent rise of ESG preferences has challenged profit maximization as the holy grail of firm operation. In this paper, we argue that the crux of the issue is that the market for ESG preferences is incomplete, so the Fisher Separation Theorem is violated, and "green" investors cannot implement their ESG preferences without interfering with the firm's profit maximization. Consider a counterfactual example, in which a perfect market for green firms exists. This market can be a financial market where green firms enjoy a lower cost of capital or a product market where green firms elicit a premium from customers. Then ESG considerations become an input to the profit maximization function because the benefits and costs of ESG are well defined in monetary terms by these markets. This example is counterfactual, and ESG preferences can be at odds with profit maximization, because of market incompleteness.

We propose crowdfunding as a possible solution. Crowdfunding matches ESG stakeholders – investors and/or customers – with firms. The crowdfunding market contracts the outputs, so firms can decide on ESG-oriented investments just like any other capital budgeting decisions in the profit maximization framework. What is new in crowdfunding is that market prices are linked to the ESG preferences of stakeholders, whereas in the conventional intermediary funding market, market prices are linked to the (non-ESG or less-ESG) preferences of financial institutions. Since stakeholders cannot commit credibly to paying more than the market price, firms would optimally choose to only cater to priced preferences. Crowdfunding allows stakeholders to financially commit to their preferences, thereby nesting ESG preferences inside profit maximization.

We empirically test this idea in the context of crowdfunded microfinance. Microfinance, and more generally financial inclusion, is an important and widely recognized ESG domain. The 2006 Nobel Peace Prize is given to the Bangladeshi economist Muhammad Yunus for pioneering the concepts of microcredit and microfinance. Moreover, the MSCI (2022) ESG Ratings include "access to finance" as one of the key ESG issues for financial firms, in which "companies are evaluated on their efforts to expand financial services to historically underserved markets".[1] Importantly, the clash

[1] Mastercard, Visa, PayPal, Bank of America, and many other financial service providers have programs related to microfinance and financial inclusion, and these programs are strong contributors to the ESG ratings of these firms. See Mastercard Labs for Financial Inclusion, Visa Foundation for Financial Inclusion, PayPal Working Capital, and Bank of America Economic Mobility & Community Development Programs.

between ESG preferences and profit maximization is clear in the microfinance context. Predatory interest rates that perpetuate the "debt trap" for borrowers have sparked wide controversy about profit-making institutions in microfinance (e.g., de Quidt, Fetzer, & Ghatak, 2018), and Muhammad Yunus himself is a strong proponent of the incompatibility between ESG preferences (uplifting the poor) and profit maximization (allowing for-profit institutions in microfinance). In this empirical context, we seek support for our idea that crowdfunding can reconcile ESG preference and profit maximization, by examining actual instances of crowdfunded microfinance, and how they affect the ESG and monetary performance of microfinance institutions (MFIs).

We find that crowdfunded microfinance exerts the crowdfunders' ESG preferences at a monetary cost, consistent with our central thesis that crowdfunding acts as a device that financially commits stakeholders to their ESG preferences. In a triple-difference framework, we find that the MFIs that crowdfund a large share of their assets have more non-performing loans, lower asset utilization rates, and higher labor costs, relative to low take-up counterparts. At the same time, the gender gap in financial inclusion narrows, primarily due to less inclusive MFIs catching up. In four regions that report below-average female borrower shares, the percentage of female borrowers increases by ten percentage points, approaching the level of female borrower shares in the other two regions (South Asia and East Asia). The negative consequences of high crowdfunding take-up (higher default rate, lower asset utilization rate, and higher labor costs) are more pronounced for the MFIs that catch up in female financial inclusion (i.e., effectively exerts the crowdfunders' ESG preferences). These financial costs are ultimately borne by the crowdfunders, who contract to lend at zero interest and absorb the potential default loss in order to commit to their ESG preferences.

The crowdfunding partnerships that we analyze are from Kiva, a crowdfunding platform for entrepreneurial loans. Figure 1 visualizes the partnership between Kiva and MFIs and how this partnership connects ESG-aware crowdfunders with necessity entrepreneurs in developing economies. As of 2023, lenders on Kiva have funded 2.4 million loans worth over $2 billion U.S. dollars, connecting two million charitable lenders to five million borrowers across 77 countries. Major destinations of Kiva loans include developing countries across Asia, Africa, and South America. Kiva partners with MFIs in these developing countries, where the MFIs make loans to local entrepreneurs, post the loans on Kiva for access to interest-free credit, and service the collection of interest and principal.

From the MFIs' perspective, the Kiva partnership incurs additional operating costs: for example, Kiva lenders require much information about the loan, leading to higher staff costs for the MFIs. Kiva lenders may also prefer borrowers with higher risk or borrowers that are harder to reach, which also weighs on the MFIs' profitability. The crowdfunders make up for these costs by funding their MFI partners at zero cost of capital, essentially waiving financial returns from their crowdfunding investment in exchange for the exertion of their ESG preferences.

Practically, Kiva enforces the crowdfunders' preferences by sending field workers to partnering MFIs and terminating partnerships with MFIs that do not align with them.[2] These institutional details distinguish Kiva from conventional financial intermediaries and charities: commercial lenders do not price ESG preferences, and charities may not enforce them nearly as closely. In Section 3, we more formally establish the preference of crowdfunders with a stylized model, which modifies Lee and Parlour (2022) and assumes that crowdfunders have a warm-glow component (Andreoni, 1990) in their utility function.

Our empirical analyses first match the MFIs that take up the financial technology with similar MFIs that do not. We use the coarsened exact matching (Iacus, King, & Porro, 2012) method to match the treated MFIs and the control MFIs on quartiles of six variables that prior research shows to predict Kiva partnership, including (i) five-year average assets growth, (ii) return on assets, (iii) percentage of female borrowers, (iv) inflation-adjusted portfolio yield, (v) 30-day portfolio at risk, and (vi) operational self-sufficiency. The treatment group ($N = 112$ MFIs) and the control group ($N = 962$ MFIs) are well matched on the outcome variables that we analyze in our regressions.

We then develop a triple-differences framework to analyze the effects of the Kiva partnership on MFI operation and financial inclusion. The triple-differences regression exploits three sources of variation in outcome variables: before and after the Kiva partnership, between the MFIs that partner with Kiva and those that do not, and between the high take-up MFIs and the low take-up counterparts. We define the take-up rate as the average percentage of an MFI's assets crowdfunded by Kiva throughout the partnership and classify an MFI as a high take-up MFI if its take-up rate is in the top quintile of the distribution, i.e., if the MFI crowdfunds over 22% of its assets. Partnering

[2]For an example of Kiva terminating partnerships with non-aligned MFIs, see https://www.kiva.org/about/where-kiva-works/partners/20. For more details on the field workers sent by Kiva, see Appendix B.

MFIs that do not meet such criteria are classified as low take-up MFIs.

Separating the high take-up and the low take-up MFIs is important in our setting because we would like to focus not on the nominal part of the partnership but on the actual action of crowdfunding. The empirical distribution of the Kiva take-up rate is heavily skewed: for many MFIs, Kiva only funds a small fraction of their assets. By running a triple-differences regression, we establish a *dosage effect* of crowdfunding and raise the hurdle for alternative explanations. For example, one may be concerned that MFIs partner with Kiva to symbolically signal their organizational health, which is already in good shape before the partnership. If such concern is true, Kiva is the outcome, not the cause, of MFIs' improvements. We reject this Kiva-as-a-trophy hypothesis with our triple-differences estimate: trophy-waving MFIs need not seek funding from Kiva. Hence, if they drive the effects, we should see larger improvements in financial inclusion among the low take-up MFIs. Our findings suggest the opposite: Kiva partnerships lead to improvements in MFI outcomes, particularly among the high take-up MFIs that receive proportionally higher Kiva fund flows.

We have four blocks of main findings. First, crowdfunded projects have higher default rates than intermediary-funded projects at similar investment levels. We show that the high take-up MFIs have more non-performing loans. For these MFIs, loans are 1.9% more likely to become overdue by a month, 3.3% more likely to be written off, and 3.8% more likely to eventually default after the Kiva partnership. The sample averages of these statistics are 4.9%, 1.5%, and 1.3%, respectively. Hence, the loan performance effects are large relative to the low baseline non-performing rates. The increase in overdue loans is transient and recovers two years after the take-up, whereas the increase in defaults and more severe delinquencies is persistent in the three-year post-intervention period. Therefore, while temporary factors may explain some of the loan overdue variations, the persistence of higher default rates suggests a more permanent shift in the MFIs' clientele riskiness, consistent with the notion of financial inclusion.

Second, crowdfunded projects incur high information production costs that are reflected in overhead expenses. Following the Kiva partnership, the high take-up MFIs become significantly more labor-intensive, and their payrolls expand substantially as a result. Unconditionally in our

sample, an average loan officer services 333 borrowers and 357 loans every year. However, conditional on the high take-up intensity, an average loan officer only services 235 borrowers and 241 loans. Assuming a constant per-officer labor supply, each loan (borrower) becomes 32% (30%) more time-consuming to service. Presumably in response to such an increase in workloads, the high take-up MFIs increase their payrolls by nearly 50%. These statistics are consistent with the prediction that Kiva loan projects require more intensive information production efforts.

Third, crowdfunded projects face high matching friction that results in low asset utilization rates. The high take-up MFIs have higher total assets, but their gross loan portfolios do not increase one-to-one with total assets. These MFIs' total assets grow by 36%, whereas their asset utilization rates (gross loan portfolio divided by total assets) drop by 17.4 percentage points. Dynamically, the asset utilization rates drop sharply in the year of the Kiva partnership for the high take-up MFIs and remain depressed two years after the partnership.

Fourth, crowdfunded projects overcome the underinvestment of projects initiated by entrepreneurs of discriminated demographics, particularly in regions where such discrimination is more pronounced. We partition our sample based on the World Bank's region definitions, distinguishing between regions with above-average and below-average female borrower shares relative to the 64% panel average. Our analysis reveals a notable catch-up effect: MFIs in the regions with historically lower female entrepreneurial participation (Africa, Eastern Europe and Central Asia, Latin America, and Middle East and North Africa) experience a significant 10.5 percentage point increase in female borrower shares following the Kiva partnership. On the other hand, for the two regions (South Asia and East Asia) in which female borrower inclusion is already high, the female inclusion impact of the Kiva partnership is limited. These findings suggest that the crowdfunding FinTech platform helps narrow the gender gap in financial inclusion, particularly in regions where women historically have less access to finance.

We perform three robustness checks. First, we vary the cutoffs used to define the high Kiva take-up MFIs. We find that the economic magnitude of the triple-difference treatment effect generally increases with the cutoff. Hence, this dosage effect indicates that it is likely the actual Kiva fund inflow, rather than the symbolic partnership, that drives our results. Second, we partition

the sample into two sub-samples, one containing only non-profit MFIs and the other containing only for-profit MFIs. Our results are not entirely driven by either type of MFI. Third, we present evidence against empire-building by the partnering MFIs (Jensen, 1989), as well as evidence against the quiet-life hypothesis (Hicks, 1935), to support our rationalization of the costs associated with financial inclusion. MFIs' operating expenses stay proportional to total assets after the crowdfunding partnerships, suggesting against aggressive, empire-building spending. MFIs with high market power are as affected by the crowdfunding partnerships as other MFIs. MFIs' motive to enjoy a quiet life thus cannot fully explain our results.

Our paper contributes to the literature on microfinance and financial inclusion. According to the definition by the World Bank (2022a), "financial inclusion means that individuals and businesses have access to useful and affordable financial products and services that meet their needs – transactions, payments, savings, credit, and insurance – delivered in a responsible and sustainable way." We show that a crowdfunding market that connects borrowers seeking financial inclusion and *altruistic* lenders could sustainably deliver financial inclusion. Unlike lenders with pecuniary (and potentially predatory) motives, altruistic lenders are aligned with underprivileged borrowers at the preference level. In this context, crowdfunding is a technology that better matches borrowers and lenders of similar preferences, despite geographic, cultural, and institutional barriers. Financial inclusion driven by synergistic preferences in our paper is akin to the financial inclusion driven by the Freedman's Savings Bank that served emancipated blacks after the American Civil War (Stein & Yannelis, 2020), but different from the competition-driven financial inclusion studied in, for example, Célerier and Matray (2019) using U.S. interstate branching deregulation between 1994 and 2005.

In our framework, financial inclusion is not an NPV calculation: we dedicate a non-pecuniary parameter in the utility function for it. Our paper is therefore related to other studies with non-standard preferences, such as Malenko and Malenko (2024) regarding ESG preferences in shareholder voting and Beason and Schreindorfer (2022) regarding disappointment-averse preferences in shaping the average U.S. equity premium. In terms of empirical design, the closest paper to ours is Luo, Ge, and Wang (2022), which shows that after partnering with Kiva, MFIs charge lower interest rates, and their operations become more self-sufficient. We discuss several methodological issues regarding

their empirical design in the Appendix and show that their results are not robust to, for example, a stacked regression approach that properly handles staggered treatment timing (Baker, Larcker, & Wang, 2022; Cengiz, Dube, Lindner, & Zipperer, 2019). Conceptually and more importantly, Kiva has a strong focus on social impact and targets non-pecuniary outcomes in its operations. We show that Kiva is indeed successful in narrowing the gender gap in financial inclusion, and studies that zoom in on pecuniary outcomes may lack comprehensiveness.

Our paper also contributes to the literature on crowdfunding. The crowdfunding market exists because the conventional intermediary funding market is incomplete. Lee and Parlour (2022) compare crowdfunding to conventional intermediary funding and show that crowdfunding can improve efficiency by funding projects with positive consumption utility but negative net present value (NPV). Conventional intermediaries would not fund these projects because consumers cannot credibly commit to paying more than the market price, and the market price implies negative NPV and no participation from the intermediaries. Theoretically, crowdfunding is a commitment device that is more effective for short-duration and low-markup projects.[3] Microfinance loans on Kiva fit this conceptual ideal of crowdfunding, as they are one-year, zero-interest entrepreneurial loans with warm-glow benefits to crowdfunders. The body of our results suggests that crowdfunding microfinance transforms it in multiple ways, with the cost primarily being pecuniary (and borne by the MFIs) and the benefits primarily non-pecuniary (and enjoyed by the crowdfunders). These results therefore rationalize the subsidy (in the form of an interest waiver plus a free credit default swap) given by the crowdfunders to the MFIs that we observe in the data.

2 Institutional Details

The policy scale of microfinance and financial inclusion is large. 1.4 billion adults around the world do not have a bank account in 2021, and many more have no formal access to credit (World Bank, 2022b). Kiva is founded in 2005 in San Francisco as a 501(c)(3) non-profit organization that aims to promote financial inclusion. It is a crowdfunded online lending platform that aims to connect

[3]Liquidity discount and the entrepreneur's market power both reduce the efficiency of crowdfunding, so short-lived and low-markup projects are more suitable for crowdfunding.

individual borrowers and charitable lenders around the world. As of 2023, lenders on Kiva have funded 2.3 million loans worth over $2 billion U.S. dollars, connecting two million charitable lenders to five million borrowers across 77 countries. As an illustration, Figure 2 plots the Kiva fund flow map for the year 2020, in which 58% of Kiva's charitable fund originates from the United States. Major destinations of Kiva fund flow include developing countries across Asia, Africa, and South America.

Kiva partners with MFIs in developing countries. The MFIs screen the borrowers, disburse the loan, post the loan on Kiva for access to interest-free credit, and service the subsequent collection of interest and principal. To become a field partner, an MFI needs to pass Kiva's on-site due diligence on its stability, governance, and risk profile. Once it becomes a Kiva partner, it can post loans online to solicit funding from Kiva lenders, who have access to several categories of information provided by the MFI.

As shown in Figure 3, Kiva lenders can see the borrower's portrait, nationality, and name; the funding progress of the loan on Kiva; and a brief description of the purpose of the loan. Lenders can click on the loans that they are interested in and explore more details on a separate web page, where they can learn about the MFI that intermediates the loan, read a short biography of the borrower, and get other information such as the repayment schedule, loan term, and whether the exchange rate risk is covered. Once the lenders decide to fund the loan, they can check out and transfer the funds to Kiva via PayPal (PayPal, in partnership with Kiva, does not charge a fee on such transfer). By default, lenders lend in $25 increments and can choose to donate some extra money (e.g., $5) to keep Kiva running.

As intermediaries, the MFIs benefit from the participation of Kiva users in three ways. First and perhaps the most obvious, Kiva users lend at zero interest. Second, MFIs can recover the principal much faster by funding their loans on Kiva. On average, Kiva users fund a loan in around 13 days, and the proceeds are transferred to the MFIs on a monthly or quarterly basis. It is thus much faster to recover the principal by posting the loan on Kiva than by waiting until its maturity, which for the average loan takes about 13 months. Third, MFIs can pass on the default risk to Kiva lenders. During the checkout phase of funding a loan, Kiva lenders need to agree to absorb the loss

in the event of default, which shields MFIs from potential default losses.[4]

While Kiva users lend charitably and do not receive interest on their loans, MFIs do charge interest from the borrowers to cover their costs and, in the cases of commercial MFIs, profit from their operations. While interest rates charged for individual loans are usually not made available, MFIs often report the percentage yield at the portfolio level, which is an aggregate proxy for the interest rates charged for each loan in the portfolio. For the MFIs that partner with Kiva, Kiva reports that the average percentage yield at the MFI level is 27% per year. In Appendix A, we compare Kiva with other fundraising platforms and review more studies related to them. In Appendix A we also discuss Kiva's non-core business models as a result of its recent development. These recent non-core business developments, such as a social impact fund, are not in our sample period that ends in 2019.

3 Model

We present a stylized model that captures the key appeal of crowdfunding. The model modifies Lee and Parlour (2022) to account for the fact that Kiva lenders do not charge interest and display behaviors similar to warm-glow giving (Andreoni, 1990). Crowdfunding – a funding mode in which the consumer is also the financier – challenges the Fisher (1930) separation theorem that firms and consumers are separate entities. To zoom in on the Fisher separation and clearly illustrate the effect of crowdfunding, we follow Lee and Parlour (2022) and abstract from other important frictions related to the interaction of principals and agents, such as moral hazard, adverse selection, and asymmetric information.

Time is discrete and there are three types of agents: an entrepreneur, a bank, and a continuum of households of unit measure. At $t = 0$, the penniless entrepreneur is endowed with a project with a fixed cost $I > 0$. The bank and the households have deep pockets and consider funding the project. All agents are risk-neutral, time discount rate is zero, and both the bank and the households have zero outside options.

[4]Kiva also recruits and sends volunteers (Kiva Fellows) to partnering MFIs. See Appendix B for more details.

Production. In each period $t = 1, 2, ...$, the project continues with probability $\delta \in (0, 1)$ and produces output of utility v when consumed. With probability $1 - \delta$ the project fails and produces nothing going forward. The expected present value of the infinite-horizon project, denoted by $\mathbf{V}$, is given by $\mathbf{V} = \Sigma_{t=1}^{\infty} \delta^t v = \frac{\delta}{1-\delta} v$. In this scenario, the first-best outcome is to fund the project whenever $\mathbf{V} > I$.

Allocation without crowdfunding. The output v is shared by the entrepreneur with bargaining power $\alpha \in [0, 1]$ and the bank with bargaining power $1 - \alpha$. The bank and the entrepreneur negotiate a constant periodic interest p, so that the bank gets p and the entrepreneur keeps $v - p$. Under generalized Nash bargaining, p is selected to maximize

$$(v - p)^{\alpha} p^{1-\alpha},$$

and the solution is

$$p = (1 - \alpha) v.$$

The project's expected present value to the bank, denoted by $\mathbf{V}^b$, is given by $\mathbf{V}^b = \Sigma_{t=1}^{\infty} \delta^t p = (1 - \alpha) \mathbf{V}$. If the entrepreneur has any bargaining power, the bank cannot extract the full surplus and values the project at $\mathbf{V}^b < \mathbf{V}$. In this case, the first-best is not achieved because the bank suboptimally rejects the project when $\mathbf{V} > I > \mathbf{V}^b$, i.e., when the project's NPV is negative. The underinvestment becomes more severe when enterpreneur's bargaining power, α, grows large.

Allocation with crowdfunding. Crowdfunding allows the households to directly fund the project. Introducing crowdfunding is beneficial if it funds any project with negative NPV ($\mathbf{V}^b - I < 0$) but positive utility ($\mathbf{V} - I > 0$). Essentially, crowdfunding is an improvement over intermediary funding if the households can internalize a larger fraction of $\mathbf{V}$ than the bank. When the households internalize the entire consumption benefit of the project, the first-best is restored and all productive projects are funded. This can be achieved via setting $\alpha = 0$ for crowdfunding, i.e., the crowdfunded

entrepreneurs produce under perfect competition and the households consume everything. Our empirical setting proposes an alternative way to model this full internalization of $\mathbf{V}$: let households derive warm glow utility $w = v$ from funding the project *without* claiming any of its output. Denote by $\mathbf{V}^c$ the samaritan households' expected present value of the project. The entrepreneur's consumption utility v is fully internalized by the household's warm glow utility w,

$$\mathbf{V}^c = \frac{\delta}{1-\delta} w = \frac{\delta}{1-\delta} v = \mathbf{V}.$$

More generally, crowdfunding is an improvement over intermediary funding for any $w \in [(1-\alpha)\, v, v]$. It is reasonable to assume that Kiva lenders have $w \geq (1-\alpha)\, v$, because they are willing to lend to borrowers typically excluded from the formal credit system. They are also unlikely to have $w > v$ (which creates overinvestment problems), because such a high level of warm glow would imply that the household derives more utility than the entrepreneur who actually consumes the output. In aggregate, as long as w does not greatly exceed v, we should predict crowdfunding to produce a net benefit by alleviating the intermediary's underinvestment problem.

Proposition 1. *If the household derives warm glow utility $w \in [(1-\alpha)\, v, v]$ from funding the entrepreneur's project, crowdfunding is an improvement over intermediary funding towards the first-best, as illustrated in Figure 4.*

3.1 Testable predictions from the model

We develop three testable predictions from the stylized model to organize our empirical analyses. First, since crowdfunding is introduced after intermediary funding, it only funds projects rejected by the bank. In these projects, entrepreneurs are likely to have higher bargaining power that suppresses the bank's valuation of the project. One relevant and widely observable metric here is the default rate of the project. We relate the observed default rate to the bargaining power parameter α to allow a more general interpretation of it.[5] Higher observed default rates indicate higher bargaining power

[5] We assume that the observed default rate is not related to δ in the model, because in the microfinance industry the project's failure may be insured by a joint liability contract or simply unverifiable (de Quidt et al., 2018). Observed defaults can also be simply due to diverting, which relates less to the projects failure probability and more to the bargaining power dynamics.

of the entrepreneur, which makes crowdfunding an appealing choice over intermediary funding.

Hypothesis 1. *Crowdfunded projects have higher default rates than intermediary-funded projects at similar investment levels.*

Next, we model the bank's prior about the entrepreneur's bargaining power and derive statistical discrimination. We demonstrate that households with warm glow utility overcome the bank's underinvestment problem. With a warm glow, crowdfunding bootstraps the financial inclusion of discriminated demographics.

In the model, the bank encounters the entrepreneur at $t = 0$ and infers the entrepreneur's bargaining power α. To deliver clear intuition, we assume that the bank observes the entrepreneur's gender and poverty status and perceives the entrepreneur to have α_m^+ (male above poverty line), α_f^+ (female above poverty line), α_m^- (male below poverty line), or α_f^- (female below poverty line). The bargaining power order of these four types of entrepreneurs is $0 \leq \alpha_m^+ < \alpha_f^+ < \alpha_m^- < \alpha_f^- \leq 1$. A microfoundation for this order is that the bank has encountered more above-poverty male entrepreneurs in the past and can exploit them more effectively. On the other hand, it has encountered fewer below-poverty female entrepreneurs, so it rejects their loans as if α_f^- is large.

The household derives warm glow utility w irrespective of α. When the entrepreneur's bargaining power α is low, the bank's valuation of the project $(1-\alpha)\,v$ is not too far below the households' evaluation $w = v$, so introducing crowdfunding has a relatively low impact. This is the case, for example, for the projects of above-poverty male entrepreneurs, who have the lowest bargaining power among four types of entrepreneurs that we model. Introducing crowdfunding has a large impact when the gap between the households' project valuation and the bank's project valuation grows large. In the stylized model, the gap is $w - (1-\alpha)\,v = \alpha v$, i.e., the entrepreneur's share of the project output. As a result, the financial inclusion effect of crowdfunding is stronger among discriminated demographics, $\alpha_f^- v > \alpha_m^+ v$. Regions in which the under-representation is more severe ($\alpha_m^+ \ll \alpha_f^-$) will see larger borrower compositional changes than regions with more equal financial access ($\alpha_m^+ \approx \alpha_f^-$), because the warm-glow utility of households are unconditional on α.

Hypothesis 2. *Crowdfunded projects overcome the underinvestment of projects initiated by entrepreneurs of discriminated demographics, particularly in regions where such discrimination is more*

pronounced.

We make a final prediction based on the institutional details of our data. Since Kiva does not have the capacity to directly engage entrepreneurs in developing economies, it relies on partnerships with local MFIs to allocate the crowdfunded capital. In these partnerships, Kiva uses zero-cost crowdfunded capital to incentivize MFIs to profile underprivileged entrepreneurs and include them in the credit system. Profiling these entrepreneurs should be more costly than profiling the average entrepreneur due to under-representation. This friction predicts larger payrolls and higher labor costs for MFIs that partner with Kiva. Matching with these entrepreneurs should also be more difficult, similarly due to under-representation. This friction predicts that for MFIs that partner with Kiva, their loan portfolios do not grow one-to-one with their assets, i.e., they utilize their assets at a lower rate.

Hypothesis 3. *Crowdfunded projects incur higher information production costs and are subject to high matching friction. As a result, MFIs that receive crowdfunded capital at a large scale have larger payrolls and lower asset utilization rates.*

4 Data

We obtain data on MFI financial performance from MIX Market, the leading data provider for fundamentals of global microfinance institutions. We collect information on partnerships between MFIs and Kiva from Kiva's API.[6] MIX and Kiva use different sets of codes to identify MFIs and no crosswalk file is readily available to merge the two datasets. Below we describe the name-matching procedure that we use to match 133 MFIs exactly, 30 MFIs fuzzily, and 69 MFIs from archival sources.

We start from a cleaned set of MFI names provided by Kiva, which contain the full names of the MFIs and/or the acronyms associated with the full names. We attempt to match either the

[6] As an advantage of our empirical setup, Kiva maintains a relatively consistent level of data transparency throughout our sample period. This is often not the case for non-charitable lending platforms such as Prosper and Lending Club. Vallée and Zeng (2019) discuss the theoretical motives (e.g., mitigating adverse selection) behind lending platforms' decisions to reduce information transparency.

full names or the acronyms to the cleaned set of MFI names provided by MIX, which similarly may record the MFIs in either their full names, their acronyms, or both. The first step returns 133 matches. Unmatched MFI names from Kiva are then each paired with the MFI name from MIX to which it has the shortest edit distance, also known as the Levenshtein distance. We manually review the pairs of names generated and retain 30 pairs, of which we have strong confidence that the two names in each pair actually refer to the same entity. We also utilize online archives of Kiva's webpages as another source of information to map MFIs from Kiva to MIX. Old versions of the profiles of Kiva partner MFIs contain links to the MFIs' profiles on the MIX platform, but the links were later removed. We manually recovered the links from archival webpages provided by the Internet Archive, a third-party online archive initiative.[7] We are able to link another 69 MFIs from Kiva to MIX with this procedure.

Less than 10% of the MFIs in the MIX dataset ever enter a Kiva partnership. To ensure that the MFIs in the treatment group and the MFIs in the control group are comparable in their ex-ante likelihood of partnering with Kiva, we perform coarsened exact matching (CEM) to match never-treated MFIs with treated MFIs in the same CEM strata. We match on six variables, each coarsened into quartiles and evaluated at the year before the Kiva partnership: five-year average assets growth, return on assets, percentage of female borrowers, inflation-adjusted portfolio yield, 30-day portfolio at risk, and operational self-sufficiency.[8] In all our regressions, we include strata-specific fixed effects to identify the intervention effect on a within-strata basis, i.e., while the treated MFIs and the control MFIs differ in their take-up of Kiva partnership, they are ex-ante similar in key characteristics that predict the partnership.

Table 1 reports summary statistics of the outcome variables that will be used in the regressions in Section 4. We tabulate the summary statistics separately for the treatment group ($N = 112$) and the control group ($N = 962$). The two groups appear to be well matched on the outcome variables that we run regressions on. We proceed with the matched sample for our regression analyses in the next section.

[7] https://archive.org/web/

[8] These characteristics are shown to predict a Kiva-MFI partnership: Dorfleitner, Oswald, and Rohe (2020) show that Kiva is more likely to partner with MFIs that lend to females, charge non-predatory interest rates, and have low-risk portfolios, and are relatively not self-sufficient.

5 Results

Starting in 2006, Kiva partners with new MFIs every year, so the crowdfunding take-ups in our sample are staggered. We address this issue with a stacked regression approach (Baker et al., 2022; Cengiz et al., 2019). Specifically, treated units are grouped into cohorts based on their timing of treatment, denoted by $d = \{2006, 2007, ..., 2015\}$. Within each cohort, treatment units are paired with control units that are similar in key characteristics (see Section 3 for details), forming a cohort-specific dataset d. For example, all MFIs that partnered with Kiva in 2006 would be paired with comparable MFIs that never partnered with Kiva, forming a dataset indexed by $d = 2006$. This cohort-specific dataset would be an MFI-by-year panel that covers the period $t = [d-3, d+2] = \{2003, 2004, 2005, 2006, 2007, 2008\}$.

After generating datasets for all treatment cohorts, all datasets are combined for a single DiD regression. This regression is standard except for the fact that it now includes *cohort-specific* time and unit fixed effects. We choose to use never-treated units as control units since plenty are available. The regression equation is

$$\begin{aligned} Y_{itd} &= \beta_1 \cdot (\text{Treat}_{id} \times \text{Post}_{td} \times \text{High_Kiva_Flow}_{id}) + \\ &\quad \beta_2 \cdot (\text{Treat}_{id} \times \text{Post}_{td}) + \text{FE}_{id} + \text{FE}_{ctd} + \epsilon_{itd}, \end{aligned} \tag{1}$$

where Y_{itd} is one of the outcome variables of MFI i, Treat_{id} is a dummy variable that evaluates to one if MFI i is a Kiva partner, Post_{td} is a dummy that takes one if $t \geq d$, $\text{High_Kiva_Flow}_{id}$ takes one if the MFI's annual average percentage of assets funded by Kiva is above 22% (i.e., within the top quintile of the Kiva take-up rate distribution), FE_{id} and FE_{ctd} are cohort-strata-specific MFI and country-year fixed effects, and ϵ_{itd} is the error term.

Our considerations in choosing the threshold are twofold. First, from the data we only observe outcome variations for the MFI as a whole, and not outcome variations for the business segment funded by Kiva. Therefore, for Kiva to actually drive the MFI-level variations and for us to observe it from the data, the Kiva take-up rate threshold cannot be low. Second, the empirical distribution of Kiva take-up rates is heavily right-skewed: the top quintile contains more variation than the other

four quintiles combined. If we conjecture that the equilibrium distribution of the take-up rates is more symmetric than what we observe (for example, by expecting Kiva to grow as the crowdfunding technology becomes more well-known), we should focus more on the observations at the right tail of the distribution.[9]

Regression (1) is a triple-difference regression that exploits three sources of variations in outcome variables: outcome variation before and after the Kiva partnership (Post_{td}), outcome variation between MFIs taking up the partnership and MFIs that do not (Treat_{id}), and outcome variation between high take-up MFIs and low take-up MFIs ($\text{High_Kiva_Flow}_{id}$). Our prediction is that $\beta_1 \neq 0$: the impact that Kiva partnership has on MFI outcomes should be even stronger among MFIs that intensively adopt the technology.

An equally important prediction is, for an unobserved confounding variable to explain our findings, it must generate outcome variation not across the take-up line (β_2; Treat_{id}) but across the intensity line (β_1; $\text{High_Kiva_Flow}_{id}$). For example, a critic may argue that Kiva is a trophy: MFIs partner with Kiva to symbolically signal their organizational health, which is already in good shape before the partnership. Therefore, the critic may argue that the Kiva partnership is the outcome, not the cause, of the MFIs' improvements. We can directly test this hypothesis in our framework: the Kiva-as-a-trophy hypothesis would predict that MFIs taking up the partnership at low intensity see the strongest improvement in the outcomes, since these are the trophy-waving, already-healthy MFIs that need not seek actual funding from Kiva. As we will report in the remaining part of this section, this is most often not the case: the empirical findings largely support our hypothesis that Kiva partnership leads to improvements in MFI outcomes *particularly* among MFIs that intensively adopt the technology and receive higher fund flows.

[9]This is akin to giving larger weights to survey responses from subgroups that are under-represented in the survey than in the population, in order to restore population representativeness. We do not establish a set of weights formally for two reasons. First, the regression already employs analytical weights from the CEM procedure. Second, while we draw from our findings to suggest that the equilibrium empirical distribution should be less skewed, without specifying the counterfactual distribution parameters we cannot pin down the weights.

5.1 FinTech Adoption and Loan Performances

We run two regressions related to the MFIs' loan performances. The first regression uses the outcome variable of 30-days portfolio at risk, defined as the value of loans overdue by more than 30 days divided by the value of the gross loan portfolio. The second regression uses the outcome variable of loan write-off ratio, the value of written-off (but potentially still recoverable) loans divided by gross loan portfolio. The third regression uses a related outcome variable, loan loss rate, defined as the value of loans written off and considered unrecoverable divided by the value of gross loan portfolio. As is well known in the literature, the sample average of these statistics is low: the average 30-days portfolio at risk is 4.9%, the average loan write-off ratio is 1.5%, and the average loan loss rate is 1.3%. We again cluster standard errors at the MFI level to account for serially correlated errors.

Table 2 reports the regression results. The high take-up MFIs have more non-performing loans than the low take-up MFIs. For these MFIs, loans are 1.9% more likely to become overdue by a month, 3.3% more likely to be written off, and 3.8% more likely to eventually default after the Kiva partnership. The dynamics plotted in Figure 6 suggest that the increase in overdue loans is transient and recovers at $t = d + 2$, i.e., two years after the take-up, whereas the increase in loan write-offs and loan defaults is more persistent and does not revert in the three-year post-intervention period. Therefore, while it is likely that loan overdues are partially driven by temporary factors (e.g., loan officer shortages, which we investigate later), the persistence of more severe delinquencies suggests a more permanent shift in the MFIs' clientele riskiness.

These results support the loan performance hypothesis from our stylized model, as illustrated in Figure 5. We bin MFIs according to their 30-day portfolio at risk with a bin width of one percentage point. The dashed line overlays the intermediary's underinvestment problem. As predicted by the model, high-takeup MFIs are more concentrated in the region of less performant loans, i.e., they pick up loans rejected by intermediaries when they receive crowdfunded capital from Kiva.

5.2 FinTech Adoption and Payrolls

To analyze how the Kiva partnership affects the workers employed by the MFIs, we choose four variables as the outcome of regression 1: the number of loans handled per year per loan officer, the number of borrowers handled per year per loan officer, the natural log of the MFI's total number of employees, and the natural log of the MFI's number of loan officers. Following the definition by the MIX, loan officers are client-facing employees who directly handle the loans. We again reported standard errors clustered at the MFI level. Results are in Table 3.

Following the Kiva partnership and relative to low take-up MFIs, the high take-up MFIs become significantly more labor intensive, and their payrolls substantially expand as a result. Unconditionally in our sample, an average loan officer services 333 borrowers and 357 loans every year. For high take-up MFIs, an average loan officer now services 235 borrowers and 241 loans, so assuming constant per-officer labor supply, each loan (borrower) now becomes 32% (30%) more time-consuming to service.

The increase in workload per borrower aligns with the institutional details we covered in Section 2. Kiva requires MFIs to produce detailed information about the borrowers and the loans before funding the loans online. The requirement includes standard information about the borrowers and the loans, as well as some information that is less standard and likely requires extra effort. For example, Kiva emphasizes that having a well-taken photo of both the borrowers and their means of subsistence (such as livestock, village stores, and farmlands) helps to accelerate the loan funding process. Even for the standard information that the MFIs may have already collected, more work is needed for translation into English and posting it online. We see, likely in response to such an increase in workloads, that MFIs substantially increase their payrolls by nearly 50% (Columns 3 and 4, Table 3; see also Figure 7 for the effect dynamics).

5.3 FinTech Adoption and Asset Utilization

We run three regressions related to asset utilization, each with a different outcome variable. The three outcomes are the natural log of total assets, the natural log of the gross loan portfolio, and

the ratio of the gross loan portfolio to total assets (i.e., the asset utilization ratio). Standard errors are clustered at the MFI level, so we account for the potential error correlation within MFIs across years.

Table 4 reports the regression results. Relative to the low take-up MFIs, the high take-up MFIs have higher total assets, but their gross loan portfolios do not increase one-to-one with total assets. These MFIs' total assets grow by an average of 36%, whereas their asset utilization ratio drops by 17.4% from the sample average of 81.5%. The under-utilization of crowdfunded capital is consistent with the prediction that crowdfunded projects are more likely to have limited capacity.

We plot the dynamic version of regression (1) in Figure 8 to see how the asset utilization ratio changes over time. For completeness, the dynamics of log assets and log gross loan portfolio are also plotted. Assets and gross loan portfolios both increase sharply in the year of the Kiva partnership and remain at high levels up to two years after the partnership for high take-up MFIs. As a result, asset utilization is persistently low throughout the post-intervention periods.

5.4 FinTech Adoption and Financial Inclusion

Kiva makes an ambitious claim on its front page to "make a loan, change a life." Does it do what it says, or fail to deliver on the promise of financial inclusion?

In the full panel of our sample, the average female borrower share for an MFI-year observation is 64%. Two regions, South Asia and East Asia, have average female borrower shares *much* higher than 64%, reporting at 85% and 72%, respectively. The other four regions, Africa, Eastern Europe and Central Asia, Latin America, and Middle East and North Africa, have average female borrower shares below this threshold. These regions are defined by the World Bank and widely used. More importantly, the wide gap between the two Asian regions and the rest of the sample makes a large room for the laggards to catch up, and opens the empirical question about whether crowdfunding FinTech aids this catch-up process.

We arrange our panel data into three samples: the full sample, an above-average sample comprising MFIs from regions (South Asia and East Asia) where female borrower share exceeds

the 64% panel average, and a below-average sample including MFIs from regions (Africa, Eastern Europe and Central Asia, Latin America, and Middle East and North Africa) with female borrower shares below 64%. We conduct the partition before matching MFIs with comparable control units, so the numbers of observations in the two subsamples do not add up to the number of observations in the full sample. For the same reason, the average female borrower shares reported below the regression in Table 5 also differ from the descriptive statistics reported above. We anchor our sample partition on the well-known World Bank region delineations, so that what goes into each of the two subsamples is not affected by the statistical jiggles.

Table 5 reports the results of the financial inclusion regression with female borrower share as the outcome variable. The first column reports full sample regression results. For the high take-up MFIs, female borrower share grows by 5.1 percentage points ($t = 1.3$) relative to the low take-up MFIs. The second column reports regression results from the above-average sample. For the MFIs in this sample, female borrower shares are already high before the intervention, so the impact of crowdfunding FinTech is negligible.

The third column reports regression results from the below-average sample. High take-up MFIs in this sample experience a 10.5-percentage-point increase ($t = 2.7$) in female borrower shares relative to low take-up MFIs after the Kiva partnership. These findings indicate a catch-up effect: as a result of the Kiva partnership, MFIs in regions with lower female entrepreneurial participation show marked improvements in their inclusivity of female borrowers. For the two regions (South Asia and East Asia) in which female borrower inclusion is already high, the impact of Kiva is limited.

5.5 Costs Without Benefits?

Regression results in the previous section suggest that the Kiva partnership has a limited effect on the financial inclusion of female entrepreneurs in regions where inclusion is already high. A follow-up question would be whether MFIs in these already inclusive regions, largely unaffected in their female borrower shares, report reduced operating efficiency after partnering with Kiva. Reduced operating efficiency and flat female borrower shares in tandem would suggest concerning cases in

which the Kiva injection of crowdfunded capital brings costly consequences and yet little purported benefit.

We show that while the negative consequences of the Kiva partnership are weaker for MFIs in already inclusive regions, these consequences are nevertheless present. Columns (4) to (6) of Table 6 report results of asset utilization regression, default rate regression, and labor cost regression using the sample of MFIs from South Asia and East Asia, i.e., the two regions with average female borrower shares above the panel sample average of 64%. We show in column (2) of Table 5 that high Kiva take-up has no effect on female entrepreneur inclusion in this sample. However, MFIs in this sample report a 19.9 percentage point lower asset utilization rate and a 2.5 percentage point higher default rate.

5.6 Imprecise Financial Inclusion Results

For completeness, we analyze, again using regression 1, other commonly reported measures of financial inclusion to see if MFIs improve on these measures after partnering with Kiva. The measures are the natural log of the number of active borrowers, the average loan size (a proxy for borrower income, assuming that borrowers only get access to loans proportional to their income), and the percentage of group borrowers. The results are reported in Table 7. These results are statistically imprecise but economically meaningful, similar to many previous results in the microfinance literature (Banerjee, Karlan, & Zinman, 2015). The effects on all three measures are correctly signed: for high take-up MFIs, the number of borrowers grows, the borrower's average income is lower, and the borrowers are less likely to be in a borrower group. The last result is consistent with the substitution between the MFIs' information production about the borrower, which intensifies due to the Kiva partnership, and their reliance on the borrower group to screen borrowers and enforce repayments. Overall, there is some evidence that crowdfunding capital from Kiva increases the MFIs' outreach to lower-income borrowers and borrowers who are not part of a borrower group. The lack of statistical precision may stem from the fact that Kiva does not pursue improvements in these dimensions as explicitly as it encourages female entrepreneur inclusion.

6 Robustness and Supplementary Checks

6.1 Different cutoffs to separate high and low Kiva take-up MFIs

In our baseline specification, we define high Kiva take-up MFIs to be the top 20% of MFIs ranked by percentages of assets funded by Kiva. We choose this cutoff for two reasons. First, considering the empirical distribution of Kiva take-up rates, this cutoff indicates that all MFIs in the high take-up group have at least 22% of assets funded by Kiva. This threshold appears more reasonable to us than, for example, using the median take-up rate of 5% as the cutoff. Since our outcome variables are measured for the whole institution and not just the part that Kiva funds, setting a high threshold helps to reduce estimation noise.

For completeness, we vary cutoffs used to define $\text{High_Kiva_Flow}_{id}$ from 12% of MFI's total assets (the sample's average and 70th percentile) to 22% of MFI's total assets (our baseline specification) at two percent increments. We repeat regression (1) with six different cutoffs $\{12\%, 14\%, 16\%, 18\%, 20\%, 22\%\}$ and nine different outcome variables, producing $6 \times 9 = 54$ estimates of β_1.[10] Figure A1 organizes these estimates by outcome variables and by cutoffs.

We have two observations. First, with a few exceptions, the triple-differences effect size generally increases with the cutoff. This dosage effect indicates that it is the actual Kiva fund inflow, rather than the nominal partnership or some unobserved confounders, that drives the results. Second, the results on payroll and labor intensity show a spike in effect size and statistical significance at the 22% cutoff. Since the top quintile of a right-skewed distribution contains the largest variation (the take-up rate varies from 22% to 73% for 19 different MFIs), observations there are relatively thin. Lowering the threshold to 18% adds another 6 MFIs to the high take-up group, so the four-percent range $[18\%, 22\%]$ is heavily over-represented in the high take-up group. We more formally make this point in Figure A2: the take-up rates are more uniformly distributed in the region $[22\%, 73\%]$, but extending the region leftward significantly skews the distribution. Setting the threshold at 22%, as we do in the baseline specification, ensures that we cover a wide range

[10] Four outcome variables $\{ln(GLP), ln(NumBorrowers), LoanPerBorrower, PercentFemale\}$ have statistically insignficant β_1 in the main specification; they also have statistically insignficant β_1 across the range of cutoffs and are not reported in Figure A1 for brevity.

of take-up rates in the high take-up group of MFIs, without some particular take-up rate values dominating the distribution.

6.2 Do non-profit MFIs differ from for-profit MFIs?

Not systematically. We partition the sample into two sub-samples, one containing only non-profit MFIs and the other containing only for-profit MFIs. The MFI profit status data is a snapshot as of 2020 (the latest MIX data release) and we cannot observe whether the MFIs switched profit status at some point in the past. With this caveat in mind, we rerun regression (1) and report the triple-differences estimates β_1 for all outcome variables in Figure A3. There are instances where the effects of large Kiva capital injection lose statistical significance at conventional levels when we partition the sample, but no instance where we observe an opposite (yet still significant) effect in the partitioned sample. The sub-sample estimates in Figure A3 indicate that while partitioning the sample by MFI profit status reduces the precision of our estimates, our results are not entirely driven by either type of MFIs.

6.3 Empire building or a quiet life?

The key appeal of crowdfunding is the alignment of preferences. To highlight this appeal, our stylized model abstracts away from important frictions such as agency costs. Empirically, agency issues likely exist between the principal (Kiva crowdfunders) and the agent (MFIs). One possibility is that with ample supply of capital, the MFIs engage in empire-building behaviors (Jensen, 1989) and spend aggressively, leading to high expense ratios. Another possibility is that with default risk transferred to Kiva lenders, the MFIs can enjoy a quiet life (Hicks, 1935) by not screening and monitoring borrowers as prudently. How much impact do these agency issues have on our results?

To provide some evidence on empire building, we test if operating expenses (scaled by total assets, total revenue, or total gross loan portfolio) are higher following the Kiva partnership. As reported in Table A2, we find no evidence across different specifications that crowdfunding leads to increased operational spending. These results suggest that while empire-building motives may

affect corporate decisions in developed economies, their impact on MFIs in developing economies is limited.

To provide some evidence on quiet lives, we partition the sample according to regional average MFI profit margins, and see if the effects of crowdfunding are weaker in regions where MFIs' market power is high. Three World Bank regions have average MFI profit margins above the sample average: East Asia (World Bank Region 2), Eastern Europe (World Bank Region 3), and Latin America (World Bank Region 4). We exclude East Asia from this analysis because female financial inclusion in this region is already high, leaving little room for further inclusion improvements. The partitioned sample therefore includes MFIs from Eastern Europe and Latin America, which enjoy a high degree of market power and large room for financial inclusion improvements. If the quiet life hypothesis of Hicks (1935) is true, in this sample we should see little impact of crowdfunding on MFI outcomes, since these MFIs are monopolistic and inclined to stay with the status quo.

The quiet life hypothesis is rejected by regression results reported in Table A3. Even for MFIs who have a fair amount of market power, crowdfunding still has impacts on their asset utilization, default rate, labor costs, and female financial inclusion. The magnitude of the effects is overall comparable to the effects reported in Table 5 and Table 6. We interpret these results as evidence against the conjecture that MFIs shirk after they transfer the default risk to Kiva lenders. These results are also consistent with institutional details: MFIs are incentivized to stay on the Kiva platform and maintain access to the low-cost crowdfunded capital, so they have reputational concerns against shirking.

7 Conclusion

We study the effect of crowdfunding on microfinance. We collect new data on a set of crowdfunding shocks to microfinance institutions (MFIs). Via crowdfunding on the Kiva platform, MFIs intermediate debt capital that targets necessity entrepreneurs in developing economies. In a triple-difference framework, we show that the MFIs that crowdfund a large share of their assets have more non-performing loans, lower asset utilization rates, and higher information production costs, relative

to low take-up counterparts. The gender gap in financial inclusion narrows, primarily due to less inclusive MFIs catching up. In four regions that report below-average female borrower shares, the percentage of female borrowers increases by ten percentage points, approaching the level of female borrower shares in the other two regions (South Asia and East Asia).

We rationalize the costly inclusion results with a model of limited commitment, in which households value inclusive projects but cannot commit to paying extra for them. When crowdfunding is introduced, households fund these projects themselves and eliminate the commitment problem. These projects have worse cash flows but are socially desired, so they are left behind by the banks and picked up by the crowdfunders.

Crowdfunding platforms are useful to crowdfunders for two purposes. First, they have an informational function that lets funders accurately match with their preferred projects (in Kiva's case, inclusive projects). Second, they have a settlement function, allowing funders to settle (and thus commit to) their payments. They are therefore distinct from an information dashboard or a money wiring service; they are a combination of both. It is plausible that incumbent intermediaries (e.g., banks and MFIs) are able to perform both functions in-house more efficiently. However, these intermediaries do not fully internalize the preferences of their principals, so the incumbent model leaves something to be desired.

Technological progress in decentralization would help to resolve agency issues without compromising as much on economies of scale, thus making the trade-off more favorable for decentralized protocols. Results in this paper indicate that crowdfunding microfinance delivers its intended benefits, but not without substantive costs stemming from information frictions and moral hazard. We hope these empirical results inform best practices for future decentralization efforts.

References

Andreoni, J. (1990). Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving. *The Economic Journal*, *100*(401), 464–477.

Baker, A. C., Larcker, D. F., & Wang, C. C. Y. (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, *144*(2), 370–395.

Balyuk, T., & Davydenko, S. (2023). Reintermediation in FinTech: Evidence from Online Lending.

Banerjee, A., Karlan, D., & Zinman, J. (2015). Six Randomized Evaluations of Microcredit: Introduction and Further Steps. *American Economic Journal: Applied Economics*, *7*(1), 1–21.

Beason, T., & Schreindorfer, D. (2022). Dissecting the Equity Premium. *Journal of Political Economy*, *130*(8), 2203–2222.

Burtch, G., Carnahan, S., & Greenwood, B. N. (2018). Can You Gig It? An Empirical Examination of the Gig Economy and Entrepreneurial Activity. *Management Science*, *64*(12), 5497–5520.

Célerier, C., & Matray, A. (2019). Bank-Branch Supply, Financial Inclusion, and Wealth Accumulation. *The Review of Financial Studies*, *32*(12), 4767–4809.

Cengiz, D., Dube, A., Lindner, A., & Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, *134*(3), 1405–1454.

Chava, S., Ganduri, R., Paradkar, N., & Zhang, Y. (2021). Impact of marketplace lending on consumers' future borrowing capacities and borrowing outcomes. *Journal of Financial Economics*, *142*(3), 1186–1208.

de Quidt, J., Fetzer, T., & Ghatak, M. (2018). Market Structure and Borrower Welfare in Microfinance. *The Economic Journal*, *128*(610), 1019–1046.

Dorfleitner, G., Oswald, E.-M., & Rohe, M. (2020). The access of microfinance institutions to financing via the worldwide crowd. *The Quarterly Review of Economics and Finance*, *75*, 133–146.

Fisher, I. (1930). The theory of interest. *New York*, *43*, 1–19.

Gafni, H., Marom, D., Robb, A., & Sade, O. (2021). Gender Dynamics in Crowdfunding (Kickstarter): Evidence on Entrepreneurs, Backers, and Taste-Based Discrimination*. *Review of Finance*, *25*(2), 235–274.

Hertzberg, A., Liberman, A., & Paravisini, D. (2018). Screening on Loan Terms: Evidence from Maturity Choice in Consumer Credit. *The Review of Financial Studies*, *31*(9), 3532–3567.

Hicks, J. R. (1935). Annual Survey of Economic Theory: The Theory of Monopoly. *Econometrica*, *3*(1), 1–20. JSTOR: 1907343

Hildebrand, T., Puri, M., & Rocholl, J. (2017). Adverse Incentives in Crowdfunding. *Management Science*, *63*(3), 587–608.

Iacus, S. M., King, G., & Porro, G. (2012). Causal Inference without Balance Checking: Coarsened Exact Matching. *Political Analysis*, *20*(1), 1–24.

Iyer, R., Khwaja, A. I., Luttmer, E. F. P., & Shue, K. (2016). Screening Peers Softly: Inferring the Quality of Small Borrowers. *Management Science*, *62*(6), 1554–1577.

Jensen, M. C. (1989). Eclipse of the Public Corporation. *Harvard Business Review*.

Lee, J., & Parlour, C. A. (2022). Consumers as Financiers: Consumer Surplus, Crowdfunding, and Initial Coin Offerings. *The Review of Financial Studies*, *35*(3), 1105–1140.

Lin, M., Prabhala, N. R., & Viswanathan, S. (2013). Judging Borrowers by the Company They Keep: Friendship Networks and Information Asymmetry in Online Peer-to-Peer Lending. *Management Science*, *59*(1), 17–35.

Lin, M., & Viswanathan, S. (2016). Home Bias in Online Investments: An Empirical Study of an Online Crowdfunding Market. *Management Science*, *62*(5), 1393–1414.

Lin, T.-C., & Pursiainen, V. (2022). Regional social capital and moral hazard in crowdfunding. *Journal of Business Venturing*, *37*(4), 106224.

Lin, T.-C., & Pursiainen, V. (2023). Gender differences in reward-based crowdfunding. *Journal of Financial Intermediation*, *53*, 101001.

Luo, X., Ge, L., & Wang, C. ((2022). Crowdfunding for Microfinance Institutions: The New Hope? *MIS Quarterly*, *46*(1), 373–400.

Malenko, A., & Malenko, N. (2024). Voting Choice. Social Science Research Network: 4390367

MSCI. (2022). MSCI ESG Ratings Key Issue Framework. Retrieved January 22, 2025, from https://www.msci.com/our-solutions/esg-investing/esg-ratings/esg-ratings-key-issue-framework

Paravisini, D., Rappoport, V., & Ravina, E. (2017). Risk Aversion and Wealth: Evidence from Person-to-Person Lending Portfolios. *Management Science*, *63*(2), 279–297.

Stein, L. C. D., & Yannelis, C. (2020). Financial Inclusion, Human Capital, and Wealth Accumulation: Evidence from the Freedman's Savings Bank. *The Review of Financial Studies*, *33*(11), 5333–5377.

Vallée, B., & Zeng, Y. (2019). Marketplace Lending: A New Banking Paradigm? *The Review of Financial Studies*, *32*(5), 1939–1982.

Wei, Z., & Lin, M. (2017). Market Mechanisms in Online Peer-to-Peer Lending. *Management Science*, *63*(12), 4236–4257.

World Bank. (2022a). Financial Inclusion. Retrieved January 21, 2025, from https://www.worldbank.org/en/topic/financialinclusion/overview

World Bank. (2022b). *The Global Findex Database 2021*. Text/HTML. 2022.

Xu, Y., & Ni, J. (2022). Entrepreneurial Learning and Disincentives in Crowdfunding Markets. *Management Science*, *68*(9), 6819–6864.

Younkin, P., & Kuppuswamy, V. (2018). The Colorblind Crowd? Founder Race and Performance in Crowdfunding. *Management Science*, *64*(7), 3269–3287.

Zhang, J., & Liu, P. (2012). Rational Herding in Microloan Markets. *Management Science*, *58*(5), 892–912.

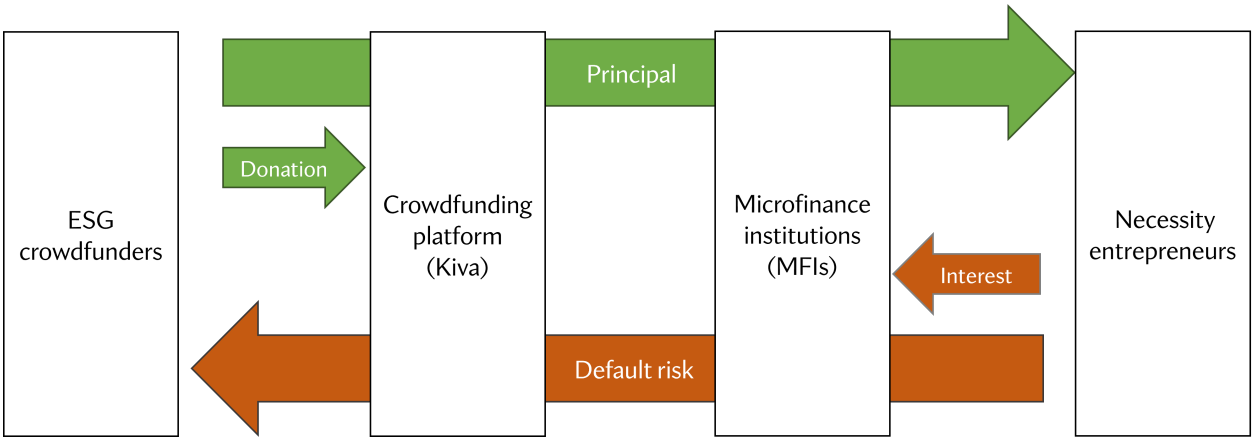


Figure 1: *Crowdfunded Microfinance.* This chart demonstrates the partnership between the crowdfunding platform, Kiva, and microfinance institutions (MFIs) in the developing economies. Starting from the left, crowdfunders with ESG preferences (e.g., the preference to improve financial inclusion or promote green agriculture) seek loan projects that align with their preferences. Kiva, the crowdfunding platform that we study, provide an online marketplace for ESG crowdfunders to match with loan projects. These loan projects are posted by MFIs in developing economies around the world, which service necessity entrepreneurs (i.e., individuals starting small businesses out of necessity) in the local areas. The necessity entrepreneurs have projects that ESG crowdfunders would like to fund for non-pecuniary reasons; the partnership between Kiva and MFIs creates a crowdfunding market for these projects.

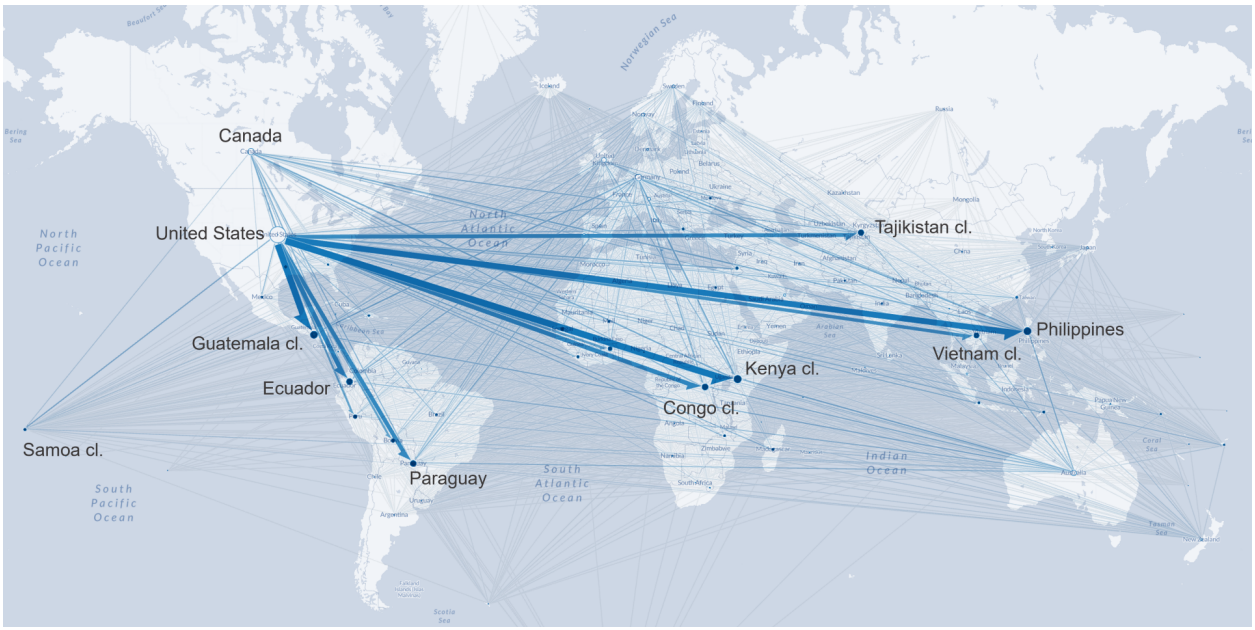


Figure 2: *Kiva Fund Flow Map in 2020.* The thickness of the arrows represents the amount of charitable fund flow between two countries. "cl." indicates that the dot on the graph is a cluster of several countries. An interactive version of this map is available here.

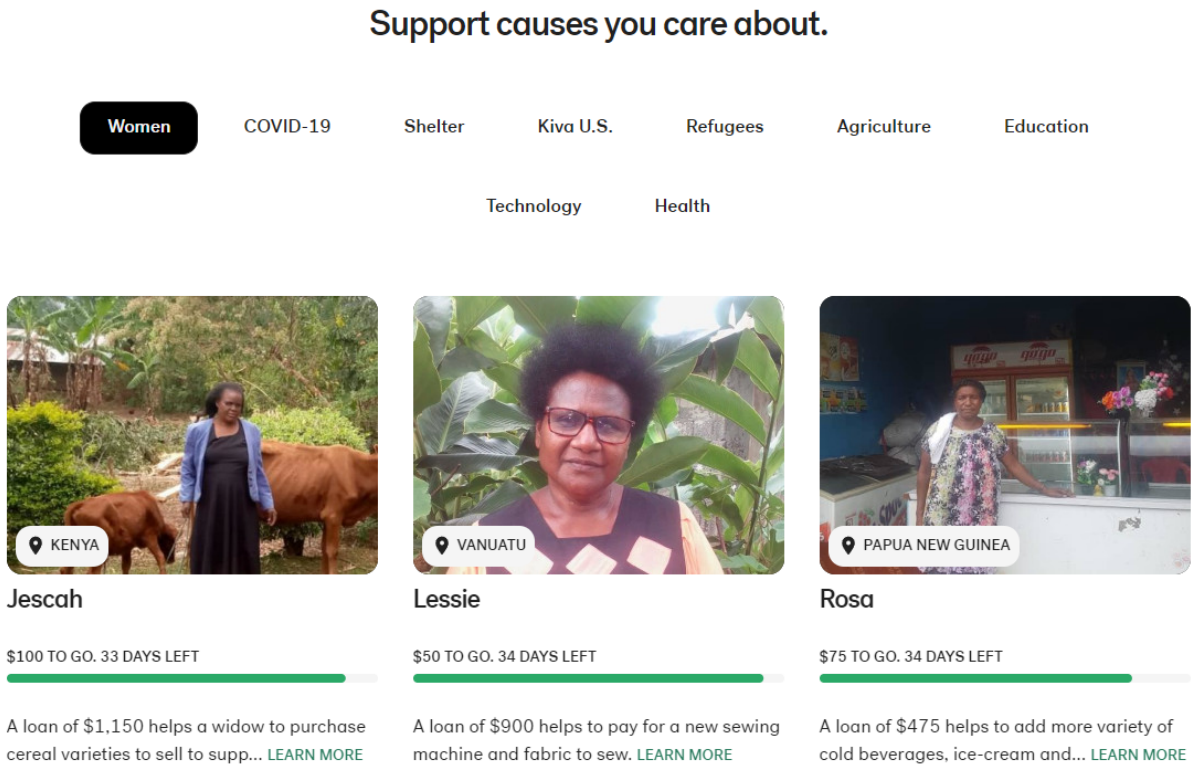


Figure 3: *A Snapshot of the Kiva Website.* Kiva lenders have access to several categories of information: the portrait, nationality, and name of borrowers, the funding progress, and a short description of the purpose of the loan. Lenders can click on the loans that they are interested in and explore more details on a separate web page, where they can learn about the MFI that intermediates the loan, read a short biography of the borrower, and get other information such as the repayment schedule, loan term, and whether the exchange rate risk is covered.

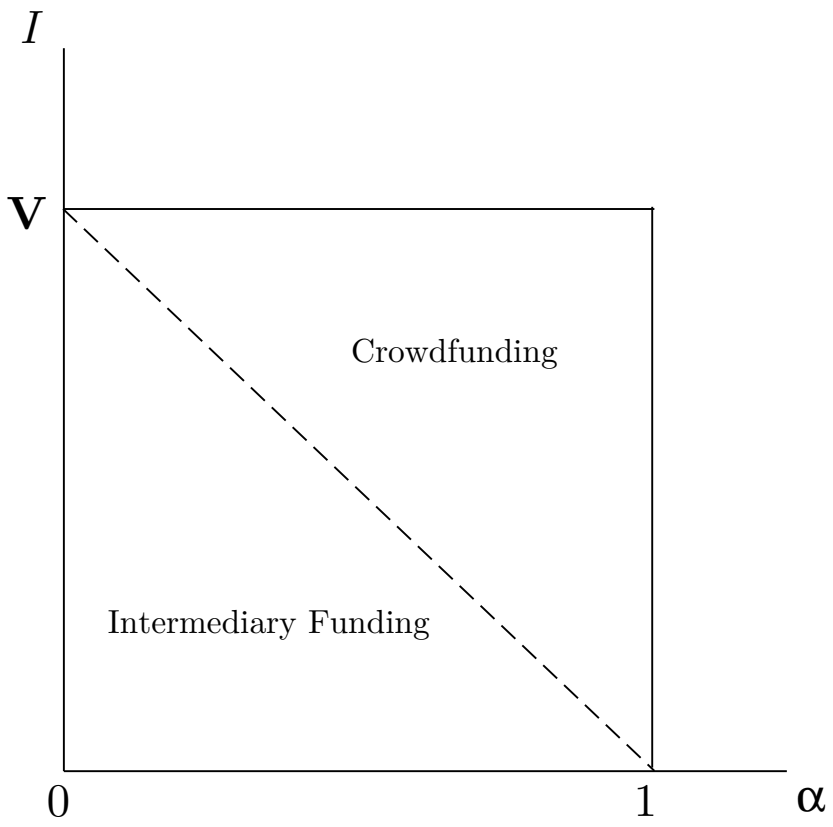


Figure 4: *A Stylized Model of Crowdfunding.* When the entrepreneur's bargaining power (α, horizontal axis) grows large, the intermediary's underinvestment problem becomes more severe, as it rejects a larger number of projects with positive utility ($\mathbf{V} > I$) but negative NPV ($I > \mathbf{V}^b = (1-\alpha)\,\mathbf{V}$). Crowdfunding is an improvement towards the first-best if it funds additional positive utility projects, which are located towards the top right and bounded by $\mathbf{V} = I$ and $\alpha = 1$.

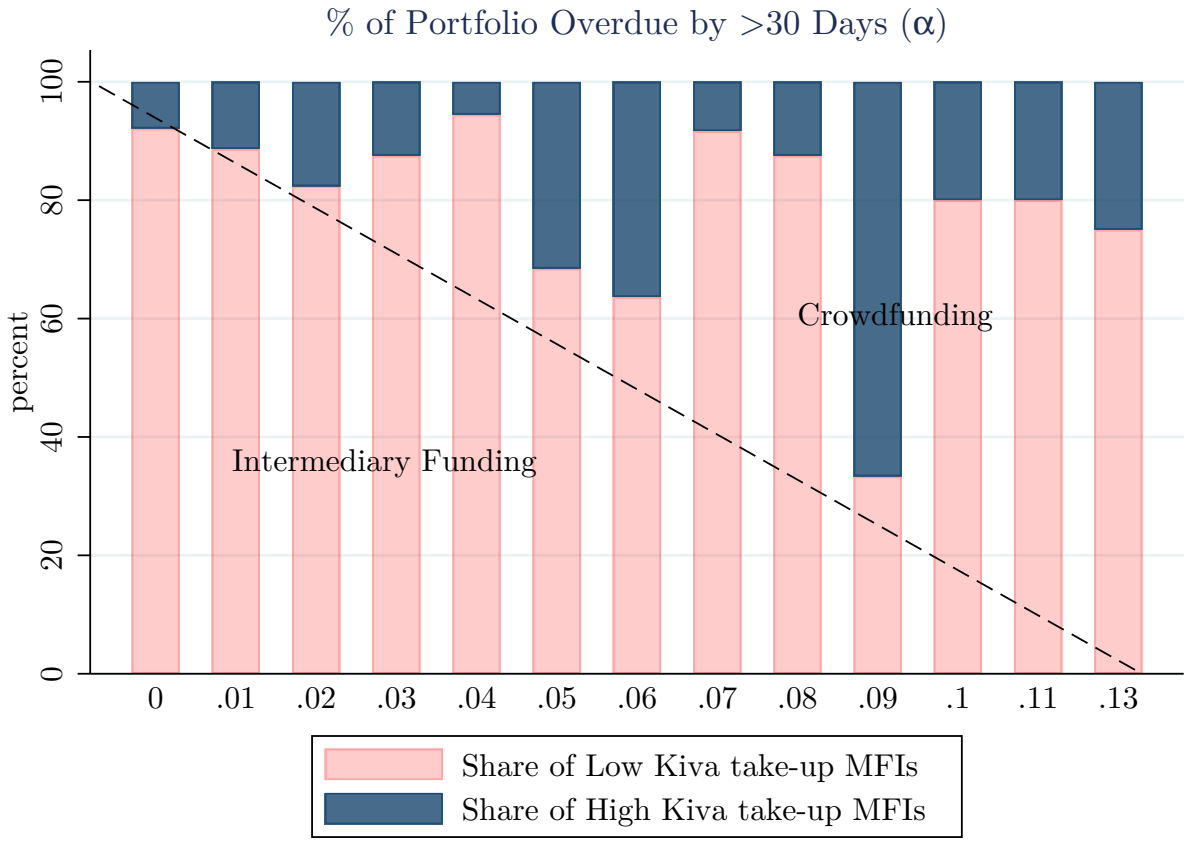


Figure 5: *Crowdfunding and Loan Performance.* We bin MFIs according to the overdue rates of their loan portfolios. Each bin is a 1% increment in overdue rates. Within each bin, the share of high Kiva take-up MFIs is colored in dark blue, and the share of low Kiva take-up MFIs is colored in light red. High Kiva take-up MFIs have a larger share of their loan portfolios funded by crowdfunders on Kiva. The dashed line overlays the intermediary's underinvestment problem as suggested by the stylized model.

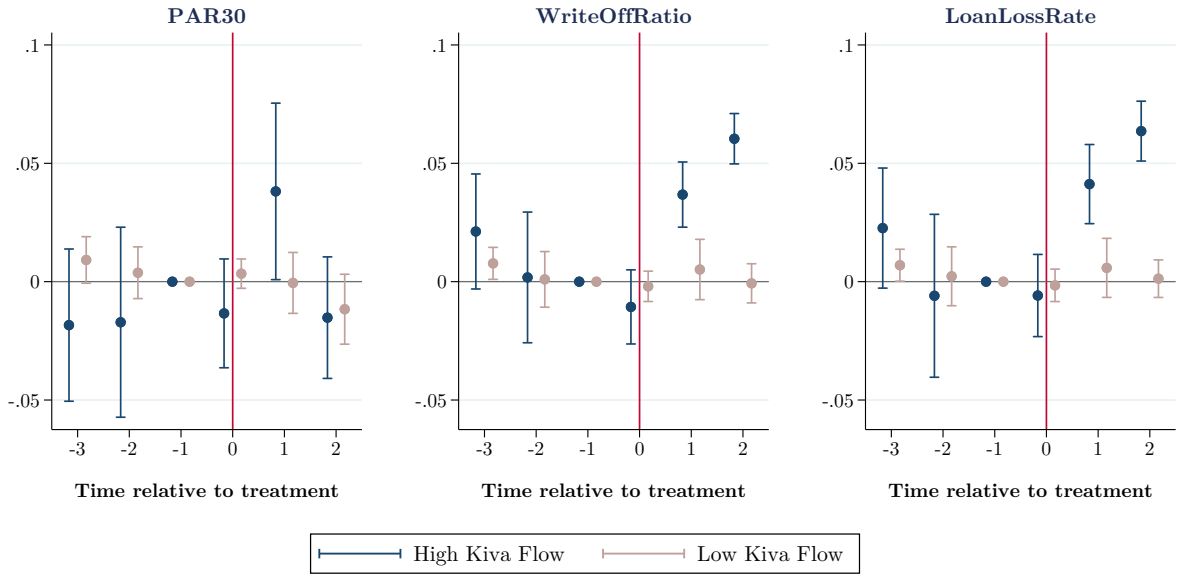


Figure 6: *Triple-differences loan performance regression.* The outcome variable is 30-days portfolio at risk in the left panel, loan write-off ratio in the middle panel, and loan loss rates in the right panel. 30-days portfolio at risk is the value of loans overdue by more than 30 days divided by gross loan portfolio. Loan write-off ratio is the value of written-off (but potentially still recoverable) loans divided by gross loan portfolio. Loan loss rate is value of unrecoverable loans divided by gross loan portfolio. The reference time period is the year before treatment. High Kiva share (plotted in dark blue) refers to the top 20% of Kiva-partnering MFIs, ranked by percentages of assets funded by Kiva. Low Kiva share (plotted in light brown) refers to the bottom 80% of Kiva-partnering MFIs.

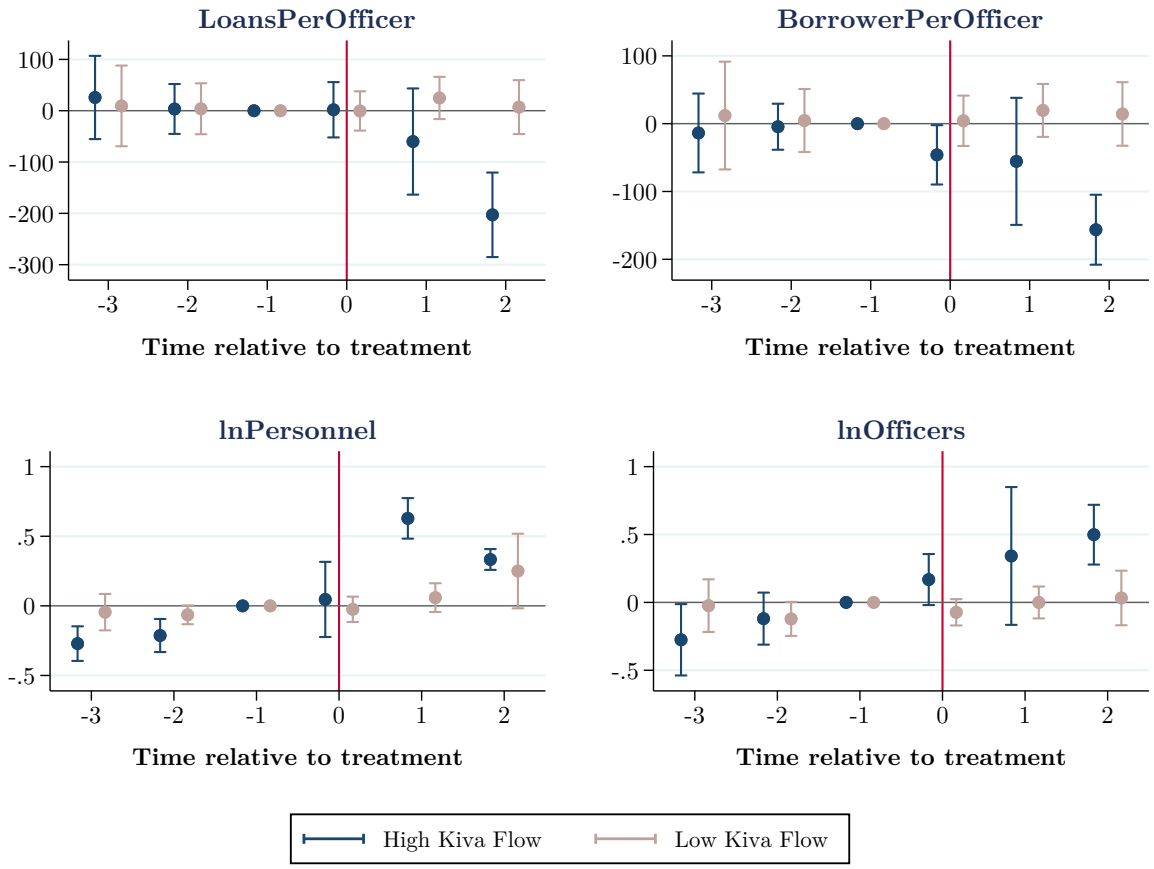


Figure 7: *Triple-differences MFI payroll regression.* The outcome variable is the number of loans serviced per loan officer in the top left panel, the number of borrowers served per loan officer in the top right panel, the natural log of the number of employees in the bottom left panel, and the natural log of the number of loan officers in the bottom right panel. The reference time period is the year before treatment. High Kiva share (plotted in dark blue) refers to the top 20% of Kiva-partnering MFIs, ranked by percentages of assets funded by Kiva. Low Kiva share (plotted in light brown) refers to the bottom 80% of Kiva-partnering MFIs.

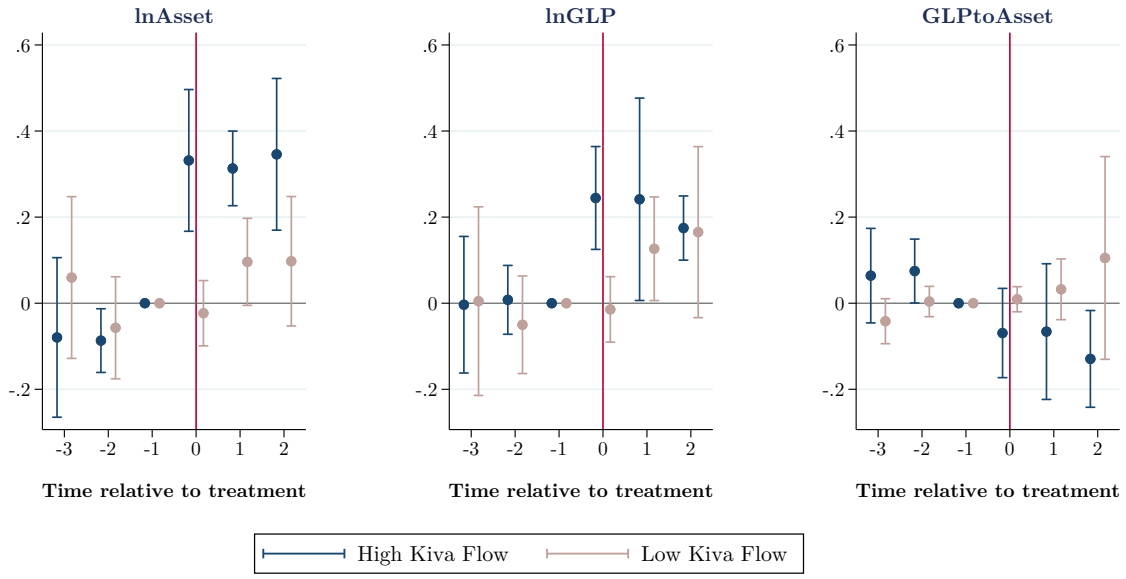


Figure 8: *Triple-differences MFI asset regression.* The outcome variable is the natural log of MFI assets in the left panel, the natural log of MFI gross loan portfolio in the middle panel, and the ratio of gross loan portfolio to assets in the right panel. The reference time period is the year before treatment. High Kiva share (plotted in dark blue) refers to the top 20% of Kiva-partnering MFIs, ranked by percentages of assets funded by Kiva. Low Kiva share (plotted in light brown) refers to the bottom 80% of Kiva-partnering MFIs.

	Treated MFIs (N=112)				Control MFIs (N=962)			
	Mean	SD	p5	p95	Mean	SD	p5	p95
lnAsset	15.88	1.37	13.53	17.95	16.48	1.89	13.57	19.79
lnGLP	15.61	1.37	13.39	17.66	16.23	1.91	13.30	19.63
GLPtoAsset	0.78	0.17	0.53	0.94	0.80	0.27	0.51	0.98
Portfolio at risk $>$ 30 days	0.04	0.05	0.00	0.12	0.05	0.06	0.00	0.16
Loan loss rate	0.01	0.03	−0.00	0.05	0.01	0.03	−0.00	0.05
Loans per loan officer	323.48	182.76	106.00	674.00	352.55	271.67	91.00	837.00
Borrowers per loan officer	312.80	180.05	107.00	657.00	328.64	243.13	88.00	750.00
lnPersonnel	4.69	1.10	2.94	6.52	4.97	1.64	2.30	7.81
lnOfficers	3.93	1.14	2.08	5.78	4.13	1.74	1.39	7.05
lnNumBorrowers	9.54	1.22	7.49	11.41	9.72	1.88	6.44	12.76
Average loan balance per borrower	751.49	867.10	91.00	2,512.00	1,419.24	1,997.90	101.00	4,976.00
Percent of female borrowers	0.74	0.22	0.36	1.00	0.67	0.24	0.32	1.00

Table 1: *Summary Statistics.* This table summarizes all the outcome variables in our regression analyses separately for the treated MFIs (i.e., MFIs that partnered with Kiva, $N = 112$) and for the control MFIs ($N = 962$). GLP: gross loan portfolio.

	(1) PAR30	(2) WriteOffRatio	(3) LoanLossRate
Treat × Post × High Kiva Flow	0.019***	0.033***	0.038***
	(3.292)	(4.761)	(7.500)
Treat × Post	−0.006	−0.002	−0.001
	(−1.484)	(−0.565)	(−0.328)
Strata-cohort-MFI FE	Yes	Yes	Yes
Strata-cohort-country-year FE	Yes	Yes	Yes
Mean of DV	0.049	0.015	0.013
R2	0.733	0.749	0.732
N	6,973	6,684	6,850

Table 2: *FinTech Adoption and Loan Performance.* Standard errors are clustered at the MFI level. t-statistics are in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1) LoansPerOfficer	(2) BorrowerPerOfficer	(3) lnPersonnel	(4) lnOfficers
Treat × Post × High Kiva Flow	−116.181***	−97.562***	0.373***	0.423**
	(−2.788)	(−2.956)	(5.741)	(2.553)
Treat × Post	6.583	7.725	0.114*	0.025
	(0.304)	(0.382)	(1.916)	(0.379)
Strata-cohort-MFI FE	Yes	Yes	Yes	Yes
Strata-cohort-country-year FE	Yes	Yes	Yes	Yes
Mean of DV	357.836	333.355	5.033	4.234
R2	0.853	0.858	0.983	0.985
N	6,802	6,827	6,949	6,841

Table 3: *FinTech Adoption and payrolls.* Standard errors are clustered at the MFI level. t-statistics are in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1) lnAsset	(2) lnGLP	(3) GLPtoAsset
Treat × Post × High Kiva Flow	0.311^{***}	0.116	-0.174^{***}
	(4.537)	(1.453)	(−2.894)
Treat × Post	0.053	0.098	0.054
	(0.816)	(1.607)	(1.174)
Strata-cohort-MFI FE	Yes	Yes	Yes
Strata-cohort-country-year FE	Yes	Yes	Yes
Mean of DV	16.434	16.206	0.815
R2	0.992	0.991	0.847
N	6,972	6,973	6,972

Table 4: *FinTech Adoption and Assets.* Standard errors are clustered at the MFI level. t-statistics are in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	DV: Female borrower share		
	(1) Full sample	(2) Above-average sample	(3) Below-average sample
Treat × Post × High Kiva Flow	0.051	−0.003	0.105***
	(1.254)	(−0.126)	(2.652)
Treat × Post	0.006	0.024	−0.010
	(0.557)	(0.886)	(−0.694)
Strata-cohort-MFI FE	Yes	Yes	Yes
Strata-cohort-country-year FE	Yes	Yes	Yes
Mean of DV	0.715	0.943	0.611
R2	0.965	0.688	0.962
N	7,123	1,519	3,829

Table 5: *FinTech Adoption and Financial Inclusion.* The first column reports regression run on the full sample. The second (third) column reports regression that only includes MFIs from regions of which the average female borrower share is above (below) the panel sample average, i.e., 64%. Regions with above-average female borrower shares are South Asia and East Asia. Regions with below-average female borrower shares are (1) Africa, (2) Eastern Europe and Central Asia, (3) Latin America, and (4) Middle East and North Africa. Region definitions are from the World Bank. Standard errors are clustered at the MFI level. t-statistics are in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	Full sample			Above-average sample			Below-average sample		
	(1) Asset utilization	(2) Default rate	(3) Loans per officer	(4) Asset utilization	(5) Default rate	(6) Loans per officer	(7) Asset utilization	(8) Default rate	(9) Loans per officer
Treat × Post × High Kiva Flow	−0.174***	0.038***	−114.881***	−0.199*	0.025***	28.206	−0.094***	0.034***	−159.017***
	(−2.942)	(7.584)	(−2.765)	(−1.714)	(2.716)	(0.478)	(−3.428)	(5.696)	(−2.861)
Treat × Post	0.054	−0.001	5.282	0.114	0.005	−27.413	0.016	−0.001	25.343
	(1.211)	(−0.485)	(0.248)	(0.973)	(1.219)	(−0.761)	(1.000)	(−0.319)	(1.089)
Strata-cohort-MFI FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Strata-cohort-country-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of DV	0.815	0.013	356.598	0.885	0.007	430.632	0.811	0.015	340.009
R2	0.848	0.734	0.854	0.882	0.615	0.810	0.877	0.793	0.895
N	7,126	7,004	6,952	1,535	1,443	1,471	3,828	3,804	3,755

Table 6: *Costly Financial Inclusion.* The first three columns report regressions run on the full sample. The second (third) three columns reports regressions that only include MFIs from regions of which the average female borrower share is above (below) the panel sample average, i.e., 64%. Regions with above-average female borrower shares are South Asia and East Asia. Regions with below-average female borrower shares are (1) Africa, (2) Eastern Europe and Central Asia, (3) Latin America, and (4) Middle East and North Africa. Region definitions are from the World Bank. Standard errors are clustered at the MFI level. t-statistics are in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1) lnNumBorrowers	(2) LoanPerBorrower	(3) PercentGroup
Treat × Post × High Kiva Flow	0.058	−107.727	−0.053
	(0.493)	(−0.542)	(−0.777)
Treat × Post	0.079	−50.133	0.089
	(1.117)	(−0.689)	(1.312)
Strata-cohort-MFI FE	Yes	Yes	Yes
Strata-cohort-country-year FE	Yes	Yes	Yes
Mean of DV	9.868	1198.925	0.300
R2	0.991	0.976	0.890
N	6,973	6,973	4,371

Table 7: *Non-significant Financial Inclusion Results.* Standard errors are clustered at the MFI level. t-statistics are in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix

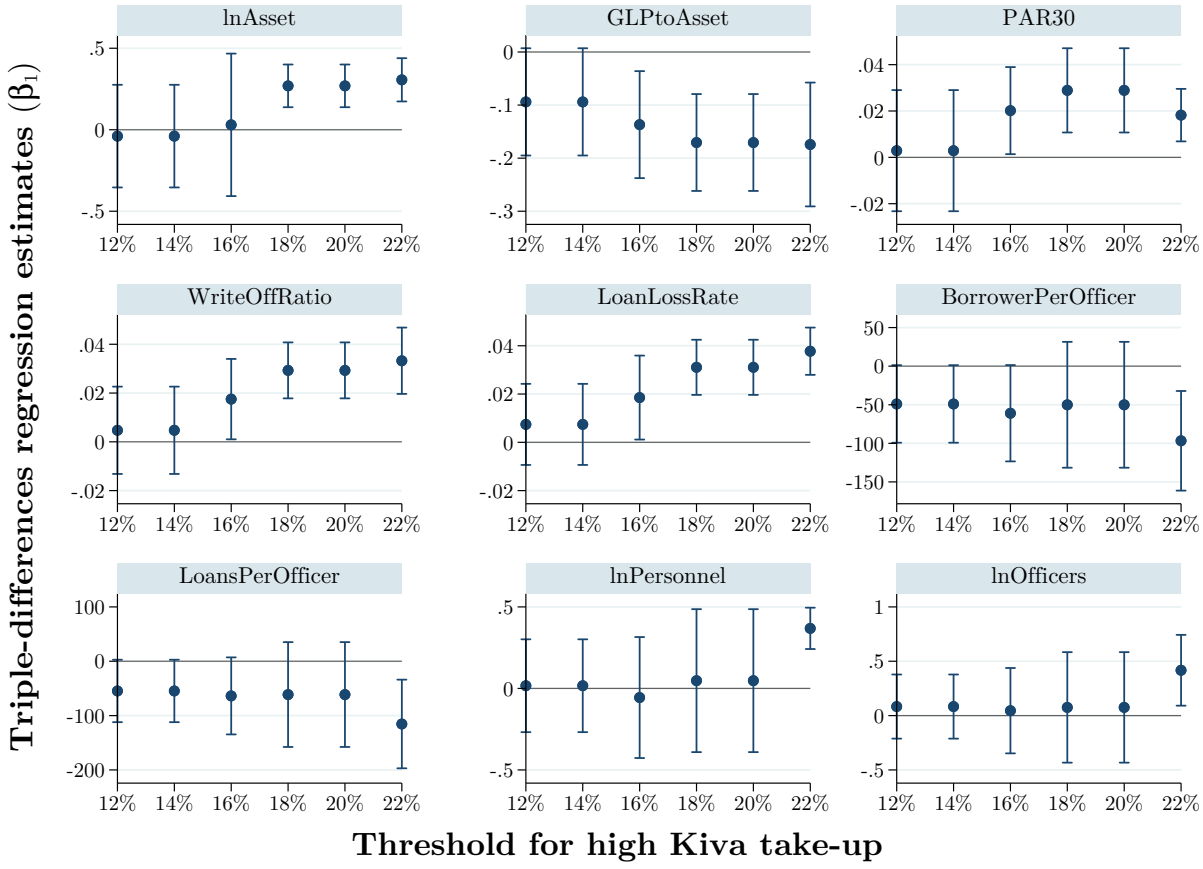


Figure A1: *Robustness to the choice of thresholds for high Kiva take-up.* Along the x-axis, we vary the threshold value that defines High_Kiva_Flow$_{id}$ from 12% of MFI's total assets to 22% of MFI's total assets. Point estimates of β_1 in (1) and corresponding 95% confidence intervals are plotted for different threshold values and for different outcome variables used in the paper.

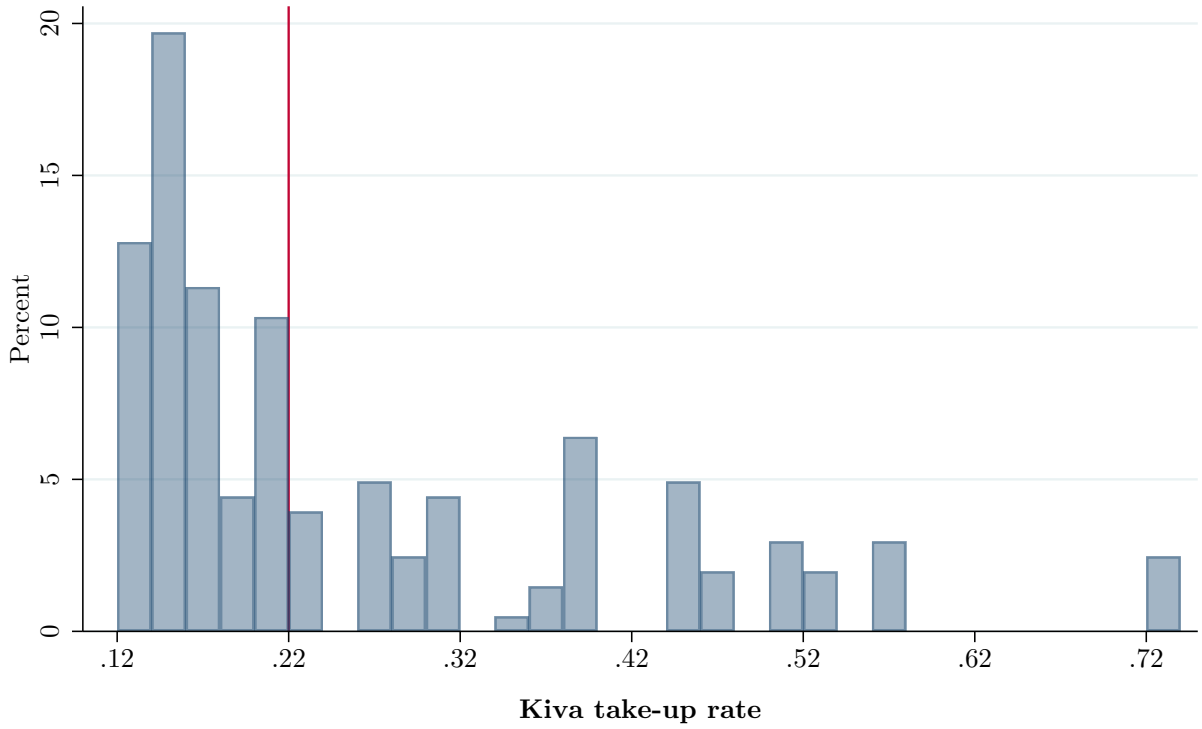


Figure A2: *Left-truncated distribution of Kiva take-up rates.* Kiva take-up rate is defined as the percentage of an MFI's assets funded by Kiva fund flows. The red vertical line marks the take-up rate of 22%, which we use as the threshold to define High_Kiva_Flow$_{id}$ in our main specification. Take-up rates below 12% (the sample average) are truncated; see Section 6.1 for details.

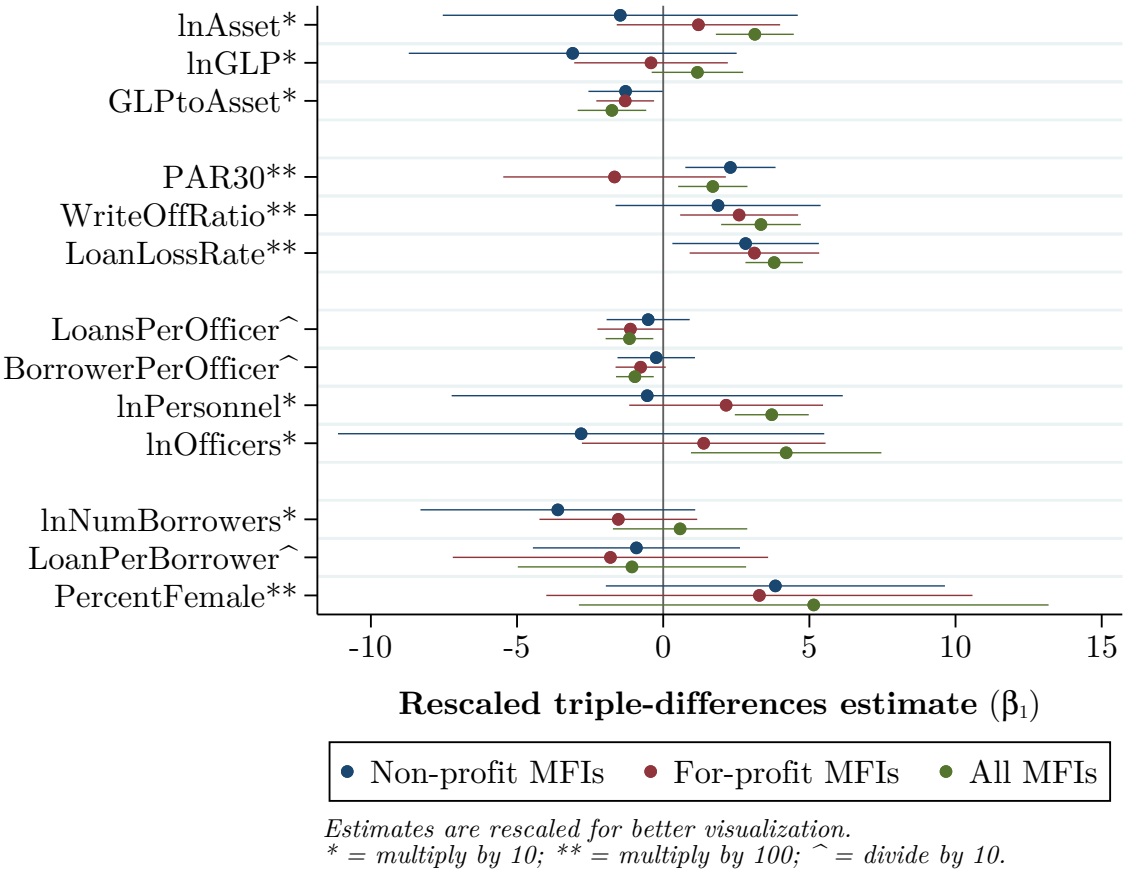


Figure A3: *Non-profit MFIs vs. for-profit MFIs.* We partition the sample into two sub-samples, one containing only non-profit MFIs and the other containing only for-profit MFIs. We rerun regression (1) and report the triple-differences estimates β_1 for all outcome variables. 95% confidence intervals are also plotted. Estimates from non-profit MFI sample are in blue; Estimates from for-profit MFI sample are in red; Estimates from the full sample (i.e., our baseline estimates) are in green, plotted here again for ease of comparison. Estimates are rescaled for better visualization.

	Non-stacked DID		Stacked DID		Triple-differences	
	(1) OSS	(2) YieldReal	(3) OSS	(4) YieldReal	(5) OSS	(6) YieldReal
Treat × Post	0.054**	−0.010	0.019	−0.004	0.017	−0.007
	(2.168)	(−1.320)	(0.662)	(−0.529)	(0.542)	(−0.975)
Treat × Post × High Kiva Flow					0.044	0.050
					(1.143)	(0.692)
Strata-cohort-MFI FE	No	No	Yes	Yes	Yes	Yes
Strata-cohort-country-year FE	No	No	Yes	Yes	Yes	Yes
MFI FE	Yes	Yes	No	No	No	No
Year FE	Yes	Yes	No	No	No	No
Mean of DV	1.119	0.253	1.187	0.224	1.187	0.224
R2	0.657	0.865	0.793	0.955	0.793	0.955
N	5,415	5,415	7,127	7,127	7,127	7,127

Table A1: *Replicating Luo, Ge, and Wang (2022).* This table replicates the two main results of Luo, Ge, and Wang (2022) on MFIs' operational self-sufficiency (OSS) and interest rate (real percentage yield). The difference-in-differences estimates in the first row for Treat × Post can be compared to those in Table 3 of Luo, Ge, and Wang (2022). Standard errors are clustered at the MFI level. t-statistics are in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1) OpExpAsset	(2) OpExpGLP	(3) OpExpRev
Treat × Post × High Kiva Flow	−0.033	0.013	−0.075*
	(−1.095)	(0.587)	(−1.715)
Treat × Post	−0.004	−0.014	−0.016
	(−0.738)	(−1.583)	(−0.663)
Strata-cohort-MFI FE	Yes	Yes	Yes
Strata-cohort-country-year FE	Yes	Yes	Yes
Mean of DV	0.145	0.185	0.713
R2	0.935	0.932	0.364
N	7,112	7,113	7,113

Table A2: *Does crowdfunding lead to empire-building?* Empire-building is proxied by aggressive spending and the increase in operating expenses. Column (1) scales operating expenses by total assets, column (2) by total revenue, and column (3) by gross loan portfolio. Standard errors are clustered at the MFI level. t-statistics are in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1) Asset utilization	(2) Default rate	(3) Loans per officer	(4) Female share
Treat × Post × High Kiva Flow	−0.062***	0.036***	−116.605**	0.084*
	(−3.086)	(6.120)	(−2.083)	(1.836)
Treat × Post	−0.000	−0.001	−2.035	0.011
	(−0.027)	(−0.333)	(−0.097)	(0.919)
Strata-cohort-MFI FE	Yes	Yes	Yes	Yes
Strata-cohort-country-year FE	Yes	Yes	Yes	Yes
Mean of DV	0.830	0.013	336.751	0.585
R2	0.840	0.731	0.900	0.959
N	3,043	3,028	2,993	3,044

Table A3: *Is crowdfunding less impactful when MFIs are living a quiet life?* In this table, we restrict the sample to MFIs in Eastern Europe (World Bank region 3) and Latin America (World Bank region 4). These two regions have high MFI market power (proxied by MFI profit margin) and low average female borrower share, so MFIs in these regions enjoy a quiet life (Hicks, 1935) while having large room for financial inclusion improvements. The null hypothesis is that MFIs with high market power stick to the status quo and exert little effort for financial inclusion, so the effect of crowdfunding is limited. Standard errors are clustered at the MFI level. t-statistics are in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A More institutional details about Kiva

A.1 Comparison of Kiva with Other Crowd-based Fundraisers

Kiva enables charitable lenders to fund borrowers around the world at zero interest, which differentiates itself from other crowdfunding or peer-to-peer (P2P) lending platforms that allow the interest rate to float and primarily serve U.S. users. Some of the commercial crowdfunding and P2P lending platforms have received more attention from the economics and finance literature in the past. We summarize the difference between Kiva and these platforms below.

A.1.1 Other Crowdfunding Platforms

The leading donation-based crowdfunding platform in the U.S. is GoFundMe. The key distinction between Kiva and GoFundMe is that Kiva facilitates loans rather than donations. The twist with Kiva loans is that Kiva lenders do not demand a risk premium, so the pecuniary payoff is at best zero. Therefore, while Kiva is not an outright donation-based crowdfunding platform, Kiva lenders' incentive compatibility does hinge on charitable motives.[11]

A.1.2 Other Peer-to-Peer Lending Platforms

The leading commercial P2P lending platforms in the U.S. are Lending Club and Prosper. Studies that use Lending Club data include Hertzberg, Liberman, and Paravisini (2018) and Paravisini, Rappoport, and Ravina (2017). Studies that use Prosper data include Hildebrand, Puri, and Rocholl (2017), Iyer, Khwaja, Luttmer, and Shue (2016), Lin, Prabhala, and Viswanathan (2013), Lin and Viswanathan (2016), Wei and Lin (2017), and Zhang and Liu (2012). Vallée and Zeng (2019) use data from robo-advisors that invest on both Lending Club and Prosper. Chava, Ganduri, Paradkar,

[11] Another popular form of crowdfunding is reward-based crowdfunding. The payoff for reward-based crowdfunders is in-kind rewards that are usually the output of the funded project. The leading reward-based crowdfunding platform in the U.S. is Kickstarter. Studies that use Kickstarter data include Burtch, Carnahan, and Greenwood (2018), Gafni, Marom, Robb, and Sade (2021), Lin and Pursiainen (2022, 2023), Xu and Ni (2022), and Younkin and Kuppuswamy (2018).

and Zhang (2021) use credit bureau data that includes credit history on both P2P lending platforms.

Kiva's key distinctions from these commercial P2P lending platforms are twofold. First, Kiva fund flow is zero interest out of the Kiva gate. The near shutdown of the pecuniary channel (no interest, small principal, and low probability of default) characterizes Kiva as a unique platform, where financial transactions are primarily governed by social considerations. Second, Kiva's fund flow is intermediated. Intermediation by MFIs around the world allows Kiva to distribute funds in underdeveloped regions, where P2P lending is not feasible in its modern flavor. Indeed, leading P2P lending platforms in the U.S. have gone through similar re-intermediation, presumably for the same reason that intermediaries, despite the moral hazard problem, have superior ability to evaluate and distribute the loans (Balyuk & Davydenko, 2023). Given Kiva's even higher demand for local intermediary expertise, intermediation has extra appeal over the P2P approach.

A.2 Kiva's non-core business developments after 2019

Our sample ends in 2019 due to data availability restrictions from MIX. It is also around 2019 that Kiva starts to experiment with alternative approaches to financial inclusion, in addition to its core model of crowdfunded intermediated loans. The expansion can be summarized into two categories. The first category is that Kiva deviates from crowdfunding and start seeking funding from institutional funders, such as Google. The second category is that Kiva deviates from intermediated loans and start experimenting with direct loans (the Kiva Zip project that later merged into the Kiva.org as Kiva U.S.) and non-loan projects (Kiva protocol, a digital identification system for Sierra Leone's seven million citizens; and Kiva Capital, an asset manager that emphasizes social impacts). For more details about these developments, see https://www.technologyreview.com/2023/08/14/1077351/microfinance-money-making/ and https://www.kiva.org/cms/stanford-social-innovation-review-kiva-reinvents-itself-case-study.fall19_final.pdf. These developments are not reflected in our sample because the MFI financial performance data provided by the World Bank is only available up to 2019. Our sample period from 2007 to 2019 represents the phase of Kiva's development of its core business, i.e., crowdfunded loans. It is in this capacity that we refer to Kiva as a crowdfunding platform.

B The Kiva Fellow Program

In addition to fund transfers, Kiva also hosts a volunteer program that sends Kiva fellows to partner MFIs. Kiva fellows travel to, live near, and provide voluntary service to Kiva's partner MFIs. The typical stay lasts for three to twelve months, during which the fellow assists the MFI to collect and post online the borrowers' profiles, the loans' journal entries, and other related information that help to engage the Kiva users online. Working closely with the partner MFIs, the Kiva fellows exchange skills and best practices with the local staff.

566 Kiva fellows have travelled to 161 MFIs from 2007 to 2012. Almost all of these fellows have completed or were pursuing a bachelor degree. 235 (42%) of them have a bachelor degree in business-related topics and 109 (19%) of them have a graduate degree, according to information disclosed in their biographies written for the fellow program. Kiva fellows are highly educated: 56 (10%) of Kiva fellows graduate from an Ivy League school, and an additional 27 (5%) of them graduate from UCLA, UC Berkeley, or Stanford. Commitment by these highly educated fellows signifies that Kiva is not merely another source of funding for MFIs, but potentially a source of knowledge exchange for them.

C Additional Discussion of Luo et al. (2022)

In terms of empirical design, the closest paper to ours is Luo et al. (2022), which shows that after partnering with Kiva, MFIs charge lower interest rates, and their operations become more self-sufficient. However, their findings suffer from two methodological issues. First, the DID regressions they use are plausibly biased due to dynamic treatment effects (Baker et al., 2022). Second, their method is subject to the reverse causality concern of "Kiva as a trophy." Specifically, it is possible that MFIs join Kiva to signal their operational success, which then becomes the cause, and not the effect, of the Kiva partnership. We address the first methodological issue by using the stacked regressions (Cengiz et al., 2019) to address the staggered partnership of Kiva. We address the second methodological issue by adopting a triple-differences framework. The benefit of a triple-difference

approach is that it identifies variations from heterogenous Kiva take-up rates, so the effects identified in our paper are driven by MFIs that *actually* crowdfunds a large fraction of their loan portfolios.

In Table A1, we show that after correcting for these methodological issues, the findings in Luo et al. (2022) are no longer statistically and economically significant. The reduction in statistical significance can largely be attributed to correcting for the staggered partnership with Kiva: once stacked regressions are used, MFIs do *not* charge lower interest rates or become more self-sufficient in their operations.