INFORMATION SYSTEMS RESEARCH
Vol. 32, No. 3, September 2021, pp. 774–785
ISSN 1047-7047 (print), ISSN 1526-5536 (online)
http://pubsonline.informs.org/journal/isre

Human–Robot Interaction: When Investors Adjust the Usage of Robo-Advisors in Peer-to-Peer Lending

Ruyi Ge,[a] Zhiqiang (Eric) Zheng,[b] Xuan Tian,[c] Li Liao[c]

[a] Department of Electronic Commerce, Shanghai Business School, Shanghai 200235, China; [b] Jindal School of Management, University of Texas at Dallas, Richardson, Texas 75080; [c] PBC School of Finance, Tsinghua University, Beijing 100083, China

Contact: gery@sbs.edu.cn (RG); ericz@utdallas.edu, https://orcid.org/0000-0001-8483-8713 (ZZ); tianx@pbcsf.tsinghua.edu.cn (XT); liaol@pbcsf.tsinghua.edu.cn (LL)

Received: January 15, 2019
Revised: December 22, 2019; June 4, 2020; August 23, 2020
Accepted: December 1, 2020
Published Online in Articles in Advance: July 20, 2021

https://doi.org/10.1287/isre.2021.1009

Abstract. We study the human–robot interaction of financial-advising services in peer-to-peer lending (P2P). Many crowdfunding platforms have started using robo-advisors to help lenders augment their intelligence in P2P loan investments. Collaborating with one of the leading P2P companies, we examine how investors use robo-advisors and how the human adjustment of robo-advisor usage affects investment performance. Our analyses show that, somewhat surprisingly, investors who need more help from robo-advisors—that is, those encountered more defaults in their manual investing—are less likely to adopt such services. Investors tend to adjust their usage of the service in reaction to recent robo-advisor performance. However, interestingly, these human-in-the-loop interferences often lead to inferior performance.

History: Hemant Jain, Balaji Padmanabhan, Paul Pavlou, and Raghu Santanam, Senior Editors; Gordan Burtch, Associate Editor. This paper has been accepted for the *Information Systems Research* Special Section on Humans, Algorithms, and Augmented Intelligence: The Future of Work, Organizations and Society.
Funding: R. Ge received financial support from the Ministry of Education in China Research Program for Humanities and Social Sciences [Grant No. 19YJC630041]. E. Zheng received financial support from the National Natural Science Foundation of China (NSFC) [Grants 71850013 and 71532004]. T. Xuan received financial support from the NSFC [Grants 71825002 and 71790591].
Supplemental Material: The online appendix is available at https://doi.org/10.1287/isre.2021.1009.

Keywords: **robo-advisor • human-in-the-loop • peer-to-peer lending • augmented intelligence**

1. Introduction

Robo-advisor (hereafter RA) is a service that provides automated, algorithm-based wealth-management advice without the use of a human financial planner.[1] Typically, these services use algorithms to help investors determine how to invest based on their risk preference, budget, and investment goals. In other words, RAs help augment investor intelligence in a personalized manner. Compared with human advisors, RAs are more accessible (being available 24/7), and they charge less (e.g., 0.25% compared with the 2–20 standard in the financial advising industry).[2] RAs also require much smaller capital outlays for receiving personal financial advice—for example, $500 for Wealthfront compared with $50,000 for Vanguard.[3] Since the first RA launched in 2008, the industry has grown rapidly. As of 2019, the three largest stand-alone RAs, Betterment, Wealthfront, and Personal Capital, boast assets under management (AUM) of approximately $16 billion, $11 billion, and $8.5 billion, respectively.[4] Recently, traditional wealth-management companies, such as Vanguard and Charles Schwab, have also started to incorporate RAs into their financial-advising services. For example, Charles Schwab's intelligent portfolio provides clients with robo-advising services for managing conventional accounts, such as 401(k), IRA, trust, and 529 plan accounts. The total AUM of the RA industry is expected to increase to $2.2 trillion by 2020 (KPMG 2016).

Most RAs in wealth management are founded based on Markowitz's portfolio-optimization theory (Friedberg 2019), creating a diversified investment portfolio with the greatest returns for each risk level (Markowitz 1952). As such, the basic inputs are typically the returns and variance–covariance matrix of asset returns. RAs then employ computer algorithms to optimize the risk–return tradeoff and recommend a diversified portfolio accordingly. Some RAs start by using sophisticated machine learning algorithms, such as random forest, neural network, and nonlinear shrinkage methods, in their optimization model (D'Acunto et al. 2019, D'Hondt et al. 2019).

Although most RAs operate in the conventional wealth-management domain and help their clients build a portfolio of traditional assets—for example, stocks, bonds, and commodities—others explore new territories such as peer-to-peer (P2P) loans. Until July 2018, more than $23 billion in loans originated in the two largest

U.S. P2P lending platforms, Prosper and Lending Club, while more than $1,080 billion in loans have been transacted on Chinese P2P lending platforms (Jiang et al. 2020). A lender (i.e., investor) on a typical P2P platform usually needs to choose among hundreds of available loans to invest in at any time. These loans have different interest rates and default probabilities, and every loan is unique with limited information (e.g., loan descriptions) for lenders to evaluate. It is, therefore, challenging for lenders to optimize their loan investment in such an environment. Given this, most mainstream P2P platforms (e.g., Lending Club, Prosper, and PPdai.com) have started providing RA services to help lenders choose loans worthy of funding. For example, a third-party company, Fastbacker, builds a robo-advising application that monitors Kickstarter projects and notifies investors when suitable projects become available. LendingRobot RAs help lenders automate the management of their accounts across multiple P2P platforms, such as Lending Club and Prosper. Some RAs even help lenders design investment strategies. For example, Lending Club collaborated with InterestRadar to offer a robo-advising service that helped lenders scan available P2P loans and assisted them in choosing appropriate investment strategies for loans that met their prespecified criteria. These RAs gained popularity among lenders quickly. For example, at PPdai, the first P2P lending platform in China, automated investments through RAs have outnumbered manual investments since it launched the RA service in 2015.

As these *intelligence-augmentation* tools become increasingly popular in people's daily lives, it is important to understand how humans and algorithms should collaborate. The nascent literature on human-in-the-loop (e.g., Dietvorst et al. 2016, Xu and Chau 2018, Fügener et al. 2019) highlights the importance of having humans engaged in designing, implementing, and refining algorithms. We draw on this body of literature in the context of human–RA interactions. Specifically, we are interested in how investors use RA services in their investments and whether having humans in the loop of RA deployment augments investment performance or not.

Researchers and practitioners, however, have little understanding on these issues. This study attempts to fill these gaps by examining the human–RA interaction through collaboration with a leading P2P lending company publicly traded on the NASDAQ.[5] Lenders there can easily access the RA service and activate it by simply clicking a specific button on the company's homepage (see Online Appendix 1 for a screenshot). Once a lender decides to use an RA, they configure their risk preference and investment amount. Lenders can turn off the RA service at any time they deem necessary.

The company provided us with the data on the complete transaction history of a random sample of lenders, including all the loans each lender funded, detailed information on each loan transacted (e.g., investment amount, date, maturity, interest rate, and payment status), and, most distinctively, the information on whether the lenders invested in the loans manually or through an RA. We observed a mix of lender populations. Some relied totally on RAs for choosing loans, some used the service occasionally, and others never tried the service at all.

This data set provides us with a unique venue to investigate the interwoven effects of investors' use of RAs and the corresponding performance of investments. Specifically, we study the following three research questions:

1. How does investors' investment performance in the past influence their RA adoption when the service becomes available?
2. How do investors adjust their usage of RAs according to the RA's investment performance?
3. How does the adjustment affect investment performance?

Taken together, the answers to these questions will help us answer the overarching question pertaining to how investors interact with RAs and whether having humans in the loop of using RA helps improve P2P investment.

We find that investors who encountered more defaults in the past are less likely to try RA when the service becomes available. RA usage is positively influenced by recent RA performance: When recent RA performance is lower, investors decrease their usage of RA immediately, and vice versa. However, such swift adjustment in RA usage often leads to worse investment performance, especially when the adjustments are frequent and substantial.

Our research makes several contributions. It represents one of the first attempts at investigating RA-augmented intelligence in P2P lending investments. It also provides the first empirical evidence demonstrating how investors' investment performance influences RA adoption. This finding can help RA marketers target certain customer segments to improve adoption rates. Moreover, as the first study on human–RA interaction, our results show that users are subject to the recency effect when evaluating RAs. They experience more losses due to being too reactive to recent RA performance. This presents a new, but negative, use case for human—artificial intelligence (AI) symbiosis, where leaving too much control to humans over when to use an RA may be counterproductive. This result reflects investors' possible misunderstanding and misuse of RAs. They may not always have proper knowledge of RA systems and may intervene counterproductively. It suggests that such RA systems need to offer more transparency in their services (Friedberg 2018), for example, by communicating with investors on their RA's objective and inner-working mechanisms. Conversely, it also suggests that a well-designed RA should anticipate the possible adjustments lenders may make and factor in such reactions in their algorithms' design.

2. Background and Research Context

2.1. Literature Review

We first briefly review the nascent body of literature on robo-advising. The scant literature largely focuses on describing the features of RAs (e.g., Lopez et al. 2015, Park et al. 2016, and Jung et al. 2017) or the IT components inside RAs (e.g., Musto et al. 2015 and Jung et al. 2018). Recently, a few studies have assessed the benefits ofs.023649.11(m).1TD-.d.9(n)v5804is11(m)-3.94-2445282.864itrst(rc)o-1.9(s.2

to these machine learning methods consist of loan characteristics (term, amount, loan description, etc.) and borrower characteristics (education, employment status, financial status, social network features, borrowing history, online behavior, mobile communication features, etc.). A sample screenshot detailing the main variables is presented in Online Appendix 2.

2. Next, building on Markowitz's portfolio-optimization approach, the RA chooses and invests in the loans that meet the lender's risk preference.

The RA service became very popular among lenders; more than half of the bids were conducted by RAs after one year of release.

3. Data Description

We obtained a random data sample of 4,374 lenders from the company with the complete history of their bids across 18 months, from January 2015 to June 2016.[6] The descriptions of the sample are presented in Table 1.

In our sample, 73% of the lenders were male, and the average lender was 37.6 years old, with 1.25 years of investment experience on the platform.[7] On average, a lender invested 251.2 renmibi (RMB) per bid and 138,555 RMB in total. The means of lenders' annualized interest rates and terms were 15.96% and 8.94 months, respectively. During the 18 months, 63% of lenders used the RA service to invest in at least one loan, and the lenders' average monthly return rate was 1%.[8]

It is noteworthy that the means of *BidAmount* and *TotalAmount* are much larger than their medians, which implies positive skewness. Therefore, we use the natural logarithm of these variables in the following analyses.

4. RA Adoption

Our first research question (RQ1) asks how investment performance in the past affects investors' RA adoption. RQ1 investigates the human–RA interaction from the adoption (i.e., first interaction) perspective as a function of investors' past performance. Past performance may affect RA adoption through two possible underlying mechanisms. First, investors' previous investment performance will affect the perceived usefulness of RAs. An investor whose past performance was inferior is more likely to count on RAs to improve their performance; that is, RAs' perceived usefulness turns higher. There has been abundant information systems literature documenting that users' perceived usefulness or performance expectancy concerning a technology increases the likelihood of technology adoption (e.g., Venkatesh et al. 2003). This suggests that investors experiencing inferior performance should be more likely to adopt RAs when the service becomes available.

On the other hand, ceteris paribus, underperforming means that the investor has encountered more defaults than others. The investor would then have a stronger perceived risk regarding P2P loans on the platform and would thus be less certain about RA performance in such cases. Prior studies have shown that investors are less likely to adopt a new technology when the perceived risk of using it is high (Featherman and Pavlou 2003). When performance risk is high—for example, the possibility of technology malfunctioning or technology not performing as designed or advertised—the technology will fail to deliver the desired benefits (Featherman and Pavlou 2003). This suggests that investors with inferior past performances would be less likely to adopt RAs because of their higher level of perceived RA performance risk.

Because these two potential effects may counteract each other, our research question sets out to answer which one dominates in our study context.

Table 1. Sample Description

Variable	Mean	S.D.	Min	Median	Max	N
Gender	0.73	0.45	0	1	1	4,370
Age	37.62	9.68	20	35	75	4,340
Experience	1.25	1.19	0	1	9	4,374
BidAmount	251.2	608.9	10	111.4	13698	4,374
TotalAmount	138,555	477,683	50	32,916	1.2e+07	4,374
InterestRate	15.96	3.75	7	16.35	23.64	4,374
Term	8.94	2.54	1	9.41	19.45	4,374
RAAdopted	0.63	0.48	0	1	1	4,374
ReturnRate	0.01	0.003	−0.03	0.01	0.02	4,374

Note. The units of *BidAmount* and *TotalAmount* are Chinese RMB.

4.1. Empirical Specifications

Our data sample began in January 2015, and the RA service launched in April 2015. In our sample, approximately 1,000 lenders had investment transactions both before and after April 2015, which provides a good setting for examining how lenders reacted to the service's launch. We find that more than 50% of the first tryouts occurred in the first month, and nearly 75% of tryouts occurred within the first three months. We examine the effect of lenders' previous investment performance on their RA adoption behavior using the following two cross-sectional models:[9]

$$\begin{aligned} Prob(RAAdopted_{i,T} = 1 \mid X) = Logit(&\alpha_0 \\ &+ \alpha_1 Previous\_Investment\_Performance_i \\ &+ \alpha_2 Previous\_Investment\_Characteristics_i \\ &+ \alpha_3 Controls_i), \end{aligned} \tag{1}$$

$$\begin{aligned} RAShare_{i,T} = \beta_0 + \beta_1 &Previous\_Investment\_Performance_i \\ &+ \beta_2 Previous\_Investment\_Characteristics_i + \beta_3 Controls_i + \varepsilon_i. \end{aligned} \tag{2}$$

We consider two alternative dependent variables, $RAAdopted_{i,T}$ and $RAShare_{i,T}$, to measure lenders' adoption behavior. $RAAdopted_{i,T}$ denotes whether a lender has ever used the RA service during a period of T months after the service becomes available; it equals one if the lender has used the service to invest in at

Table 2. The Effect of Previous Investment Performance on RA Adoption

	Panel A: Logit specification		Panel B: Tobit specification	
	(1)	(2)	(3)	(4)
Variable	$RAAdopted_{i,T=1}$	$RAAdopted_{i,T=3}$	$RAShare_{i,T=1}$	$RAShare_{i,T=3}$
$ReturnRate_i$	−52.979	−119.361	−15.934	−27.746
	(76.178)	(79.433)	(20.159)	(21.655)
ln(#*Default*)$_i$	−0.499**	−0.487**	−0.106**	−0.180***
	(0.210)	(0.208)	(0.049)	(0.060)
$InterestRate_i$	−0.080	−0.042	−0.025	−0.016
	(0.062)	(0.066)	(0.017)	(0.018)
$Term_i$	0.306***	0.265***	0.087***	0.088***
	(0.041)	(0.039)	(0.010)	(0.010)
ln(*BidAmount*)$_i$	−0.078	−0.057	0.029	0.011
	(0.127)	(0.126)	(0.032)	(0.036)
ln(*TotalAmount*)$_i$	0.195***	0.136**	0.005	0.012
	(0.059)	(0.058)	(0.016)	(0.017)
Lender characteristics	Controlled	Controlled	Controlled	Controlled
Observations	924	984	924	984
R^2	0.077	0.066	0.083	0.057

$^{**}p < 0.05$; $^{***}p < 0.01$.

least one loan during the period, and zero otherwise. $RAShare_{i,T}$ is the proportion of RA bids among all the bids a lender invested during the period, capturing the intensity of a lender's RA usage. In both models, i indexes lender; T equals one or three, standing for one month or three months after the RA launch (i.e., April 2015). In other words, if $T = 1$, we calculate $RAAdopted_i$ and $RAShare_i$ based on the data from May 2015; if $T = 3$, we calculate $RAAdopted_i$ and $RAShare_i$ using the three-month data from May to July 2015.[10]

Previous_Investment_Performance$_i$ takes two measures:[11] $ReturnRate_i$ and ln(#*Default*$_i$), representing lenders' average monthly return rate and (the natural logarithm of) the number of defaulted loans that lenders encountered before the launch of RA, respectively. The vector *Previous_Investment_Characteristics*$_i$ includes the average interest rate, the average terms of a lender's investment (i.e., $InterestRate_i$ and $Term_i$), and the natural logarithm of the lender's bid amount and the total amount (i.e., ln(*BidAmount*$_i$) and ln(*TotalAmount*$_i$)) before the RA launch. Because our data sample began in January 2015, *Previous_Investment_Performance*$_i$ and *Previous_Investment_Characteristics*$_i$ are calculated based on January, February, and March 2015 data. The vector *Controls* contains lender characteristic variables, including *gender*, *age*, and *experience*.

4.2. Results

In Table 2, panel A reports the results from the above logit specification. The coefficients for $ReturnRate_i$ are insignificant, whereas those for ln(#*Default*$_i$) are significant. Specifically, columns (1) and (2) suggest that when #*Default*$_i$ increases by 1%, the odds of $RAAdopted_{i,T=1}$ decreases by 39.3% (odds ratio = 0.607), and the odds of $RAAdopted_{i,T=3}$ decreases by 38.5% (odds ratio = 0.615). These results indicate that a lender experiencing a higher level of loan defaults is less likely to try the RA service. Columns (3) and (4) in panel B report the Tobit regression results with the dependent variable $RAShare_{i,T}$. The results show that ln(#*Default*$_i$) exhibits a significant and negative effect on $RAShare_{i,T}$, which is consistent with the results of panel A.

Taken together, the results in Table 2 suggest that investors' past investment performance affects their adoption of RAs. In other words, a human's own past performance may well be in the loop regarding the first interaction (adoption) decision with RA services. Interestingly, it is the number of defaulted bids, rather than bid return rates, that influence lender adoption behavior significantly, possibly because #*Default* conveys a clearer and more straightforward risk message, as opposed to *ReturnRate*. Loan defaults are painful, salient events for investors. According to prospect theory (Tversky and Kahneman 1974), salient instances affect people's assessments of the probability of an event occurring the most. Lenders experiencing more defaulted loans are more likely to perceive the P2P market to be risky and, thus, tend to rely more on their own judgment rather than an RA's, echoing the findings of Featherman and Pavlou (2003), who show that risk perceptions exert a negative impact on the use of e-services.

4.3. Robustness Checks

4.3.1. Alternative Explanation.

One potential alternative explanation is that investors' capability, rather than investment performance, drives investors' RA adoption. However, we do not directly observe investors' hidden abilities. To alleviate this concern, we replace ln(#*Default*) with ln(#*Default_Ultimate*). The former records whether a loan was defaulted before the RA launch, while the latter

Table 3. The Effect of Investors' Capability on RA Adoption

	(1)	(2)	(3)	(4)
Variable	$RAAdopted_{T=}$			

is a forward-looking metric capturing whether a loan ultimately defaults. Obviously, the latter is a more accurate proxy for investors' capabilities. This renders a different result: The coefficients for ln(#*Default_Ultimate*) are not statistically significant (see Table 3), whereas those of ln(#*Default*) remain significant. This test alleviates the concern that the capability of investors is the more likely driver behind RA adoption.

4.3.2. Coarsened Exact Matching.

We then use the coarsened exact matching (CEM) approach to alleviate the above endogeneity concern further. CEM coarsens each covariate into meaningful bins, matches observations based on these bins, and then retains the covariates' original values for analysis (Blackwell et al. 2009). Compared with some other matching methods, such as propensity score matching, CEM can generate matched data sets with lower imbalance (Iacus et al. 2012). To make full use of the data, we use the $T = 3$ data set, which has more observations than $T = 1$ for the matching procedure. We divide lenders into two groups: One group of lenders encountered no defaults (i.e., control group) before the RA launch. The other group of lenders encountered at least one default (i.e., treatment group) before the RA launch. We match these two groups with two sets of covariates, previous investment characteristics and lender characteristics, with a CEM procedure. In total, 126 lenders are matched. The logit regression results of the covariates before and after matching are shown in Table 4.

Then, we use the matched samples to regress lenders' adoption outcome on the treatment. Table 5 shows that the treatment exerts significant and negative effects on *RAAdopted* and *RAShare*, consistent with the main model's results.

5. Adjustment of RA Usage

Our second research question (RQ2) studies how investors adjust their RA usage based on recent RA performance. RQ2 looks at investor interaction with RAs during the phase of using RA services in investments as a function of past RA performance. The investor is allowed to enable or disable RA services at any point in time.

We explain how RA usage is adjusted from the lens of the recency effect, which is the tendency of an individual to recall or emphasize the most recent events. This effect was first discovered in cognitive science (Deese and Kaufman 1957, Murdock 1962) and then applied in finance (e.g., Cushing and Ahlawat 1996, Arnold et al. 2000, and Pompian 2011). Pompian (2011) points out that a manifestation of the recency effect among investors explains their misuse of investment-performance records for mutual funds. Investors tend to analyze a small data sample, such as the fund performance of recent periods, and then make investment decisions based on such recent experiences without paying attention to the cyclical nature of asset class returns. RA services in P2P lending are designed to select suitable loans from all the listed loans on the platform, so as to build a portfolio that meets an investor's long-term risk and return objectives. However, P2P lending platforms typically release the performance of RA investments to investors monthly. Thus, it is interesting to investigate whether RA users are subject to the recency effect, adjusting their RA usage mainly based on RAs' recent and short-term performance.

5.1. Empirical Specification

In order to examine how recent RA investment performance influences investors' usage of the RA service, we construct a one-year panel starting from May 2015 (immediately after RA became available), with which improve their investment performance. However, it is not clear whether such interference pays off. As pointed out by Pompian (2011), the recency effect can cause investors to make suboptimal decisions as a result of relying on historical data samples that are too small to ensure accuracy, which may inadvertently end up in losses. Our investigation of RQ2 reveals that investors make adjustments based on RAs' monthly performance. RA services in P2P lending typically focus on relatively long-term returns (Ludwig 2020), such as annual returns. Evaluating RA performance and adjusting RA usage based on monthly data may thus be suboptimal. We set out to answer RQ3 by examining the impact of RA usage adjustment on the return rate of loans.

6.1. Empirical Specifications

133gi0-e 6(t)04n06-51o305(33in2n9(4-v2)n9(7(st)21(6-me)n9)n.2-298(4)am)17(s)ou49-th.2(t31

translates into a decrease of nearly 235 RMB in annual return for an average investor, when she increases RA usage by 10% in its coefficient of variation, holding her mean and standard deviation of *RAShare* constant at 42% and 43%, respectively. For the second sample, a one-unit increase in *RAShare_Std* decreases the average monthly return rate by 0.1%—that is, a 1.2% decrease in the average annual return rate. Hence, it seems better to let the algorithms do the work; having humans (lenders) in the decision loop in terms of enabling or disabling RA services can be counterproductive. RAs aim to achieve long-term portfolio optimization concerning risk and return, which means there is a long-term mean that the RA targets. A bad loan is just a small deviation from the long-term mean. Manually adjusting the usage of RAs too frequently and substantially may inadvertently disrupt the stochastic process of RA performance, leading to inferior performance.

6.3. Robustness Checks

6.3.1. Coarsened Exact Matching. Here, we apply a CEM approach to strengthen identification. We first divide lenders into two groups based on *RAShare_CoV*. Lenders in the top 20% of *RAShare_CoV* form the treatment group (n = 241), and lenders in the bottom 20% of *RAShare_CoV* form the control group (n = 238). We utilize the CEM method to eliminate the difference caused by the covariates between the two groups and then regress the outcome variable on the treatment. We match the two groups with the two sets of covariates, investment characteristics and lender characteristics, with two CEM procedures. In the first procedure, 134 lenders are matched. The matched samples are mostly balanced, except for *experience*. In the second procedure, 59 lenders are matched. The matched samples are completely balanced. Table 10 tabulates the logit regression results of covariates before and after matching. We then use the matched samples to regress lenders' *ReturnRate* on the treatment. Table 11 shows that the treatment has significant and negative effects on *ReturnRate*, consistent with the main models' results.

6.3.2. Causal Forest. As a sensitivity analysis to further solidify identification, we applied a causal forest approach to estimate the treatment effect (Wager and Athey 2018). Causal forest has been widely used to estimate and infer heterogeneous treatment effects (e.g., Davis and Heller 2017 and Luo et al. 2019). In our case, however, we are not attempting to estimate heterogeneous treatment effects. Instead, we use this approach as an alternative matching method to CEM, where the samples falling into each leaf are considered homogeneous. We treat *RAShare_CoV* adjustments. Potentially, investors may decrease RA shares when RAs did not do well in the previous period or, conversely, increase RA shares when RAs per-

[5] The nondisclosure agreement we signed with the company requires us to ensure anonymity of the name of the company.

[6] The company randomly selected these samples according to the last two digits of lenders' user IDs. Each user ID is generated based on lenders' registration sequences. The random sample accounts for 0.3% of the entire lender population in 2015.

[7] The distributions of gender and age in our sample are similar to the population statistics released by the platform in 2016. A few lenders did not report their genders or ages, making the number of records (N in Table 1) of the two variables slightly smaller than that of others.

[8] The monthly return rate is calculated based on the equation specified by the platform—that is, ReturnRate = (Σinterests obtained in the focal month − Σprincipal losses in the focal month)/Σprincipals that were lent out in the focal month.

[9] We also model the hazard of adoption as a function of previous performance. The results, presented in the online appendix, are consistent with the main models.

[10] We also estimate Equations (1) and (2) with a larger adoption window, $T = 6$, which covers nearly 90% tryouts after the launch of the RA service. The results are consistent with the results of $T = 1$ and $T = 3$. The details are presented in the online appendix.

[11] We use these two performance measures because $ReturnRate_i$ and $\#Default_i$ are directly displayed on the monthly report provided by the platform to lenders. However, $\%Default_i$ is not provided by the

Pompian M (2011) *Behavioral Finance and Wealth Management: How to Build Investment Strategies that Account for Investor Biases* (John Wiley & Sons, Hoboken, NJ).

Schanke S, Burtch G, Ray G (2021) Estimating the impact of 'humanizing' customer service chatbots. *Inform. Systems Res.*, ePub ahead of print May 24, https://doi.org/10.1287/isre.2021.1015.

Schierz PG, Schilke O, Wirtz BW (2010) Understanding consumer acceptance of mobile payment services: An empirical analysis. *Electronic Commerce Res. Appl.* 9(3):209–216.

Srivastava SC, Chandra S, Theng YL (2010) Evaluating the role of trust in consumer adoption of mobile payment systems: An empirical analysis. *Commun. Assoc. Inform. Systems* 27:561–588.

Tversky A, Kahneman D (1974) Judgment under uncertainty: Heuristics and biases. *Science* 185(4157):1124–1131.

Venkatesh V, Morris M, Davis G, Davis F (2003) User acceptance of information technology: Toward a unified view. *MIS Quart.* 27(53):425–478.

Wager S, Athey S (2018) Estimation and inference of heterogeneous treatment effects using random forests. *J. Amer. Statist. Assoc.* 113(523):1228–1242.

Wang J, Ipeirotis PG, Provost F (2017) Cost-effective quality assurance in crowd labeling. *Inform. Systems Res.* 28(1):137–158.

Xin D, Ma L, Liu J, Macke S, Song S, Parameswaran A (2018) Accelerating human-in-the-loop machine learning: Challenges and opportunities. *DEEM'18 Proc. Second Workshop Data Management End-to-End Machine Learn.* (Association for Computing Machinery, New York), 1-4.

Xu J, Chau M (2018) Cheap talk? The impact of lender-borrower communication on peer-to-peer lending outcomes. *J. Management Inform. Systems* 35(1):53–85.

Yin J, Luo J, Brown SA (2021) Learning from crowdsourced multilabeling: A variational Bayesian approach. *Inform. Systems Res.* Forthcoming.

Zhou T (2013) An empirical examination of continuance intention of mobile payment services. *Decision Support Systems* 54(2):1085–1091.