Recommending teams promotes prosocial lending in online microfinance

Wei Ai[a], Roy Chen[b], Yan Chen[a,c,1], Qiaozhu Mei[a], and Webb Phillips[d]

[a]School of Information, University of Michigan, Ann Arbor, MI 48109; [b]Department of Economics, National University of Singapore, Singapore 117570; [c]Department of Economics, School of Economics and Management, Tsinghua University, Beijing 100084, China; and [d]Kidaptive, Mountain View, CA 94043

This paper reports the results of a large-scale field experiment designed to test the hypothesis that group membership can increase participation and prosocial lending for an online crowdlending community, Kiva. The experiment uses variations on a simple email manipulation to encourage Kiva members to join a lending team, testing which types of team recommendation emails are most likely to get members to join teams as well as the subsequent impact on lending. We find that emails do increase the likelihood that a lender joins a team, and that joining a team increases lending in a short window (1 wk) following our intervention. The impact on lending is large relative to median lender lifetime loans. We also find that lenders are more likely to join teams recommended based on location similarity rather than team status. Our results suggest team recommendation can be an effective behavioral mechanism to increase prosocial lending.

social identity | charitable giving | microfinance | field experiment | recommender systems

Understanding strategies to increase prosocial behavior has important policy implications. Charities have explored various mechanisms to increase giving, such as seed money, matching gifts, and peer pressure (1). In comparison, an underexplored class of mechanisms uses group membership and intergroup competition (2, 3) to increase both participation and giving amounts. Compared with price-based strategies, such as matching gifts and rebates, empirical analysis of naturally occurring data indicates that identity-based mechanisms have longer-lasting effects (4). Our research explores two questions through a large-scale field experiment on a crowdlending community with a natural group structure (teams). First, which types of team recommendations are most likely to motivate lenders to join teams? Second, once they join a team, what is the subsequent impact on lending?

Our research is conducted at Kiva.org, a crowdlending community created to help micro and small enterprises in developing countries, which often lack access to the formal banking sector. Specifically, Kiva partners with local microfinance institutions to match individual lenders with low-income entrepreneurs in developing countries as well as selected cities within the United States. Through Kiva's platform, anyone can make a zero-interest loan of $25 or more to support an entrepreneur. Since its inception in 2005, Kiva has increased its membership significantly. However, although many lenders join Kiva for prosocial motives, they do not participate fully. Indeed, 36% of them have never made a single loan, and many others do not come back to Kiva after making their first loan (5). Kiva's challenge is not unique, as many online contribution communities struggle with the issue of how to sustain member engagement and contributions.

To increase member engagement, some online communities have created group structures. For example, in 2008, Kiva instituted a lending teams program, a system through which lenders can create teams or join existing teams of other lenders. Once a team is created, it appears on Kiva's team leaderboard, which sorts teams by the total loan amounts designated to them by their team members. Since 2008, more than 38,957 Kiva teams have been created based on lender group affiliations such as organizations, geographic location, religious affiliation, or sports interests. Of note, many of the highly ranked teams are identity based, such as the "Atheists" and the "Kiva Christians." Each team has a dedicated forum where team members can coordinate their lending activities, ask and answer questions, and set goals for the team.

The use of groups to increase charitable contributions has intuitive appeal, but its success is difficult to measure with naturally occurring field data because of sample selection bias. For example, lenders who join teams might simply be those who are more active in general (4). To establish the causal relationship between group membership and prosocial lending, we use a randomized field experiment that enables us to combine the control of a laboratory experiment with the external validity of a field study (6, 7).

Our approach is inspired by the economic theory of social identity (2, 8) as well as the development of big data analytics in computer science. Research on social identity has consistently found that people derive their sense of identity from groups (9, 10). This group identity can be used to increase voluntary contribution and improve coordination among team members in the laboratory (11–16). Building on these findings, we conduct a large-scale randomized field experiment to evaluate the effectiveness of team recommendation as a behavioral mechanism for increasing participation among Kiva members. Our approach enables us to synthesize the predictive accuracy of machine learning with the causal inference of economic theory and field experiments (17).

Significance

With three billion people subsisting on the equivalent of $2.50 per day, alleviating poverty is one of the most urgent challenges facing the world today. One solution to this problem has been to encourage the growth of small enterprises through microlending. A successful innovation is represented by Kiva.org, which matches citizen lenders with low-income entrepreneurs in developing countries. To increase prosocial lending, we use a large-scale field experiment and machine-learning methods to recommend lending teams to lenders. We find that lenders who join a team contribute significantly more compared with those who do not. Our results suggest team recommendation can be an effective and low-cost behavioral mechanism to increase charitable contributions.

Fig. 1. Proportion of lenders joining teams in each experimental condition. This figure presents the proportion of lenders who join a lending team in each experimental condition after our email intervention. Location-based recommendations exhibit a higher proportion of lenders joining recommended teams (67.96%), compared with lending history similarity (42.31%) or leaderboard (44.37%)-based recommendations ($P < 0.01$, proportion of t tests). Similar results are observed when we focus on lenders who open our email (*Right*).

Literature Review

Our study builds upon findings from three streams of literature: charitable giving, advertising and recommender systems, and social identity. The charitable-giving literature has uncovered several motivations and mechanisms for people to voluntarily give to charity (1). In addition to the neoclassical preferences for public goods (18), people might derive a "warm glow" from the amount they give, which increases giving (19, 20). People also respond positively to mechanisms that decrease the price of giving, such as tax subsidies (21), matching gifts, or rebates (22, 23). Sequential giving mechanisms (24–26), which use leadership gifts to transmit information or signal the value of the public good, have been shown to increase giving in the laboratory and field (27, 28). Closely related to our study, researchers have shown both theoretically and experimentally that people might give because they care about their social image (29, 30), peer pressure (31), or social pressure (32). In our context, when lenders join a team, team members can activate several of these mechanisms, such as leadership giving and social pressure, by posting messages on the team forum (4).

Our research is also related to the advertising literature. Recent field experiments show that advertising content, especially when it appeals to intuition, significantly affects demand (33). More generally, personalized recommendations based on various machine-learning algorithms have increased consumer adoption of recommended items, and have thus been widely used by e-commerce sites (34, 35). Instead of recommending items, such as products, our study recommends lending teams to Kiva users.

Last, our study builds upon social identity theory (2, 3) and recent experimental research that uncovers the positive effects of group identity on voluntary contribution and coordination in the laboratory (11–16, 36) and the field (37). Our team recommendation approach extends social identity research to the realm of behavioral mechanism design at scale.

Methods

In our study, we use a lender's likelihood of joining a team to recommend teams based on both homophily and status. Homophily refers to the tendency to associate with similar others (38, 39). As such, we recommend teams to lenders based on their similarity to the existing members of those teams. In our study, we use two different measures of homophily: location similarity and loan history similarity. The former is based on the number of lenders in a team who share the same location as the target lender, whereas the latter is based on how often the lenders have lent to the same borrowers. In addition to homophily, we recommend teams based on status (40), using the top three teams on the Kiva leaderboard as the high-status teams.

Using a 3 × 2 factorial design (Table S1), we vary our recommendation algorithms along one factor based on lender–team location similarity, loan

history similarity, or team status. Along the other factor, we vary whether our recommendation rationale is explained to the lender. The computer science literature suggests that providing an explanation can increase the acceptance of a recommendation (34, 41). By varying whether a lender receives an explanation, we can obtain a better understanding of whether a factor impacts the effectiveness of the recommender system. We also include a control condition where we do not contact lenders (no contact) and a placebo condition where we email lenders to make them aware that there are lending teams on Kiva without providing any specific recommendations (teams exist) to control for any contact effect. The text of the email is completely identical across treatments, except for the variables that change across treatments (Fig. S1 and *SI Methods*).

To study the causal effects of team recommendations on the likelihood of joining a team and increasing contributions, we use a group of 69,802 lenders who have made at least two loans in the past 6 mo but have

Table 1. Treatment effects on the likelihood of joining teams: Probit regressions

Dependent variable: Joined a Team

	(1)	(2)	(3)
	All Users	Opened and No-Contact	Opened
Team-Exist	0.0045***	0.0155***	
	(0.002)	(0.003)	
Location-Explanation	0.0094***	0.0256***	0.0145***
	(0.002)	(0.002)	(0.004)
Location-NoExplanation	0.0062***	0.0189***	0.0050
	(0.002)	(0.002)	(0.004)
History-Explanation	0.0070***	0.0212***	0.0083**
	(0.002)	(0.002)	(0.004)
History-NoExplanation	0.0061***	0.0182***	0.0039
	(0.002)	(0.003)	(0.004)
Leaderboard-Explanation	0.0062***	0.0185***	0.0043
	(0.002)	(0.002)	(0.004)
Leaderboard-NoExplanation	0.0063***	0.0197***	0.0062
	(0.002)	(0.002)	(0.004)
Number of subjects	64,800	29,055	20,371

Notes: (*i*) SEs in parentheses. (*ii*) Significant at the: ****10%, **5%, and ***1% levels. Marginal effects reported, calculated at the mean level of the covariates. (*iii*) In the first model, the decision to join a team is regressed on the seven treatment dummies for all lenders in our sample ($n = 64,800$). (*iv*) The second model uses the same specifications but is restricted to the lenders who opened their emails or were not contacted ($n = 29,055$). (*v*) The third model is restricted to lenders who were sent emails and opened them ($n = 20,371$). Applying a multiple-hypothesis testing correction (42) yields the same significance levels as above, except for the "History-Explanation" variable in column (3), which becomes insignificant at the 10% level.

Fig. 2. Treatment effects on the likelihood of joining teams. This figure presents the treatment effects on the likelihood that a lender joins a lending team (Table 1). When we focus on all lenders (lines with red triangle), we find that every treatment significantly increases the likelihood of joining a team compared with the control condition. When focusing on lenders who open our email (lines with green circle), we find that the homophily-based recommendations with an explanation also significantly increase the likelihood of joining a team, compared with the teams-exist condition. Explanations increase the likelihood of joining a team for only the location-based recommendations (All: $P = 0.02$; Lenders who open our email: $P = 0.01$; Wald tests).

never joined a team. We then randomly assign each lender to one of eight experimental conditions with equal probability. Pairwise Kolmogorov–Smirnov tests based on observable characteristics verify that our randomization works (Table S2).

We send each lender in our treatment groups an email from Kiva. After excluding lenders whose emails bounced and those who made their accounts private, we have a total of 64,800 lenders whom we intend to treat (henceforth "All"; Fig. S2). Of these lenders, we find that one-third

($n = 20,371$) open our email, constituting our treated subsample (henceforth "Opened"). We then track the team joining and lending behavior of each lender for the next 2 mo. *SI Methods* includes a detailed description of our experimental procedure and email scripts. Anonymized data will be available from the open Inter-university Consortium for Political and Social Research data repository. Our research protocol was approved by the University of Michigan institutional review board (HUM00050208), which exempted us from obtaining informed consent.

Table 2. Choice model: Conditional logit regressions

	Dependent variable: Joined a Team							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	No-Contact	Team-Exist	Loc.-Exp	Loc.-NoExp	Hist.-Exp	Hist.-NoExp	Lead.-Exp	Lead.-NoExp
Location similarity (percentile)	1.03***	1.02***	1.02***	1.05***	1.02***	1.01***	1.02***	1.02***
	(0.011)	(0.005)	(0.006)	(0.017)	(0.005)	(0.005)	(0.007)	(0.005)
History similarity (percentile)	0.99**	1.00	1.01	1.02***	1.01	1.01	0.99**	0.99**
	(0.006)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.006)	(0.005)
Top-10 team	13.13***	13.77***	0.92	1.26	7.99***	13.99***	18.32***	6.03***
	(6.237)	(5.225)	(0.292)	(0.481)	(2.997)	(5.262)	(8.936)	(2.450)
Team size (percentile)	1.01	1.00	1.00	0.99	1.00	1.01**	1.00	1.02**
	(0.010)	(0.009)	(0.009)	(0.008)	(0.006)	(0.007)	(0.008)	(0.009)
Recommended			83.57***	39.16***	122.51***	192.96***	8.83***	8.23***
			(28.286)	(15.027)	(36.817)	(66.952)	(2.748)	(2.926)
Number of teams	491	491	491	491	491	491	491	491
Number of subjects	35	61	105	74	80	72	72	74

Notes: (*i*) SEs in parentheses, clustered at the subject level. (*ii*) Significant at the *10%, **5%, and ***1% levels. Odds ratios reported. Whether the subjects join teams is regressed against the two similarity measures (coded as the percentile of the measure for each subject–team pair), whether the team is one of the top teams, the team size, and whether or not the team was recommended through the experiment. This regression is performed separately for each treatment. Although location similarity and recommendations always significantly increase the likelihood of joining a team, the effects of the other variables depend on the treatment. When the teams are recommended based on either lending history (columns 5 and 6) or the leaderboard (columns 7 and 8), giving no explanation for the recommendation increases the importance of team size. A location recommendation (columns 3 and 4) causes subjects to ignore the top-10 teams. The leaderboard recommendation also decreases the degree to which subjects pay attention to lending history.

Fig. 3. Effects of team membership on prosocial lending. This figure reports the results of our two-stage least-squares instrumental variable regression coefficients (Table 3), indicating the effects of joining a lending team on contributions for the 1-d (left red bar) and 7-d (middle red bar) window. The median Kiva lender's lifetime contributions ($25) is plotted to provide a benchmark (green bar).

Results

We first examine what types of recommendations are most effective in increasing team membership. Fig. 1 presents the proportion of lenders who join a lending team in each treatment after our email intervention, for both all lenders (*Left*) and those who open our emails (*Right*). For both groups, lenders who receive a location similarity explanation are most likely to join a team, accounting for 3% of the group who open their emails. This participation rate is comparable to that in other charitable-giving field experiments using mailing campaigns (23, 28).

We next conduct a regression analysis (Table 1 and Fig. 2) and find that every treatment leads to a significantly higher likelihood of joining a team, compared with the no-contact control condition, for both the all-lenders (column 1) and opened-email (column 2) groups ($P < 0.01$). Of those who open their emails, lenders in the location similarity with explanations treatment are more likely to join a team compared with those in the teams-exist condition ($P < 0.01$). These results are robust to a multiple-hypothesis testing correction (42).

We next explore which types of teams lenders are most likely to join by examining the characteristics of teams joined by our lenders. Table 2 displays the results of eight conditional logit specifications with odds ratios reported, with one specification per treatment. In our regressions, we use whether each lender joined each team as our dependent variable, and location similarity, loan history similarity, team status, team size, and experimenter recommendation as our independent variables.

The results for our control and teams-exist conditions (columns 1 and 2) show that lenders are more likely to join teams with higher location similarity and status. The odds of a lender joining a team whose location similarity is 1 percentile higher is 3% higher, whereas the odds of a lender joining a top-10 team is 13 times higher than those of joining a non–top-10 team. Lending history actually has a small negative effect on the odds of joining a team in the control and no effect in the teams-exist condition. We also find that team size has no impact on lenders' choices. These findings show that lenders value both homophily and status when deciding to join a team. It is also noteworthy that location and status information are easily found on Kiva's website, whereas lending histories are more difficult to locate.

Interestingly, we find that the provision of a location similarity recommendation mitigates the influence of team status, leading lenders to join recommended teams (columns 3 and 4) or teams with higher history similarity (column 4). By contrast, our recommendations based on loan history similarity (columns 5 and 6) do not substantially change how lenders choose their teams. Finally, recommendations based on team status (columns 7 and 8) seem to decrease the importance of lending history.

Finally, we study whether joining a team increases prosocial lending. To address any potential endogeneity issues caused by self-selection, we use the random treatment assignment in our experiment, namely, whether the lender received an email, as an instrumental variable for joining a team. Fig. 3 and Table 3 display the results of our two-stage least-squares instrumental variable regression. In the first stage, we find that the "Email" variable, denoting whether a lender received an email, is not a weak instrument for joining a team, with an F statistic of 23.55. Next, for this instrument to satisfy the exclusion restriction, it must be the case that an email does not directly affect lending except through increasing the likelihood that a lender joins a team. This might occur if contacting the lenders regarding Kiva reminds them of Kiva's existence, prompting them to lend. However, because our previous field experiment on Kiva has shown that

Table 3. Difference-in-differences regressions of average daily lending amount (2SLS)

	IV first stage	IV second stage: Average amount			OLS		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		1 d	7 d	30 d	1 d	7 d	30 d
Email	0.0053***						
	(0.001)						
Join Team		298.5579***	55.9145***	10.2310	5.2565***	0.5662*	0.5166***
		(72.283)	(21.058)	(7.318)	(0.755)	(0.337)	(0.134)
Constant	0.0045***	-2.6593***	-0.9359***	-0.2357***	0.0066	-0.4328***	-0.1474***
	(0.001)	(0.670)	(0.195)	(0.068)	(0.072)	(0.032)	(0.013)
Observations	64,800	64,800	64,800	64,800	64,800	64,800	64,800

Notes: (*i*) SEs in parentheses. (*ii*) Significant at the: *10%, **5%, and ***1% levels. The endogenous variable, whether a lender joins a team ("Join Team"), is instrumented with whether a lender receives an email in the experiment ("Email"). As the results of a two-stage least-squares instrumental variable (IV) regression, the coefficients on the "Join Team" variable in columns 2–4 give local average treatment effects, or the effects on the subset of lenders who only join a team because of our email ("compliers"). The different columns give different window sizes in a difference-in-differences setting. The effect is significant up to a week after we send the email. Ordinary least-squares (OLS) estimates are also displayed in columns 5–7 for comparison. The difference between the IV and OLS estimates is due to the difference in the local average treatment effects (given by the IV regressions), which only gives the effect on compliers, and average treatment effects (given by the OLS estimates, although with potential selection bias), which gives the effect on all subjects. There are a large number of lenders who do not join any team in our sample, and the effects of our treatment on these subjects are not captured by the IV estimates.

simply contacting the lenders does not affect lending (4), we conclude that the instrument satisfies the exclusion restriction.

This regression uses a difference-in-differences approach. For three different window sizes, the dependent variable in each second-stage regression is the difference in total loan amounts t days before and after our treatment, where t is the window size. Thus, the coefficients on the "Join Team" variable indicate how much more lenders who join teams give than those who do not join teams after the treatment, controlling for the same difference before the treatment. The results of this regression show that joining a team significantly increases lending. However, it is important to note that, because these estimates are derived from an instrumental variables regression, they give the local average treatment effect, not the average treatment effect (43). Therefore, the estimates apply only to lenders who would join a team if prompted by an email.

This effect is also insignificant beyond 1 wk. One possible reason for the lack of an observed long-term effect is that lenders may wait until initial loans are repaid before lending again, a process which may take 12 to 18 mo. However, even the 1-wk effect ($392) is more than 15 times the lifetime contribution of the median Kiva lender ($25), indicating that team membership is effective in increasing member contributions on those lenders who would join a team because of our email.

Discussion

This paper reports the results of a large-scale field experiment designed to test the hypothesis that team membership can increase participation and lending for an online crowdlending community, Kiva. We find that emails increase the likelihood that a lender joins a team, and that joining a team increases lending in an 1-wk window following the decision to join. Although this experiment does not explore the mechanism through which joining a team increases giving, our prior empirical analyses and field experiment point to two mechanisms at work (4). First, joining a team increases information sharing about specific borrowers on the team forum, which reduces team members' search costs and increases their lending. Second, joining a team increases the pressure to help improve the team's ranking on the Kiva leaderboard. Therefore, effective teams share information and coordinate their loans to reduce search costs, and emphasize team competition through goal setting. Our results suggest that recommending teams to members of an online lending community based on homophily is an effective mechanism to engage community members and increase their contributions.

1. Andreoni J, Payne AA (2013) Charitable giving. *Handbook Publ Econ* 5:1–50.
2. Akerlof GA, Kranton RE (2000) Economics and identity. *Q J Econ* 115(3):715–753.
3. Tajfel H, Turner J (1979) An integrative theory of intergroup conflict. *The Social Psychology of Intergroup Relations*, eds Worchel S , Austin W (Brooks/Cole, Monterey, CA), pp 94–109.
4. Chen R, Chen Y, Liu Y, Mei Q (2017) Does team competition increase pro-social lending? Evidence from online microfinance. *Game Econ Behav*, 10.1016/j.geb.2015.02.001.
5. Liu Y, Chen R, Chen Y, Mei Q, Salib S (2012) I loan because...: Understanding motivations for pro-social lending. *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining* (ACM, New York), pp 503–512.
6. Harrison GW, List JA (2004) Field experiments. *J Econ Lit* 42(4):1009–1055.
7. Chen Y, Konstan J (2015) Online field experiments: A selective survey of methods. *J Econ Sci Assoc* 1(1):29–42.
8. Akerlof GA, Kranton RE (2010) *Identity Economics: How Our Identities Shape Our Work, Wages, and Well-Being* (Princeton Univ Press, Princeton).
9. Chen Y, Li SX (2009) Group identity and social preferences. *Am Econ Rev* 99(1):431–457.
10. Cohn A, Fehr E, Maréchal MA (2014) Business culture and dishonesty in the banking industry. *Nature* 516:86–89.
11. Bornstein G, Gneezy U, Nagel R (2002) The effect of intergroup competition on group coordination: An experimental study. *Game Econ Behav* 41:1–25.
12. Eckel CC, Grossman PJ (2005) Managing diversity by creating team identity. *J Econ Behav Organ* 58(3):371–392.
13. Charness G, Rigotti L, Rustichini A (2007) Individual behavior and group membership. *Am Econ Rev* 97:1340–1352.
14. Croson R, Marks M, Snyder J (2008) Groups work for women: Gender and group identity in the provision of public goods. *Negot J* 24(4):411–427.
15. Chen R, Chen Y (2011) The potential of social identity for equilibrium selection. *Am Econ Rev* 101(6):2562–2589.
16. Charness G, Holder P, Saunders D (2015) Charity in the laboratory: Matching funds, individual competition, and team competition. Working paper (University of California, Santa Barbara, CA). Available at https://static1.squarespace.com/static/575338d64c2f85618c6c01a8/t/58477aa9e6f2e130d85c8ac6/1481079469608/charity_november19_2016.pdf. Accessed December 6, 2016.
17. Kleinberg J, Ludwig J, Mullainathan S, Obermeyer Z (2015) Prediction policy problems. *Am Econ Rev* 105(5):491–95.
18. Bergstrom T, Blume L, Varian H (1986) On the private provision of public goods. *J Publ Econ* 29(1):25–49.
19. Andreoni J (1989) Giving with impure altruism: Applications to charity and Ricardian equivalence. *J Polit Econ* 97(6):1447–1458.
20. Andreoni J (1990) Impure altruism and donations to public goods: A theory of warm-glow giving. *Econ J* 100(401):464–477.
21. Bakija J, Heim B (2011) How does charitable giving respond to incentives and income? New estimates from panel data. *Natl Tax J* 64(2):615–650.
22. Eckel CC, Grossman PJ (2003) Rebate versus matching: Does how we subsidize charitable giving matter? *J Publ Econ* 87(3-4):681–701.
23. Karlan D, List JA (2007) Does price matter in charitable giving? Evidence from a large-scale natural field experiment. *Am Econ Rev* 97(5):1774–1793.
24. Andreoni J (1998) Toward a theory of charitable fund-raising. *J Polit Econ* 106(6):1186–1213.
25. Marx LM, Matthews SA (2000) Dynamic voluntary contribution to a public project. *Rev Econ Stud* 67(2):327–358.
26. Vesterlund L (2003) The informational value of sequential fundraising. *J Publ Econ* 87(3-4):627–657.
27. Potters J, Sefton M, Vesterlund L (2005) After you—endogenous sequencing in voluntary contribution games. *J Publ Econ* 89(8):1399–1419.
28. List JA, Lucking-Reiley D (2002) The effects of seed money and refunds on charitable giving: Experimental evidence from a university capital campaign. *J Polit Econ* 110(1):215–233.
29. Andreoni J, Bernheim BD (2009) Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects. *Econometrica* 77(5):1607–1636.
30. Ariely D, Bracha A, Meier S (2009) Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *Am Econ Rev* 99(1):544–555.
31. Meer J, Rosen HS (2011) The ABCs of charitable solicitation. *J Publ Econ* 95(5-6):363–371.
32. DellaVigna S, List JA, Malmendier U (2012) Testing for altruism and social pressure in charitable giving. *Q J Econ* 127(1):1–56.
33. Bertrand M, Karlan D, Mullainathan S, Shafir E, Zinman J (2010) What's advertising content worth? Evidence from a consumer credit marketing field experiment. *Q J Econ* 125(1):263–306.
34. Ricci F, Rokach L, Shapira B (2011) *Introduction to Recommender Systems Handbook* (Springer, New York).
35. Jordan MI, Mitchell TM (2015) Machine learning: Trends, perspectives, and prospects. *Science* 349(6245):255–260.
36. Chen Y, Li SX, Liu TX, Shih M (2014) Which hat to wear? Impact of natural identities on coordination and cooperation. *Game Econ Behav* 84:58–86.
37. Erev I, Bornstein G, Galili R (1993) Constructive intergroup competition as a solution to the free rider problem: A field experiment. *J Exp Soc Psychol* 29(6):463–478.
38. McPherson M, Smith-Lovin L, Cook JM (2001) Birds of a feather: Homophily in social networks. *Annu Rev Sociol* 27(1):415–444.
39. Girvan M, Newman MEJ (2002) Community structure in social and biological networks. *Proc Natl Acad Sci USA* 99(12):7821–7826.
40. Shayo M (2009) A model of social identity with an application to political economy: Nation, class, and redistribution. *Am Polit Sci Rev* 103:147–174.
41. Herlocker JL, Konstan JA, Riedl J (2000) Explaining collaborative filtering recommendations. *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work* (ACM, New York), pp 241–250.
42. List JA, Shaikh AM, Xu Y (2016) Multiple hypothesis testing in experimental economics. Working paper 21875 (NBER, Cambridge, MA). Available at www.nber.org/papers/w21875. Accessed December 1, 2016.
43. Imbens GW, Angrist JD (1994) Identification and estimation of local average treatment effects. *Econometrica* 62(2):467–475.

Supporting Information

Ai et al. 10.1073/pnas.1606085113

SI Methods

We first determine the experimental design and our subject pool. Then, in collaboration with Kiva, we implement our study by sending out mass emails with our recommendations.

Experimental Design. Our experiment consists of six treatments (three types of recommendations, with and without explanation), a control condition (no contact), and a placebo condition (teams exist). Table S1 displays our experimental design.

Although lenders in the control condition were not contacted during the experiment, for each treatment, we sent one of five email messages. Each email consists of three parts. Part 1 is common to all treatments and the placebo,

> "Hi [FirstName], Since you're such an awesome Kiva lender, we wanted to let you know about a fun feature of the Kiva experience: Kiva Lending Teams! Lending Teams are self-organized groups around shared interests—location, alumni orgs, social causes, you name it. You can connect with other lenders, discover loans you might be interested in, and track your collective impact."

Likewise, each email ends with Part 3,

"[Or] Check out the thousands of [other] lending teams to find the right one for you. Thanks for being a part of the Kiva community and making a difference around the world."

Although the text of emails sent to lenders in the placebo ("teams exist") condition consists of parts 1 and 3, lenders in the six treatments also received one of the following in the second part of the email:

1. Leaderboard with explanation treatment (Leaderboard-Explanation):

   "Some of the most popular teams are: [TEAMS]."

2. Location similarity with explanation treatment (Location-Explanation):

   "Other lenders who live near you enjoy being a part of these teams: [TEAMS]."

3. Loan history similarity with explanation treatment (History-Explanation):

   "Based on your past lending, people who have made similar loans enjoy being a part of these teams: [TEAMS]."

4. Recommendations without explanations treatments (Leaderboard-NoExplanation, Location-NoExplanation, History-NoExplanation)

   "Here are a few teams you may want to check out: [TEAMS]."

A sample email from the Location-Explanation treatment is included in Fig. S1. We now explain our recommendation algorithms.

Recommendations Based on Team Status. The simplest recommendation strategy is to recommend teams that are ranked highly on the team leaderboard. Kiva provides several leaderboards that rank teams based on either the total loan amount attributed to the team or the number of team members, in the most recent month or all time. For the experiment, we use the default leaderboard that lenders see when they visit the Kiva Team page, the all-time total amount lent.

Note that every lender receives the same recommendations under this strategy. The three teams we recommend to the

lenders are "Atheists, Agnostics, Skeptics,...," "Kiva Christians," and "Guys holding fish."

Recommendations Based on Location Similarity. The goal of this algorithm is to recommend the most popular teams in a lender's local area. This is motivated by the fact that there are many location-based teams on Kiva and by the conclusion of our previous work that the maximum location similarity between a lender and all of the teams is partially correlated with whether the lender has joined a team (4). This also reflects the results of an online data mining competition we ran with doctoral students at the University of Michigan using the Kiva API data hosted on https://inclass.kaggle.com/c/predict-new-team-memberships-on-kiva. The following algorithm, written by the first author, is the one that performed best in that competition. We calculate the location similarity between two lenders u and v as $l_{uv} \in \{0, 1, 2\}$ (4). If the two lenders are from different countries, $l_{uv} = 0$. If two lenders are from the same city, $l_{uv} = 2$. The condition for $l_{uv} = 1$ includes the following two cases: (*i*) if the two lenders are not in the same city but in the same state in the United States or Australia, or the same province in Canada, or (*ii*) if they are from the same country other than the United States, Australia, or Canada. This is because there are significantly more lenders on Kiva from the United States, Australia, or Canada than from any other country.

The location similarity of a team t in the neighborhood of a lender u is calculated as the sum of the location similarities between that lender and all lenders in that team. That is, $L(u, t) = \sum_{v \in T} l_{uv}$, where T denotes the set of lenders belonging to team t. For every lender, we rank all teams by the location similarity of these teams and recommend the three highest-ranked teams. For these recommendations, we exclude the three teams highest on the leaderboard: "Atheists, Agnostics, Skeptics,...," "Kiva Christians," and "Guys holding fish," for two reasons. First, the Atheists and Christians are outliers in that they overwhelm all other teams in size. Consequently, they often appear as winners of location-similarity–based recommendations. Second, to differentiate between status-based and homophily-based recommendations, we exclude all three teams.

Recommendations Based on Loan History Similarity. We also construct a recommender system based on the loan history of a lender. This is motivated by the homophily conjecture that lenders who lend to similar borrowers share similar interests and are thus more likely to join the same teams.

Borrowers on Kiva are registered in 80 countries from eight geographical regions (Oceania, Asia, etc.). They loan to facilitate 149 types of activities, which are further categorized into 15 sectors. Let S_u be a set of loans made by a user u and S_t be a set of loans that are attributed to a team t. The relevance of the team to the user is scored by the following function:

$$Relevance(u, t) = \sum_{i \in S_u} \sum_{j \in S_t} [f_g(i, j) + f_a(i, j)], \quad \textbf{[S1]}$$

where $f_g(i, j)$ equals 2 if the two loans i and j are from the same country, 1 if they are from two different countries in the same region, and 0 if they are not from the same region; $f_a(i, j)$ equals 2 if the two loans i and j are for the same activities, 1 if they are for different activities in the same sector, and 0 if they are not for activities in the same sector.

Note that the relevance score as defined in Eq. **S1** favors large teams that have made many loans. We further normalize the

score by taking into account the total number of loans made by each team. That is,

$$Normalized\_Relevance(u, t) = \frac{Relevance(u, t)}{|S_t| + 50}. \quad \textbf{[S2]}$$

Given a user who has not joined a team, we calculate the normalized relevance score for every team and recommend the three top-scoring teams to that user. For consistency with the recommendations based on location similarity, we also exclude the top three teams on the leaderboard, "Atheists, Agnostics, Skeptics,...," "Kiva Christians," and "Guys holding fish," for these recommendations.

Subject Pool. Based on the Kiva privacy policy and the information need of our recommendation algorithms, we select lenders for our experiment based on the following criteria:

- Their pages and loans are set to public in their account settings.
- They allow marketing emails in their communication settings.
- They have never joined a team.
- They provide location information in their profile.
- They have made at least two nonpromotion loans in the past 6 mo.

This gives us 69,845 users.

We then assign each user to one of the treatments, the placebo, or the control condition using stratified randomization. The stratified random assignment is based on the total loan amount by each lender before the experiment. We want to ensure that the most active Kiva lenders are not all concentrated into one treatment, so we rank the lenders based their total loan amounts, taking the top eight lenders and randomly assigning them to different conditions. We then repeat this for each group of eight lenders, proceeding down the ranked list. Between assigning lenders to conditions and running the experiment, 43 users joined a team and were dropped from our sample. This yields a final sample of 69,802 users. The size of the sample and population is summarized with a Venn diagram in Fig. S2.

Before running the experiment, we run pairwise Kolmogorov–Smirnov tests of the equality of distributions based on the user statistics to verify that our randomization produces balanced treatments across observable characteristics. The results of these tests show that the number of loans, average amount per loan, balance, average loan terms for fundraising or repayment, and autolending settings do not differ significantly at the 10% level between any treatments. Thus, the Kolmogorov–Smirnov tests do not reject the hypothesis that these values are drawn from the same distribution. We summarize the lending and location statistics of each treatment in Table S2.

Experimental Procedure. The experiment was conducted when the first author undertook an internship at Kiva.org in 2014.

We conduct the experiment in 2014, with Kiva sending out 61,077 emails to lenders in our sample (all except those in the no-contact control condition) within 1 d. After excluding lenders whose emails bounced and those who switched their pages to private and reincluding the lenders from our no-contact control group, we have a total of 64,800 lenders whom we intend to treat (henceforth "All"). Of these lenders, 20,371 open our email, constituting our treated subsample (henceforth "Opened"). We follow the team joining and lending behavior of each participant for 2 mo.
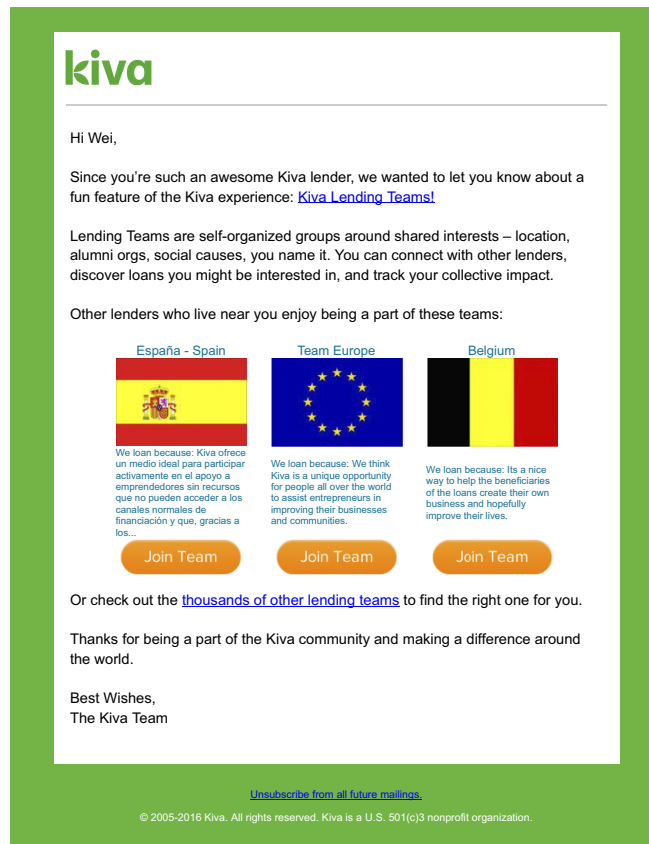
Fig. S1. Email screenshot. This email is an example from the Location-Explanation treatment. [© 2005–2016 Kiva. All rights reserved. Kiva is a US 501(c)3 nonprofit organization.]
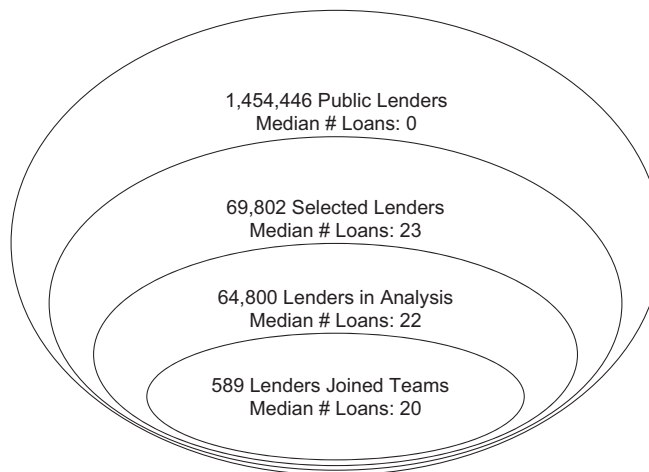

Fig. S2. Sample and population comparison. The number of lenders and median number of loans of all public users, those who are selected as participants, those whose data are used in our analyses, and those who joined at least one team during our experiment.

Table S1. Summary of experimental treatments

		Explanation of recommender algorithm	
		Explanation	No Explanation
Recommendation algorithm	Location	Location-Explanation	Location-NoExplanation
	Loan History	History-Explanation	History-NoExplanation
	Leaderboard	Leaderboard-Explanation	Leaderboard-NoExplanation
Control		No Contact	
Placebo		Teams Exist	

Table S2. Lending statistics of each treatment during 6 mo before experiment

		Lending statistics (average)			
Experimental condition	No. of users	Amount loaned	No. loans	Repayment term	Account balance
No-Contact	8725	184.29	6.07	18.50	36.24
Teams-Exist	8725	181.15	5.96	18.33	35.89
Location-Explanation	8726	181.34	6.04	18.45	35.22
Location-NoExplanation	8726	182.68	6.02	18.32	37.13
History-Explanation	8726	181.54	5.93	18.29	37.89
History-NoExplanation	8725	181.78	5.94	18.38	35.62
Leaderboard-Explanation	8723	182.14	6.05	18.40	34.37
Leaderboard-NoExplanation	8726	195.83	6.51	18.28	37.89

Note: Pairwise Kolmogorov–Smirnov tests comparing each experimental condition with the other yield $P > 0.10$ for each observable characteristic. "Amount loaned" and "Account balance" are in US dollars, whereas "Repayment term" is in months.