Search and Negotiation with Biased Beliefs in Consumer Credit Markets*

Erik Berwart Sean Higgins Sheisha Kulkarni Santiago Truffa

December 17, 2025

Abstract

How do biased beliefs about the distribution of interest rates affect search, negotiation, and loan terms in consumer credit markets? Motivated by a model of sequential search and negotiation with potentially biased beliefs, we conducted a randomized controlled trial with 112,063 loan seekers in collaboration with Chile's financial regulator. We first elicited beliefs about the interest rate distribution, then showed treated participants a price comparison tool that we built using administrative data on the universe of consumer loans merged with borrower characteristics. The tool shows loan seekers a conditional distribution of interest rates based on similar loans obtained recently by similar borrowers. We find that most consumers thought interest rates were lower than they actually were and also underestimated price dispersion, and the price comparison tool caused them to update their beliefs. The price comparison tool did not cause people to search or apply at more institutions, but it did cause them to be 39% more likely to negotiate with their lender, to receive 13% more offers and 11% lower interest rates, and to be 5% more likely to take out a loan. *JEL Codes:* G51, D83, D14, D18.

*Berwart: Comisión para el Mercado Financiero, eberwart@cmfchile.cl. Higgins (corresponding author): Department of Finance, Kellogg School of Management, Northwestern University, 2211 Campus Dr., Evanston, IL 60208, sean.higgins@kellogg.northwestern.edu. Kulkarni: Department of Finance, McIntire School of Commerce, University of Virginia, sk7nn@virginia.edu. Truffa: Department of Economics and Finance, ESE Business School, Universidad de los Andes, struffa.ese@uandes.cl. We gratefully acknowledge funding from the Think Forward Initiative (TFI), ING Bank; Financial Institutions and Markets Research Center, Avi Nash Fund, and Guthrie Center for Real Estate Research, Northwestern Kellogg; Household Finance Small Grant Program, Alfred P. Sloan Foundation and NBER; Sparkassenstiftung für Internationale Kooperation; University of Virginia; Fondecyt iniciacion: 11200010; Digital Credit Observatory and Lab for Inclusive Fintech (LIFT), UC Berkeley; and Consumer Protection Research Initiative, Innovations for Poverty Action. We are very grateful to Bronson Argyle, Deniz Aydın, Claire Célérier, Ben Charoenwong, Isis Durmeyer, John Grigsby, Chris Hansman, Clément Mazet-Sonilhac, Monica Paiella, Guosong Xu, and Constantine Yannelis for discussing our paper. We thank conference and seminar participants at AIEA, Bank of Canada, BI Norwegian Business School, Boulder Summer Conference on Consumer Financial Decision Making, Cambridge Finance Conference, Canadian Economics Association, CEPR European Workshop on Household Finance, CEPR WEFIDEV Workshop on Finance and Development, CFPB, Chicagoland Development Economics Meeting, Central Bank of Chile, Columbia Workshop in New Empirical Finance, Commonwealth Finance Workshop, Conference on Digital Experimentation, CSIO/TSE Joint Workshop on Industrial Organization, CUHK, EFA, Emory University, Erasmus University, EUROFIDAI-ESSEC, FCA, FDIC, HEC Montréal, ISB, IPA–GPRL Researcher Gathering, Junior Household Finance Seminar, Lab for Inclusive FinTech at UC Berkeley, MFA, MIT, MoFiR Singapore, NBER Summer Institute Household Finance, New York University, Northwestern University, Paris December Finance Conference, Peking University, Philadelphia Federal Reserve, PUC Chile, SFS Cavalcade North America, Stanford SITE Financial Regulation Workshop, Stanford IFDM Conference, Universidad de los Andes, University of Chile, University of Houston, University of Oxford, University of Virginia, University of Washington Summer Finance Conference, Wayne State University, and WFA for helpful comments and discussions. We thank Arturo Charleston, Vicente Corral, Andrés Cruz, Matías Fuentes, Jora Li, Nicolas Min, Erick Molina, Limin Peng, Carlos Restituyo, Agustín Valenzuela Pidal, and Francisco Villarroel for their excellent research assistance. This RCT was pre-registered in the AEA RCT registry under trial number 8553 (https://www.socialscienceregistry.org/trials/8553). IRB approvals: NBER #22_099, #23_154, #24_138, Northwestern University #STU00213001, and Universidad de los Andes #CEC201946.

1 Introduction

Consumer credit markets feature large amounts of *within-consumer* price dispersion (Stango and Zinman, 2016; Ponce, Seira, and Zamarripa, 2017). Even though many consumers pay substantial costs by borrowing at higher rates than they could, this price dispersion can persist in equilibrium if consumers engage in limited search and negotiation (Stahl, 1989; Hortaçsu and Syverson, 2004; Allen, Clark, and Houde, 2014).[1] Why, then, do consumers not search or negotiate more? The existing literature has focused on *costs* that prevent search or negotiation, including time and travel costs, high rejection rates, cognitive effort to compare complex offers, and the costs of gathering additional quotes to use in a negotiation.[2] We focus instead on the expected *benefits*, and test whether biased beliefs about the interest rate distribution constrain search and negotiation.

Motivated by a model of sequential search and negotiation with potentially biased beliefs, we conducted a randomized controlled trial (RCT) in close collaboration with Chile's financial regulator, the Comisión para el Mercado Financiero (CMF). The RCT, conducted with 112,063 Chileans searching for loans, tests how correcting biased beliefs about the distribution of interest rates affects search, negotiation, and loan terms in consumer credit markets. We both measure beliefs about the interest rate distribution and show loan seekers in the treatment group a price comparison tool designed to correct biased beliefs. We built the price comparison tool using administrative data from CMF on the universe of consumer loans merged with borrower characteristics. The tool shows participants the conditional distribution of interest rates based on similar loans obtained recently by similar borrowers.

In a standard sequential search model (without negotiation), biased beliefs would affect the perceived benefits of search and thus the consumer's reservation rate. If consumers underestimate the first moment of the distribution of interest rates (as we find most consumers do), they will search more than is optimal because they will overestimate the expected benefit from another draw from the distribution of rates. If consumers underestimate the second moment of the distribution (as we also find that most consumers do), they will search less than is optimal because they underestimate the expected benefit from another draw.

Biased beliefs can also affect negotiation. In a model where consumers negotiate with lenders and the lender observes a signal about the consumer's beliefs about the interest rate distribution

[1] Consumers incur a substantial cost to borrow at higher rates than they could obtain with more search or negotiation. In the US auto loan market, the average borrower pays $488 more in present value for a $17,000 car (Argyle, Nadauld, and Palmer, 2023). In mortgage markets, Woodward and Hall (2012) estimate that borrowers pay $1,000 more for a $100,000 mortgage, and Bhutta, Fuster, and Hizmo (2024) estimate that borrowers pay $6,250 more for a median $250,000 mortgage.

[2] On physical search costs across branches, see Allen, Clark, and Houde (2013) and Argyle, Nadauld, and Palmer (2023). On high rejection rates leading to higher per-offer search costs for less-creditworthy borrowers, see Agarwal, Grigsby, Hortaçsu, Matvos, Seru, and Yao (2024). On the cognitive effort required to compare offers, see Galenianos and Gavazza (2022). For evidence on the costs of negotiating in mortgage markets, see Allen and Li (2025).

during the negotiation, we show that a consumer's reservation rate—which is a function of their beliefs and their search cost—is no longer a sufficient statistic for search. Consumers who underestimate the second moment of the interest rate distribution will (successfully) negotiate less than if they had accurate beliefs. The model's prediction for how biased beliefs about the first moment affect negotiation, however, is non-monotonic. Thus, in a sequential search model with negotiation, correcting consumers' underestimates of the second moment should increase negotiation while not necessarily increasing search, and the effects of correcting consumers' underestimates of the first moment depend on how biased their beliefs about the first moment were. In particular, for consumers who overestimated or substantially underestimated the first moment of the distribution, the tool will not have a treatment effect on negotiation, while it will increase the probability of successfully negotiating for those who somewhat underestimated the first moment.

We test these predictions from the model in an RCT where we recruited participants through Google ads targeted to people searching for keywords related to consumer loans in Chile. After participants clicked on the Google ad and consented to participate in the study, we collected their contact information and national ID numbers, which we use to track participants' future loan outcomes in administrative data.[3] We then had them fill out a baseline survey that randomized whether we asked them their beliefs about the distribution of interest rates (which we refer to as the "elicit beliefs" treatment). After the baseline survey, we cross-randomized whether we showed them a price comparison tool, a simple tool showing our estimate of the cost savings from search in pesos, or a control video. We then elicit beliefs again from those assigned to the elicit beliefs treatment, to measure the extent to which participants update their beliefs after seeing the price comparison or simple tool. To measure real outcomes on search, negotiation, and loan terms, we use a combination of administrative data from CMF and follow-up phone surveys.

We first document that the majority of participants have biased beliefs about both the first and second moments of the interest rate distribution. While there is significant heterogeneity in beliefs, the vast majority underestimated both the interest rate they would get on the loan they took out, as well as the dispersion in rates. We measure whether participants underestimated the rate they would get by comparing actual interest rates obtained by participants after the study (according to administrative data) to their beliefs about the rate they would obtain. Nearly three-quarters (73.2%) thought they would obtain an interest rate lower than what they actually obtained. Their estimates of dispersion in the interest rates banks would offer them are also much lower than suggested by administrative data: specifically, 74.8% of participants underestimate dispersion in the interest

[3]Chile's national ID number, or *rol único tributario* (RUT), is commonly used in everyday life. For example, people are asked to give their national ID numbers when they check out at the grocery store. The data merge using national ID number was completed in accordance with the terms of the collaboration with CMF and with the study's IRB protocols.

rates banks would offer them compared to administrative data.[4]

We then test the effects of a price comparison tool designed to correct biased beliefs on (i) beliefs about interest rates, (ii) search behavior, (iii) negotiation, and (iv) loan terms, including whether they were offered a loan, the terms of the offer, and whether they took out the loan. The price comparison tool was built using administrative data on loan and borrower characteristics for the universe of consumer loans in Chile, i.e., data from over 1.8 million loans to approximately 1.2 million borrowers over two years. The tool shows treated participants the conditional distribution of interest rates that similar borrowers obtained for similar loans over the previous six months.

Immediately after seeing the price comparison tool, simple tool, or control video, participants assigned to the elicit beliefs treatment were again asked their beliefs about the distribution of interest rates. When treated with the price comparison tool, participants update and report expecting to receive a 16.2 pp *higher* interest rate on the loan they obtain, or a 54.9% increase compared to the control mean posterior belief of a 29.2% expected annual interest rate. The price comparison tool also led participants to increase their expectation of how much price dispersion they face in the market by 15.9 pp, or 68% relative to the control mean posterior of 23.2 pp dispersion in annual interest rates.[5]

To measure effects on search behavior, negotiation, and loan terms, we combine the administrative data on originated loans with rich data on participants' search histories that we collected in a follow-up phone survey conducted with a subset of 6,441 participants. The price comparison tool did not affect the number of institutions at which participants searched for information (which includes "soft search" such as visiting a bank website or branch to get a sense of the interest rate the bank would offer without formally applying), the number of institutions at which they applied for a loan, nor the specific institutions at which they searched. While the lack of a treatment effect on search could be due to offsetting effects of updating inaccurate beliefs about the first and second

[4]Following the macroeconomic uncertainty literature (Coibion, Georgarakos, Gorodnichenko, Kenny, and Weber, 2024), our preferred measure is the range between the highest rate a bank would offer them and the lowest rate a bank would offer them, due to its simplicity. This measure performed better in piloting than more complicated measures of the distribution. Consistent with our piloting, in the inflations expectations literature eliciting a more detailed distribution leads to higher survey dropout (Weber, D'Acunto, Gorodnichenko, and Coibion, 2022), which is a particular concern in our setting given that our participants take the survey online and are not professional survey respondents unlike in some other studies. Following Coibion, Georgarakos, Gorodnichenko, Kenny, and Weber (2024) we also ask participants what percent of loan offers they think are above the midpoint of the distribution (i.e., the midpoint between the lowest and highest rates they reported to us) to capture potential asymmetry in participants' perceived distribution.

[5]We avoid using the term "prior" to refer to their belief prior to seeing the tool, as some may have already updated their beliefs—for example by visiting a bank's website, seeing a bank ad, or getting an offer from a bank—before participating in our study. To distinguish the belief elicited before seeing the price comparison tool, simple tool, or control video and the belief elicited afterwards, we refer to the latter as a "posterior belief." Although we winsorize responses to these interest rate questions at the 95th percentile (as prespecified), the distributions are nevertheless quite skewed and the medians are substantially lower than the means. The control median posterior of the annual rate people expect to get is 18%, while the control median posterior of dispersion is 10.7 pp.

moments of the distribution, even among subsamples that underestimated the second moment and did not underestimate the first moment we do not find a statistically significant impact on search. Instead, as highlighted by our model, the lack of an effect on search is likely due to the possibility of negotiating, which breaks the sufficient statistic relationship between a consumer's reservation rate and search.

Despite not affecting the number of institutions at which participants searched or applied, the price comparison tool led them to obtain 13.1% more offers and 11.9% lower interest rate offers (measured in the follow-up survey), and to be 4.7% more likely to take out a loan (measured in the administrative data). Increased negotiation explains how borrowers obtained more loan offers and more favorable terms: the price comparison tool increases the probability of negotiating by 39%. In a subsequent survey we conducted to gather more data on negotiation, we find that much negotiation occurs prior to receiving a formal offer, which explains why negotiating also affects the probability of receiving an offer: in many cases the lender first informally offers an interest rate, then only if terms are agreed on does the lender issue a formal offer.

We next test two key predictions of our model. In the model, lenders make initial offers that are conditional on borrower characteristics and their desired loan characteristics, without observing the consumer's beliefs. The lender then receives a signal about the consumer's beliefs during the negotiation. The lender can then issue a new take-it-or-leave-it offer at a lower price, but must incur a cost to do so, and thus only does so if the consumer's beliefs indicate that they are likely to accept the new offer.

Assuming that seeing the price comparison tool leads borrowers to partially update their beliefs, our model has two key predictions. First, the price comparison tool should cause consumers who underestimated dispersion to negotiate more (but not necessarily search more). Second, the effect of the price comparison tool on negotiation should be non-monotonic in biased beliefs about the first moment of the interest rate distribution. Intuitively, consumers who overestimated the first moment will not benefit from updating their beliefs closer to the truth, as the lender will assume based on the consumer's beliefs that the consumer will accept the initial offer. Consumers who vastly underestimated the first moment also will not benefit, as their partially updated beliefs will still be too far below the true distribution, and thus the lender will infer that it cannot profitably lend to them, and will not incur the cost of negotiating. Meanwhile, consumers who somewhat underestimated the first moment will successfully negotiate more, as their updated beliefs make it such that the lender can still profitably lend to them at a lower, negotiated rate, so the threat of walking away leads the lender to lower the rate.

Our results align with these two predictions. First, the effect of the price comparison tool on negotiation is concentrated among those who underestimated dispersion, increasing their probability of negotiating by 78% (statistically significant at the 1% level). Meanwhile, the price comparison

tool did not have an effect on negotiation for those who did not underestimate dispersion. Second, the treatment effect on negotiation is non-monotonic in how biased beliefs are about the first moment of the distribution. There is no treatment effect for those who overestimated the first moment, nor for those who vastly underestimated it, while the effect on negotiation is concentrated among those who somewhat underestimated the first moment.

Finally, cross-randomizing whether we elicited beliefs about the interest rate distribution led participants to search at 0.13 (3.9%) more institutions and to obtain 9.6% lower interest rates than participants who were not asked their beliefs. Unlike the price comparison tool treatment, eliciting beliefs does not have an effect on negotiation or the number of offers received: instead, participants asked their beliefs obtained lower rates by searching more.

Our paper makes three main contributions. First, we contribute to the literature documenting high within-borrower price dispersion in consumer credit markets, the resulting high interest rate costs incurred by borrowers who do not search much, and the constraints to search. High price dispersion faced by the same borrower (or, in some studies, observationally similar borrowers) is a hallmark of consumer credit markets around the world (Zinman, 2015), including markets for credit cards (Stango and Zinman, 2016; Ponce, Seira, and Zamarripa, 2017), consumer loans (Karlan and Zinman, 2019; Cuesta and Sepúlveda, 2021), auto loans (Argyle, Nadauld, and Palmer, 2023), and mortgages (Allen, Clark, and Houde, 2014; Gurun, Matvos, and Seru, 2016; Guiso, Pozzi, Tsoy, Gambacorta, and Mistrulli, 2022; Coen, Kashyap, and Rostom, 2024). Due to this price dispersion, not searching much leads borrowers to pay substantially higher interest costs (Argyle, Nadauld, and Palmer, 2023; Bhutta, Fuster, and Hizmo, 2024).

Existing studies have focused primarily on the role of search *costs* in preventing search, while assuming—due to data limitations—that consumers have correct beliefs about the distribution of prices from which they are drawing, and thus know the *benefits* of search. These search costs include physical search costs across branches (Allen, Clark, and Houde, 2013; Argyle, Nadauld, and Palmer, 2023), high rejection rates (Agarwal, Grigsby, Hortaçsu, Matvos, Seru, and Yao, 2024), and the cognitive effort of comparing offers (Galenianos and Gavazza, 2022)—especially since financial products are often complex (Célérier and Vallée, 2017; Kulkarni, Truffa, and Iberti, 2025) and can include shrouded costs (Campbell, Jackson, Madrian, and Tufano, 2011; Stango and Zinman, 2014; Ferman, 2016; Alan, Cemalcilar, Karlan, and Zinman, 2018). We show that individuals have biased beliefs about the interest rate distribution they face and test the effects of a price comparison tool designed to correct biased beliefs.[6] While we do not find effects of

[6]In other contexts beyond consumer financial markets, Arteaga, Kapor, Neilson, and Zimmerman (2022) and Agte, Allende, Kapor, Neilson, and Ochoa (2024) study how beliefs affect search in the education market, Bandiera, Bassi, Burgess, Rasul, Sulaiman, and Vitali (2025) study this in the labor market, Jäger, Roth, Roussille, and Schoefer (2024) study how beliefs affect job search and negotiation intentions, and Caldwell, Haegele, and Heining (forthcoming) show that workers signal their beliefs to employers by providing their salary expectations before they receive a job offer.

correcting biased beliefs on search, we do find effects on negotiation and interest rates obtained, and show that this can be rationalized when negotiation is added to a model of sequential search with biased beliefs.

Second, we contribute to the literature on the importance of negotiation in markets with price dispersion. Differing negotiation leverage across consumers can cause persistent interest rate dispersion in equilibrium (Allen, Clark, and Houde, 2014). Theoretical models often posit that fixed costs constrain negotiation (Rubinstein, 1982), and there is empirical evidence that this is true (Backus, Blake, Larsen, and Tadelis, 2020). Allen, Clark, and Houde (2013, 2019) and Allen and Li (2025) model negotiation and competition in mortgage markets, and assume that negotiation requires searching for additional quotes.[7] However, negotiating may merely require accurate information on the price distribution (Grennan and Swanson, 2020). We introduce a novel way that beliefs can affect negotiation by allowing for biased beliefs: in our model, the lender learns about the consumer's beliefs about the interest rate distribution during the negotiation. Thus, correcting biased beliefs can affect the probability that the consumer negotiates successfully and receives a lower interest rate. Consistent with this, we find that the price comparison tool causes borrowers to be more likely to negotiate and to obtain lower interest rates.

Third, we contribute to the literature on the effects of beliefs in financial markets. At a macroeconomic level, biased expectations drive credit cycles (Bordalo, Gennaioli, and Shleifer, 2018; Bordalo, Gennaioli, Shleifer, and Terry, forthcoming). At a microeconomic level, the effects of beliefs on financial decision-making has been studied on both the assets and liabilities sides of the household balance sheet, as well as within firms. On the household assets side, individuals who have experienced low stock market returns are more pessimistic about future stock market returns and less likely to participate in the stock market (Malmendier and Nagel, 2011). Increased expectations about house price growth cause increases in real estate investments (Armona, Fuster, and Zafar, 2019; Liu and Palmer, 2023). On the household liabilities side, experiencing inflation leads households to expect more inflation in the future and to borrow more using fixed-rate mortgages (Malmendier and Nagel, 2016). Correcting biased beliefs about the dollar costs of payday loans reduces loan demand among first-time borrowers (Bertrand and Morse, 2011), while more experienced payday borrowers have more accurate beliefs but are willing to pay to avoid repeat borrowing (Allcott, Kim, Taubinsky, and Zinman, 2022). Within firms, biased CFO expectations about earnings growth affect corporate investment (Gennaioli, Ma, and Shleifer, 2016). We show how experimentally-induced changes in beliefs about the distribution of interest rates affect search, negotiation, and loan terms in the market for consumer loans.

[7] Beyond mortgage markets, Argyle, Nadauld, Palmer, and Pratt (2021) show evidence consistent with the importance of consumers obtaining lower interest rates by negotiating in auto loan markets.

2 Model

We present a simple theoretical framework to illustrate how biased beliefs affect search and negotiation, and to generate empirical predictions that motivate the design of our RCT and that we will take to the data.

2.1 Model Assumptions

Search A consumer i searches for loans sequentially from a continuum of banks, paying a search cost $s_i \sim F(\cdot)$ to search at each additional bank. When a consumer meets a bank j, the bank can provide a loan at cost $c_{ij} \sim G(\cdot)$ if the consumer and bank agree to originate the loan. Let $H_i(\cdot)$ denote the distribution of interest rates that consumer i could obtain across banks, which is conditional on the consumer's characteristics and the characteristics of the loan they are looking for. Let $\hat{H}_i(\cdot)$ denote the consumer's *perceived* distribution of interest rates that banks would offer them, based on their potentially biased beliefs. The consumers' type, denoted by θ_i, describes how much their perceived distribution of interest rates $\hat{H}_i(\cdot)$ differs from the actual distribution $H_i(\cdot)$.[8]

Negotiation The interest rate is decided in a two-stage bargaining game. The bank begins by posting a loan rate r_{ij}^1, which can be conditional on observable characteristics of the consumer (and the characteristics of the loan they are looking for), but not on the consumer type θ_i as this is not yet observed by the bank. If the consumer is unsatisfied with the first-round offer, they initiate a negotiation with the bank. During this negotiation, the bank learns about the consumer's beliefs about the interest rate distribution (i.e., the consumer's type θ_i), but the consumer's search cost s_i and the bank's marginal cost c_{ij} remain private information even during the negotiation. The intuition is that if for example $r_{ij}^1 = 26\%$, the bank learns about the consumer's beliefs about rates based on how they counter (e.g., a consumer who responds "how about 5% instead?" has very different beliefs from a consumer who responds "I know other banks would offer me 20%.")

After the negotiation, the bank decides whether to lower the loan rate. If the bank proposes a second-round rate $r_{ij}^2 < r_{ij}^1$ to the consumer, the bank incurs a fixed cost $d_j \sim K(\cdot)$, e.g., for the additional paperwork. In contrast, we model negotiation as costless to consumers; otherwise, the bank would be able to infer something about the consumer's search cost based on whether the consumer negotiates—rather than the search cost remaining private information—which would likely make the model intractable. The consumer then decides to accept or reject the second-round offer. After rejecting the second-round offer, the consumer will have to continue searching. Otherwise, when the consumer and the bank make a deal, the bank will provide the consumer with

[8] For simplicity, we assume that consumers are naïve in the sense that they do not realize that their beliefs are biased, nor do they update their beliefs based on the draws they receive while searching.

a loan at the first- or second-round rate r_{ij}, depending on at which round the consumer accepts the offer, at the predetermined cost c_{ij}.

Treatment Our treatment shows consumers the true distribution of interest rates $H_i(r)$ that the borrower faces, conditional on consumer characteristics and the characteristics of their desired loan. We assume that treatment leads to a *partial* correction of biased beliefs. Specifically, after seeing the true distribution $H_i(r)$, the consumer's type changes from θ_i to $\lambda_i\theta_i$ for some $\lambda_i \in (0,1)$.

2.2 Model Predictions

We characterize the equilibrium in Appendix A, and now turn to the predictions from the model that we will take to the data. The proofs of the predictions are also in Appendix A. We drop the i and j subscripts for notational convenience.

In a standard sequential search model where we add biased beliefs but do not add negotiation, we show that the consumer's reservation rate (determined by their search cost and beliefs about the interest rate distribution they face) is a sufficient statistic for the number of institutions at which the consumer searches and the rate they ultimately obtain.

Prediction 1. *In a sequential search model with biased beliefs but without negotiation, a consumer's equilibrium search result is completely characterized by the consumer's reservation rate.*

Specifically, if a consumer underestimated the first moment of the distribution (as we find most consumers do), the price comparison tool will increase their reservation rate, thus decreasing the number of institutions searched and increasing the interest rate they obtain. If, on the other hand, the consumer underestimated the second moment (as we also find most consumers do), the price comparison tool will reduce their reservation rate, thus increasing the number of institutions searched and lowering the interest rate they obtain.

After negotiation is introduced into the model, however, this sufficient statistic property of the reservation rate breaks down, and the number of institutions searched might not change as a result of changes in beliefs.

Prediction 2. *In a sequential search model with biased beliefs* ***and negotiation****, a consumer's equilibrium search result is* ***not*** *completely characterized by the consumer's reservation rate.*

Thus, we should no longer necessarily expect an effect of correcting consumers' beliefs on search.

We next turn to the predicted treatment effects of our price comparison tool, which depend on the consumer's biased beliefs. For expositional convenience and to take these predictions to the data, we characterize the consumer's type θ as their bias about the first and second moment of the interest rate distribution, $\theta = (\Delta\mu, \Delta\sigma^2)$.

Prediction 3. *If a consumer is biased about the first moment of the interest rate distribution, i.e., $\theta = (\Delta\mu, 0)$ for some $\Delta\mu \neq 0$, then the treatment effects of the price comparison tool on negotiation and the interest rate obtained depend on the direction and magnitude of the bias. In particular:*

1. *There exist $L_1, L_2 > 0$ such that when $\Delta\mu < -L_1$ or $\Delta\mu > L_2$, there is no treatment effect on the probability of a successful negotiation or the interest rate obtained.*

2. *For $\Delta\mu$ in some subset of $(-L_1, L_2)$, treatment increases the probability of a successful negotiation and reduces the interest rate obtained.*

Prediction 4. *If a consumer is biased about the second moment of the distribution, and in particular expects a lower variance than the actual distribution, i.e, $\theta = (0, \Delta\sigma^2)$ for some $\Delta\sigma^2 < 0$, then the price comparison tool increases the consumer's probability of a successful negotiation with the bank and reduces the interest rate they obtain.*

3 Institutional Context

3.1 Chilean Consumer Loans

Consumer loans are a popular credit product offered by banks in Chile: 43% of Chilean households have outstanding consumer credit, with consumer loans and credit cards being equally popular forms of obtaining credit from banks (Banco Central de Chile, 2021). Consumer loans are typically uncollateralized, have fixed interest rates, and are paid in equal monthly installments up until the loan matures.

According to administrative data from the CMF on the universe of consumer loans obtained between November 2021 and February 2024 (N = 1,863,087 consumer loans), the mean and median annual interest rates are 25.8% and 23.9%, respectively. The median loan amount is $4,488 USD, and the median maturity is 3 years. Based on our survey data, consumer loans are most commonly used to pay down other higher-interest debt (23.7% of borrowers), purchase or repair a car (16.3%), invest in their business (10.7%), make home improvements (5.2%), and purchase consumer durables (4.1%).

Unlike in the US and many other countries, Chilean credit bureaus do not report continuous credit scores; rather, they report binary flags of whether people have defaulted on prior loans. In 2012, the government passed legislation requiring a one-off deletion of information on default in response to the financial shock that many households experienced due to a large earthquake; Liberman, Neilson, Opazo, and Zimmerman (2018) study the effects of this policy.

3.2 Regulation in the Consumer Credit Market

Chile has a number of regulatory conditions that must be fulfilled when consumers are offered a loan. In 2011, the Chilean parliament defined a new credit term, *carga anual equivalente* (CAE), which functions as the Chilean equivalent of the annual percentage rate (APR) and must include fees. By law, both the CAE and the interest rate include the costs of all services inherent to loan operations. In addition, the CAE must include any additional costs of the loan, such as any insurance included with the loan (e.g., insurance that will pay off the outstanding balance if the borrower becomes unemployed or incapacitated). As part of the same legislation, borrowers have to be shown a universal credit contract (potentially in addition to another loan being offered) that represents a standardized "plain vanilla" contract and does not include insurance or various other types of fees that are sometimes included in consumer loan offers.[9]

In 2012, a new law mandated that formal loan offers must be made through a standardized disclosure sheet in which the CAE is prominently displayed in large bold numbers in the upper right-hand corner. Additional fees included in the CAE were also itemized in the disclosure sheet. Kulkarni, Truffa, and Iberti (2025) study the impacts of both the 2011 and 2012 regulations.

Finally, in 2013, the Chilean government lowered interest rate caps on consumer loans. Price ceilings have been in place for consumer loans since 1981, but the 2013 law substantially lowered this cap. The maximum interest rate on consumer loans is conditional on the loan terms, and is defined as 1.5 times the "current interest rate," where the current interest rate is calculated as a volume-weighted average of interest rates on originated consumer loans (conditional on loan characteristics). This law also expanded the interest rate caps to not only consumer loans but also other financial products such as credit cards. Cuesta and Sepúlveda (2021) study the effects of this law on both access to credit and interest rates.

3.3 Search for Consumer Loans

Chileans search for loans a number of ways: 92.3% of our follow-up survey respondents visited at least one bank website during their search, 37.3% used a mobile banking app, 33.9% visited a branch in person, 32.9% communicated with a bank by email, and 26.5% communicated with a bank by phone. Soft search (i.e., searching for information without formally applying for a loan) plays an important role: while control participants formally applied to 1.3 institutions on average, they searched across 3.4 institutions.

In our baseline and follow-up surveys, we asked participants what features of a loan were most important to them to better understand search behavior. Figure A.1a shows that in our baseline

[9]This legislation applied to all consumer loans with loan amounts below approximately $40,000 USD. Thus, it applied to nearly all consumer loans.

survey (when participants were looking for a loan), the three most important features of the loan were all functions of the interest rate: 25.1% of participants reported that the total loan cost was the most important feature, 21.7% reported monthly payment, and 19.9% reported the interest rate or APR. Similarly, in our follow-up survey, the most common reason for choosing a particular lender was a lower interest rate, with 44.4% of participants giving this answer (Figure A.1b).

Another potentially important feature is the probability of being approved for a loan, as consumer credit markets feature high rates of rejection and consumers need to "search for approval" (Agarwal, Grigsby, Hortaçsu, Matvos, Seru, and Yao, 2024). In our context, approval rates conditional on formally applying for a loan are 51.1%. Furthermore, 48.7% of survey respondents reported that the bank gave them some indication of whether their application would be approved before or without formally applying. In our baseline survey, 15% of borrowers named getting approved for the loan as the most important feature of the loan for which they were searching. In our follow-up survey, 28.2% of borrowers chose a particular lender because they were quickly approved by that institution, while 19.2% did so because that was the only offer they received.[10]

Less important features that participants reported during the baseline survey when they were searching for a loan included whether the bank branch was nearby (the most important feature for 9.4% of participants) and whether it is a bank in which they already had an account (7.8%). In the follow-up survey, when choosing a loan the less important features included whether the loan payment could be automatically deducted from payroll (5.9%), whether they were a client of that bank (5.7%), trust in the institution (4.5%), and getting approved for a higher loan amount (3.4%).

We also asked participants what strategy they employed while searching for their loan to better understand the prevalence of sequential vs. simultaneous search (De Los Santos, Hortaçsu, and Wildenbeest, 2012) and of searching for approval (Agarwal, Grigsby, Hortaçsu, Matvos, Seru, and Yao, 2024). We find that both sequential and simultaneous search are common (Figure B.2): 60.3% of participants reported having a target interest rate (consistent with sequential search), while 42% said they planned to search at a target number of banks or until receiving a target number of offers (consistent with simultaneous search). Searching for approval was also common, with 69.4% of participants reporting that they planned to stop searching after they were approved by one institution. These survey questions were not mutually exclusive, and based on the responses it appears that loan seekers implement a combination of strategies.

3.4 Online Tools

Because our intervention is an online tool that provides information about the distribution of interest rates a borrower faces conditional on their characteristics and the characteristics of the loan

[10]Being quickly approved and only receiving one offer were the second- and third-most common reasons for choosing a particular lender. Reasons were not mutually exclusive, as participants could name more than one.

they are looking for, we briefly describe other online tools available in the Chilean consumer credit market. We describe two types of tools: (i) tools provided by particular banks on their websites and (ii) third-party comparison platforms. We scraped data from as many of these websites as possible—conditional on the loan and borrower characteristics of our RCT participants—in order to quantify how accurate the information on these websites is compared to the loans participants actually received in the administrative data. We present results from this exercise in Section 6 and describe the details of our procedure in Appendix C. In short, neither bank websites nor third-party comparison websites provide accurate information.

Bank Websites Prospective borrowers can get interest rate quotes from bank websites, usually through online tools provided by the bank that are known in Chile as "simulators." Nearly all (93.2%) of our participants used at least one bank simulator while looking for a loan.

We identified twelve banks that have consumer loan simulators on their websites. The simulators ask for a range of inputs (Table B.1, panel A). The most common inputs requested by these tools are loan amount and maturity (requested by all banks). All but one bank request the user's national ID number, but we show in Appendix C that the interest rate numbers shown do not vary based on the ID number the user enters. Five out of twelve banks ask for the user's income, and none ask for the user's neighborhood, or *comuna*, despite this being an important predictor of interest rates used by banks in their algorithms.

Third-Party Comparison Websites There are two main third-party comparison websites for consumer loans, but only 12.2% of participants reported using such a tool when searching for a loan. One is provided by a private company and the other is run by a different government agency. Table B.1, panel B, describes the inputs required by these two comparison websites. In both of these tools, consumers input their desired loan size and maturity and receive quotes for loans from different institutions. However, neither asks for any borrower characteristics, and thus the interest rate quotes they provide are not conditional on borrower characteristics.

4 Experimental Design

4.1 Participant Recruitment

Figure 1 shows the design of the RCT and the funnel of participant recruitment. We recruited 112,063 participants to the RCT from November 2021 to June 2023 (see Appendix D for a detailed timeline of the experiment). We targeted Google ads from the CMF to people who searched for keywords related to consumer loans in Chile. Our Google ads campaign included 4,107,376 ads

served from November 2021 to June 2023, and 18.5% of people searching for keywords related to consumer loans in Chile were served our ad. Figure B.3 shows an example of one of the Google ads included in our campaign. Those who clicked on the ads were taken to a landing page with a description of our study and informed consent to participate. The following page asked for their national ID number—which is commonly given out in Chile (e.g., for rewards programs at the grocery store)—and their contact information including email address and phone number. We then conducted a baseline survey prior to showing the price comparison tool, simple tool, or a control video to the participant. Immediately after seeing the treatment, the participant was asked additional survey questions.

The ads we served were clicked 612,945 times, i.e., 14.9% of ads were clicked. From these clicks, 112,063 (18.3%) consented to participate and continued taking the baseline survey long enough to randomly be asked or not asked the questions on their expectations about the distribution of interest rates banks would offer them and how much they would search; this is our sample for measuring the effects of eliciting beliefs. Many consumers abandoned the survey during this module: 46,051 consumer loan seekers (41.1% of those who consented) continued taking the baseline survey long enough to reach the module where we randomized whether they saw the price comparison tool, simple tool, or control video; this is our sample for measuring the impact of the price comparison and simple tools.[11]

4.2 Elicit Beliefs Treatment

After obtaining their national ID number and contact information, participants completed modules on sociodemographic characteristics and other financial products that they currently have or loans they had in the past. We then randomly assigned 75% of participants to be asked questions about their expectations of (i) the lowest interest rate a bank could offer them, (ii) the highest interest rate a bank could offer them, (iii) the fraction of offers that would have an interest rate above the midpoint between the lowest and highest rates, (iv) the rate they expected the first bank where they searched to offer, (v) the rate they expect the second bank where they searched to offer, and (vi) the rate they expected to get on the loan they ultimately took out. The first three of these are borrowed

[11] Despite the smaller sample size of 46,051 for measuring the effects of the price comparison tool and simple tool, compared to the sample size of 112,063 for measuring the effect of the elicit beliefs treatment, the research design is internally valid. We do not randomize participants into one of the price comparison tool, simple tool, or control arms until they reach that module of the online survey. As a result, we can simply remove those who do not make it to the tool treatment module from the sample for estimating the effect of the tools, and still have balance across these treatment arms (both in theory and, as we show, in practice). Because we use a cross-randomized design, the results on the effects of the tools are a weighted average of the effect for the subsample who also received the elicit beliefs treatment and the subsample who did not receive the elicit beliefs treatment, relative to a control group in which the same proportions did and did not receive the elicit beliefs treatment; thus, if there is an interaction effect between the tool treatments and the elicit beliefs treatment, it would be reflected in the effect of the tool treatments that we estimate (Muralidharan, Romero, and Wüthrich, 2023).

from the macroeconomic expectations literature (Coibion, Georgarakos, Gorodnichenko, Kenny, and Weber, 2024). In addition, we asked them at how many banks they would search, at which bank they would search first, and at which bank they would search second. We did not ask any of these expectations questions to a randomly selected 25% of the sample in order to test whether these survey questions have a treatment effect, motivated by evidence from Zwane, Zinman, Van Dusen, Pariente, Null, Miguel, Kremer, Karlan, Hornbeck, Giné, Duflo, Devoto, Crepon, and Banerjee (2011) and Stango and Zinman (2014), and indeed we find that asking these questions led people to search more and obtain loans with lower interest rates. After viewing either the price comparison tool, simple tool, or a control video, we again asked the 75% of participants assigned to the elicit beliefs treatment the same interest rate expectation and search questions to test whether their expectations were affected by treatment.

4.3 Price Comparison Tool and Simple Tool Treatments

Price Comparison Tool Our price comparison tool (Figure 2a) showed participants a conditional distribution of interest rates that similar borrowers had received for similar loans over the past six months. We built the tool using administrative data on loan characteristics merged with borrower characteristics for the universe of consumer loans in Chile, i.e., data from over 1.8 million loans to 1.2 million borrowers over two years. Appendix E provides more details on how the histograms were created. We refreshed the data every month to show the previous six months of interest rate data, based on tests we conducted to determine the optimal time period of data to show (where we traded off showing accurate information vs. having sufficient data points underlying the histogram shown to each participant). Appendix E.1 provides more detail on this trade-off and describes the rationale behind why we showed participants data on loans from the last six months.

Walking through each component of the price comparison tool shown in Figure 2a, the borrower and loan characteristics that the participant already answered in the baseline survey were loaded automatically in the top panel of the tool ("1. Verify that your data are correct"), but these values could be modified by the user. The second panel of the tool ("2. Look at the information") showed the user the distribution of interest rates that similar borrowers had obtained for similar loans in the past six months. We conducted focus groups to test a prototype of the tool. Based on the findings from these focus groups, in order to make the histogram understandable to consumers that may not be familiar with interpreting data from graphs and histograms, participants could hover over the histogram's bars to see a tool-tip that explained what that bar indicated. Specifically, the tool-tip told the participant the number of loans that had that interest rate, gave a cumulative distribution function interpretation of the bar (what percent of loans had an interest rate at or below that rate), and converted the interest rate to a monthly and total loan cost in Chilean

pesos based on the loan amount and maturity entered by the participant. In addition, we created a tutorial video that the user could watch to better understand how to use the tool.

In the third panel of the comparison tool ("3. Compare the impact of different interest rates for your wallet"), we compared two interest rates in the histogram to show the participant the implications of these different rates for their monthly and total loan costs (in Chilean pesos). The inclusion of this part of the price comparison tool was inspired by research that it is important to translate differences in APRs into dollar costs (Bertrand and Morse, 2011), that borrowers target monthly payments rather than interest rates (Argyle, Nadauld, and Palmer, 2020), and more broadly that market participants are more perceptive to dollars rather than percentages (Shue and Townsend, 2021). By default, the highest and lowest interest rates were compared, but the user could drag the two triangle markers on the x-axis of the histogram in order to change which interest rates were compared. Alternatively, participants could manually enter interest rate values to see how they would translate into costs. Participants experimenting with this feature should see the concrete consequences of the market's price dispersion, i.e., how they may pay substantially different costs for their loan depending on the interest rate they obtain.

One concern is that the dispersion shown in the tool could reflect both within-borrower dispersion and across-borrower dispersion based on characteristics that are unobservable to us but observable to lenders. To test this, we compare the following for each consumer who reported receiving more than one offer in our follow-up survey. We first calculate the within-borrower standard deviation of interest rates in the survey, which we know reflects within-borrower dispersion since these rates are from offers received by the same borrower. We then conduct 1,000 simulations where we randomly draw an equal number of interest rates per simulation as the number of offers the consumer reported receiving—we draw from the consumer's conditional distribution in the administrative data (i.e., the same distribution we show them in the price comparison tool if they are in that treatment arm). We then calculate the average—across the 1,000 simulations—of the standard deviation in interest rate offers "obtained" by that borrower from their conditional distribution. If the conditional distribution in administrative data includes substantial across-borrower dispersion, the simulated standard deviation should be higher than the within-borrower standard deviation reported in the survey. In contrast, the standard deviations are almost identical, with the standard deviation in administrative data 0.01 pp *lower*. We conclude that the dispersion shown in the price comparison tool is an accurate representation of within-borrower dispersion.

Simple Tool on Benefits of Search Participants in the simple tool treatment arm viewed a simpler tool that provided the user with just two numbers on the estimated benefits of search (Figure 2b). This treatment was designed to be simpler and avoid the information overload that might be present in the price comparison tool. The borrower and loan characteristics that the participant

already answered in the baseline survey were again loaded automatically in the top panel of the tool ("1. Verify that your data are correct"). The bottom panel ("2. Look at the information") told the following to the borrower: "Using real data from loans granted to people similar to you, we estimate that shopping at 1 additional bank would lower your monthly payment by $X and the total cost of your loan by $Y, on average." The number of additional banks could be modified using a drop-down menu, which the participant could use to determine the expected benefits of searching at up to five additional banks (i.e., of searching at up to six total banks relative to searching at just one bank).

To estimate the amount they could save in Chilean pesos, we used the conditional distribution corresponding to that participant's characteristics and the characteristics of the loan they were searching for, and simulated consumer searches across one to six banks. We then averaged across these simulated searches to calculate how much the participant could expect to save on average. The "More details" link provided the participant with a description of how we calculated the expected savings. Appendix E.2 provides more detail on the calculation of search benefits.

Control Video The control video was a 1 minute and 35 second long animated video created by the CMF describing key credit terms. The video was designed to provide information related to loans that would *not* be useful for search. The video defined what a lender and debtor are, what a loan contract is and what is included in it, and key loan terms like maturity and principal. Figure B.4 shows a screenshot from our control video. In all treatment arms including the control video, the participants were required to stay on the treatment module page for one minute prior to clicking "Next" to proceed to the following module.

5 Data

In this section, we describe our administrative and survey data sets. Appendix F provides more detail on data cleaning, links to the survey questionnaires, and a timeline of the surveys.

5.1 Administrative Loan Data

We use administrative data on the universe of consumer loans from 2015–2024 from the Chilean financial regulator, the CMF. We observe the following borrower characteristics that banks use to determine whether to offer loans: age, marital status, gender, income, and neighborhood of residence.[12] Importantly, credit bureaus in Chile do not report continuous credit scores, but rather

[12]Note that if applicants already have other products at the bank where they are applying for a loan, the bank might also use that information in its lending decision, and we do not observe these bank-specific data.

report binary flags if the borrower has defaulted on prior loans; thus, the interest rates that banks offer are not conditional on continuous credit scores. As for loan characteristics, we see each loan's amount, interest rate, and maturity, as well an anonymized code for the lender. We use these data in our construction of the conditional distribution of interest rates for both the price comparison tool and the simple tool on the benefits of search.

By obtaining participants' national ID number, we are able to merge their treatment status and survey responses with future administrative data to measure treatment effects on the loans they obtain after participating. In total, 21,522 out of 112,063 participants from our RCT took out a consumer loan between the time they participated in our RCT and one year later (Figure B.5). Of these, 8,988 participants among the sample size of 46,051 participants for measuring the effect of the price comparison tool and simple tool took out a consumer loan within one year of participating.

We also use the administrative data to compare participants in our RCT who took out consumer loans to the universe of consumer loan borrowers in Chile. Figure B.6 shows that borrowers in our RCT are—unsurprisingly—not perfectly representative of the overall population of borrowers in Chile; nevertheless, there is a large amount of overlap in the distributions of characteristics of borrowers in our RCT and the overall population of borrowers in Chile.[13] The two groups are relatively similar on gender (all borrowers: 38% women vs. RCT sample: 37.4%), the percentage who live in the capital Metropolitan Region (all borrowers: 51.5% vs. RCT sample: 50.4%). First-time borrowers (defined as those who did not have a previous consumer loan in the administrative data prior to the RCT) make up 37.5% of the RCT sample compared to 41% of the overall population of borrowers. Borrowers in our RCT are relatively better-off than the overall population of borrowers—as the distribution of annual income for RCT borrowers is shifted right of that of all borrowers—though there is extensive overlap in the support of the distributions. The variable with the starkest differences between the RCT sample and the overall population of borrowers is age, where participants in our RCT are younger than the general population of borrowers (median age of all borrowers: 38 vs. RCT sample: 34). These differences are unsurprising considering the online nature of our recruitment process.

The distributions of loan terms (interest rate, loan amount, and maturity) obtained by our RCT participants and all borrowers in Chile also exhibit differences but have a large degree of overlap (Figure B.7). In general, borrowers in our RCT obtain slightly larger, longer-maturity, lower-interest rate loans. For example, the average loan maturity in our sample is 37 months as compared to 34 months in the overall population.

[13] We exclude borrowers who participated in our RCT from the "all borrowers" group (i.e., overall population of borrowers in Chile) in order to compare two mutually exclusive groups.

5.2 Baseline Survey

The baseline survey was conducted online after participants who searched for keywords related to loans clicked on a Google ad from the financial regulator and consented to participate in the study. In addition to the questions about beliefs for those assigned to the elicit beliefs treatment (described in Section 4.2), we asked participants about their sociodemographic characteristics and detailed questions regarding their existing banking relationships and other financial products.

We also asked participants questions to determine how they form beliefs about interest rates. Specifically, we asked them if they had ever obtained a quote for a consumer loan from a bank website, if they had seen an ad for a consumer loan advertising an interest rate, or if someone they know had told them what interest rate they got for a consumer loan. If they answered yes to any of these questions, we asked them how long ago this was and what interest rates were given by the bank website, advertisement, or person they know. We then asked whether they had searched for a loan before, and if so how long ago it was, how many offers they had received, and the range of interest rates of those offers. Finally, we asked questions on financial literacy, behavioral biases (e.g., financial procrastination), and a set of simple questions used to measure cognitive ability, which are all related to search and the formation of beliefs (D'Acunto, Hoang, Paloviita, and Weber, 2023).

Table 1 reports means for characteristics from the baseline survey, and tests for balance between the 75% of participants who were randomly assigned to the elicit beliefs treatment—i.e. were asked questions on their expectations about the distribution of interest rates and how much they would search—vs. the 25% of participants who were not asked these questions. Participants in our sample are roughly 36 years old on average with an average monthly income of 1,125,959 pesos (1,142 USD at market exchange rates). Participants have a wide range of education: 3.7% did not complete high school, 36% completed only high school, 21.3% completed a 2-year post-secondary program (equivalent to an associate's degree), and 39% completed a 5-year degree program or higher (equivalent to a bachelor's degree). As for financial experience, 67.8% of our participants had a bank account, and 70.2% had taken out a loan.

As expected due to randomization, our sample of 112,063 participants who were randomized to either receive or not receive the elicit beliefs questions is balanced: the p-value of the omnibus F-test regressing the elicit beliefs dummy on all baseline survey characteristics is 0.463 (Table 1). Furthermore, only one variable—the probability of having a loan already at baseline—is not balanced, as could be expected by chance: those assigned to the elicit beliefs treatment are 0.6 pp less likely to have a prior loan (significant at the 5% level).

Table 2 tests for balance across the price comparison tool, simple tool, and control arms.[14] The

[14]The loan characteristics variables included in Table 2 are not included in Table 1 because they are asked in the same module as the elicit beliefs treatment, and the elicit beliefs treatment caused some participants to stop participat-

sample size in this table, 46,051 participants, is smaller than that of Table 1 because of participant attrition between the module in which we randomized whether we elicited beliefs and the module in which we randomized assignment to one of the tool or control arms. We again find that the sample is balanced across treatment arms. The p-value for our omnibus F-test of whether characteristics jointly predict the price comparison tool treatment is 0.279, and that for the simple tool treatment is 0.207. The only variable that has a statistically significant difference between the treatment arms and the control arm is having a bank account: participants are 1.6 and 1.3 pp more likely to have a bank account in the price comparison tool and simple tool arms compared to the control group (statistically significant at the 5% level).

5.3 Endline Phone Survey

We surveyed participants via phone at least six months after they participated in the RCT. We attempted to contact 39,713 participants (35.4% of our 112,063 sample), and ultimately collected 6,441 completed surveys, for a 15.5% response rate. Table B.2 shows that response rates are balanced across both the elicit beliefs and tool treatments.

The primary objective of the endline phone survey was to collect rich data on participants' search histories. Search data are poorly captured in most administrative data sets: in many administrative data sets including the CMF data, only originated loans are recorded. Even if all applications were recorded, true search behavior also involves informal quote requests, or even inquiring about the probability of approval at a particular lender (Agarwal, Grigsby, Hortaçsu, Matvos, Seru, and Yao, 2024). For each bank at which consumers searched for information, we ask detailed questions about how they searched (e.g., using the bank's website or mobile banking app, going to a branch in person, emailing, calling by phone), whether they informally received any information about their probability of acceptance or an estimate of the interest rate they would receive, whether they formally applied, whether they were accepted or rejected, what loan terms they were formally offered if accepted, whether they negotiated the offer, the loan terms they were offered after this negotiation, and whether they took out the loan.

5.4 WhatsApp Survey

We collected WhatsApp surveys to ask participants more detailed questions about the timing of negotiation, given our findings that the price comparison tool increased negotiation. We attempted to survey 15,510 participants and received 2,344 responses, for a response rate of 15%.

ing in the survey. (Note that participants who abandoned the survey during the elicit beliefs module are still tracked in the administrative data and included in the sample to estimate the effect of eliciting beliefs.)

6 Results

6.1 Participants Underestimate Rate They Will Obtain and Dispersion

Comparing participants' expectations about the distribution of interest rates with administrative data, we find that prior to viewing the tool, most users thought interest rates were lower than they actually were, and also underestimated price dispersion.

Figure 3 compares the interest rates that participants report expecting to receive on the loan they take out to the interest rates we observe in administrative data for the loan they actually obtained subsequently (restricting to those who did take out a loan after participating in the RCT). It shows that participants have inaccurate beliefs on the interest rate that they will ultimately receive on their loan: 73.2% of borrowers think they will receive an interest rate *lower* than the rate they later receive. Conditional on underestimating, these borrowers think they will get a loan that is 14.3 pp lower than the rate they ultimately receive; the average bias including those who are accurate or overestimate is –6.9 pp.

Figure 4 shows the distribution of the difference between a participant's beliefs on dispersion—measured as the highest rate they think a bank would offer them minus the lowest rate they think a bank would offer them—and the corresponding measure in the administrative data (i.e., the highest minus lowest rates we observe in the conditional interest rate distribution, conditional on that participant's characteristics and the characteristics of the loan they are looking for). We use administrative data for similar borrowers and loans over the past six months—i.e., the same data we would show the borrower if assigned to the price comparison tool arm. Unlike in the case of beliefs about the loan they will obtain, for dispersion we cannot compare to offers they actually received, as these would only be a subset of draws from the full distribution of interest rates if they were to search across all banks. We find that the majority of participants (74.8%) *underestimate* price dispersion, but also that there is a long right tail of participants who substantially *overestimate* dispersion.

Why are beliefs biased? We first test whether biased beliefs about the first moment of the distribution may be due to changes in interest rates over time. We find that even though interest rates were indeed changing over time during the 19 months of our RCT, changes in interest rates do not explain biased beliefs. We then ask which consumer characteristics, if any, predict biased beliefs. Finally, we explore whether other sources of information to which people have access lead them to have biased beliefs. We find that various other sources of information indeed provide downward-biased estimates of interest rates, which may explain the bias we observe.

On average, interest rates increased over the 19 months that we conducted the experiment, with the median monthly interest rate among the universe of loans in administrative data ranging from 19.6% to 26.7% over this time period. If consumers are slow to update their beliefs, this

could explain why they underestimate the first moment of the interest rate distribution. To test whether consumers are slow to update, we regress beliefs about the rate a consumer will obtain on the monthly median interest rate during the month in which that consumer participated in the RCT. We find that when median consumer loan interest rates are 1 pp higher, beliefs about the rate people expect to obtain are 1.3 pp higher (Table B.3), indicating that people do update as interest rates change over time and that underestimates of the rate people will obtain are not caused by people being slow to update.

Alternatively, consumers might have accurate beliefs about the first moment of the distribution but take a while to obtain a loan. Since interest rates were increasing on average over the course of our experiment, we may then misattribute the difference between their belief and the rate they obtained to a bias about current interest rates, when it is instead due to underestimating future changes in interest rates and/or the amount of time before they would obtain a loan. We show, however, that even when we restrict the sample to the first quartile of the number of days between participating in the RCT and obtaining a loan—those who obtained a loan within 21 days of participating—the bias in the first moment looks very similar, with 74.9% underestimating rates (Figure B.8). Furthermore, while the median interest rates are rising over the course of our sample, consumers' interest rates expectations are persistently lower than the level of rates in the market but do reflect the upward trend in median rates (Figure B.9), visually confirming the regression results from Table B.3. We conclude that consumers' underestimates of the first moment of the interest rate distribution cannot be explained by changes to interest rates over time.

Another possibility is that we are overestimating dispersion in the administrative data if the conditional distribution combines within-borrower dispersion and across-borrower dispersion based on characteristics unobservable to us but observable to banks. However, we already showed in the exercise described in Section 4.3—comparing within-borrower dispersion reported in the survey by the same borrower across offers they received to simulations from their corresponding conditional distribution in the administrative data—that the dispersion we measure in administrative data is an accurate representation of within-borrower dispersion.

What predicts biased beliefs? We estimate a series of regularized regression models using elastic net (Friedman, Hastie, and Tibshirani, 2010) in Appendix H. Table H.1 presents the coefficient estimates from these models. Individuals with more-biased beliefs tend to be younger and have lower incomes, less education, and less experience with financial products (i.e., bank accounts and loans). Furthermore, those looking for a smaller loan amount also tend to have more-biased beliefs.

We next assess what sources of information people use to form their beliefs, as well as whether these sources provide accurate information. In our survey data, 41.2% of participants report having seen advertisements by banks, 43.8% used tools on bank websites known as "simulators" that provide an estimated interest rate to the consumer, 12.2% used third-party comparison websites,

and 22.8% asked friends and family about interest rates. To assess whether bank advertisements, bank websites, third-party comparison websites, or family and friends might cause people to have biased beliefs, we compare rates our participants *would* have seen in each of these contexts with the rate they ultimately received in our administrative data (Figure 5).

For advertisements, we cannot observe all bank advertisements but can observe those that banks place on Google. We randomly sample combinations of search terms that led people to the Google ads for our experiment and neighborhoods of participants in our experiment. We then conduct automated Google searches using that keyword and specifying the location from which the search originated, and scrape the resulting first page of Google results—including both Google ads and regular Google search results. More details are provided in Appendix G. The difference between interest rates that are shown in Google ads or search results and the rates people actually obtain are heavily negatively skewed: 74.2% of rates shown on Google were lower than what participants ultimately received from the same bank in our administrative data (Figure 5a).

For bank websites and third-party comparison websites, we use a script to input participants' characteristics and the characteristics of the loan they are looking for into the interest rate simulator on each bank's website and the main third-party comparison website and scrape the resulting loan terms. More details are provided in Appendix C. Bank simulators tend to show inaccurate rates, as the rate a bank website showed can be as much as 27 pp lower or 20 pp higher than the rate they ultimately receive from that bank. While there is substantial noise in the quotes from bank websites, they are not biased in one direction or the other: 50% of participants would have been shown an interest rate that is lower than the rate they ultimately received (Figure 5b). As for the third-party comparison website, the difference between the rate shown and the rate the same consumer obtained from the same bank is also negatively skewed, with 74.1% of quotes being lower than the rate the borrower ultimately received. These results suggest that banks have an incentive to provide attractive quotes to borrowers in a context where the borrower is still deciding which bank to apply for an offer from, but that they can subsequently bait-and-switch the customer and offer them a higher rate when providing a formal loan offer (Figure 5c).

Finally, only fifteen of our participants responded in the follow-up survey that they received information from friends and family, reported the interest rates that those friends and family told them, and also received a loan in the administrative data to compare. For this small sample of participants the difference in rates between what friends and family told them and what they ultimately received is also negatively skewed, with 68.8% being lower than the rate the borrower actually received (Figure B.10). This could be due to selection in which friends and family are willing to divulge the rate they obtained, depending on whether they think it is was a good or bad rate.

6.2 Price Comparison Tool Leads to Large Updates in Beliefs

After seeing the price comparison tool, did participants revise their expectations about the distribution of interest rates? To test this, we estimate the following specification:

$$Posterior_i - Prior_i = \beta_1 \mathbb{1}(\text{Simple Tool})_i + \beta_2 \mathbb{1}(\text{Price Comparison Tool})_i + \lambda_{b(i)} + \varepsilon_i, \quad (1)$$

where $Prior_i$ is the interest rate expectation participant i reported prior to seeing the tool or control video and $Posterior_i$ is the interest rate expectation they reported after seeing it. In our main specification, interest rates are annualized and measured in percentage points (e.g., an expected interest rate of 18% per year would be coded as 18). The treatment dummies $\mathbb{1}(\text{Simple Tool})_i$ and $\mathbb{1}(\text{Price Comparison Tool})_i$ equal one if the participant was assigned to that treatment arm and zero otherwise, and $\lambda_{b(i)}$ are bin density fixed effects. The bin density fixed effects are deciles of the number of observations in the tool that were shown or would have been shown to the participant, to control for the fact that people in higher-population neighborhoods or with more borrowers with similar characteristics would see more observations in the price comparison tool and might infer that there is more dispersion than those seeing fewer observations.[15]

Table 3 shows the results. On average, the price comparison tool caused participants to increase their beliefs about the rate they would obtain by 16.2 pp, or 54.9% relative to the control mean posterior of 29.2%.[16] Comparing posteriors to priors, treated participants' expectations about the entire distribution shift rightward. They update their expectation about the lowest interest rate a bank would offer them by 10.9 pp and their expectation about the highest interest rate a bank would offer them by 30.4 pp. Their expectations about dispersion also increase by 16 pp compared to a control mean posterior of 23.2 pp of dispersion, an increase of 68%. Tables B.4 and B.5 show that the same conclusions hold if we use levels of posteriors as the dependent variable, with or without controlling for priors on the right-hand side. Tables B.6 and B.7 show the same pattern when we log-transform beliefs about interest rates.

One concern is that the increased expectations about dispersion are due to scale effects around an increased first moment of the distribution, given that neither the standard deviation nor our preferred measure of dispersion are scale-invariant. To test whether the effects on treatment on expectations about dispersion are driven entirely by a scale effect, we create a normalized measure of dispersion where we divide the highest minus lowest rate that a bank would offer by the midpoint

[15] For those in the simple tool and control groups, the bin density fixed effect is based on how many observations are in the price comparison tool histogram that *would have been shown* to the participant had they been assigned to the price comparison tool arm.

[16] Although we winsorize responses to these interest rate questions at the 95th percentile, the distributions are nevertheless quite skewed and the medians are substantially lower than the means. The control median posterior of the annual rate people expect to get is 18%.

between the highest and lowest rate, which is a scale-invariant measure. Table B.8 estimates the results of the treatment on this normalized measure of dispersion. Even with this scale-invariant measure of dispersion, we find that our comparison tool treatment increases expectations about dispersion (statistically significant at the 1% level). We conclude that the effect of treatment on participants' expectations about dispersion is not just a scaling effect from increasing their expectations of the first moment of the distribution.

Appendix I.1 tests for heterogeneous treatment effects of the price comparison tool on interest rate beliefs using the machine-learning methodology proposed by Chernozhukov, Demirer, Duflo, and Fernández-Val (2025). We reject the null hypothesis of no heterogeneity in treatment effects of the tool for belief updating about the rate participants expect to obtain and the highest rate a bank would offer them, but we fail to reject the null hypothesis of no heterogeneity for the lowest rate a bank would offer them and dispersion (Table I.1). We find that the same characteristics that predict more-biased beliefs also predict a larger treatment effect of the tool on beliefs about both the expected rate and the highest rate. In particular, the treatment effect of the tool is larger for participants who are younger and have lower incomes, less education, and less experience with financial products (i.e., bank accounts and loans), as well as those looking for smaller and shorter-maturity loans (Table I.3).

In contrast, the simple tool quantifying the benefits of search but providing no direct information on the distribution of interest rates hardly affected priors about interest rates: the coefficient of the effect of the simple tool on expectations about the interest rate the participant will obtain is less than 1 pp, while the coefficient on dispersion is very close to 0 at 0.01 pp, and neither is statistically significant (Table 3).

6.3 Effects of Tool on Search, Negotiation, and Loan Terms

Based on the findings from the previous two subsections, we revisit predicted treatment effects from our model. There are three key insights from the model in terms of the effects of the price comparison tool (partially) correcting biased beliefs on treated participants' search and negotiation. First, negotiation breaks the sufficient statistic relationship between reservation rates and search, so we should not necessarily expect an effect of shifting beliefs on search (Predictions 1 and 2). Second, the effect of the tool on those with biased beliefs about the first moment is non-monotonic: it should have no effect on negotiation for those who overestimated or vastly underestimated the first moment, but should increase negotiation for those who somewhat underestimated the first moment (Prediction 3). Third, the tool should lead those who underestimated dispersion to negotiate more (Prediction 4).

We run the following regression:

$$y_i = \alpha + \beta_1 \mathbb{1}(\text{Simple Tool})_i + \beta_2 \mathbb{1}(\text{Price Comparison Tool})_i + \varepsilon_i, \tag{2}$$

where y_i are search, negotiation, and loan term outcomes for participant i. The treatment dummies $\mathbb{1}(\text{Simple Tool})_i$ and $\mathbb{1}(\text{Price Comparison Tool})_i$ equal one if the participant was assigned to that treatment arm and zero otherwise.

Table 4 shows the effect of our treatments on search behavior, negotiation, and loan outcomes. The price comparison tool did not lead people to search at more institutions or to formally apply for loans at more institutions. Furthermore, Figure B.11 plots cumulative distribution functions (CDFs) of the number of institutions searched by treatment arm and shows that the null average treatment effect on search is not masking offsetting effects in different parts of the distribution.

We find that despite not searching or applying at more institutions, the tool makes consumers 3.7 pp (39%) more likely to negotiate and to receive 0.069 (13.1%) more offers in the survey data. In a subsequent WhatsApp survey we conducted to further understand the negotiation mechanism, we find that much of this negotiation happens prior to the bank issuing a formal offer, which explains the effect on the number of formal offers received. Participants also receive 11.9% lower (post-negotiation) interest rate offers, are 3.6 pp (9.7%) more likely to take out a loan, and have 10.5% lower interest rates on the loans they take out according to the survey data.

In administrative data, the tool makes borrowers 0.9 pp (4.7%) more likely to take out a loan, which is a smaller treatment effect compared to the survey data but nevertheless statistically significant at the 5% level. According to our survey data, borrowers who did not take out a loan overwhelmingly did not make the purchase or investment for which they were seeking the loan.

We do not find a treatment effect on interest rates in administrative data. The lack of a treatment effect in the administrative data appears to be due to a selection effect: because the price comparison tool causes some people who would not have otherwise obtained a loan in the absence of treatment to obtain one, and because only consumers who obtain a loan can be included in the interest rate regression, this regression suffers from selection. In particular, if the "compliers" who only obtain a loan if they are treated are less-creditworthy borrowers who thus obtain higher interest rates than the "always-takers" who would take out a loan regardless of treatment status, this selection effect could offset a negative treatment effect of the tool on the always-takers. There are two pieces of evidence that, taken together, suggest that this indeed drives the lack of a treatment effect on interest rates in the administrative data and the discrepancy compared to the effect in survey data.

First, we use machine learning methods from Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) to attempt to isolate participants most likely to be always-takers and most likely to

be compliers, and we estimate treatment effects on interest rates for these two groups.[17] Details are included in Appendix J. We find that we are able to effectively isolate likely always-takers and compliers: for the decile of participants most likely to be always-takers we observe a precise zero treatment effect of the tool on the probability of taking out a loan, while we observe that the tool causes a 6.1 pp increase in the probability of taking out a loan for the decile of participants most likely to be compliers. We then estimate treatment effects on interest rates for these two groups, finding an estimated 5.6% decrease in interest rates for always-takers and a 3.4% increase in interest rates among those likely to be compliers (Table J.1).[18]

Second, in administrative data, the tool does not cause a selection effect—specifically, there is a precise null treatment effect on the probability of taking out a loan—for those who obtain a loan very shortly after (within one day of) participating in the experiment. This group likely already had a loan offer in hand when they participated in the RCT. Given the lack of a selection effect for this group, we expect and indeed find a negative treatment effect on interest rates for this group, and that treatment effect is –9.5%, similar in magnitude to the effect we find in the survey data (Figure 6). Finally, for this selection effect to explain the discrepancy between the treatment effects estimated in administrative and survey data, it would need to be the case that participants who obtained loans sooner after participating are overrepresented in the survey data, which we find is indeed the case (Figure B.12). This leads our estimated treatment effects in survey data to be more heavily weighted towards those who obtained a loan shortly after participating, who are always-takers whom the tool caused to obtain lower interest rates.

Appendix I.2 tests for heterogeneous treatment effects of the price comparison tool on search, negotiation, and loan terms, again using machine learning (Chernozhukov, Demirer, Duflo, and Fernández-Val, 2025). Across all of the tests on search and negotiation behavior and loan terms, we find no detectable heterogeneity in treatment effects (Table I.4). This suggests that the null effects on search are not masking positive effects on search for some individuals and negative effects for others. Furthermore, it suggests that treatment effects of the tool on negotiation, the probability of taking out a loan, and interest rates do not differ by the characteristics of the participants, but rather tend to be spread evenly across participants.

The simple tool that showed only the estimated interest cost savings from search did not have

[17] Medina and Pagel (2025) use a similar conceptual approach. In particular, they use machine learning to estimate predicted individual treatment effects of a nudge on saving, then isolates the group that is predicted to have the highest treatment effect. For this group, they then estimate treatment effects on a separate outcome—borrowing—to study whether experimentally-induced saving causes borrowing.

[18] As shown in Table J.1, the difference in treatment effects on interest rates between the likely compliers and likely always-takers is statistically significant at the 5% level. While we can reject that the treatment effects on interest rates for the two groups are the same, and while the point estimate for compliers is negative and that for always-takers is positive, the confidence interval for each group's point estimate does include zero, so we treat this evidence as suggestive.

effects on any search, negotiation, or loan outcomes (Table 4), which is not surprising given that the simple tool did not lead people to update their priors about the interest rate distribution.

Next, we turn to two key predictions from the model about heterogeneous effects of the tool on negotiation, based on consumers' biased beliefs about the first and second moments of the interest rate distribution that they face. The first prediction is that the effect of the tool on negotiation is non-monotonic in how biased the belief is about the first moment. In particular, there should be no effect of the tool for those who (prior to seeing the tool) overestimated the first moment of the distribution. There should also be no effect for those who vastly underestimated the first moment, as they will still underestimate it by enough after partially updating that the bank will determine that it likely cannot profitably lend to that consumer based on their beliefs, and thus will not incur the cost of lowering their initial offer. In contrast, for those who somewhat underestimated rates, the tool should make them more likely to negotiate by bringing their belief into the range where the bank will determine that it can profitably lend and thus will engage in negotiation. This is exactly what we find: those who overestimated or substantially (by more than 30 pp) underestimated rates do not negotiate more in response to the tool, while those who somewhat (by less than 30 pp) underestimated rates are more likely to negotiate (Figure 7).

The second prediction is that the tool will increase negotiation for consumers who underestimated the second moment of the distribution. We separately estimate specification (2) for those who underestimated the second moment and for all others in Table 5.[19] For those who underestimated dispersion, the tool led to a 7.9 pp (78.2%) increase in negotiation. It did not, however, lead to an increase in search, consistent with the insight from our model that introducing negotiation breaks the sufficient statistic relationship between reservation rates and search. For those who did not underestimate dispersion, on the other hand, there is no effect on negotiation, with a small and not statistically significant point estimate.

6.4 Eliciting Beliefs Leads to More Search and Lower Rates

Randomizing whether we elicited priors about the interest rate distribution and the number of institutions at which participants intended to search—rather than asking these questions to all participants—was motivated by evidence from Zwane, Zinman, Van Dusen, Pariente, Null, Miguel, Kremer, Karlan, Hornbeck, Giné, Duflo, Devoto, Crepon, and Banerjee (2011) and Stango and Zinman (2014) that survey questions can have treatment effects on real-world behavior and outcomes.

[19] We define underestimating dispersion as beliefs about dispersion prior to treatment being at least 1 pp lower than the observed dispersion in the administrative data. The breakdown of our participants' accuracy in beliefs about both rates and dispersion are presented in Figure B.13. The majority of participants from whom we elicited interest rate beliefs underestimate both the levels of rates and dispersion in rates.

We run the following regression:

$$y_i = \alpha + \beta \mathbb{1}(\text{Elicit Beliefs})_i + \varepsilon_i, \tag{3}$$

where y_i are search, negotiation, and loan term outcomes for participant i. The treatment dummy $\mathbb{1}(\text{Elicit Beliefs})_i$ equals one if the participant was assigned to the elicit beliefs treatment and zero otherwise.

Table 6 shows the effect of eliciting participants' beliefs on search, negotiation, and loan outcomes. We find that merely asking these questions led consumers to search at 0.13 more institutions, or a 3.9% increase compared to the control mean of 3.357 institutions, but had no effect on negotiation. This increased search led borrowers to obtain 7.1% lower interest rate offers and 9.6% lower interest rates on the loans they took out compared to participants who were not asked these questions, according to survey data. The effect on interest rates is also statistically significant in administrative data (at the 5% level)—as expected given that the elicit beliefs treatment does not cause a selection effect (i.e., an increase in the probability of taking out a loan). The treatment effect on interest rates in administrative data is lower in magnitude, however, suggesting 1.2% lower interest rates.

7 Conclusion

We document that consumers have inaccurate beliefs about both the first and second moment of the distribution of interest rates that lenders would offer them. About 73% of participants thought they would obtain an interest rate lower than what they actually received and about 75% underestimated the dispersion of interest rate offers they could get. The presence of inaccurate beliefs suggests that consumers are likely to have different search and negotiation behavior than that predicted by models where agents are assumed to perfectly know the distribution of rates from which they are drawing.

We designed a price comparison tool to correct inaccurate beliefs and test their effects on search, negotiation, and loan terms using an RCT. The tool showed participants a histogram of interest rates that borrowers with similar characteristics obtained on similar loans in the last six months. We built the tool in collaboration with Chile's financial regulator using administrative data on 1.8 million loans from 1.2 million borrowers over two years. We recruited participants online through Google ads and both surveyed and treated them online. We measure outcomes in administrative data merged with RCT participants using their national ID numbers and a follow-up phone survey we conducted to collect rich data on their search and negotiation behavior and loan outcomes.

We find that the price comparison tool caused participants to update their beliefs about both moments of the interest rate distribution. In terms of real effects of these changes in beliefs, the tool led participants to receive 13.1% more offers and 11% lower interest rates, by negotiating 39% more. It also made them 4.7% more likely to take out a loan. Thus, we find that negotiation is an important action in consumer credit markets, in addition to search (which has received substantially more attention in the literature).

While correcting biased beliefs can help consumers negotiate successfully for lower interest rates, we find that it is extremely difficult for consumers to correct biased beliefs themselves. In particular, rates shown by Google, bank websites, and third-party comparison tools are all biased or very noisy. Getting accurate interest rate quotes requires actually applying at a lender, which involves substantial search costs. Thus, there is a role for a financial regulator to require that lenders report loan-level information that includes both loan terms and borrower characteristics, as Chile's financial regulator does, and to synthesize and provide this information to consumers, as we did in this experiment.

References

Agarwal, Sumit, John Grigsby, Ali Hortaçsu, Gregor Matvos, Amit Seru, and Vincent Yao (2024). "Searching for Approval." *Econometrica* 92(4), 1195–1231.

Agte, Patrick, Claudia Allende, Adam Kapor, Christopher Neilson, and Fernando Ochoa (2024). "Search and Biased Beliefs in Education Markets." Working Paper.

Alan, Sule, Mehmet Cemalcilar, Dean Karlan, and Jonathan Zinman (2018). "Unshrouding: Evidence from Bank Overdrafts in Turkey." *The Journal of Finance* 73(2), 481–522.

Allcott, Hunt, Joshua Kim, Dmitry Taubinsky, and Jonathan Zinman (2022). "Are High-Interest Loans Predatory? Theory and Evidence from Payday Lending." *The Review of Economic Studies* 89(3), 1041–1084.

Allen, Jason, Robert Clark, and Jean-François Houde (2013). "The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry." *American Economic Review* 104(10), 3365–3396.

Allen, Jason, Robert Clark, and Jean-François Houde (2014). "Price Dispersion in Mortgage Markets." *The Journal of Industrial Economics* 62(3), 377–416.

Allen, Jason, Robert Clark, and Jean-François Houde (2019). "Search Frictions and Market Power in Negotiated-Price Markets." *Journal of Political Economy* 127(4), 1550–1598.

Allen, Jason and Shaoteng Li (2025). "Dynamic Competition in Negotiated Price Markets." *The Journal of Finance* 80(1), 561–614.

Argyle, Bronson, Taylor Nadauld, and Christopher Palmer (2023). "Real Effects of Search Frictions in Consumer Credit Markets." *The Review of Financial Studies* 36(7). Ed. by Gregor Matvos, 2685–2720.

Argyle, Bronson, Taylor Nadauld, Christopher Palmer, and Ryan Pratt (2021). "The Capitalization of Consumer Financing into Durable Goods Prices." *The Journal of Finance* 76(1), 169–210.

Argyle, Bronson S, Taylor D Nadauld, and Christopher J Palmer (2020). "Monthly Payment Targeting and the Demand for Maturity." *The Review of Financial Studies* 33(11). Ed. by Itay Goldstein, 5416–5462.

Armona, Luis, Andreas Fuster, and Basit Zafar (2019). "Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment." *The Review of Economic Studies* 86(4), 1371–1410.

Arteaga, Felipe, Adam J Kapor, Christopher A Neilson, and Seth D Zimmerman (2022). "Smart Matching Platforms and Heterogeneous Beliefs in Centralized School Choice." *The Quarterly Journal of Economics* 137(3), 1791–1848.

Athey, Susan and Guido Imbens (2016). "Recursive Partitioning for Heterogeneous Causal Effects." *Proceedings of the National Academy of Sciences* 113(27), 7353–7360.

Backus, Matthew, Thomas Blake, Brad Larsen, and Steven Tadelis (2020). "Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions*." *The Quarterly Journal of Economics* 135(3), 1319–1361.

Banco Central de Chile (2021). "Encuesta Financiera de Hogares 2021: Principales Resultados." *Report. Joint FAO/WHO Expert Committee on Nutrition*.

Bandiera, Oriana, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali (2025). "The Search for Good Jobs: Evidence from a Six-Year Field Experiment in Uganda." *Journal of Labor Economics* 43(3), 885–935.

Bertrand, Marianne and Adair Morse (2011). "Information Disclosure, Cognitive Biases, and Payday Borrowing." *The Journal of Finance* 66(6), 1865–1893.

Bhutta, Neil, Andreas Fuster, and Aurel Hizmo (2024). "Paying Too Much? Borrower Sophistication and Overpayment in the US Mortgage Market." *SSRN Electronic Journal*.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018). "Diagnostic Expectations and Credit Cycles." *The Journal of Finance* 73(1), 199–227.

Bordalo, Pedro, Nicola Gennaioli, Andrei Shleifer, and Stephen J Terry (forthcoming). "Real Credit Cycles." *American Economic Review*.

Caldwell, Sydnee, Ingrid Haegele, and Jörg Heining (forthcoming). "Bargaining and Inequality in the Labor Market." *Quarterly Journal of Economics*.

Campbell, John Y., Howell E. Jackson, Brigitte C. Madrian, and Peter Tufano (2011). "Consumer Financial Protection." *Journal of Economic Perspectives* 25(1), 91–114.

Célérier, Claire and Boris Vallée (2017). "Catering to Investors Through Security Design: Headline Rate and Complexity." *The Quarterly Journal of Economics* 132(3), 1469–1508.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val (2025). "Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, With an Application to Immunization in India." *Econometrica* 93(4), 1121–1164.

CMF (2024). "SISTEMA CONTABLE (Instrucciones Generales)."

Coen, Jamie, Anil K. Kashyap, and May Rostom (2024). "Price Discrimination and Mortgage Choice." Working Paper.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, Geoff Kenny, and Michael Weber (2024). "The Effect of Macroeconomic Uncertainty on Household Spending." *American Economic Review* 114(3), 645–677.

Cuesta, José Ignacio and Alberto Sepúlveda (2021). "Price Regulation in Credit Markets: A Trade-off between Consumer Protection and Credit Access."

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber (2023). "IQ, Expectations, and Choice." *The Review of Economic Studies* 90(5), 2292–2325.

De Los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest (2012). "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." *American Economic Review* 102(6), 2955–2980.

Ferman, Bruno (2016). "Reading the Fine Print: Information Disclosure in the Brazilian Credit Card Market." *Management Science* 62(12), 3534–3548.

Friedman, Jerome, Trevor Hastie, and Robert Tibshirani (2010). "Regularization Paths for Generalized Linear Models via Coordinate Descent." *Journal of Statistical Software* 33(1).

Galenianos, Manolis and Alessandro Gavazza (2022). "Regulatory Interventions in Consumer Financial Markets: The Case of Credit Cards." *Journal of the European Economic Association* 20(5), 1897–1932.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer (2016). "Expectations and Investment." *NBER Macroeconomics Annual* 30, 379–431.

Gertler, Paul, Sean Higgins, Aisling Scott, and Enrique Seira (2025). "Using Lotteries to Attract Deposits." Tech. rep. w31529. Cambridge, MA: National Bureau of Economic Research.

Grennan, Matthew and Ashley Swanson (2020). "Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining." *Journal of Political Economy* 128(4), 1234–1268.

Guiso, Luigi, Andrea Pozzi, Anton Tsoy, Leonardo Gambacorta, and Paolo Emilio Mistrulli (2022). "The Cost of Steering in Financial Markets: Evidence from the Mortgage Market." *Journal of Financial Economics* 143(3), 1209–1226.

Gurun, Umit G., Gregor Matvos, and Amit Seru (2016). "Advertising Expensive Mortgages." *The Journal of Finance* 71(5), 2371–2416.

Hortaçsu, A. and C. Syverson (2004). "Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds." *The Quarterly Journal of Economics* 119(2), 403–456.

Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer (2024). "Worker Beliefs About Outside Options." *The Quarterly Journal of Economics* 139(3), 1505–1556.

Karlan, Dean and Jonathan Zinman (2019). "Long-Run Price Elasticities of Demand for Credit: Evidence from a Countrywide Field Experiment in Mexico." *The Review of Economic Studies* 86(4), 1704–1746.

Kulkarni, Sheisha, Santiago Truffa, and Gonzalo Iberti (2025). "Removing the Fine Print: Disclosure, Standardized Products, and Consumer Outcomes." Unpublished Manuscript.

Liberman, Andres, Christopher Neilson, Luis Opazo, and Seth Zimmerman (2018). "The Equilibrium Effects of Information Deletion: Evidence from Consumer Credit Markets." National Bureau of Economic Research Working Paper No. 25097.

Liu, Haoyang and Christopher Palmer (2023). "Implicit Extrapolation and Beliefs Channel of Investment Demand."

Malmendier, Ulrike and Stefan Nagel (2011). "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?*." *The Quarterly Journal of Economics* 126(1), 373–416.

Malmendier, Ulrike and Stefan Nagel (2016). "Learning from Inflation Experiences *." *The Quarterly Journal of Economics* 131(1), 53–87.

Medina, Paolina C. and Michaela Pagel (2025). "Does Saving Cause Borrowing? Implications for the Coholding Puzzle." *The Journal of Finance* 80(5), 2689–2738.

Muralidharan, Karthik, Mauricio Romero, and Kaspar Wüthrich (2023). "Factorial Designs, Model Selection, and (Incorrect) Inference in Randomized Experiments." *The Review of Economics and Statistics*, 1–44.

Ponce, Alejandro, Enrique Seira, and Guillermo Zamarripa (2017). "Borrowing on the Wrong Credit Card? Evidence from Mexico." *American Economic Review* 107(4), 1335–1361.

Rubinstein, Ariel (1982). "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50(1), 97–109.

Shue, Kelly and Richard R. Townsend (2021). "Can the Market Multiply and Divide? Non-Proportional Thinking in Financial Markets." *Journal of Finance* 76(5), 2307–2357.

Stahl, Dale O. (1989). "Oligopolistic Pricing with Sequential Consumer Search." *The American Economic Review* 79(4), 700–712.

Stango, Victor and Jonathan Zinman (2014). "Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Bank Overdraft Fees." *The Review of Financial Studies* 27(4), 990–1030.

Stango, Victor and Jonathan Zinman (2016). "Borrowing High versus Borrowing Higher: Price Dispersion and Shopping Behavior in the U.S. Credit Card Market." *Review of Financial Studies* 29(4), 979–1006.

Topkis, Donald M (1998). *Supermodularity and Complementarity*. Princeton University Press.

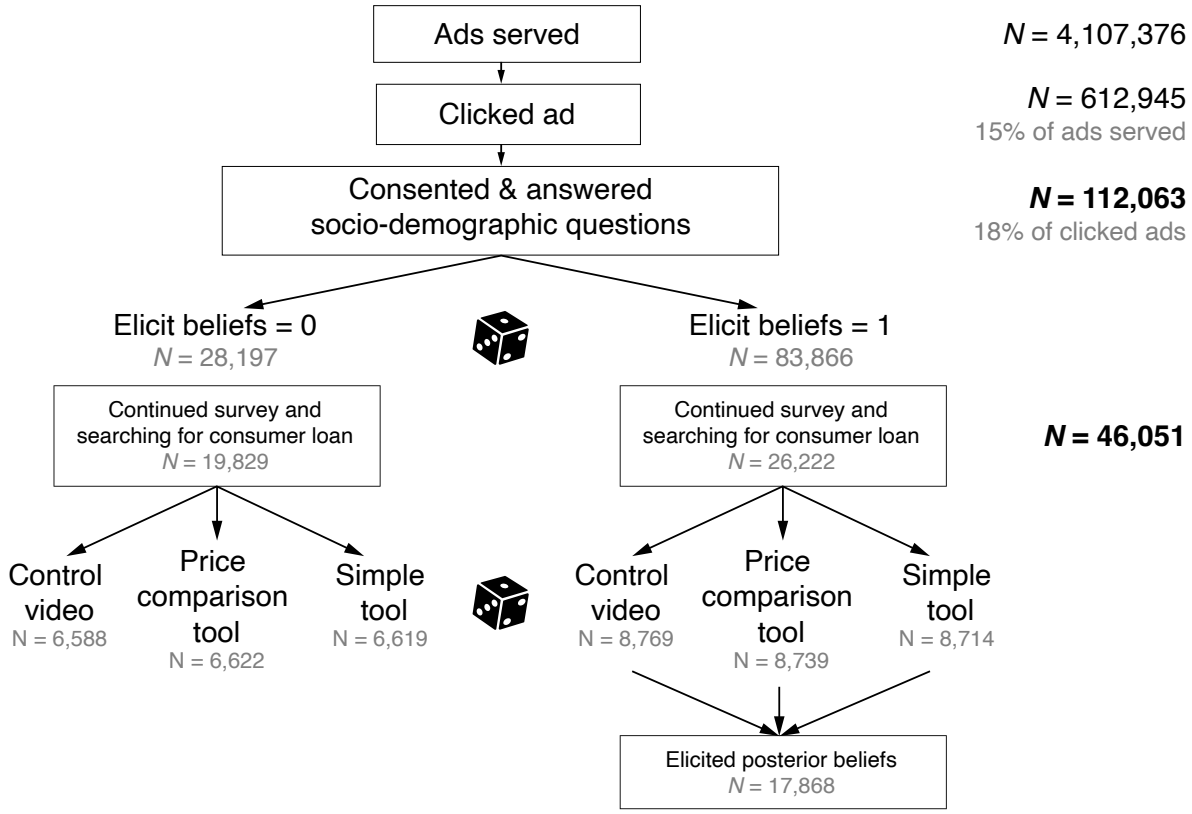
Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion (2022). "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications." *Journal of Economic Perspectives* 36(3), 157–184.

Woodward, Susan E and Robert E Hall (2012). "Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence." *American Economic Review* 102(7), 3249–3276.

Zinman, Jonathan (2015). "Household Debt: Facts, Puzzles, Theories, and Policies." *Annual Review of Economics* 7(1), 251–276.

Zwane, Alix Peterson, Jonathan Zinman, Eric Van Dusen, William Pariente, Clair Null, Edward Miguel, Michael Kremer, Dean S. Karlan, Richard Hornbeck, Xavier Giné, Esther Duflo, Florencia Devoto, Bruno Crepon, and Abhijit Banerjee (2011). "Being Surveyed Can Change Later Behavior and Related Parameter Estimates." *Proceedings of the National Academy of Sciences* 108(5), 1821–1826.

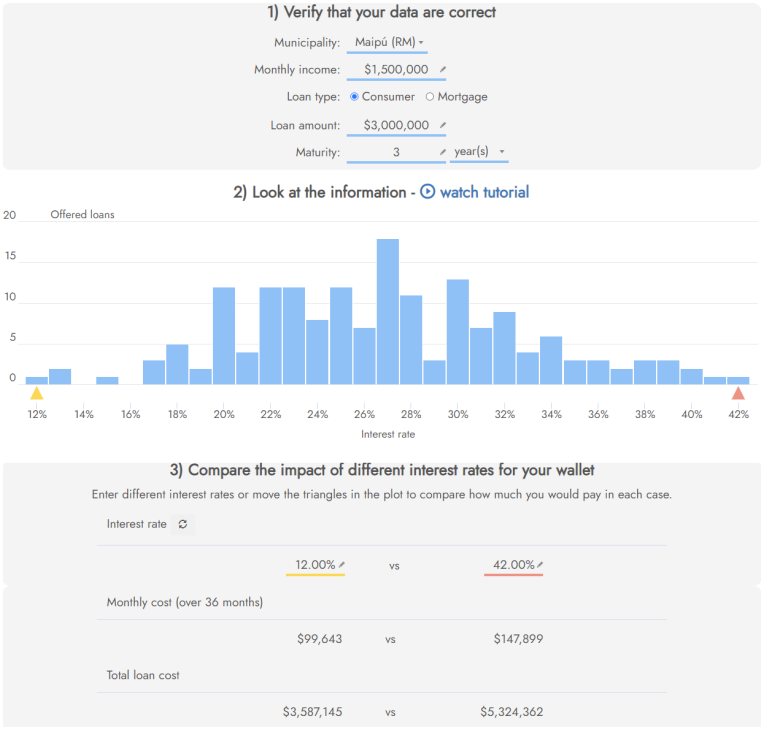
Figure 1: Experimental Design



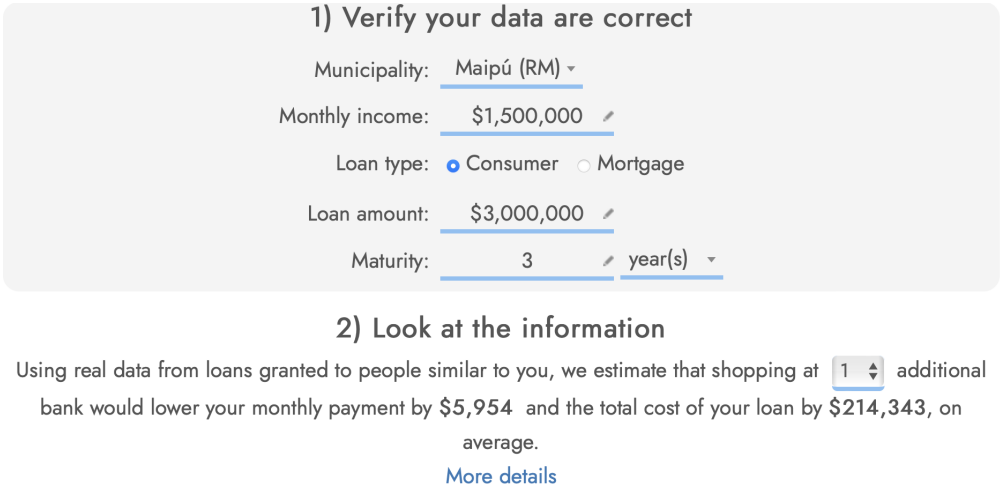
This figure shows the progression of the participants through our study. They are randomized at two key points: when they are assigned either "Elicit beliefs = 0" or "Elicit beliefs = 1" and subsequently when they are cross-randomized to one of our three treatment arms: the control video, price comparison tool, or simple tool.

Figure 2: Price Comparison Tool and Simple Tool Treatments

(a) Interest Rate Price Comparison Tool

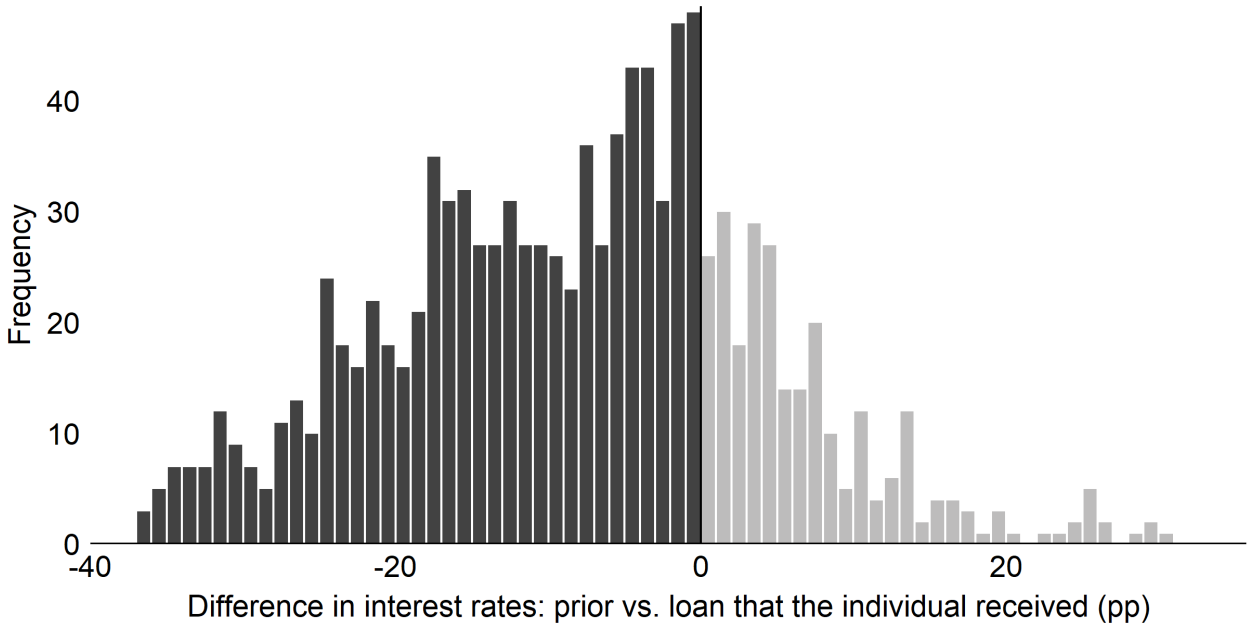


(b) Simple Tool



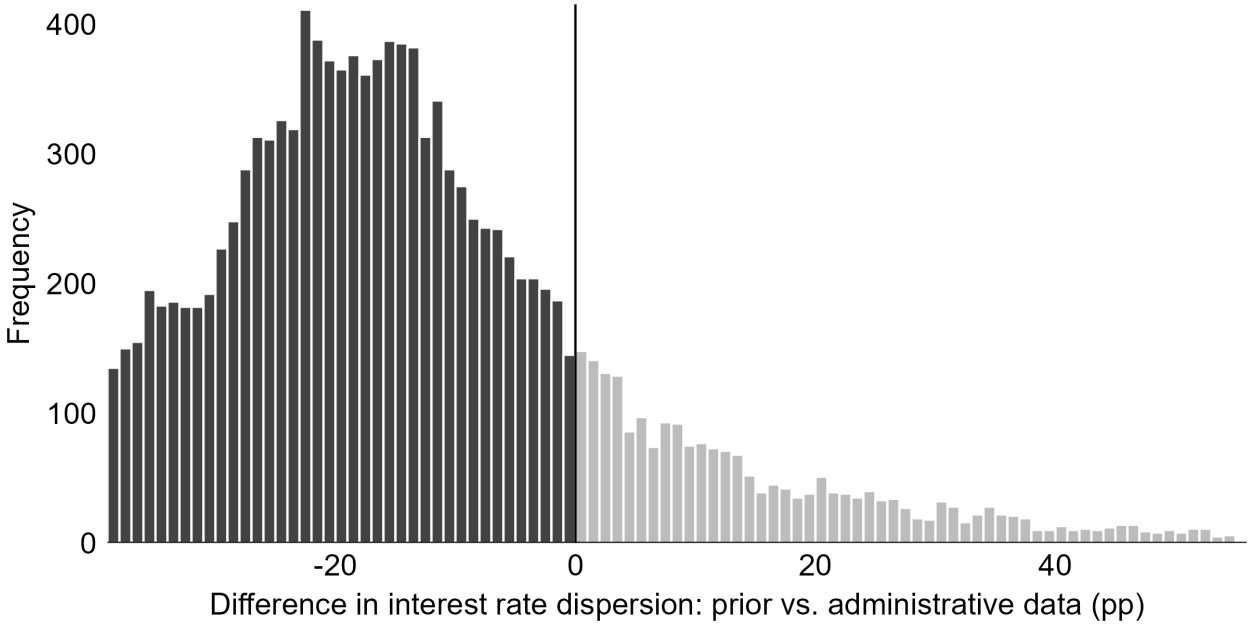
This figure shows a screenshot of English translations of our price comparison tool (panel a) and simple tool highlighting the benefits of search (panel b). For both tools, consumers had already entered the borrower and loan characteristics in the top panel of the tool in our baseline survey; this information is automatically populated for them. Participants can also change this information, in which case the tool is automatically refreshed to show the corresponding data. For the price comparison tool, participants can hover over the histogram bars for more information that helps them interpret and understand the information in the histogram. Participants can also move the triangles along the x-axis to see the implications on monthly and total loan costs. More details are in Appendix E.1. For the simple tool, participants can select from a drop-down menu the number of additional banks they plan to search (up to six banks). The simple tool then displays the amount of money they could save on the monthly and total cost by searching at that many additional banks. More details are in Appendix E.2.

Figure 3: Difference in Interest Rates Between Beliefs and Loan the Individual Received (pp)



This figure shows that most participants underestimate the interest rate they ultimately obtain. The figure is a histogram of the difference between a participant's beliefs about the interest rate they will get on the loan they take out (as reported in our baseline survey) and the interest rate they ended up receiving on the loan they took out in our administrative data. We restrict to the subset of participants in the control group who took out a loan after participating and compare the belief they had reported in the baseline survey to the interest rate they obtained in the administrative data. For participants who obtained more than one loan after participating, we restrict to the first loan they obtained after participating. We remove observations beyond the 5th and 95th percentiles in the graph for legibility. The number of observations is 1,198. The percentage of people who underestimated the rate they would receive, i.e., the percentage of the sample in the negative portion of the histogram, is 73.2%.

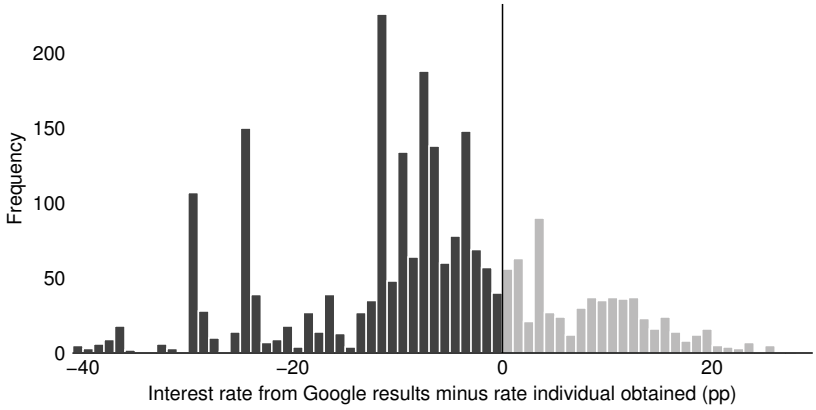
Figure 4: Difference in Interest Rate Dispersion Between Beliefs and Administrative Data (pp)



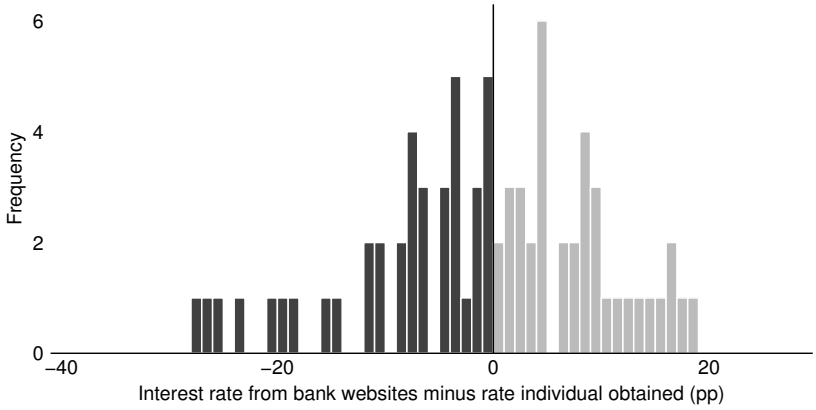
This figure shows that most participants underestimate dispersion. The figure is a histogram of the difference between a participant's prior expectations about the dispersion in interest rates that a bank could offer them, measured as the highest rate a bank could offer them minus the lowest rate a bank could offer them, compared to the dispersion we observe based on their characteristics in the administrative data (i.e., the dispersion they would have seen in the price comparison tool if assigned to that treatment arm). We remove observations beyond the 5th and 95th percentiles in the graph for legibility. The number of observations is 14,890. The percentage of people who underestimated the dispersion, i.e., the percentage of the sample in the negative portion of the histogram, is 74.8%.

Figure 5: Difference in Interest Rates Between Sources and Loan the Individual Obtained (pp)

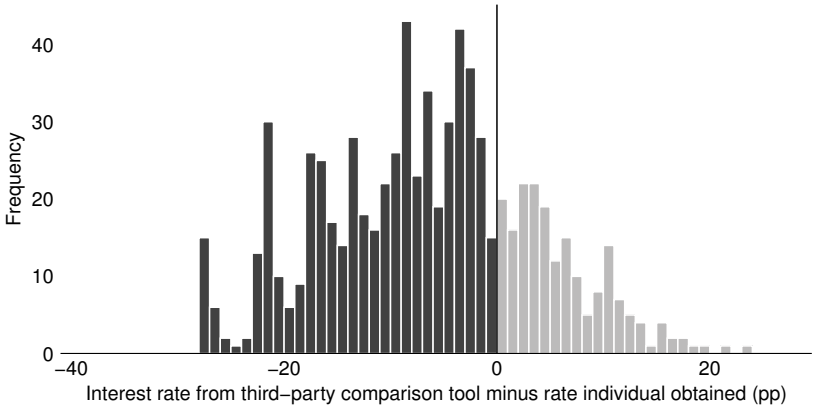
(a) Google Search Results



(b) Bank Websites



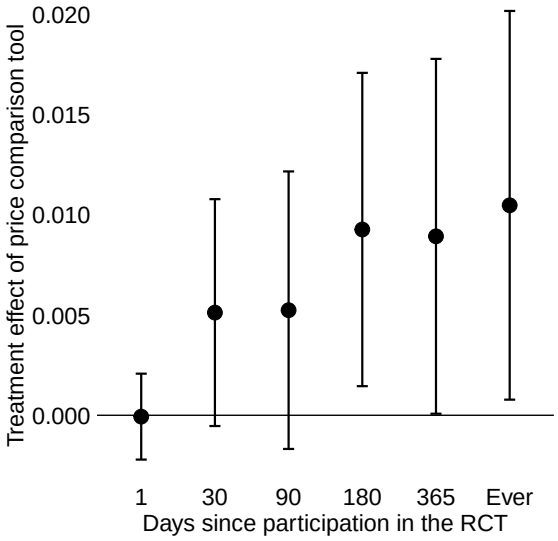
(c) Third-Party Comparison Tool



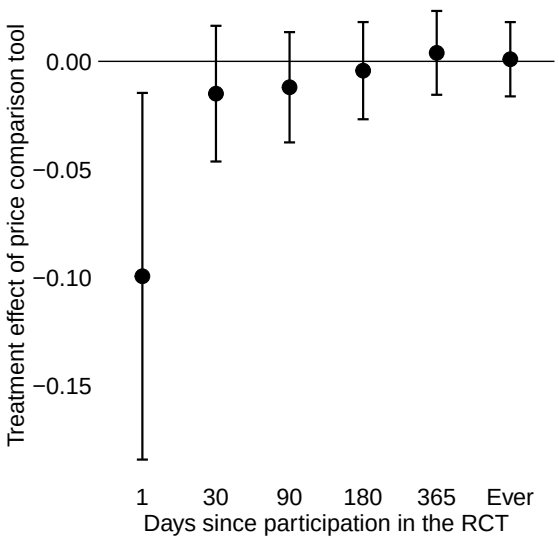
This figure shows histograms of differences between interest rates shown by various sources and the actual interest rate received by participants in the administrative data. Panel (a) shows a histogram of the difference between the interest rate a participant would have seen searching for loan keywords on Google from the bank where they obtained a loan and the rate they actually received from that bank in the administrative data. There are 2,493 observations, of which 74.2% are negative. Panel (b) shows a histogram of the difference between the interest rate a participant would have seen using the bank website of the bank where they obtained a loan and the rate they received from that bank in the administrative data. There are 76 observations, of which 50% are negative. Panel (c) shows a histogram of the difference between the interest rate a participant would have seen on the most popular third-party comparison tool for the bank where they obtained a loan and the rate they received from that bank in the administrative data. There are 749 observations, of which 74.1% are negative. We restrict the analysis to loans obtained during months that overlap with the months in which we scraped data, to ensure that differences are not driven by changes in interest rates over time. See Appendices C and G for more detail.

Figure 6: Effect of Price Comparison Tool over Time (Administrative Data)

(a) Probability of taking a loan

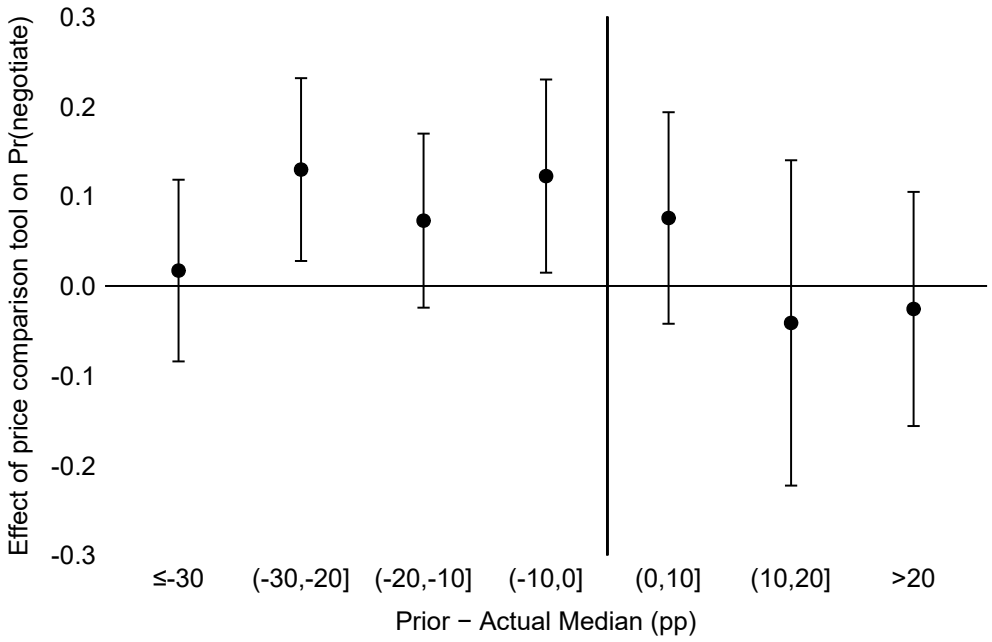


(b) Log interest rate



This figure shows estimates from specification (2) over time relative to when consumers participate in the RCT. For each number of days since participation in the RCT, we measure treatment effects considering only loans obtained within that number of days after participation. RCT = randomized controlled trial.

Figure 7: Heterogeneous Effects of Tool on Negotiation by Beliefs about First Moment



This figure shows heterogeneous treatment effects of the price comparison tool on the probability of negotiating, by how biased participants' beliefs about the first moment were. It shows results from specification (2) estimated separately by bins of the degree to which participants over- or underestimated the first moment of the interest rate distribution. The x-axis measures the participant's belief measured prior to being treated and the observed median of the interest rate distribution they face in administrative data, in percentage points. Bins are each 10 percentage points wide.

Table 1: Balance of Pre-Treatment Characteristics by Elicit Beliefs Treatment

	Elicit Beliefs = 0 Mean	Elicit Beliefs	N
	(1)	(2)	(3)
Personal characteristics			
Age	35.939***	-0.106	112,063
	(0.059)	(0.068)	
log(Income)	13.625***	0.001	109,665
	(0.007)	(0.008)	
Incomplete high-school	0.037***	-0.001	108,809
	(0.001)	(0.001)	
Complete high-school	0.358***	0.003	108,809
	(0.003)	(0.003)	
Complete 2-year program	0.214***	-0.002	108,809
	(0.002)	(0.003)	
Complete 5-year program or higher	0.391***	0.000	108,809
	(0.003)	(0.003)	
Financial products			
Bank account	0.677***	0.002	106,220
	(0.003)	(0.003)	
Any loan	0.707***	-0.006**	107,127
	(0.003)	(0.003)	
Omnibus F-statistic		0.979	112,063
		[0.463]	
Number of participants by arm	28,197	83,866	112,063

This table tests the balance of pre-treatment characteristics by elicit beliefs treatment for the full sample. We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta \mathbb{1}(\text{Elicit Beliefs})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Elicit Beliefs})_i$ is a dummy indicating whether participant i was assigned to the elicit beliefs treatment. Column (1) shows α which is the mean for the $\mathbb{1}(\text{Elicit Beliefs})_i = 0$ group. Column (2) shows β which is the difference in means between the $\mathbb{1}(\text{Elicit Beliefs})_i = 1$ and $\mathbb{1}(\text{Elicit Beliefs})_i = 0$ groups. Column (3) shows the number of observations in each regression, which can change across covariates due to missing values when the respondent did not answer that question in the baseline survey. For the omnibus F-test to test whether the covariates jointly predict treatment, we run the following regression: $\mathbb{1}(\text{Elicit Beliefs})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$. The omnibus F-statistic is a test of $\gamma_1 = \cdots = \gamma_K = 0$. To retain the full sample in the omnibus F-test, if a participant did not answer a particular question, we create a dummy variable indicating whether the variable was missing, replace the missing value with zero, and include the missing value dummy as an additional X^k covariate in the regression. The p-value of the omnibus F-statistic is included in square brackets. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values). "Any loan" is a dummy variable equal to 1 if the participant had any of the following types of loan: consumer loan, mortgage, auto loan, credit card, or cash advance. Unlike in Table 2, we cannot include characteristics of the loan they are searching for in the balance tests since these questions were asked in the same module as the elicit beliefs treatment (rather than a prior module), and the elicit beliefs treatment affected whether participants continued in the survey. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Balance of Pre-Treatment Characteristics by Tool Treatment Arm

		Difference relative to control mean			
	Control Mean	Price Comparison Tool	Simple Tool	Joint test F-stat	N
	(1)	(2)	(3)	(4)	(5)
Personal characteristics					
Age	35.773***	-0.145	0.057	1.616	46,051
	(0.082)	(0.116)	(0.116)	[0.199]	
log(Income)	13.460***	0.000	0.004	0.06	44,978
	(0.010)	(0.014)	(0.014)	[0.942]	
Incomplete high-school	0.041***	0.001	0.002	0.426	44,615
	(0.002)	(0.002)	(0.002)	[0.653]	
Complete high-school	0.425***	-0.008	-0.007	1.068	44,615
	(0.004)	(0.006)	(0.006)	[0.344]	
Complete 2-year program	0.222***	0.006	0.005	0.865	44,615
	(0.003)	(0.005)	(0.005)	[0.421]	
Complete 5-year program or higher	0.312***	0.000	0.000	0.002	44,615
	(0.004)	(0.005)	(0.005)	[0.998]	
Financial products					
Bank account	0.618***	0.016***	0.013**	4.566**	43,272
	(0.004)	(0.006)	(0.006)	[0.01]	
Any loan	0.668***	0.002	0.006	0.526	43,675
	(0.004)	(0.006)	(0.006)	[0.591]	
Desired loan characteristics					
log(Loan Amount)	14.737***	0.020	0.017	0.883	43,775
	(0.012)	(0.017)	(0.017)	[0.413]	
log(Maturity (years))	1.320***	-0.003	0.009	1.334	40,920
	(0.005)	(0.007)	(0.008)	[0.263]	
Omnibus F-statistic					
Price Comparison Tool		1.179			30,718
		[0.279]			
Simple Tool			1.277		30,690
			[0.207]		
Number of participants by arm	15,357	15,361	15,333		46,051

This table tests the balance of pre-treatment characteristics across treatment arms for the sample of consumer loan seekers who continued in the baseline survey long enough to reach the module in which they were assigned to one of the tool treatment arms or the control group. We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta_1 \mathbb{1}(\text{Simple Tool})_i + \beta_2 \mathbb{1}(\text{Price Comparison Tool})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Simple Tool})_i$ and $\mathbb{1}(\text{Price Comparison Tool})_i$ are dummies indicating whether participant i was assigned to the simple tool or price comparison tool arms, respectively. Column (1) shows α which is the mean for the control group. Column (2) shows β_2 which is the difference in means between the price comparison tool treatment arm and the control group. Column (3) shows β_1 which is the difference in means between the simple tool treatment arm and the control group. Column (4) shows the F-statistic from a joint test of $\beta_1 = \beta_2 = 0$. Column (5) shows the number of observations in each regression, which can change across covariates due to missing values when the respondent did not answer that question in the baseline survey. For the omnibus F-test to test whether the covariates jointly predict treatment, we run the following regression for the participants assigned to simple tool and control groups and for those assigned to the price comparison tool and control groups, separately: $\mathbb{1}(\text{Treatment})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$, where $\mathbb{1}(\text{Treatment})_i$ is either $\mathbb{1}(\text{Simple Tool})_i$ or $\mathbb{1}(\text{Price Comparison Tool})_i$. The omnibus F-statistic is a test of $\gamma_1 = \cdots = \gamma_K = 0$. To retain the full sample in the omnibus F-test, if a participant did not answer a particular question, we create a dummy variable indicating whether the variable was missing, replace the missing value with zero, and include the missing value dummy as an additional X^k covariate in the regression. The p-values of the joint test and omnibus F-statistics are included in square brackets. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values). "Desired loan characteristics" refer to characteristics of the loan the participant is searching for. "Any loan" is a dummy variable equal to 1 if the participant had any of the following types of loan: consumer loan, mortgage, auto loan, credit card, or cash advance. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Effects of Price Comparison Tool and Simple Tool on Beliefs

	Expected rate (1)	Lowest rate (2)	Highest rate (3)	Dispersion (4)
Simple Tool	0.70	0.84**	−0.19	0.01
	(0.43)	(0.35)	(0.79)	(0.66)
Price Comparison Tool	16.18***	10.89***	30.35***	15.93***
	(1.18)	(0.93)	(2.24)	(1.45)
Observations	6,817	6,760	6,661	6,272
Control Mean Posterior	29.22	22.65	47.45	23.18
Control Median Posterior	18	12	25	10.68
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows the effect of the simple tool and price comparison tool on beliefs. It shows results from specification (1). Each column shows β_1 and β_2 for one of the following outcome variables: (1) the interest rate the participant expects to get on the loan they take out; (2) the lowest interest rate the participant expects a bank could offer them; (3) the highest interest rate the participant expects a bank could offer them; (4) the difference between the highest and the lowest interest rates the participant expects a bank could offer them. The sample sizes for each column differ based on the number of participants who responded to the corresponding questions. The bin density fixed effects (FEs) are deciles of how many loans are in the price comparison tool histogram that would have been shown to the participant based on their borrower and loan characteristics, had they been assigned to the price comparison tool arm. If any of the participant's loan amount, maturity, or income is outside of the support of the administrative data, the bin density variable would be a missing value; for these individuals, we create an additional fixed effect category that we include in the regression, i.e., we include an additional dummy equal to 1 if the bin density is a missing value. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values) by treatment arm. Tables B.4-B.8 show alternative specifications including using the prior as a control for the posterior rather than subtracting the prior on the left-hand side, using the posterior on the left-hand side without controlling for the prior, taking the natural logarithm of the interest rate expectations, taking the natural logarithm of the interest rate expectations without controlling for priors, and using a normalized measure of dispersion to test whether the treatment effect on dispersion is solely due to a scaling effect. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Real Effects of Price Comparison Tool and Simple Tool

	Survey Data							Administrative Data	
	N of inst. searched (1)	N of inst. applied (2)	N of offers (3)	Pr(negotiate) (4)	Log interest rate offered (5)	Pr(take loan) (6)	Log interest rate taken (7)	Pr(take loan) (8)	Log interest rate taken (9)
(Intercept)	3.450***	1.121***	0.531***	0.097***	3.302***	0.369***	3.213***	0.190***	3.174***
	(0.048)	(0.037)	(0.022)	(0.009)	(0.049)	(0.015)	(0.052)	(0.003)	(0.007)
Simple Tool	0.053	0.018	0.019	0.013	0.000	0.013	−0.031	0.006	0.005
	(0.071)	(0.052)	(0.032)	(0.013)	(0.074)	(0.021)	(0.072)	(0.005)	(0.010)
Price Comparison Tool	0.017	0.025	0.069**	0.037***	−0.127**	0.036*	−0.111*	0.009**	0.004
	(0.071)	(0.051)	(0.033)	(0.014)	(0.062)	(0.021)	(0.065)	(0.005)	(0.010)
Observations	3,283	3,167	3,147	3,114	555	3,143	364	46,051	8,988

This table shows the effect of the simple tool and price comparison tool on search, negotiation, and loan terms using follow-up survey and administrative data. It shows results from specification (2). The outcomes and samples in each column are as follows, and the sample in all columns excludes survey respondents who did not reach the module in which the tool treatments were assigned. Column (1) is a count variable of the number of institutions at which participants report searching for a consumer loan; the sample excludes participants who did not know or refused to answer how many institutions they searched. Column (2) is a count variable of the number of institutions at which participants applied for a loan at any time after participating in the RCT; compared to column (1), this sample also excludes participants who did not know or refused to answer whether they applied for a loan at all of the institutions where they searched. Column (3) is a count variable of the number of approved loan applications at any time after participating in the RCT; compared to column (2), the sample also excludes participants who did not know or refused to answer whether the loan was approved for all their loan applications. Column (4) is a dummy variable equal to 1 if the participant tried to negotiate with at least one institution where they applied; compared to column (3), the sample also excludes participants who did not know or refused to answer whether they negotiated for all their loan offers. Column (5) is the natural logarithm of the interest rate offered post-negotiation; compared to column (3), the sample excludes searches that did not generate an offer and offers for which the participant did not report the offered interest rate. Column (6) is a dummy variable equal to 1 if the participant reported taking out a loan at any time after participating in the RCT; compared to column (3), the sample also excludes participants who did not know or refused to answer whether they took out the loan for all loan offers they received. Column (7) is the natural logarithm of the reported interest rate obtained; compared to column (6), the sample excludes offers that the participant did not take and offers for which the participant did not report the interest rate. For columns (5) and (7), each observation is a loan offer or loan taken. Column (8) is a dummy variable equal to 1 if the participant obtained a consumer loan within 1 year after participating in the RCT according to administrative data from the CMF. Column (9) is the natural logarithm of the interest rate on the loan the participant took out according to administrative data from the CMF; compared to column (8), the column excludes those who did not take out a loan in administrative data within 1 year after participating. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values) by treatment arm. Tables B.9 and B.10 show balance tests for the surveyed subsample and subsample receiving loans in the administrative data in this table, respectively. Heteroskedasticity-robust standard errors are reported in parentheses for columns at the individual level, while standard errors clustered at the individual level are reported for columns at the loan offer level or loan level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Heterogeneous Effects of Tools by Beliefs about Dispersion

	Underestimated dispersion				All others			
	N of inst. searched	N of inst. applied	N of offers	Pr(negotiate)	N of inst. searched	N of inst. applied	N of offers	Pr(negotiate)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(Intercept)	3.562***	1.188***	0.622***	0.101***	3.391***	1.034***	0.444***	0.078***
	(0.086)	(0.068)	(0.041)	(0.017)	(0.084)	(0.056)	(0.034)	(0.014)
Simple Tool	0.015	0.013	−0.065	0.031	0.289**	0.109	0.104**	0.007
	(0.123)	(0.099)	(0.061)	(0.026)	(0.134)	(0.088)	(0.053)	(0.021)
Price Comparison Tool	0.084	0.033	0.050	0.079***	0.003	0.035	0.068	0.028
	(0.140)	(0.096)	(0.062)	(0.028)	(0.124)	(0.084)	(0.054)	(0.022)
Observations	965	939	935	925	1,063	1,026	1,021	1,013

This table shows the heterogeneous effect of the price comparison tool on search and negotiation using follow-up survey data. It shows results from specification (2) estimated separately for those who underestimate dispersion in columns (1)–(4) and all others in columns (5)–(8). Underestimating dispersion is defined as beliefs about dispersion being at least 1 pp lower than the observed dispersion faced by that consumer in the administrative data. The outcomes and samples in each column are as follows, and the sample in all columns excludes survey respondents who did not reach the module in which the tool treatments were assigned. Column (1) is a count variable of the number of institutions at which participants report searching for a consumer loan; the sample excludes participants who did not know or refused to answer how many institutions they searched. Column (2) is a count variable of the number of institutions at which participants applied for a loan at any time after participating in the RCT; compared to column (1), this sample also excludes participants who did not know or refused to answer whether they applied for a loan at all of the institutions where they searched. Column (3) is a count variable of the number of approved loan applications at any time after participating in the RCT; compared to column (2), the sample also excludes participants who did not know or refused to answer whether the loan was approved for all their loan applications. Column (4) is a dummy variable equal to 1 if the participant tried to negotiate with at least one institution where they applied; compared to column (3), the sample also excludes participants who did not know or refused to answer whether they negotiated for all their loan offers. Columns (5)–(8) repeat the outcomes from columns (1)–(4). Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values) by treatment arm. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Real Effects of Eliciting Beliefs

	Survey Data							Administrative Data	
	N of inst. searched (1)	N of inst. applied (2)	N of offers (3)	Pr(negotiate) (4)	Log interest rate offered (5)	Pr(take loan) (6)	Log interest rate taken (7)	Pr(take loan) (8)	Log interest rate taken (9)
(Intercept)	3.357***	1.192***	0.579***	0.111***	3.553***	0.360***	3.469***	0.195***	3.174***
	(0.040)	(0.033)	(0.021)	(0.008)	(0.035)	(0.012)	(0.041)	(0.002)	(0.005)
Elicit Beliefs	0.130***	−0.031	−0.003	0.011	−0.073*	0.001	−0.101**	−0.004	−0.012**
	(0.048)	(0.038)	(0.024)	(0.010)	(0.042)	(0.015)	(0.048)	(0.003)	(0.006)
Observations	5,774	5,565	5,525	5,465	1,241	5,516	724	112,063	21,522

This table shows the effect of the elicit beliefs treatment on search, negotiation, and loan terms using follow-up survey and administrative data. It shows results from specification (3). The outcomes and samples in each column are as follows. Column (1) is a count variable of the number of institutions at which participants report searching for a consumer loan; the sample excludes participants who did not know or refused to answer how many institutions they searched. Column (2) is a count variable of the number of institutions at which participants applied for a loan at any time after participating in the RCT; compared to column (1), this sample also excludes participants who did not know or refused to answer whether they applied for a loan at all of the institutions where they searched. Column (3) is a count variable of the number of approved loan applications at any time after participating in the RCT; compared to column (2), the sample also excludes participants who did not know or refused to answer whether the loan was approved for all their loan applications. Column (4) is a dummy variable equal to 1 if the participant tried to negotiate with at least one institution where they applied; compared to column (3), the sample also excludes participants who did not know or refused to answer whether they negotiated for all their loan offers. Column (5) is the natural logarithm of the interest rate offered post-negotiation; compared to column (3), the sample excludes searches that did not generate an offer and offers for which the participant did not report the offered interest rate. Column (6) is a dummy variable equal to 1 if the participant reported taking out a loan at any time after participating in the RCT; compared to column (3), the sample also excludes participants who did not know or refused to answer whether they took out the loan for all loan offers they received. Column (7) is the natural logarithm of the reported interest rate obtained; compared to column (6), the sample excludes offers that the participant did not take and offers for which the participant did not report the interest rate. For columns (5) and (7), each observation is a loan offer or loan taken. Column (8) is a dummy variable equal to 1 if the participant obtained a consumer loan within 1 year after participating in the RCT according to administrative data from the CMF. Column (9) is the natural logarithm of the interest rate on the loan the participant took out according to administrative data from the CMF; compared to column (8), the column excludes those who did not take out a loan in administrative data within 1 year after participating. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values) by treatment arm. Tables B.11 and B.12 show balance tests for the surveyed subsample and subsample receiving loans in the administrative data in this table, respectively. Heteroskedasticity-robust standard errors are reported in parentheses for columns at the individual level, while standard errors clustered at the individual level are reported for columns at the loan offer level or loan level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix A Model Details and Proofs

This appendix characterizes the equilibrium of the model and includes proofs of the predictions in Section 2.

A.1 Equilibrium Behavior

Consumer behavior Since the negotiation process is costless to consumers and they expect a possibly lower rate after the negotiation, the consumer will always choose not to accept the first-round offer immediately and to instead attempt to negotiate. Thus, in the model, the consumer always negotiates and our predictions can be thought of as predictions about whether the bank engages in negotiation by offering a strictly lower second-round rate $r^2 < r^1$. In the second round, the consumer then makes decisions as in a standard sequential search model, i.e., by comparing the benefit of accepting the offer and the cost of searching further.

A consumer who has an offer in hand from bank j with interest rate r (after attempting to negotiate) will search at another bank if and only if their search cost is lower than the expected gain of getting an additional draw by searching at an additional institution, given their potentially biased beliefs:

$$s \leq \int_{-\infty}^{r} (r - \tilde{r}) d\hat{H}(\tilde{r}; \theta).$$

Then the consumer's optimal strategy can be described by a reservation rate r^* defined as the rate at which the consumer is indifferent between searching at another bank or accepting the loan offer in hand. Thus a consumer with search cost s and beliefs $\hat{H}$ has a reservation rate r^* that satisfies

$$\begin{aligned} s &= \int_{-\infty}^{r^*} (r^* - \tilde{r}) d\hat{H}(\tilde{r}; \theta) \\ &= \hat{H}(\tilde{r}; \theta)(r^* - \tilde{r})\Big|_{-\infty}^{r^*} + \int_{-\infty}^{r^*} \hat{H}(\tilde{r}; \theta) d\tilde{r} \\ &= \int_{-\infty}^{r^*} \hat{H}(\tilde{r}; \theta) d\tilde{r}. \end{aligned}$$

This determines the reservation rate as a function of the consumer's search cost s and belief type θ, i.e., $r^*(s, \theta)$.

Since $F(\cdot)$ is the distribution of search costs, the distribution of reservation rates is[20]

$$\Phi(r; \theta) = \mathbb{P}_s[r^*(s, \theta) \leq r] = \mathbb{P}\left[s \leq \int_{-\infty}^{r} \hat{H}(\tilde{r}; \theta) d\tilde{r}\right] = F\left(\int_{-\infty}^{r} \hat{H}(\tilde{r}; \theta) d\tilde{r}\right).$$

[20] We use subscripts of $\mathbb{P}$ and $\mathbb{E}$ to indicate over which variable(s) these probabilities and expectations are calculated.

This describes the aggregate behavior of consumers with type θ.

Bank behavior The bank's behavior in the two-stage bargaining game is solved backwards.

After the negotiation, the bank learns about the consumer's type θ and knows the distribution $\Phi(r;\theta)$ of reservation rates for consumers with type θ. The intuition is that the reservation rate is a function of search costs *and* beliefs, so after learning the consumer's type θ the bank knows the distribution of reservation rates that corresponds to consumers of that type, but does not know the consumer's exact reservation rate because the search cost remains private information.

When deciding whether to lower the offer rate, the bank has to compare the cost d with the benefits of enlarging its market share (i.e., increasing the probability that the consumer accepts the offer) by offering a lower second-round rate.

Following the analysis of Agarwal, Grigsby, Hortaçsu, Matvos, Seru, and Yao (2024), who model sequential search with rejection but without biased beliefs or negotiation, the residual demand faced by the bank is

$$q(r,\theta) = \int_r^{+\infty} \frac{d\Phi(\tilde{r},\theta)}{H(\tilde{r};\theta)} = \int_r^{+\infty} \frac{1}{H(\tilde{r};\theta)} dF\left(\int_{-\infty}^{\tilde{r}} \hat{H}(\hat{r};\theta)d\hat{r}\right).$$

The bank then attempts to set the rate through the profit maximization problem

$$\max_{r\leq r^1} \pi(r;\theta,c) = q(r,\theta)(r-c),$$

where r^1 is the first-round offered rate. The bank lowers the rate if and only if

$$\max_{r\leq r^1} \pi(r;\theta,c) - \pi(r^1;\theta,c) \geq d.$$

We define the negotiation as successful if the bank lowers the rate and the consumer accepts the bank's offer.

Let the monopoly interest rate $r^{2\star}$ be the unconditionally optimal solution to the profit maximization problem

$$r^{2\star}(\theta,c) = \arg\max_r \pi(r;\theta,c),$$

and let the optimal threshold $r^{2\circ}$ be the largest solution $r < r^1$ to

$$\pi(r;\theta,c) - d = \pi(r^1;\theta,c).$$

Then the optimal second-round rate is

$$r^2(\theta,c,d,r^1) = \begin{cases} r^{2\star}(\theta,c), & r^{2\star}(\theta,c) < r^{2\circ}(\theta,c,d,r^1), \\ r^1, & r^{2\star}(\theta,c) \geq r^{2\circ}(\theta,c,d,r^1). \end{cases}$$

Later, we also write $r^2(\theta,c,d)$ as an abbreviation for $r^2(\theta,c,d,r^1(c,d))$. The first-round rate r^1 is then chosen as

$$r^1(c,d) = \arg\max_{r^1} \mathbb{E}_\theta \left[\left(\pi(r^{2\star};\theta,c) - d \right) \mathbb{1}[r^{2\star} < r^{2\circ}] + \pi(r^1;\theta,c)\mathbb{1}[r^{2\star} \geq r^{2\circ}] \right].$$

Thus, the optimal first-round rate is determined by the distribution of the consumer's type and the costs of financing this consumer c and of lowering the interest rate in a negotiation d.

Equilibrium search and interest rates obtained The consumer accepts the bank's offer if and only if

$$r^2(\theta,c,d) \leq r^*(s,\theta). \tag{4}$$

Therefore, the probability of a successful negotiation is

$$\mathbb{P}_{c,d}[r^2(\theta,c,d) \leq r^*(s,\theta)].$$

The expected number of banks at which the consumer will search is given by

$$\frac{1}{\mathbb{P}_{c,d}[r^2(\theta,c,d) \leq r^*(s,\theta)]}. \tag{5}$$

The expected interest rate they will ultimately obtain is given by

$$\mathbb{E}_{c,d}[r^2(\theta,c,d) | r^2(\theta,c,d) \leq r^*(s,\theta)]. \tag{6}$$

Both of these equilibrium outcomes are determined not only by the consumer's reservation rate $r^*(s,\theta)$ but also by the consumer's type θ.

A.2 Sufficient Statistic Predictions

In this section, we include proofs for Predictions 1 and 2 from Section 2.

Prediction 1. *In a sequential search model with biased beliefs but without negotiation, a consumer's equilibrium search result is completely characterized by the consumer's reservation rate.*

Proof. When there is no negotiation, the bank can no longer observe the consumer's belief type θ, as this is only revealed during the negotiation. Thus, the bank believes that the consumer has a distribution of reservation rates

$$\bar{\Phi}(\tilde{r}) = \mathbb{E}_{s,\theta}[\Phi(\tilde{r}; s, \theta)].$$

The residual demand faced by the bank is then given by

$$\bar{q}(r) = \int_r^{+\infty} \frac{d\bar{\Phi}(\tilde{r})}{\bar{H}(\tilde{r})}.$$

Therefore, the bank will offer the interest rate to maximize its profit:

$$r^\star(c) = \arg\max_r \hat{\pi}(r; c) = \arg\max_r \bar{q}(r)(r - c).$$

This indicates that the bank's offered rate is solely determined by their cost of providing a loan, c, and the aggregate distribution of reservation rates. As a result, a consumer accepts the current offer if and only if

$$r^\star(c) \leq r^*(s, \theta).$$

The expected number of institutions at which the consumer will search is given by

$$\frac{1}{\mathbb{P}_c[r^\star(c) \leq r^*(s, \theta)]}$$

and the expected interest rate they will ultimately obtain by

$$\mathbb{E}_c[r^\star(c) | r^\star(c) \leq r^*(s, \theta)].$$

The consumer's search cost s and type θ only enter these equilibrium quantities through their affect on r^*. In other words, for a given reservation rate $r^*(s, \theta)$, these quantities are uniquely determined and independent of the search cost s or the consumer's belief type θ. □

Prediction 2. *In a sequential search model with biased beliefs **and negotiation**, a consumer's equilibrium search result is **not** completely characterized by the consumer's reservation rate.*

Proof. A consumer with reservation rate r^* accepts the current offer by the bank if and only if (4) holds. The consumer's type θ enters into this expression not only through its affect on the reservation rate $r^*(s, \theta)$ but also through the second-round rate $r^2(\theta, c, d)$. This can also be seen by noting that the equilibrium expected number of banks at which the consumer will search in

(5) and the expected interest rate they will ultimately obtain in (6) are functions of $r^2(\theta, c, d)$. Intuitively, since the bank can learn the consumer's type during the negotiation, the bank's profit function in the second round $\pi(r;\theta,c)$ is computed based on the perceived residual demand faced by the bank, $q(r,\theta)$, which has incorporated the bank's knowledge about the consumer's potentially biased beliefs. □

A.3 Lemmas for Treatment Effect Predictions

We next build up the lemmas required for the treatment effect predictions, and state technical assumptions needed for the lemmas to hold.

Biased beliefs about the first moment We begin with analyzing the effect of biases about the first moment of the interest rate distribution. Throughout this section, we fix $\Delta\sigma^2 = 0$ and take $\theta = (\Delta\mu, 0)$. Hence, the parameter θ only captures first-order biases in the consumer's belief.

The main results depend on the following two distribution assumptions.

Assumption 1. The distribution of search costs $F(\cdot)$ has a continuous, log-concave, and decreasing density function.

Examples of distributions satisfying Assumption 1 include exponential distributions and half-normal distributions.

Assumption 2. The cumulative distribution function of the consumer's belief $\hat{H}(r;\theta)$ is log-supermodular in (r,θ).

Assumption 2 indicates that when θ increases, the rate at which $\hat{H}(r;\theta)$ decreases (in the sense of first-order stochastic dominance, FOSD) is slower for larger r.

Lemma 1. *Under Assumptions 1 and 2, if $\hat{H}(r;\theta)$ increases in θ in the sense of FOSD, then the density function $\phi(r;\theta)$ is also log-supermodular in (r,θ), i.e., the distribution of reservation rates $\Phi(r;\theta)$ increases in θ in the sense of a monotone-likelihood-ratio (MLR) shift.*

Proof. Note that

$$\log\phi(r;\theta) = \log f\left(\int_{-\infty}^{r}\hat{H}(\tilde{r};\theta)d\tilde{r}\right) + \log\hat{H}(r;\theta).$$

It holds that

$$\frac{\partial^2\log\phi(r;\theta)}{\partial r\partial\theta} = (\log f)''\cdot\hat{H}(r;\theta)\int_{-\infty}^{r}\hat{H}_\theta(\tilde{r};\theta)d\tilde{r} + (\log f)'\cdot\hat{H}_\theta(r;\theta) + \frac{\partial^2\log\hat{H}(r;\theta)}{\partial r\partial\theta}.$$

Since $F(\cdot)$ has a log-concave and decreasing density, $(\log f)'' < 0$ and $(\log f)' > 0$. The FOSD increase of $\hat{H}$ in θ implies $\hat{H}_\theta(x;\theta) \leq 0$ as well as $\int_{-\infty}^{r} \hat{H}_\theta(\tilde{r};\theta)d\tilde{r} \leq 0$. Also, the log-supermodularity of $\hat{H}(r;\theta)$ indicates that the last term is positive. Therefore, each term in the expression above is positive, so $\partial^2 \log \phi(r;\theta)/\partial r \partial \theta$ is also positive. This indicates that $\phi(r;\theta)$ is also log-supermodular in (r,θ).

Meanwhile, when $\hat{H}(r;\theta)$ increases in θ in the sense of FOSD, $\int_{-\infty}^{r} \hat{H}(\tilde{r};\theta)d\tilde{r}$ also decreases in θ for all r. Thus, $\Phi(r;\theta)$ decreases in θ for all r. The log-supermodularity of $\phi(r;\theta)$ is thus equivalent to an MLR increase of $\Phi(r;\theta)$ in θ. □

Lemma 2. *If the distribution of reservation rates $\Phi(r;\theta)$ increases in θ in the MLR sense, then the monopoly interest rate $r^{2\star}$ also increases in θ.*

Proof. First, it is true that the residual demand function $q(r,\theta)$ is also log-supermodular in (r,θ). Note that

$$\begin{aligned}
\frac{\partial^2}{\partial r \partial \theta} \log q(r;\theta) &= \frac{\partial}{\partial \theta} \frac{q_r(r;\theta)}{q(r;\theta)} \\
&= \frac{q_{r\theta}(r;\theta)q(r;\theta) - q_r(r;\theta)q_\theta(r;\theta)}{(q(r;\theta))^2} \\
&= \frac{1}{(q(r;\theta))^2}\left(-\frac{\phi_\theta(r;\theta)}{H(r)}\int_r^{\bar{r}} \frac{\phi(\tilde{r};\theta)}{H(\tilde{r})}d\tilde{r} + \frac{\phi(r;\theta)}{H(r)}\int_r^{\bar{r}} \frac{\phi_\theta(\tilde{r};\theta)}{H(\tilde{r})}d\tilde{r}\right) \\
&= \frac{1}{(q(r;\theta))^2}\int_r^{\bar{r}} \frac{\phi(r;\theta)\phi_\theta(\tilde{r};\theta) - \phi(\tilde{r};\theta)\phi_\theta(r;\theta)}{H(r)H(\tilde{r})}d\tilde{r} \\
&= \frac{1}{(q(r;\theta))^2}\int_r^{\bar{r}} \frac{\phi(r;\theta)\phi(\tilde{r};\theta)}{H(r)H(\tilde{r})}\left(\frac{\phi_\theta(\tilde{r};\theta)}{\phi(\tilde{r};\theta)} - \frac{\phi_\theta(r;\theta)}{\phi(r;\theta)}\right)d\tilde{r} \\
&= \frac{1}{(q(r;\theta))^2}\int_r^{\bar{r}} \frac{\phi(r;\theta)\phi(\tilde{r};\theta)}{H(r)H(\tilde{r})}\left(\int_r^{\tilde{r}} \frac{\partial^2}{\partial\theta\partial x}\log\phi(x;\theta)dx\right)d\tilde{r} > 0.
\end{aligned}$$

Next, the log-supermodularity of $q(r;\theta)$ implies a higher monopoly rate $r^{2\star}$. Note that the monopoly rate is

$$\begin{aligned}
r^{2\star} &= \arg\max_r q(r;\theta)(r-c) \\
&= \arg\max_r \log q(r;\theta) + \log(r-c).
\end{aligned}$$

Since we have established the log-supermodularity of $q(r;\theta)$, the theorm from Topkis (1998) applies. Thus, $r^{2\star}$ is monotonically increasing with a higher θ (i.e., an FOSD increase in Φ). □

Next, we note that an extremely deviated prior (i.e., a large bias about the first moment of the interest rate distribution) leads to negotiation failure.

Lemma 3. *If the mean of the reservation rate distribution is either too low or too high, the bank will not lower the interest rate in the second round of the negotiation, i.e.,* $r_2 = r_1$.

Proof. Let $\mu_\Phi(\theta) = \mathbb{E}_r[\Phi(r;\theta)]$ be the mean of the reservation rate distribution. First, consider the case where $\mu_\Phi(\theta)$ is too low. Note that for $r \geq c$, the demand function

$$q(r;\theta) \leq \frac{1}{H(c;\theta)}(1-\Phi(r;\theta)).$$

By the uniform shift assumption, there exists $R > c$ such that for every $r > R$, $\Phi(r;\theta)$ increases monotonically when $\mu_\Phi(\theta)$ decreases. Therefore, for $r > R$,

$$\pi(r;\theta,c) = q(r;\theta)(r-c) \leq (1-\Phi(R;\theta))(R-c).$$

Thus, for sufficiently large $R(\theta,c)$, the profit $\pi(r;\theta,c) \leq d$ for all $r > R(\theta,c)$. Next, by Markov's inequality,

$$1-\Phi(c;\theta) \leq \frac{1}{c}\mu_\Phi(\theta).$$

Therefore, for $c \leq r \leq R(\theta,c)$,

$$q(r;\theta) \leq \frac{1}{cH(c;\theta)}\mu_\Phi(\theta),$$

so

$$\pi(r;\theta,c) \leq \frac{R-c}{cH(c;\theta)}\mu_\Phi(\theta).$$

For sufficiently low $\mu_\Phi(\theta)$, the profit $\pi(r;\theta,c) \leq d$. In conclusion, we have proven that for $\mu_\Phi(\theta)$ too low, the profit never exceeds d. Therefore, the bank will never find it profitable to lower the interest rate.

Second, consider the case where $\mu_\Phi(\theta)$ is too high. Since the monopoly interest rate $r^{2\star}(\theta,c)$ increases with the mean of the reservation rate distribution, it will eventually be larger than $r^1(c,d)$ for sufficiently high $\mu_\Phi(\theta)$. The bank is only allowed to lower the interest rate, so the bank's optimal strategy is to keep the rate unchanged. □

Combining these two lemmas, we conclude that the optimal interest rate provided by the bank is weakly increasing with the mean of the reservation rate distribution except for extremely low means.

We next consider the effects of biased beliefs on the aggregate behaviors of a large quantity of

bank-consumer pairs with belief $\hat{H}$. When faced with a consumer of belief $\hat{H}$, the probability of a successful negotiation and the expected value of the offered interest rate are respectively $\mathbb{P}[r^{2\star} < r^{2\circ}]$ and $\mathbb{E}[r^2]$. This requires us to take expectations of each individual-bank pair's quantities over all distributions F, G, and K.

Lemma 4. *Fix the variance of $\hat{H}$. The probability of a successful negotiation is a single-peaked function of the bias in the first moment $\Delta\mu$. Moreover, there exist $L_1, L_2 > 0$ such that when $\Delta\mu < -L_1$ and $\Delta\mu > L_2$, it is impossible for the negotiation to be successful.*

Proof. Following the proof of Lemma 3, the effects of extreme biases are uniform over parameters other than the biased beliefs. □

Biased beliefs about the second moment We next analyze the effect of consumer biases about the second moment of the interest rate distribution. Throughout this section, we fix $\Delta\mu = 0$ and take $\theta = (0, \Delta\sigma^2)$. Hence, the parameter θ only captures second-order biases in the consumer's beliefs.

We first show that expecting a lower variance leads to a higher reservation rate.

Lemma 5. *If $\hat{H}(\cdot)$ increases in the sense of second-order stochastic dominance (SOSD), then $\Phi(\cdot)$ increases in the sense of FOSD.*

Proof. In this case, the quantity

$$\int_{-\infty}^{\tilde{r}} \hat{H}(r)dr$$

decreases for each $\tilde{r} \in \mathbb{R}$. Also, the distribution $F(\cdot)$ is increasing, so $\Phi(\tilde{r})$ also decreases for each $\tilde{r} \in \mathbb{R}$. This implies that $\Phi(\cdot)$ increases in the sense of FOSD. □

To proceed in the second-order case, we have to strengthen the conclusions of Lemma 5 by imposing an assumption on the distribution of reservation rates.

Assumption 3. The distribution of reservation rates $\Phi(r;\theta)$ belongs to a location family of distributions $F(x;\theta) = F_0(x - \mu(\theta))$, where $\mu(\cdot)$ is strictly increasing and F_0 is log-concave.

This assumption is satisfied, for example, when the distribution of reservation rates is a normal distribution with a fixed variance. Under this assumption, we are able to establish the following conclusion.

Lemma 6. *Under Assumption 3, if $\Phi(\cdot)$ increases in the sense of FOSD, the monopoly interest rate $r^{2\star}$ increases.*

Proof. When Assumption 3 holds, an FOSD increase in $\Phi(\cdot)$ is equivalent to a MLR shift in $\Phi(\cdot)$. To see this, note that

$$\frac{\partial^2 \log \phi(r;\theta)}{\partial r \partial \theta} = \frac{\partial^2 \log f_0(r-\mu(\theta))}{\partial r \partial \theta} = -\frac{\partial^2 \log f_0}{\partial x^2}\bigg|_{r-\mu(\theta)} \mu'(\theta).$$

Therefore, for log-concave distribution F_0 and strictly increasing μ, this is strictly positive, and thus $\Phi(r;\theta)$ experiences an MLR shift.

With this result, we then apply Lemma 2 to conclude the proof. □

We next discuss the impact of a second-order deviation from the accurate prior.

Lemma 7. *Fix the mean of $\hat{H}$. As the variance of $\hat{H}$ decreases from the accurate prior, the probability of a successful negotiation decreases, and the final interest rate increases.*

Proof. As the variance of $\hat{H}$ decreases, suggesting that $\hat{H}$ SOSD H, the distribution of reservation rates shifts to the right. Since the variance of $\hat{H}$ is smaller than that of the accurate prior and the mean is correct, there is no chance for this scenario to fall into extreme cases. According to the prior lemmas, the probability of a successful negotiation decreases, and the final interest rate increases. □

A.4 Treatment Effect Predictions

Prediction 3. *If a consumer is biased about the first moment of the interest rate distribution, i.e., $\theta = (\Delta\mu, 0)$ for some $\Delta\mu \neq 0$, then the treatment effects of the price comparison tool on negotiation and the interest rate obtained depend on the direction and magnitude of the bias. In particular:*

1. *There exist $L_1, L_2 > 0$ such that when $\Delta\mu < -L_1$ or $\Delta\mu > L_2$, there is no treatment effect on the probability of a successful negotiation or the interest rate obtained.*

2. *For $\Delta\mu$ in some subset of $(-L_1, L_2)$, treatment increases the probability of a successful negotiation and reduces the interest rate obtained.*

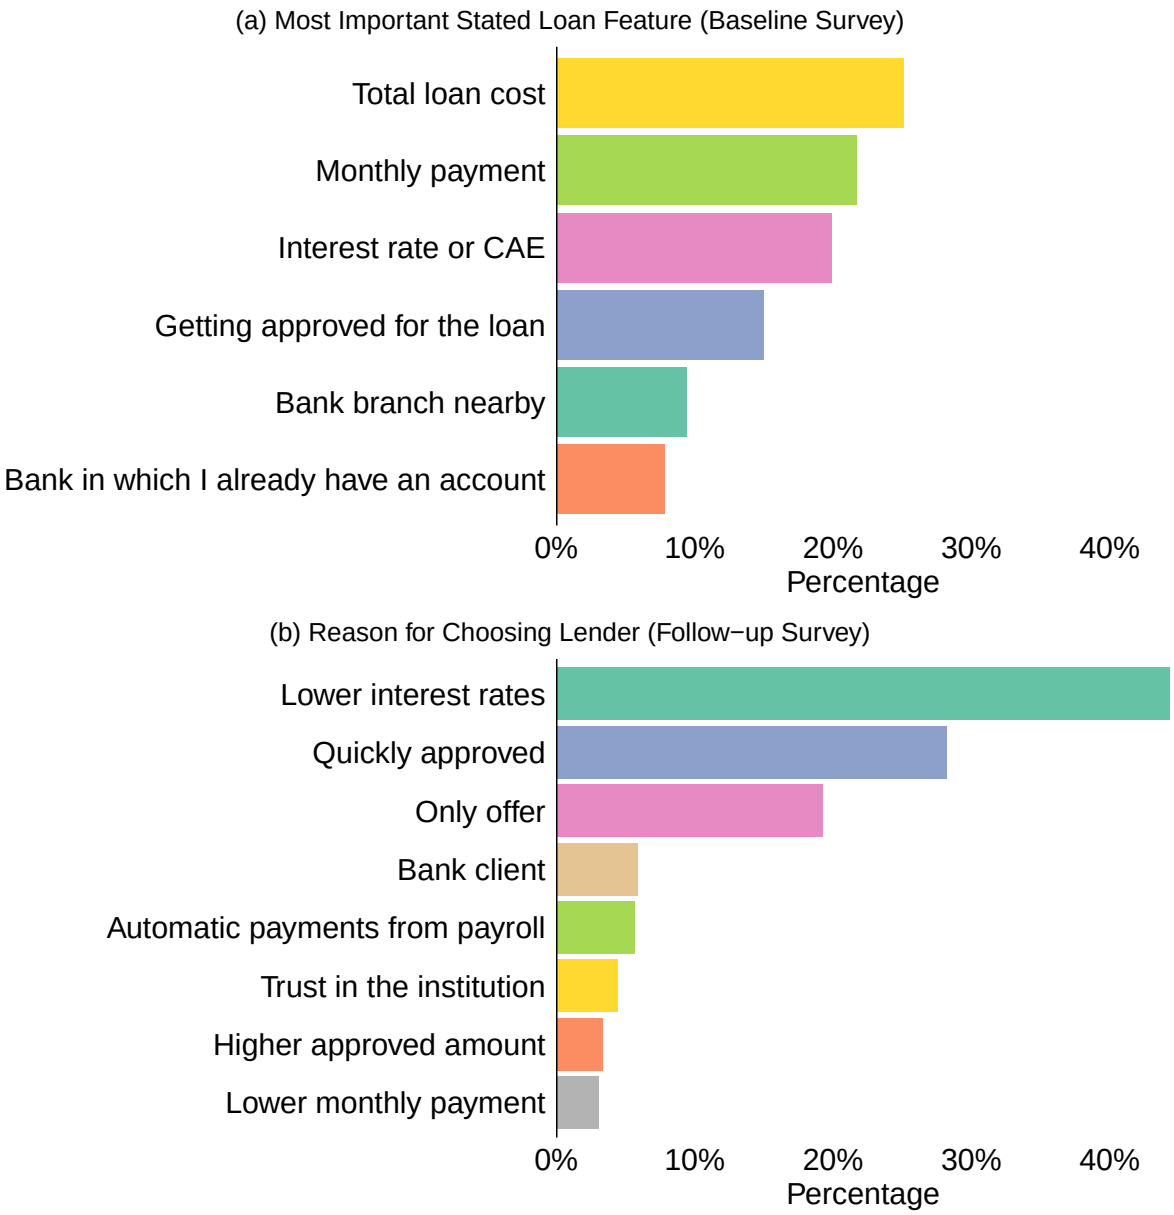
Proof. This is a direct result of Lemma 4. □

Prediction 4. *If a consumer is biased about the second moment of the distribution, and in particular expects a lower variance than the actual distribution, i.e, $\theta = (0, \Delta\sigma^2)$ for some $\Delta\sigma^2 < 0$, then the price comparison tool increases the consumer's probability of a successful negotiation with the bank and reduces the interest rate they obtain.*

Proof. This is a direct result of Lemma 7. □

Appendix B Additional Figures and Tables

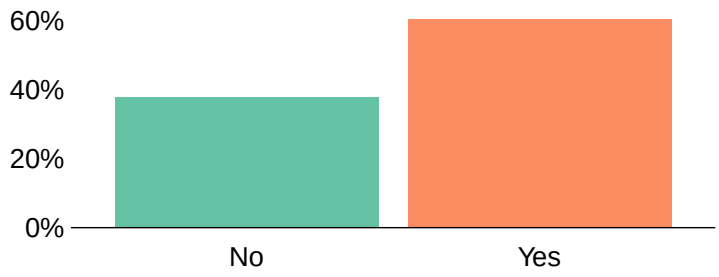
Figure B.1: Stated Importance of Loan Features



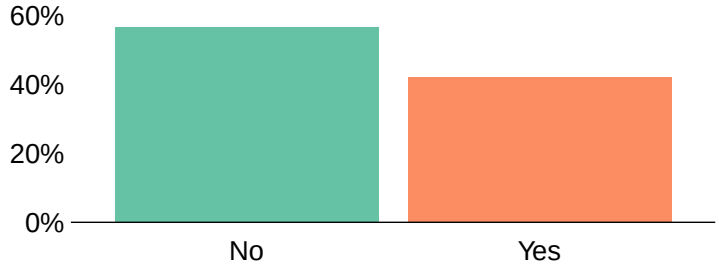
This figure shows the most important features of a loan reported in the baseline survey and the reason they chose to accept a particular loan offer in the follow-up survey. Panel (a) shows results from the baseline survey, conducted when participants were searching. It shows the reasons that participants ranked as most important in response to the question "What are the most important features of the loan you are looking for?" Panel (b) shows responses in our follow-up survey for the subset of participants who took out a loan. It shows responses to the question "Why did you take the loan from {Bank X} compared to offers you saw or received from other banks?" CAE refers to the *carga anual equivalente* which is analogous to an annualized percentage rate (APR).

Figure B.2: Sequential Search, Simultaneous Search, and Searching for Approval

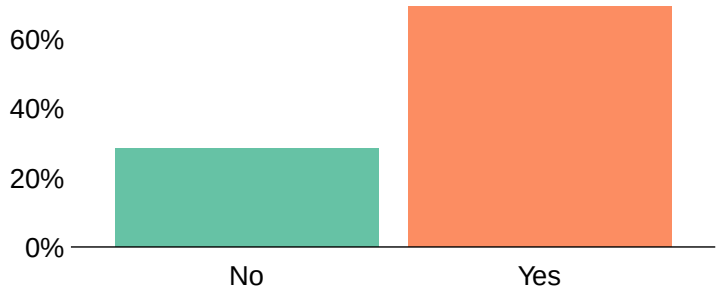
(a) Sequential Search: Target Interest Rate



(b) Simultaneous Search: Target Number of Offers or Banks



(c) Searching for Approval: Search Until One Offer



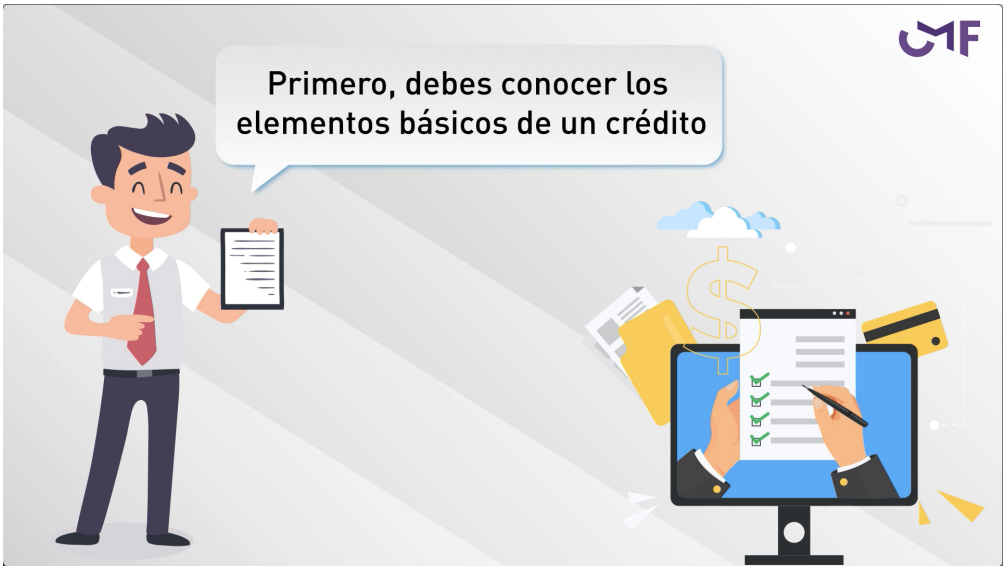
This figure shows the results of asking participants in our follow-up phone survey questions about their search strategy. We asked four yes/no questions that underlie the three panels: (a) "Did you plan to search until you reached a target interest rate and then stop searching?"; (b) "Did you have a target number of offers you would like to receive from financial institutions to stop searching?" and "Did you have a target number of financial institutions from which you wanted to obtain information about loans?"; (c) "Did you expect to search until a financial institution approved your application and then take a loan from that institution?" For each panel, we report the proportions that respond yes or no to the questions, and for panel (b), we report the number of participants who answered "yes" to either of the two questions.

Figure B.3: Sample Google Advertisement for Participant Recruitment



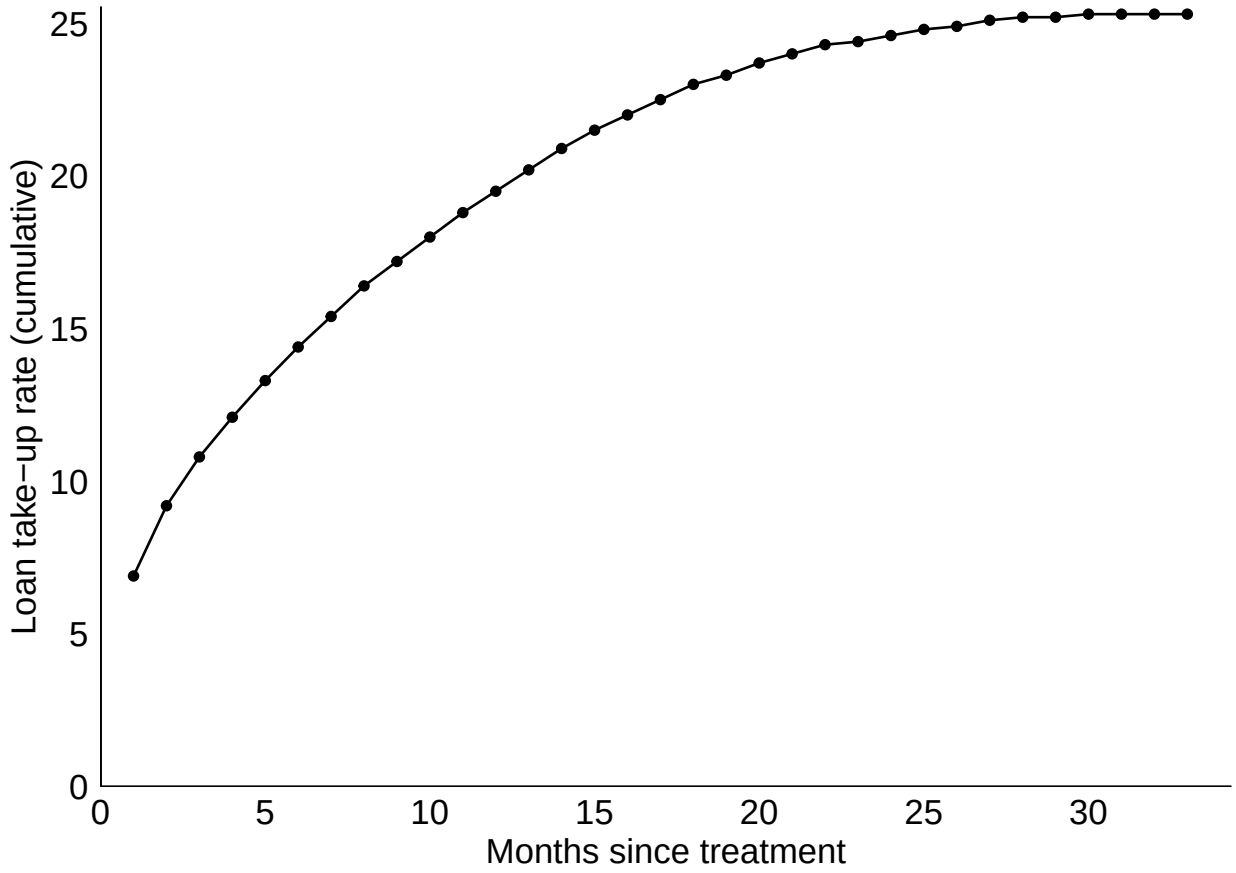
This figure shows an English translation of one of our Google advertisements that we targeted to people searching for keywords related to consumer loans in Chile to recruit them as participants in the RCT.

Figure B.4: Screenshot of Control Video



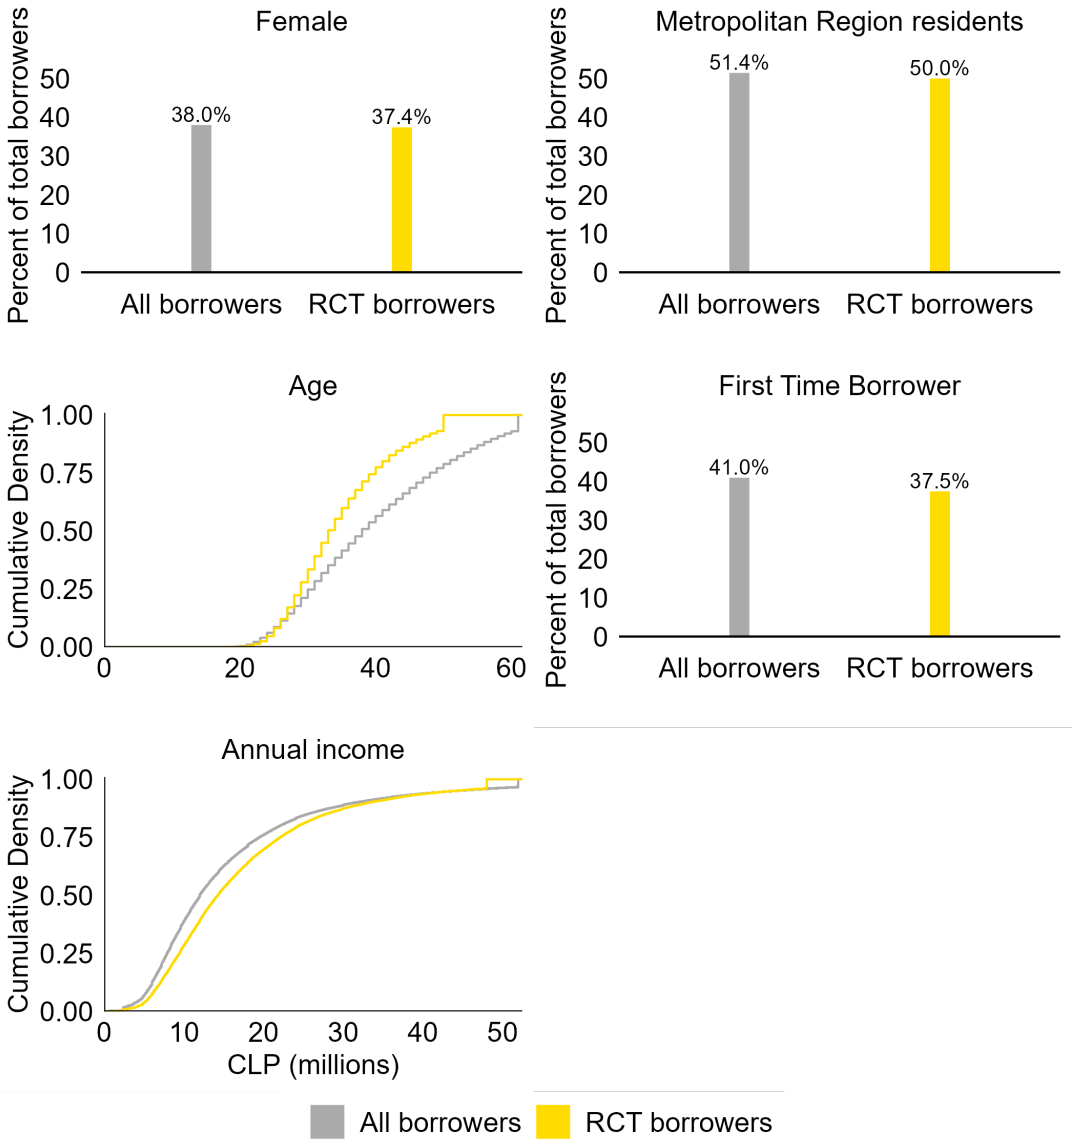
This figure shows a screenshot of the animated video shown to the control group. The video lasts 1 minute and 35 seconds and was developed by the Comisión Mercado Financiero (CMF) to provide basic loan terminology, but not provide information that would affect search or negotiation.

Figure B.5: Participant Loan Take-Up Rate Since Treatment



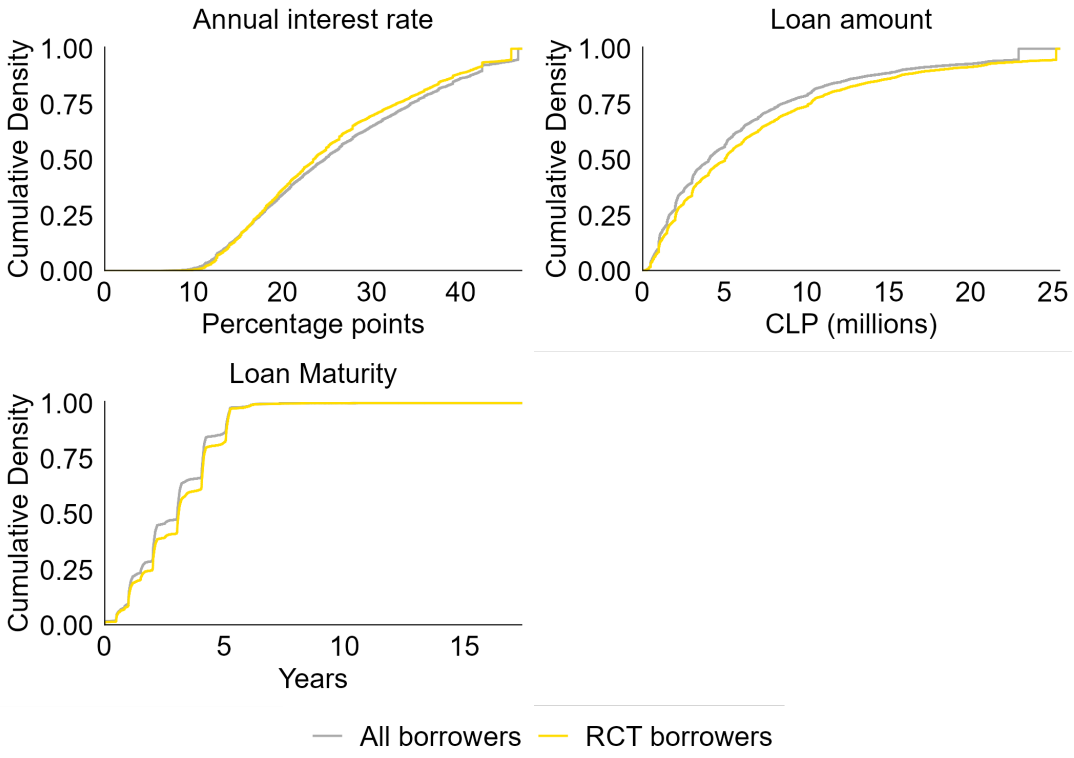
The figure shows the cumulative loan take-up rates of the 46,051 participants who were assigned to either our control video, price comparison tool, and simple tool. Overall, 11,666 of our RCT participants ended up taking out a consumer loan. We define loan take up as the participants having a loan in our administrative data on bank consumer loans.

Figure B.6: External Validity: Borrower Characteristics



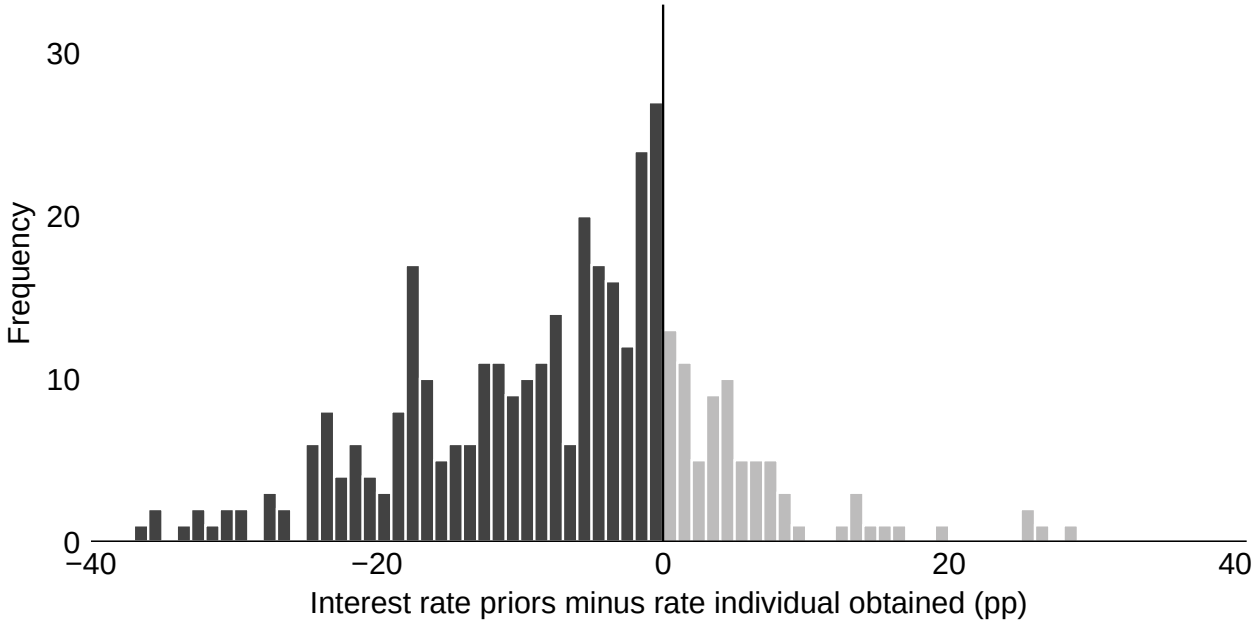
This figure compares borrower characteristics for participants in our RCT who subsequently obtained a loan in the administrative data (yellow bars), to those of all borrowers in the administrative data excluding RCT participants (gray bars). The sample is constructed from 27,130 loans taken out by RCT borrowers and 1,454,216 loans taken out by other borrowers from November 2021 to June 2024.

Figure B.7: External Validity: Loan Terms



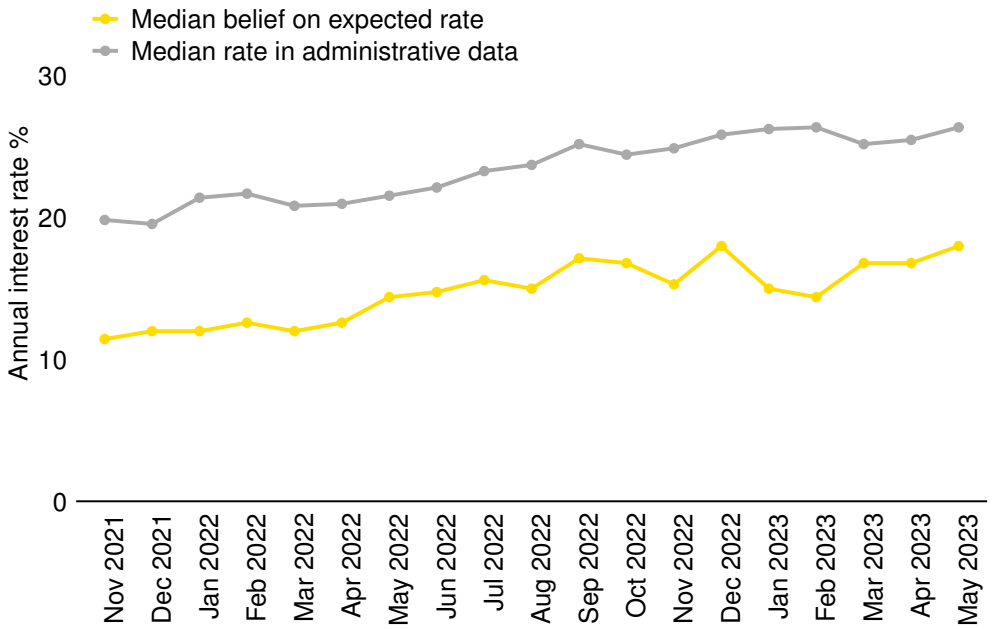
This figure compares loan terms for loans obtained by participants in our RCT who subsequently obtained a loan in the administrative data (yellow bars), to those of loans obtained by all borrowers in the administrative data excluding RCT participants (gray bars). The sample is constructed from 27,130 loans taken out by RCT borrowers and 1,454,216 loans taken out by other borrowers from November 2021 to June 2024.

Figure B.8: Difference in Interest Rates Between Beliefs and Rate the Individual Received (pp), Restricted to First Quartile of Time Between Participation and Obtaining Loan



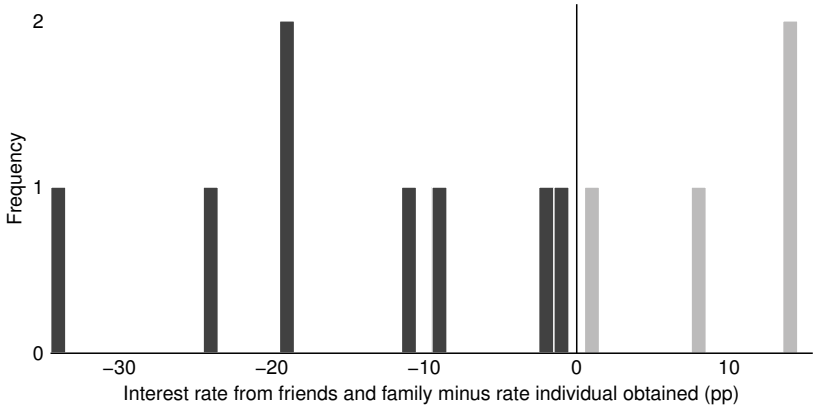
This figure is equivalent to Figure 3 but restricted to the first quartile of the number of days between participating in the RCT and obtaining a loan. It shows that even among participants who obtain a loan shortly after participating (within 22 days, which corresponds to the 25th percentile), most underestimate the interest rate they ultimately obtain. The figure is a histogram of the difference between a participant's beliefs about the interest rate they will get on the loan they take out (as reported in our baseline survey) and the interest rate they ended up receiving on the loan they took out in our administrative data. We restrict to the subset of participants in the control group who took out a loan within 22 days after participating and compare the belief they had reported in the baseline survey to the interest rate they obtained in the administrative data. For participants who obtained more than one loan, we restrict to the first loan they obtained after participating. We remove observations beyond the 5th and 95th percentiles in the graph for legibility. The number of observations is 394. The percentage of people who underestimated the rate they would receive, i.e., the percentage of the sample in the negative portion of the histogram, is 74.9%.

Figure B.9: Interest Rates Over Time



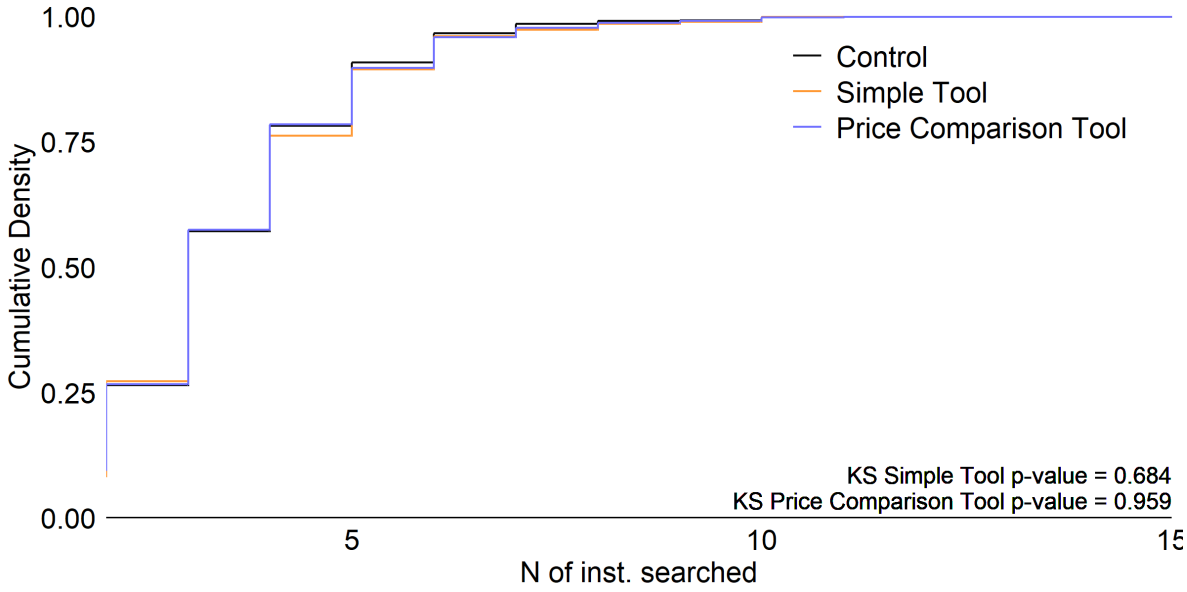
This figure shows the median annual interest rate from administrative data and the median belief about the expected annual interest rate reported by respondents in the baseline survey. For each month, we compute the median across all observed consumer loans (in administrative data) and across all respondents' beliefs (in survey data). Months are based on the response date for survey data and the loan operation date for administrative data.

Figure B.10: Difference in Interest Rates Shared by Friends and Family and Obtained



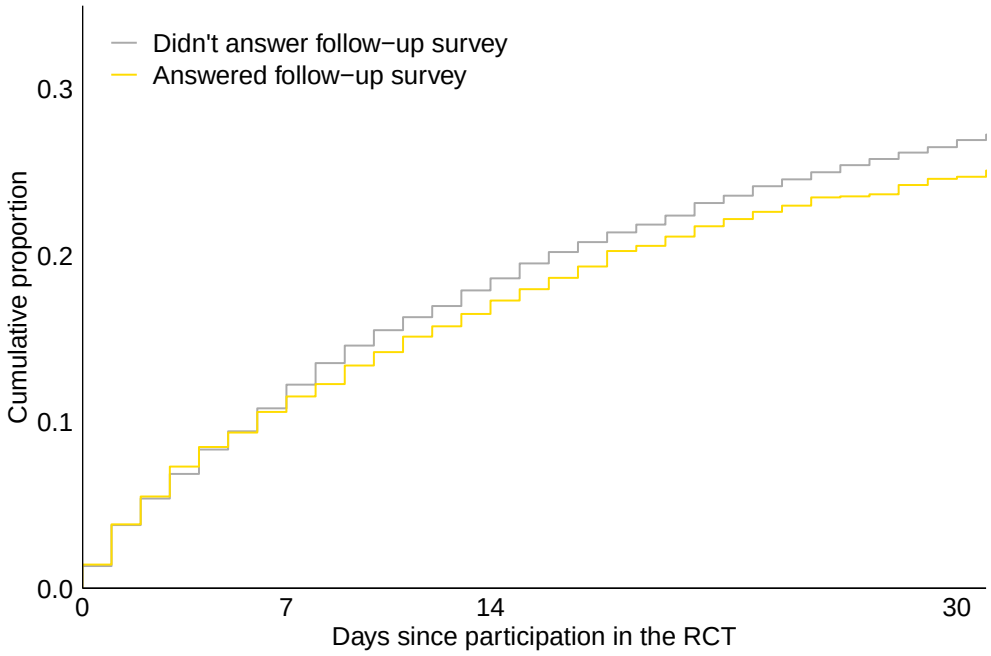
This figure shows a histogram of differences between interest rates shared by family and friends (from survey data) and the actual interest rate received by participants in the administrative data. We remove observations beyond the 5th and 95th percentiles in the graph for legibility. There are 15 observations, 68.8% of which are negative.

Figure B.11: Cumulative Distribution Function of Number of Institutions Searched



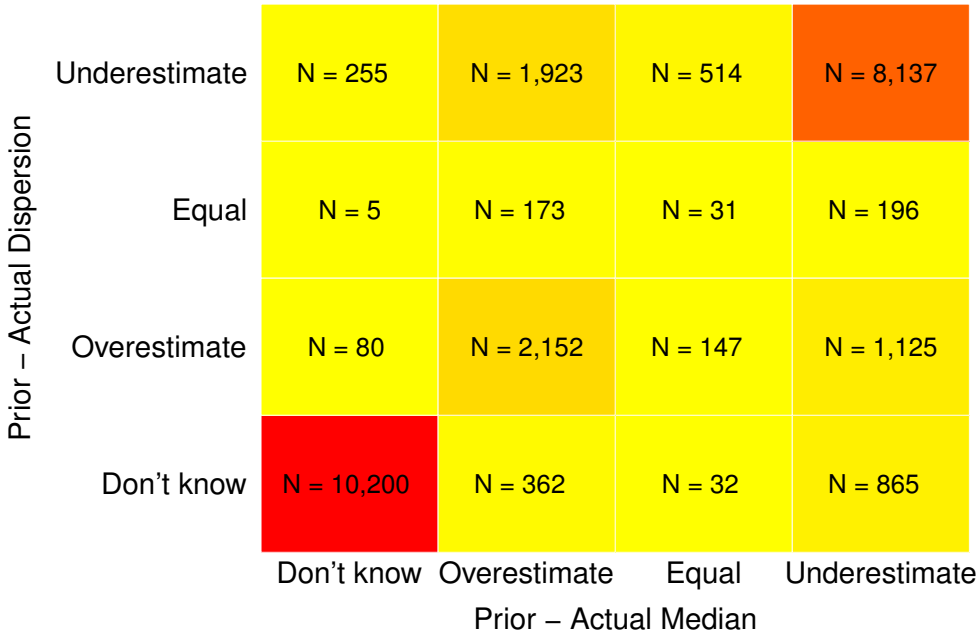
This figure shows the cumulative distribution function (CDF) of the number of institutions searched by treatment arm. N =3,283. The Kolmogórov-Smirnov (KS) test is estimated comparing each tool treatment to the control group and is calculated by Monte Carlo simulations with 10,000 replications.

Figure B.12: Timing of Loan Take-Up by Survey Response



This figure shows the proportion of people who obtained a loan in the administrative data within a certain number of days of participating in the RCT (up to 30), conditional on obtaining a loan. Results are shown separately by those who responded to and did not respond to the follow-up survey, among the subset we attempted to survey.

Figure B.13: Belief Heterogeneity



This figure shows participants' beliefs about interest rate dispersion and the interest rate they expected to receive, compared to actual values observed in administrative data, conditional on each participant's characteristics. *Prior – actual median* is the difference between a participant's beliefs about the interest rate they will get on the loan they take out (as reported in our baseline survey) and the median rate we observe based on their characteristics in the administrative data (i.e., the median they would have seen in the price comparison tool if assigned to that treatment arm). *Prior – actual dispersion* is the difference between a participant's beliefs about the dispersion in interest rates that a bank could offer them—measured as the highest rate minus the lowest rate they believed they could be offered—and the actual dispersion we observe based on their characteristics in the administrative data. We allow for a tolerance of ± 1 percentage point when defining equality, as it is unlikely participants would report exact matches with the administrative data. Each cell reports the number of participants in each category. Colors reflect cell frequencies, with red indicating the cell with the highest numbers of observations. The sample includes participants who made it far enough in the baseline survey to be assigned to a tool treatment arm or to the control group.

Table B.1: Bank Website and Third-Party Comparison Tool Inputs

	Borrower characteristics						Loan characteristics			
	Name (1)	RUT (2)	Document number (3)	Income (4)	Phone number/email (5)	Comuna (6)	Loan amount (7)	Maturity (8)	First payment date (9)	Insurance options (10)
Panel A: Bank websites										
Bank 1		Y		Y			Y	Y	Y	Y
Bank 2		Y					Y	Y	Y	Y
Bank 3		Y		Y			Y	Y	Y	Y
Bank 4	Y	Y		Y	Y		Y	Y	Y	Y
Bank 5		Y		Y			Y	Y	Y	Y
Bank 6							Y	Y		Y
Bank 7		Y					Y	Y	Y	Y
Bank 8		Y	Y		Y		Y	Y	Y	Y
Bank 9		Y								
Bank 10		Y					Y	Y	Y	Y
Bank 11	Y	Y		Y	Y		Y	Y		Y
Bank 12		Y			Y		Y	Y		
Panel B: Comparison Tools										
Third party comparison tool							Y	Y		
Public comparison tool							Y	Y		

This table shows what inputs are required by each bank website's consumer loan simulator and each third-party comparison tool as of April 3, 2024. Bank names are anonymized. Column (1), "Name", refers to the name of the person searching for information. Column (2), "RUT", refers to the *rol único tributario*, the national ID number in Chile. Column (3), "Document number", refers to the serial number on the national identity card which is distinct from the national ID number or RUT. Column (6), "Comuna", is a geographic area analogous to a neighborhood; we include this column to emphasize that no banks or third-party comparison tools request this information, despite it being an important predictor of interest rates used by banks in their algorithms. Column (9), "First payment date", refers to a field asking the consumer when they would like to make their first payment, which sometimes allows a specific day and other times a specific month beginning in the month after the current month.

Table B.2: Follow-Up Survey Response

	Pr(answer the survey)	
	(1)	(2)
(Intercept)	0.154***	0.150***
	(0.004)	(0.003)
Simple Tool	−0.004	
	(0.006)	
Price Comparison Tool	−0.006	
	(0.006)	
Elicit Beliefs		0.003
		(0.004)
Observations	21,007	37,664

This table tests for differential response rates to the follow-up survey by tool treatment status and by elicit beliefs treatment status. It uses specifications (2) and (3), where y_i is a dummy variable equal to 1 if participant i responded to the follow-up survey. The sample is restricted to participants whom we attempted to contact in the follow-up survey, and column (1) is further restricted to participants who made it far enough in the baseline survey to be assigned to a tool treatment arm or the control group. We define answering the survey as successfully completing (i.e., reaching the end of) the survey. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.3: Correlation between Beliefs and Monthly Median Rate in Administrative Data

	Expected rate	Lowest rate	Highest rate	Dispersion
	(1)	(2)	(3)	(4)
Median Rate$_t$	1.33***	1.29***	3.25***	1.67***
	(0.15)	(0.12)	(0.31)	(0.19)
Observations	16,015	15,875	15,618	15,045

This table shows the correlation between beliefs and the monthly median rate during the month in which a consumer participated. Coefficients are from a regression of the participant i's belief (measured in percentage points) against the median rate during the month t in which participant i participated in the RCT. The monthly median rate is calculated using the universe of consumer loans in administrative data for each month. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.4: Effects of Tools on Beliefs, Controlling for Priors

	Expected rate (1)	Lowest rate (2)	Highest rate (3)	Dispersion (4)
Prior	0.80***	0.80***	0.68***	0.52***
	(0.02)	(0.02)	(0.02)	(0.02)
Simple Tool	−0.26	0.44	1.07	−1.96**
	(0.86)	(0.70)	(1.46)	(0.80)
Price Comparison Tool	18.94***	14.36***	39.16***	20.61***
	(1.55)	(1.23)	(2.93)	(1.78)
Observations	6,817	6,760	6,661	6,272
Control Mean Posterior	29.22	22.65	47.45	23.18
Control Median Posterior	18	12	25	10.68
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows the effect of the simple tool and price comparison tool on beliefs. We estimate the following specification: $Posterior_i = \theta Prior_i + \beta_1 \mathbb{1}(\text{Simple Tool})_i + \beta_2 \mathbb{1}(\text{Price Comparison Tool})_i + \lambda_{b(i)} + \varepsilon_i$, where $Prior_i$ is the interest rate expectation participant i reported prior to seeing one of the tools or the control video, $Posterior_i$ is the interest rate expectation they reported after seeing it, and $\lambda_{b(i)}$ are bin density fixed effects. Each column shows θ, β_1, and β_2 for one of the following outcome variables: (1) the interest rate the participant expects to get on the loan they take out; (2) the lowest interest rate the participant expects a bank could offer them; (3) the highest interest rate the participant expects a bank could offer them; (4) the difference between the highest and the lowest interest rates. The sample sizes for each column differ based on the number of participants who responded to the corresponding questions. The bin density fixed effects (FEs) are deciles of how many loans are in the price comparison tool histogram that would have been shown to the participant based on their borrower and loan characteristics, had they been assigned to the price comparison tool arm. If any of the participant's loan amount, maturity, or income is outside of the support of the administrative data, the bin density variable would be a missing value; for these individuals, we create an additional fixed effect category that we include in the regression, i.e., we include an additional dummy equal to 1 if the bin density is a missing value. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values) by treatment arm. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.5: Effects of Tools on Beliefs, without Subtracting or Controlling for Priors

	Expected rate	Lowest rate	Highest rate	Dispersion
	(1)	(2)	(3)	(4)
Simple Tool	−1.01	−0.02	−0.88	−2.95***
	(1.19)	(0.95)	(2.00)	(0.98)
Price Comparison Tool	22.13***	17.22***	43.87***	23.38***
	(1.83)	(1.48)	(3.38)	(1.93)
Observations	7,792	7,640	7,533	7,321
Control Mean Posterior	29.91	22.94	48.12	23.72
Control Median Posterior	17.88	12	25	10
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows the effect of the simple tool and price comparison tool on beliefs. We estimate the following specification: $Posterior_i = \beta_1 \mathbb{1}(\text{Simple Tool})_i + \beta_2 \mathbb{1}(\text{Price Comparison Tool})_i + \lambda_{b(i)} + \varepsilon_i$, where $Posterior_i$ is the interest rate expectation participant *i* reported prior to seeing one of the tools or the control video, and $\lambda_{b(i)}$ are bin density fixed effects. Each column shows β_1 and β_2 for one of the following outcome variables: (1) the interest rate the participant expects to get on the loan they take out; (2) the lowest interest rate the participant expects a bank could offer them; (3) the highest interest rate the participant expects a bank could offer them; (4) the difference between the highest and the lowest interest rates the participant expects a bank could offer them. The sample sizes for each column differ based on the number of participants who responded to the corresponding questions. The bin density fixed effects (FEs) are deciles of how many loans are in the price comparison tool histogram that would have been shown to the participant based on their borrower and loan characteristics, had they been assigned to the price comparison tool arm. If any of the participant's loan amount, maturity, or income is outside of the support of the administrative data, the bin density variable would be a missing value; for these individuals, we create an additional fixed effect category that we include in the regression, i.e., we include an additional dummy equal to 1 if the bin density is a missing value. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values) by treatment arm. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.6: Effects of Tools on Beliefs in Logs, Controlling for Priors

	log(Expected rate) (1)	log(Lowest rate) (2)	log(Highest rate) (3)	log(Dispersion) (4)
log(Prior)	0.695***	0.701***	0.684***	0.578***
	(0.010)	(0.010)	(0.010)	(0.013)
Simple Tool	−0.038*	−0.008	−0.041*	−0.091***
	(0.023)	(0.023)	(0.024)	(0.032)
Price Comparison Tool	0.315***	0.273***	0.367***	0.335***
	(0.028)	(0.027)	(0.029)	(0.038)
Observations	6,817	6,760	6,661	6,272
Control Mean Posterior (Levels)	29.22	22.65	47.45	23.18
Control Median Posterior (Levels)	18	12	25	10.68
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows the effect of the simple tool and price comparison tool on beliefs. We estimate the following specification: $\log(Posterior_i + 1) = \theta \log(Prior_i + 1) + \beta_1 \mathbb{1}(\text{Simple Tool})_i + \beta_2 \mathbb{1}(\text{Price Comparison Tool})_i + \lambda_{b(i)} + \varepsilon_i$, where $Prior_i$ is the interest rate expectation participant i reported prior to seeing one of the tools or the control video, $Posterior_i$ is the interest rate expectation they reported after seeing it, and $\lambda_{b(i)}$ are bin density fixed effects. To account for any rate expectations of 0 (which do occur in the data), the transformation $\log(y_i + 1)$ was applied to the annualized interest rates in levels (where, for example, an 18% expected annual interest rate would be coded as 18). Each column shows θ, β_1, and β_2 for one of the following outcome variables: (1) the interest rate the participant expects to get on the loan they take out; (2) the lowest interest rate the participant expects a bank could offer them; (3) the highest interest rate the participant expects a bank could offer them; (4) the difference between the highest and the lowest interest rates. The sample sizes for each column differ based on the number of participants who responded to the corresponding questions. The bin density fixed effects (FEs) are deciles of how many loans are in the price comparison tool histogram that would have been shown to the participant based on their borrower and loan characteristics, had they been assigned to the price comparison tool arm. If any of the participant's loan amount, maturity, or income is outside of the support of the administrative data, the bin density variable would be a missing value; for these individuals, we create an additional fixed effect category that we include in the regression, i.e., we include an additional dummy equal to 1 if the bin density is a missing value. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values) by treatment arm. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.7: Effects of Tools on Beliefs in Logs, without Subtracting or Controlling for Priors

	log(Expected rate) (1)	log(Lowest rate) (2)	log(Highest rate) (3)	log(Dispersion) (4)
Simple Tool	−0.057*	−0.031	−0.066*	−0.128***
	(0.033)	(0.032)	(0.035)	(0.039)
Price Comparison Tool	0.407***	0.376***	0.459***	0.398***
	(0.034)	(0.034)	(0.036)	(0.043)
Observations	7,792	7,640	7,533	7,321
Control Mean Posterior (Levels)	29.91	22.94	48.12	23.72
Control Median Posterior (Levels)	17.88	12	25	10
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows the effect of the simple tool and price comparison tool on beliefs. We estimate the following specification: $\log(Posterior_i + 1) = \beta_1 \mathbb{1}(\text{Simple Tool})_i + \beta_2 \mathbb{1}(\text{Price Comparison Tool})_i + \lambda_{b(i)} + \varepsilon_i$, where $Posterior_i$ is the interest rate expectation participant i reported prior to seeing one of the tools or the control video, and $\lambda_{b(i)}$ are bin density fixed effects. To account for any rate expectations of 0 (which do occur in the data), the transformation $\log(y_i + 1)$ was applied to the annualized interest rates in levels (where, for example, an 18% expected annual interest rate would be coded as 18). Each column shows β_1 and β_2 for one of the following outcome variables: (1) the interest rate the participant expects to get on the loan they take out; (2) the lowest interest rate the participant expects a bank could offer them; (3) the highest interest rate the participant expects a bank could offer them; (4) the difference between the highest and the lowest interest rates the participant expects a bank could offer them. The sample sizes for each column differ based on the number of participants who responded to the corresponding questions. The bin density fixed effects (FEs) are deciles of how many loans are in the price comparison tool histogram that would have been shown to the participant based on their borrower and loan characteristics, had they been assigned to the price comparison tool arm. If any of the participant's loan amount, maturity, or income is outside of the support of the administrative data, the bin density variable would be a missing value; for these individuals, we create an additional fixed effect category that we include in the regression, i.e., we include an additional dummy equal to 1 if the bin density is a missing value. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values) by treatment arm. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.8: Effects of Tools on Beliefs about Normalized Dispersion

	Normalized Dispersion (1)
Simple Tool	−0.02
	(0.01)
Price Comparison Tool	0.03***
	(0.01)
Observations	6,272
Control Mean Posterior	0.67
Control Median Posterior	0.67
Bin Density FEs	Yes

This table shows the effect of the simple tool and price comparison tool on beliefs. It shows results from specification (1). Normalized dispersion is measured as the highest rate minus the lowest rate divided by the midpoint of the highest rate and lowest rate. The bin density fixed effects (FEs) are deciles of how many loans are in the price comparison tool histogram that would have been shown to the participant based on their borrower and loan characteristics, had they been assigned to the price comparison tool arm. If any of the participant's loan amount, maturity, or income is outside of the support of the administrative data, the bin density variable would be a missing value; for these individuals, we create an additional fixed effect category that we include in the regression, i.e., we include an additional dummy equal to 1 if the bin density is a missing value. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values) by treatment arm. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.9: Balance by Tool Treatment Arm for Survey Subsample in Table 4

		Difference relative to control mean			
	Control Mean	Price Comparison Tool	Simple Tool	Joint test F-stat	N
	(1)	(2)	(3)	(4)	(5)
Personal characteristics					
Age	36.533***	-0.374	0.094	0.669	3,253
	(0.304)	(0.426)	(0.434)	[0.512]	
log(Income)	13.546***	-0.001	0.006	0.012	3,200
	(0.033)	(0.051)	(0.049)	[0.988]	
Incomplete high-school	0.026***	-0.002	-0.001	0.053	3,176
	(0.005)	(0.007)	(0.007)	[0.948]	
Complete high-school	0.381***	-0.027	-0.019	0.88	3,176
	(0.015)	(0.021)	(0.021)	[0.415]	
Complete 2-year program	0.205***	0.037**	0.012	2.119	3,176
	(0.012)	(0.018)	(0.018)	[0.12]	
Complete 5-year program or higher	0.388***	-0.008	0.008	0.288	3,176
	(0.015)	(0.021)	(0.021)	[0.749]	
Financial products					
Bank account	0.648***	0.019	0.026	0.83	3,120
	(0.015)	(0.021)	(0.021)	[0.436]	
Any loan	0.698***	0.031	-0.003	1.868	3,147
	(0.014)	(0.020)	(0.020)	[0.155]	
Desired loan characteristics					
log(Loan Amount)	14.981***	0.033	0.016	0.153	3,083
	(0.041)	(0.059)	(0.058)	[0.859]	
log(Maturity (years))	1.361***	0.008	0.000	0.065	2,945
	(0.019)	(0.027)	(0.027)	[0.937]	
Omnibus F-statistic					
Price Comparison Tool		0.973			2,150
		[0.481]			
Simple Tool			1.34		2,194
			[0.169]		
Number of participants by arm	1,091	1,059	1,103		3,253

This table tests the balance of pre-treatment characteristics across treatment arms for the survey subsample in Table 4. We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta_1 \mathbb{1}(\text{Simple Tool})_i + \beta_2 \mathbb{1}(\text{Price Comparison Tool})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Simple Tool})_i$ and $\mathbb{1}(\text{Price Comparison Tool})_i$ are dummies indicating whether participant i was assigned to the simple tool or price comparison tool arms, respectively. Column (1) shows α which is the mean for the control group. Column (2) shows β_2 which is the difference in means between the price comparison tool treatment arm and the control group. Column (3) shows β_1 which is the difference in means between the simple tool treatment arm and the control group. Column (4) shows the F-statistic from a joint test of $\beta_1 = \beta_2 = 0$. Column (5) shows the number of observations in each regression, which can change across covariates due to missing values when the respondent did not answer that question in the baseline survey. For the omnibus F-test to test whether the covariates jointly predict treatment, we run the following regression for the participants assigned to simple tool and control groups and for those assigned to the price comparison tool and control groups, separately: $\mathbb{1}(\text{Treatment})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$, where $\mathbb{1}(\text{Treatment})_i$ is either $\mathbb{1}(\text{Simple Tool})_i$ or $\mathbb{1}(\text{Price Comparison Tool})_i$. The omnibus F-statistic is a test of $\gamma_1 = \cdots = \gamma_K = 0$. To retain the full sample in the omnibus F-test, if a participant did not answer a particular question, we create a dummy variable indicating whether the variable was missing, replace the missing value with zero, and include the missing value dummy as an additional X^k covariate in the regression. The p-values of the joint test and omnibus F-statistics are included in square brackets. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values). "Desired loan characteristics" refer to characteristics of the loan the participant is searching for. "Any loan" is a dummy variable equal to 1 if the participant had any of the following types of loan: consumer loan, mortgage, auto loan, credit card, or cash advance. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.10: Balance by Tool Treatment Arm for Subsample Obtaining Loans in Table 4

		Difference relative to control mean			
	Control Mean	Price Comparison Tool	Simple Tool	Joint test F-stat	N
	(1)	(2)	(3)	(4)	(5)
Personal characteristics					
Age	35.082***	0.042	0.245	0.711	8,988
	(0.156)	(0.219)	(0.220)	[0.491]	
log(Income)	13.905***	0.019	0.022	0.777	8,865
	(0.014)	(0.019)	(0.019)	[0.46]	
Incomplete high-school	0.007***	0.002	0.001	0.26	8,832
	(0.002)	(0.002)	(0.002)	[0.771]	
Complete high-school	0.243***	-0.009	-0.015	0.978	8,832
	(0.008)	(0.011)	(0.011)	[0.376]	
Complete 2-year program	0.205***	0.011	0.015	1.023	8,832
	(0.008)	(0.011)	(0.011)	[0.359]	
Complete 5-year program or higher	0.544***	-0.003	0.000	0.039	8,832
	(0.009)	(0.013)	(0.013)	[0.962]	
Financial products					
Bank account	0.863***	0.016*	0.004	1.775	8,844
	(0.006)	(0.009)	(0.009)	[0.17]	
Any loan	0.882***	0.000	0.003	0.069	8,875
	(0.006)	(0.008)	(0.008)	[0.934]	
Desired loan characteristics					
log(Loan Amount)	15.428***	0.054*	0.041	1.779	8,604
	(0.021)	(0.030)	(0.030)	[0.169]	
log(Maturity (years))	1.425***	0.039**	0.021	3.268**	8,373
	(0.011)	(0.015)	(0.015)	[0.038]	
Omnibus F-statistic					
Price Comparison Tool		1.149			5,982
		[0.305]			
Simple Tool			0.661		5,928
			[0.825]		
Number of participants by arm	2,922	3,060	3,006		8,988

This table tests the balance of pre-treatment characteristics across treatment arms for the subsample obtaining loans in the administrative data in Table 4, column (9). We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta_1 \mathbb{1}(\text{Simple Tool})_i + \beta_2 \mathbb{1}(\text{Price Comparison Tool})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Simple Tool})_i$ and $\mathbb{1}(\text{Price Comparison Tool})_i$ are dummies indicating whether participant i was assigned to the simple tool or price comparison tool arms, respectively. Column (1) shows α which is the mean for the control group. Column (2) shows β_2 which is the difference in means between the price comparison tool treatment arm and the control group. Column (3) shows β_1 which is the difference in means between the simple tool treatment arm and the control group. Column (4) shows the F-statistic from a joint test of $\beta_1 = \beta_2 = 0$. Column (5) shows the number of observations in each regression, which can change across covariates due to missing values when the respondent did not answer that question in the baseline survey. For the omnibus F-test to test whether the covariates jointly predict treatment, we run the following regression for the participants assigned to simple tool and control groups and for those assigned to the price comparison tool and control groups, separately: $\mathbb{1}(\text{Treatment})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$, where $\mathbb{1}(\text{Treatment})_i$ is either $\mathbb{1}(\text{Simple Tool})_i$ or $\mathbb{1}(\text{Price Comparison Tool})_i$. The omnibus F-statistic is a test of $\gamma_1 = \cdots = \gamma_K = 0$. To retain the full sample in the omnibus F-test, if a participant did not answer a particular question, we create a dummy variable indicating whether the variable was missing, replace the missing value with zero, and include the missing value dummy as an additional X^k covariate in the regression. The p-values of the joint test and omnibus F-statistics are included in square brackets. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values). "Desired loan characteristics" refer to characteristics of the loan the participant is searching for. "Any loan" is a dummy variable equal to 1 if the participant had any of the following types of loan: consumer loan, mortgage, auto loan, credit card, or cash advance. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.11: Balance by Elicit Beliefs Treatment for Survey Subsample in Table 6

	Elicit Beliefs = 0 Mean	Elicit Beliefs	N
	(1)	(2)	(3)
Personal characteristics			
Age	36.822***	-0.307	5,729
	(0.251)	(0.294)	
log(Income)	13.589***	0.035	5,624
	(0.032)	(0.037)	
Incomplete high-school	0.028***	0.000	5,592
	(0.004)	(0.005)	
Complete high-school	0.348***	-0.014	5,592
	(0.012)	(0.014)	
Complete 2-year program	0.210***	-0.001	5,592
	(0.010)	(0.012)	
Complete 5-year program or higher	0.414***	0.015	5,592
	(0.013)	(0.015)	
Financial products			
Bank account	0.682***	0.009	5,491
	(0.012)	(0.014)	
Any loan	0.738***	-0.016	5,538
	(0.011)	(0.013)	
Omnibus F-statistic		0.959	5,729
		[0.482]	
Number of participants by arm	1,563	4,166	5,729

This table tests the balance of pre-treatment characteristics by elicit beliefs treatment for the survey subsample in Table 6. We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta \mathbb{1}(\text{Elicit Beliefs})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Elicit Beliefs})_i$ is a dummy indicating whether participant i was assigned to the elicit beliefs treatment. Column (1) shows α which is the mean for the $\mathbb{1}(\text{Elicit Beliefs})_i = 0$ group. Column (2) shows β which is the difference in means between the $\mathbb{1}(\text{Elicit Beliefs})_i = 1$ and $\mathbb{1}(\text{Elicit Beliefs})_i = 0$ groups. Column (3) shows the number of observations in each regression, which can change across covariates due to missing values when the respondent did not answer that question in the baseline survey. For the omnibus F-test to test whether the covariates jointly predict treatment, we run the following regression: $\mathbb{1}(\text{Elicit Beliefs})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$. The omnibus F-statistic is a test of $\gamma_1 = \cdots = \gamma_K = 0$. To retain the full sample in the omnibus F-test, if a participant did not answer a particular question, we create a dummy variable indicating whether the variable was missing, replace the missing value with zero, and include the missing value dummy as an additional X^k covariate in the regression. The p-value of the omnibus F-statistic is included in square brackets. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values). "Any loan" is a dummy variable equal to 1 if the participant had any of the following types of loan: consumer loan, mortgage, auto loan, credit card, or cash advance. Unlike in Table 2, we cannot include characteristics of the loan they are searching for in the balance tests since these questions were asked in the same module as the elicit beliefs treatment (rather than a prior module), and the elicit beliefs treatment affected whether participants continued in the survey. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.12: Balance by Elicit Beliefs Treatment for Subsample Obtaining Loans in Table 6

	Elicit Beliefs = 0 Mean (1)	Elicit Beliefs (2)	N (3)
Personal characteristics			
Age	35.179***	0.012	21,522
	(0.110)	(0.127)	
log(Income)	14.040***	0.004	21,268
	(0.009)	(0.010)	
Incomplete high-school	0.007***	0.000	21,213
	(0.001)	(0.001)	
Complete high-school	0.207***	0.002	21,213
	(0.006)	(0.006)	
Complete 2-year program	0.200***	-0.002	21,213
	(0.005)	(0.006)	
Complete 5-year program or higher	0.585***	-0.001	21,213
	(0.007)	(0.008)	
Financial products			
Bank account	0.888***	0.007	21,238
	(0.004)	(0.005)	
Any loan	0.888***	-0.002	21,303
	(0.004)	(0.005)	
Omnibus F-statistic		0.401	21,522
		[0.956]	
Number of participants by arm	5,503	16,019	21,522

This table tests the balance of pre-treatment characteristics across treatment arms for the subsample obtaining loans in the administrative data in Table 4, column (9). We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta \mathbb{1}(\text{Elicit Beliefs})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Elicit Beliefs})_i$ is a dummy indicating whether participant i was assigned to the elicit beliefs treatment. Column (1) shows α which is the mean for the $\mathbb{1}(\text{Elicit Beliefs})_i = 0$ group. Column (2) shows β which is the difference in means between the $\mathbb{1}(\text{Elicit Beliefs})_i = 1$ and $\mathbb{1}(\text{Elicit Beliefs})_i = 0$ groups. Column (3) shows the number of observations in each regression, which can change across covariates due to missing values when the respondent did not answer that question in the baseline survey. For the omnibus F-test to test whether the covariates jointly predict treatment, we run the following regression: $\mathbb{1}(\text{Elicit Beliefs})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$. The omnibus F-statistic is a test of $\gamma_1 = \cdots = \gamma_K = 0$. To retain the full sample in the omnibus F-test, if a participant did not answer a particular question, we create a dummy variable indicating whether the variable was missing, replace the missing value with zero, and include the missing value dummy as an additional X^k covariate in the regression. The p-value of the omnibus F-statistic is included in square brackets. Continuous variables are winsorized at the 95th percentile (and 5th percentile if they take negative values). "Any loan" is a dummy variable equal to 1 if the participant had any of the following types of loan: consumer loan, mortgage, auto loan, credit card, or cash advance. Unlike in Table 2, we cannot include characteristics of the loan they are searching for in the balance tests since these questions were asked in the same module as the elicit beliefs treatment (rather than a prior module), and the elicit beliefs treatment affected whether participants continued in the survey. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix C Bank Websites and Comparison Websites

In our follow-up survey data, 43.8% of participants report using tools on bank websites known as "simulators" that provide an interest rate estimate and 12% report using third-party comparison websites (aggregators) during their search. Thus, these channels are likely a way that some consumers form their beliefs about loan interest rates. We investigate whether the simulator tools on banks' own websites and the main third-party comparison tool provide accurate information. If not, these sources potentially contribute to participants' biased beliefs. We scraped data from seven banks' consumer loan simulators and the main third-party comparison website. For each participant in our RCT, we obtained data from their baseline survey of individual and desired loan characteristics. We then ran a script that fed these inputs into each website (including bank simulators and the third-party comparison website) and scraped the output. Next, we compared the rates participants would have seen on these websites with the rates they actually received in the administrative data.

C.1 Description of Bank Websites

Many Chilean banks provide a "simulator" on their websites, which allows visitors to see what interest rate they could expect to receive on a loan. Consumers input their personal information along with desired loan amount, terms, and other details. The simulator then shows the consumer the interest rate they can expect to obtain for this loan, as well as the total loan cost and monthly cost. The input variables required by each simulator are shown in Table B.1, panel A. All bank websites require information on loan amount and maturity. All but one bank also require the consumer's RUT (national ID number), but in tests that we show below, we find that the interest rates that banks show in these simulators typically do not vary based on the RUT that is entered. Five out of twelve bank websites require the consumers' income as an input. On the other hand, none of them require the users to enter their neighborhood of residence (*comuna*), despite this being an important variable that banks use to price loans.

We were able to scrape data from seven of the twelve bank simulators. The simulators that we were not able to scrape were due to firewalls, returning errors when attempting to obtain a quote, and requiring the user to have a bank account already with that bank or to pay for the quote.

C.2 Description of Comparison Websites

There are two main third-party or government-run comparison websites providing estimated loan terms from multiple banks, also known as aggregators. One is provided by a private company, using information reported to the third-party comparison tool by banks, which report the rates they

would offer for different loan types. Banks may have an incentive, however, to report downward-biased quotes to comparison websites as a bait-and-switch technique, as putting lower rates on comparison websites can direct traffic to them over other banks. Table B.1, panel B, describes the inputs required by this comparison website as well as a comparison website run by a different government agency. We scraped the third-party comparison website but not the other government comparison website as the latter was down for several months when we were running our scrapers.

C.3 Obtaining Data from Bank and Comparison Websites

We use the loan and consumer characteristics of each consumer-loan seeker in the baseline survey as input to the simulators, thereby replicating what our survey respondents would see should they use these tools. For identification-related inputs, such as RUT (national ID number) and contact information, we use random fake RUT numbers generated by adapting the code at `https://codepen.io/alisteroz/pen/KEoqgQ` for Python. To test whether the outputs shown by the bank websites depend on the RUT entered, we conducted tests where we held all inputs fixed except RUT. In these tests, we set the other characteristics such as loan amount, maturity, and income to the median values observed in our data and hold them constant. We set the test size to 100 observations and tested the five bank websites we scraped where IDs were requested as an input. As shown by Figure C.1, despite occasional variations in interest rates for different RUTs for two rates, the annualized interest rates remain largely identical across a random sample of RUTs. Our data collection period spanned from September 28th, 2023 to October 9th, 2023.

Similarly, four bank websites (three of which we could successfully scrape) require the phone number as an input. We conduct a similar test of whether the interest rates shown by the bank depend on the phone number (e.g., the bank might use the phone number's area code and condition the interest rate on where the consumer lives) by randomly generating phone numbers and again testing 100 observations where other inputs are held fixed. Figure C.2 shows that interest rates do not differ by area code for any of the three banks that require phone number as an input.

Many simulators provide users with the flexibility to select their preferred grace period (i.e. difference between loan origination and first payment date) and insurance options. These choices do not influence the interest rate of the loan, but they impact the CAE (APR) and the total loan cost. Since we did not ask about the preferred grace period or insurance options in the baseline survey (as many respondents would not have known how to respond to these questions), we extract a range of CAEs (APRs) that the user might have seen based on different inputs. In particular, we choose the grace period and insurance option that would either minimize or maximize the CAE (APR) and total cost of the loan, holding other inputs constant. For example, opting for no grace period and declining all insurance resulted in the lowest APR and total loan cost, while choosing

the longest grace period and all available insurance yielded the highest APR and total loan cost. Nevertheless, because we observe interest rate (rather than CAE/APR) in the administrative data, the interest rate is the more relevant output that we scrape, and the interest rate is not affected by the choice of grace period or insurance.

We obtain the following simulated loan outcomes for each participant in our RCT from each scraped bank website: monthly interest rate, equivalent annual charge (carga anual equivalente, or CAE, which is analogous to an APR), and total cost of the loan (costo total del crédito, or CTC).

C.4 Comparison of Websites' Rates and Received Rates

To compare the rate an individual in our RCT would have seen on bank and comparison websites to the rate that they actually received in the bank administrative data, we begin by matching the interest rates we scraped from these websites that correspond to what an individual RCT participant would have seen to the interest rates of the loans that these individuals actually received. First, we restrict our sample to the 30,979 people in our RCT who took out a loan. Next, for each of these participants, we keep one unique consumer loan from the administrative data, which is the first loan taken after treatment based on the date that the individual participated in the RCT and the date that the loan was taken out. In the rare case that the individual took two loans on the same day (0.55% of the sample), we take the largest loan taken on the first day after treatment that any loans were taken out by that individual. Then, for each individual who took up a consumer loan in the administrative banking data, we match interest rates the individual would have seen on the bank and comparison websites—based on the loan amount, loan maturity, and income they reported in the baseline survey—with the interest rate they obtained in the administrative data.

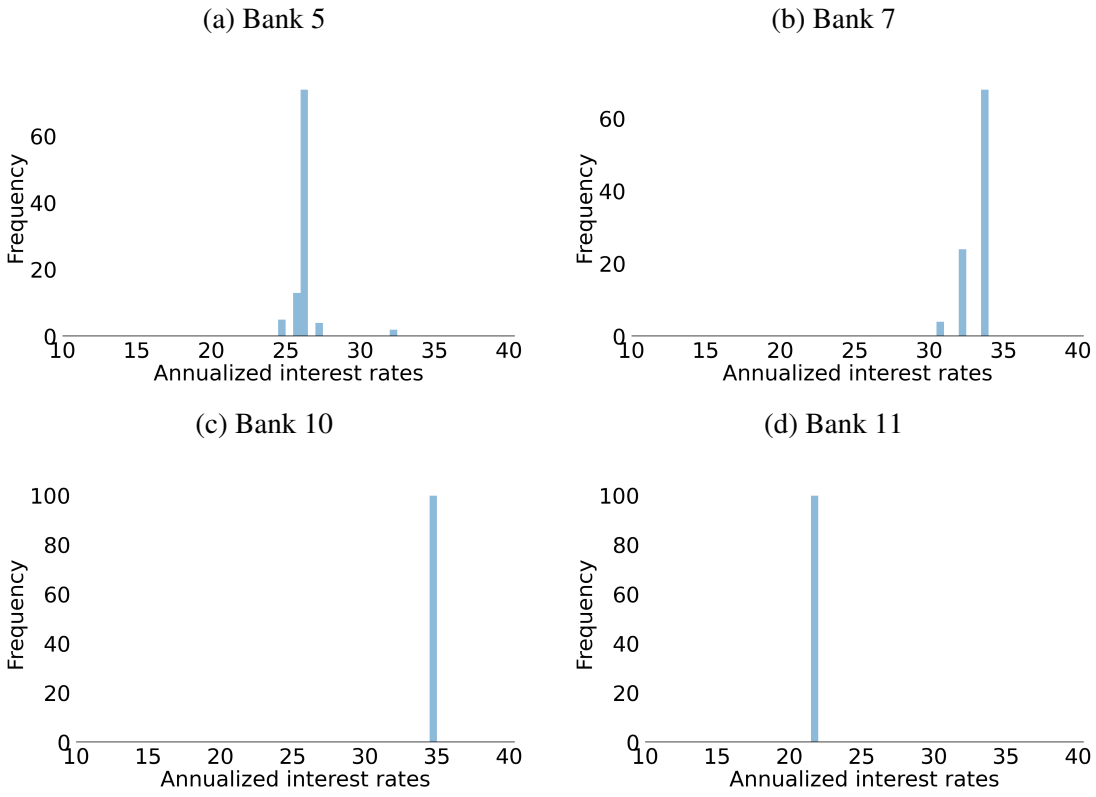
We merge at the consumer and bank level using encrypted bank identifiers in the administrative data: for example, if the individual took a loan from Bank A, we merge the interest rate they obtained with what they would have seen on Bank A's website, and what the comparison websites would have shown them for a loan from Bank A. If for a given loan, there are multiple matched interest rates quotes from different sources of the same bank, e.g., one from the third-party comparison tool and the other from the bank's website, we keep all quotes. The matched sample has 14,354 observations. We further restrict to loans in the administrative data that were obtained during the same months in which we ran the scrapers, i.e., September–October 2023. This limits the size of the sample but ensures that any observed differences between rates from bank and third-party comparison websites and rates obtained in administrative data are not due to changes in interest rates over time.

The interest rate quotes shown by banks and comparison websites are highly inaccurate (Figures 5b and 5c). This section documents possible explanations for this inaccuracy. The first and

most compelling explanation is that these websites do not ask the user for key inputs: none ask for the comuna of residence, and only three out of seven ask for income, both of which are significant predictors of interest rate. Thus, they do not provide quotes conditional on all the relevant borrower characteristics that influence the interest rate. Second, the Chilean credit bureau does not provide a continuous credit score; instead, they provide a binary flag for whether a borrower has defaulted on a loan in the past. This is a severe credit event and only happens if the borrower has missed three payments and judicial proceedings have been initiated against them. For the borrower, this flag effectively shuts them out of credit markets. Banks are able to create a proxy for credit risk by creating an average provision score across all banks reporting to the CMF. Each bank sets aside a certain fraction of the loan as revenues in case the borrower misses a payment or defaults as part of their risk management procedures (CMF, 2024). Borrowers are unaware of this number and while banks could pre-populate borrowers' risk scores by RUT in their simulators, in practice they do not.

An alternative explanation could be that the loan that the participants ultimately took out might have different terms (e.g., loan amount, maturity) than the terms the participant was initially looking for. To test this, we plot the observed difference between rates participants would have observed on the bank simulator and the rate they took the loan out at the same bank on the difference in loan size and maturity between the desired terms they reported on the baseline survey and the terms of the actual loan they obtained. The scatter plots are presented in Figure C.3 for banks and Figure C.4 for the third-party comparison tool. Loan size differences are presented in panel a and maturity differences in panel b. The majority of points are clustered along the vertical line at zero, indicating little difference in the other loan terms (loan amount or maturity). However, there is still substantial variation in interest rates received for a given value of the difference in other loan terms. We regress the difference in interest rates on the difference in loan terms and find that the R^2 of these regressions are 0.173 for loan amount and 0.265 for maturity for banks, and 0.086 for loan amount and 0.060 for maturity for the third-party comparison website. This suggests that the differences between rates seen on bank or third-party comparison websites and the rates participants actually obtained cannot primarily be explained by participants obtaining loans with different loan amount or maturity than they were initially searching for.

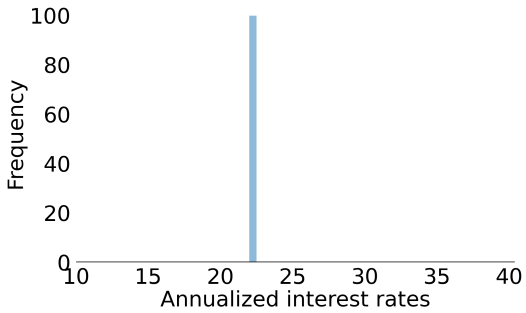
Figure C.1: Test of Whether Bank Simulator Terms Vary by RUT



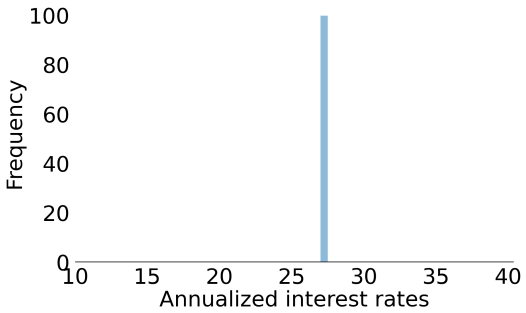
This figure shows the annualized interest rates from bank websites, derived from a test where we varied only the RUT while maintaining other inputs constant. Specifically, we standardized inputs such as income, loan amount, and loan maturity by using the median values from the baseline sample: a monthly income of $740,000 pesos, a loan amount of $2,500,000 pesos, and a maturity period of 3 years. The x-axis represents the annualized interest rates calculated by each bank's simulator tool. Each panel includes 100 observations.

Figure C.2: Test of Whether Bank Simulator Terms Vary by Phone Area Code

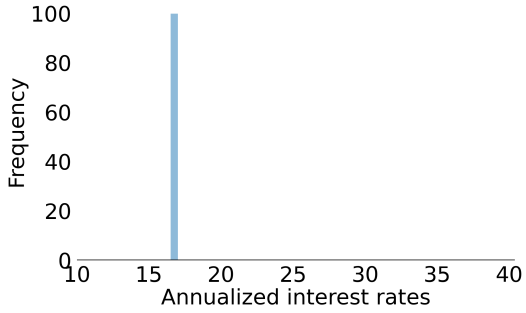
(a) Bank 11



(b) Bank 12



(c) Bank 8

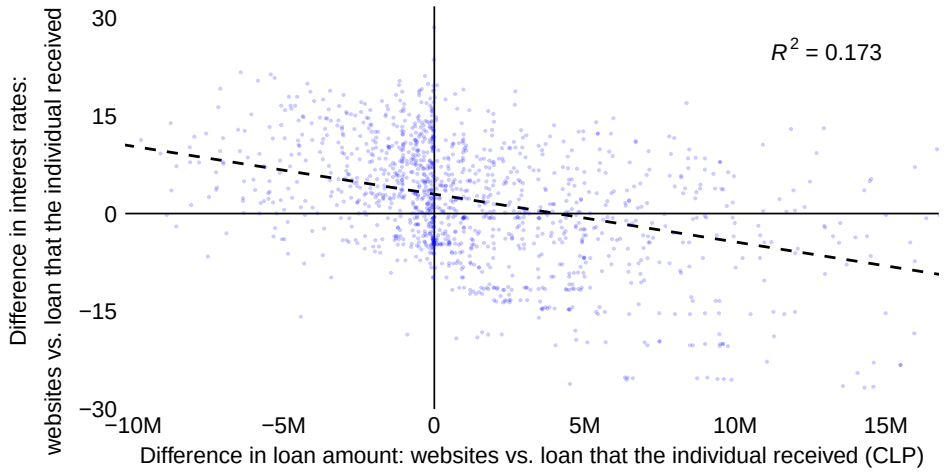


This figure shows the annualized interest rates from bank websites, derived from a test where we varied only the phone number area code while maintaining other inputs constant. Specifically, we standardized inputs such as income, loan amount, and loan maturity by using the median values from the baseline sample: a monthly income of $740,000 pesos, a loan amount of $2,500,000 pesos, and a maturity period of 3 years. The x-axis represents the annualized interest rates calculated by each bank's website simulation. Each simulation consists of 100 observations.

Figure C.3: Difference in Interest Rates Between Bank Websites and Rate the Individual Received

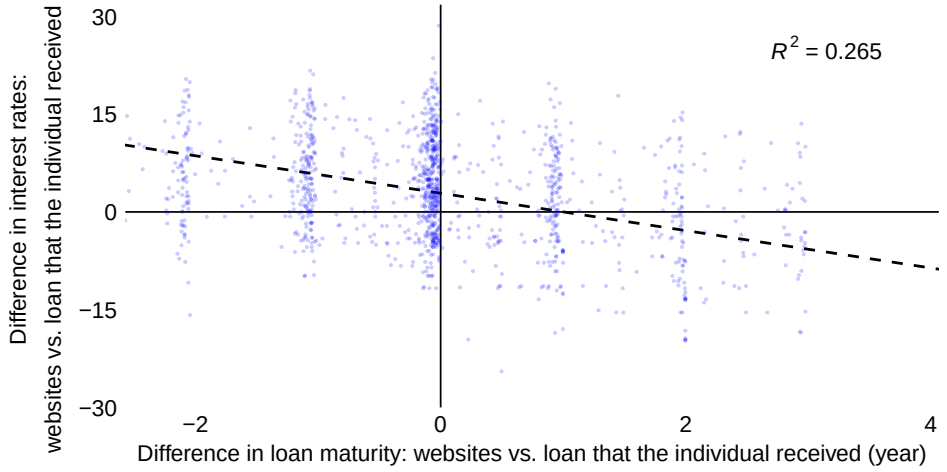
Difference in Interest Rates vs. Difference in Loan Amount

(a) Difference in Loan Amount



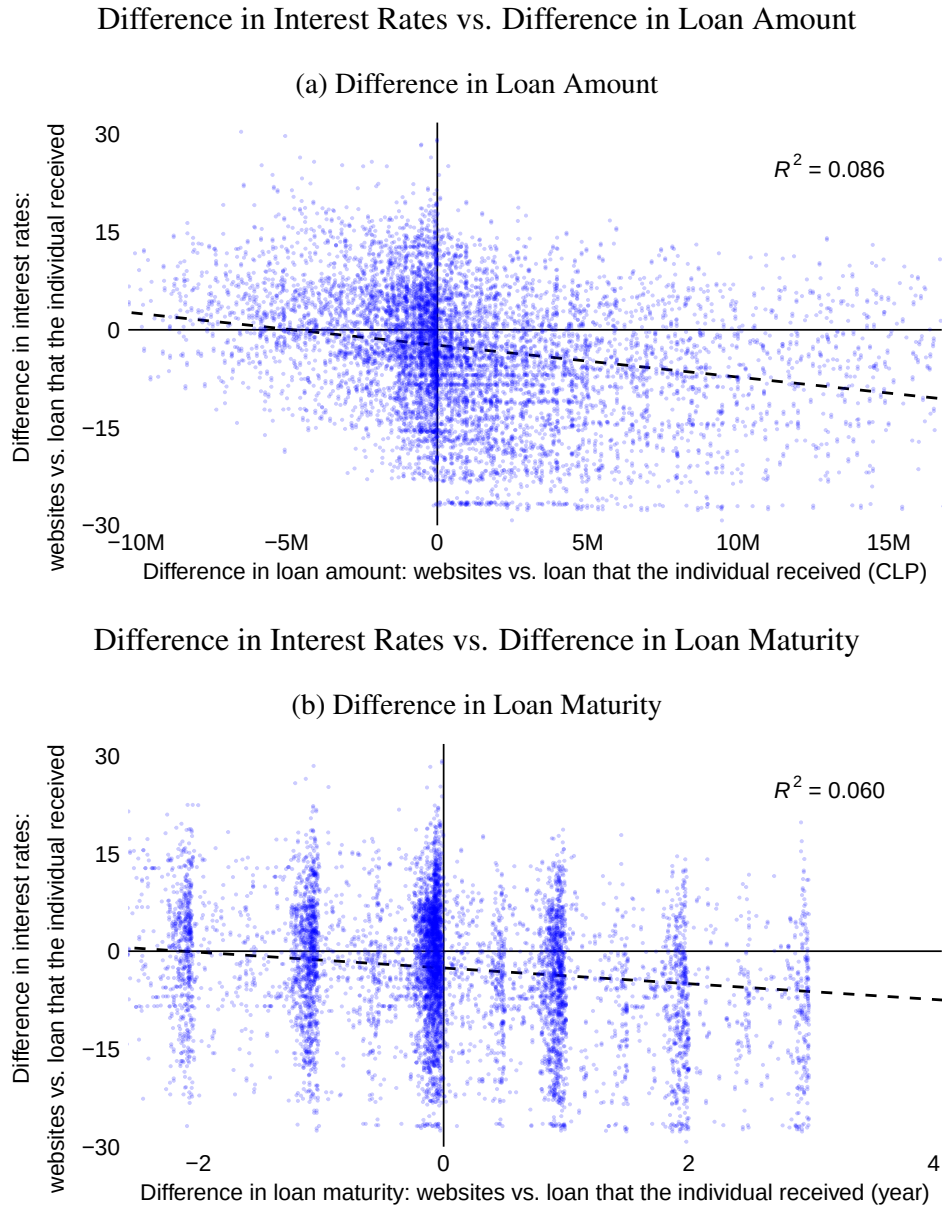
Difference in Interest Rates vs. Difference in Loan Maturity

(b) Difference in Loan Maturity



This figure shows the correlation of the difference in interest rates with the difference in loan amount (panel a) or the difference in loan maturity (panel b). Interest rates on bank websites are the rates displayed on consumer loan simulation websites, given an individual's baseline survey inputs. The difference in interest rates is calculated as: interest rates on websites - interest rates of the loan that the individual received. The difference in loan amount (maturity) is calculated as: loan amount (maturity) from the baseline survey - loan amount (maturity) of the loan that the individual received. The variables on the x-axis, difference in loan amount (maturity), are winsorized at both the 5th and the 95th percentiles within each treatment group. The dashed line is the line of best fit from a linear regression of the difference in interest rates on the winsorized difference in loan amount or loan maturity. The R^2 at the top-right corner of the plot corresponds to the line of best fit. The number of observations is 1,659. For legibility, the bottom and top 5th percentiles of difference in loan amount are excluded from panel a, and the bottom 5th and top 10th percentiles of difference in loan maturity are excluded from panel b.

Figure C.4: Difference in Interest Rates Between Third-Party Tool and Rate the Individual Received



This figure shows the correlation of the difference in interest rates with the difference in loan amount (panel a) or the difference in loan maturity (panel b). Interest rates on bank websites are the rates displayed on consumer loan simulation websites, given an individual's baseline survey inputs. The difference in interest rates is calculated as: interest rates on websites - interest rates of the loan that the individual received. The difference in loan amount (maturity) is calculated as: loan amount (maturity) from the baseline survey - loan amount (maturity) of the loan that the individual received. The variables on the x-axis, difference in loan amount (maturity), are winsorized at both the 5th and the 95th percentiles within each treatment group. The dashed line is the line of best fit from a linear regression of the difference in interest rates on the winsorized difference in loan amount or loan maturity. The R^2 at the top-right corner of the plot corresponds to the line of best fit. The number of observations is 12,695. For legibility, the bottom and top 5th percentiles of difference in loan amount are excluded from panel a, and the bottom 5th and top 10th percentiles of difference in loan maturity are excluded from panel b.

Appendix D Experimental Timeline

Pilot From November 2019 to March 2020, we conducted a randomized pilot of the loan price comparison tool. The pilot tested our online ads campaign, survey, and a prototype of the interactive tool. In our pilot, 8,916 Chilean users searching for keywords related to consumer loans clicked our ads.

Focus Groups In October and November 2020 we ran a series of focus groups and surveys, in which 81 Chilean participants provided us with qualitative comments on the user experience, design, and perceived usefulness of the price comparison tool and survey questionnaire. The results from these focus groups and individual interviews prompted us to add more guidance, tips, and directions to our tool, as some participants had trouble understanding the graphs and data.

Google ads testing From January to July 2021 we ran Google ads tests to determine the most effective ads.

RCT After making a series of improvements to the interactive price comparison tool based on our pilot experiment and focus groups, we launched the RCT in November 2021. We subsequently made monthly updates to the tool so that it always showed data based on the previous six months. One issue that arose during the RCT was that from January 24, 2022 to March 10, 2022, due to a bug in our code for the baseline survey, we mistakenly showed *all* participants the questions eliciting prior beliefs. As the timing of participation in the RCT is not random, we exclude these 4,979 observations from the regressions testing the effect of eliciting priors (Table 6), but keep these observations in regressions testing the effects of the tools. Thus, our experiment actually included 117,042 rather than 112,063 participants, with the difference representing those affected by this bug.

Appendix E Price Comparison and Simple Tool Construction

E.1 Price Comparison Tool

We detail how we processed the data from administrative, loan-level data from the CMF into histograms that participants saw in the price comparison tool treatment arm.

First, loans are merged with borrower characteristics at the time of loan application, also reported to the CMF. We then bin all loan-borrower observations within a neighborhood into income quartiles, loan amount quartiles, and one-year maturity bins, where the quartile cutoffs for both

income and loan amount vary by comuna. In order to ensure that the loan histograms do not contain identifiable information, we impose the following two conditions for a participant to view a histograms:

1. There are at least five data points in total and at least two distinct interest rates.
2. Both the participant's income and loan amount fall within extended ranges of the data to be shown to them (with extended ranges calculated as 75% of the minimum value to 125% of the maximum value, for both variables).

If the above two conditions are met, then the histogram is displayed as in Figure 2a, with the caveat that data points more than five standard deviations away from the interest rate mean are removed as extreme outliers. The share of comunas that had enough data to show comuna-level histograms was 87.5%. The share of participants that saw comuna-level histograms was 82.2%.

If one of these conditions fails, then the geographic range is expanded to contiguous comunas. Existing loans in the comuna and all contiguous comunas are again binned into income quartiles, loan amount quartiles, and one-year maturity bins. If the above two conditions are met for the histogram that corresponds to the original comuna and its contiguous comunas, a histogram is returned. A message also accompanies the histogram specifying that there was insufficient data for the participant's comuna, and that they are seeing information that includes neighboring comunas. The share of comunas that did not have enough data for comuna-level histograms but did had enough data to show first-degree neighbor comuna-level histograms was 10%. The share of participants that saw first-degree neighboring comuna histograms was 4.6%.

If one of the conditions still fails, then the geographic range is expanded to second-degree contiguous comunas and the steps above are repeated. The share of comunas that had enough data to show second-degree neighbor comuna-level histograms but not enough data for histograms at a smaller geographic level was 2%. The share of participants that saw second-degree comuna matching was 0.6%.

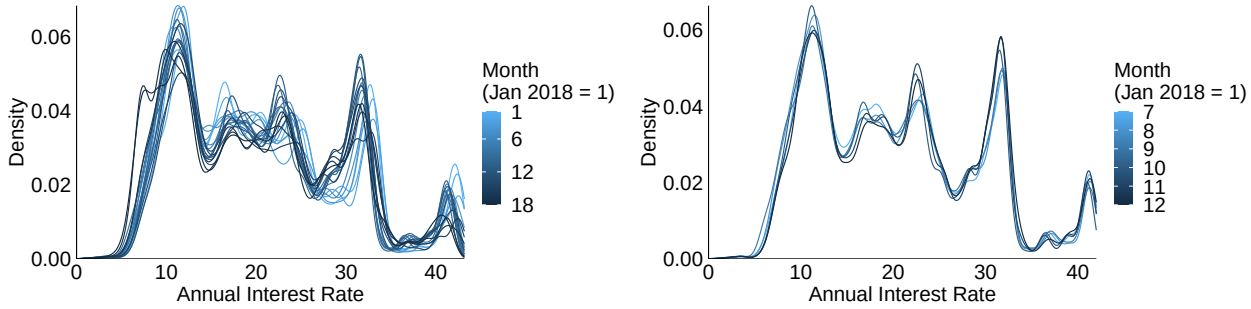
If a histogram still could not be created for a participant based on the above conditions, then data for the entire region (equivalent in Chile to a US state) is used with a corresponding message that the information displayed is at the regional level. The share of comunas that required regional data was 0.5%. The share of participants that saw regional data was 8.4%.

Length of Past Data Shown In order to provide consumers with personalized data based on loans actually obtained by similar consumers, we necessarily have to use "historical" data that goes a certain distance back in time. (For example, if we only used data from the past month, for many consumers there would not be enough observations within each cell defined by consumer and loan characteristics to show the distribution of rates.) Thus, we face a trade-off between how

recent the data used by the tool are—which is more relevant if the distribution of interest rates changes over time—and how many data points we show each consumer. We determined that using 18 months of data goes too far into the past given that the distribution of interest rates does change over time. In Figure E.1a below, we show the distribution of interest rates for consumer loans in each of the 18 months between January 2018 and June 2019.

On the other hand, we determined that the distribution of interest rates is relatively stable over six months, as shown in Figure E.1b. Furthermore, using data from the last six months still provides sufficient observations within each cell to show consumers a distribution of prices faced by similar consumers for similar loans. We thus refreshed the data underlying the tool each month so that the tool always showed the most recent six months of data.

Figure E.1: Distribution of Interest Rates over Time



(a) 18 Months

(b) 6 Months

This figure shows how the distribution of consumer loan interest rates changes over time. Panel (a) displays the distribution of interest rates over the 18 months from January 2018 to June 2019. Panel (b) shows the distribution of interest rates for the most recent six months from July 2018 to December 2019. It demonstrates that the overall distribution of interest rates is relatively stable within this shorter timeframe, allowing for a more accurate and up-to-date comparison for consumers.

E.2 Simple Tool

We use loan-level CMF data to estimate the benefits of searching at more banks. First, we subset the data to originated loans in a given comuna, income quartile, loan size quartile, and maturity bin over the last six months. This is equivalent to the data the participant would have seen if they were assigned to the price comparison tool. Within a given bin, there are J banks that have originated L loans to borrowers. We randomly draw an interest rate ℓ_0 from bank j and consider it the participant's "first offer." We then draw another quote ℓ_1 from the remaining $J-1$ banks. We then use the desired maturity and loan size to calculate the monthly interest payments of loan ℓ_0 and ℓ_1. We then consider two cases:

1. If $\ell_1 < \ell_0$, we calculate the difference in monthly payments, and calculate the difference in total loan costs as the present value of the difference in monthly payments over the life of

the loan.

2. If $\ell_1 \geq \ell_0$, we set the "benefit of searching at one additional bank" to zero.

We then repeat this drawing of two sequential quotes 1,000,000 times. We estimate the "benefit of searching at one additional bank" to be the average reduction in monthly cost and in total loan cost across these simulations. Thus, the benefit of searching is always a non-negative number, though zeros are included in the average.

To find the benefit of searching at $n \in \{2,3,4,5\}$ additional banks, we simulate the process for drawing n loans from $J-n$ banks 1,000,000 times and take the mean. All benefits are calculated by comparing the lowest interest rate ℓ_n from any draw $n > 0$ to the interest rate from the initial draw ℓ_0, and if $\ell_n \geq \ell_0$ for all $n > 0$ within a particular simulation, the benefit of searching is coded as zero in tht simulation.

We repeat this procedure for all bins determined by loan and borrower characteristics. As in the loan price comparison tool, if there are fewer than five loans fewer than two unique interest rates in the bin, we expand the bin to include comunas that border our reference comuna. If there are still less than five loans or two unique rates, we expand the bin again to include comunas that border the bordering comunas to the reference comuna.

Appendix F Data Details

F.1 Administrative Data

Annualizing Rates Survey participants could provide their interest rates received in the follow-up phone survey as an annual or a monthly percentage rate. About 34% of participants report annual rates, and 66% report monthly rates which we annualize. However, there were cases where participants appear to have responded with a monthly rate but stated that it was an annual rate. Figure F.1a shows the distribution of annualized interest rates based on participants' follow-up survey responses. About 32% of respondents gave annual rates below 10 percentage points; furthermore, the distribution of annualized reported rates is bimodal, and appears to combine one distribution of rates below 10% and another of rates above 10%. In contrast, the true distribution of annual interest rates in administrative data is unimodal and nearly all rates are above 10%. To correct for this inaccurate reporting, we assume that reported annual interest rates of 10% or lower are implausibly low, and should be recoded as monthly rates that are then converted to annual rates. Figure F.1b shows that the distribution of our corrected annualized interest rate distribution is much closer to the distribution of administrative interest rates.

Winsorizing We winsorize continuous variables at 95th percentile for variables that are bounded below at 0 and winsorize at the 5th and 95th percentile for variables that can take negative values. For variables measured after treatment, we winsorize within treatment arm to avoid truncating a true treatment effect. We winsorize the final variable used in our analysis; for example, if the outcome of interest is the difference between prior and posterior beliefs, we first take the difference of the raw variables and then winsorize the difference.

F.2 Survey Data

Survey Questionnaires English translations of our survey questionnaires are available at the following links:

1. Baseline survey: `https://seankhiggins.com/assets/pdf/pricecomp_rct/baseline_survey_en.pdf`
2. Follow-up phone survey: `https://seankhiggins.com/assets/pdf/pricecomp_rct/followup_survey_en.pdf`
3. Follow-up Whatsapp survey: `https://seankhiggins.com/assets/pdf/pricecomp_rct/whatsapp_survey_en.pdf`

The original Spanish versions of the questionnaires are also available:

1. Baseline survey: `https://seankhiggins.com/assets/pdf/pricecomp_rct/baseline_survey_es.pdf`
2. Follow-up phone survey: `https://seankhiggins.com/assets/pdf/pricecomp_rct/followup_survey_es.pdf`
3. Follow-up Whatsapp survey: `https://seankhiggins.com/assets/pdf/pricecomp_rct/whatsapp_survey_es.pdf`

Timeline The baseline survey was conducted as we recruited participants to our RCT, from November 2021 to June 2023. The follow-up phone survey was conducted from January 2023 to March 2024.[21] We surveyed participants via phone at least six months after they participated in the RCT. We collected WhatsApp surveys to ask more detailed questions about the timing of negotiation from April 2024 to May 2025.

[21]We also conducted "recovery surveys" from February 2024–March 2024 for some questions that we had not initially included in the survey but realized we should add while surveys were ongoing. During these recovery surveys, we called back those who had already responded to the survey to ask a small set of additional questions. We had originally intended to conclude the follow-up survey in March 2024. We then received last-minute additional funding to continue the surveys in April 2024, but most of the surveyors had already found other jobs since the surveys were planned to end in March 2024. The survey firm hired new surveyors for the April 2024 surveys, but determined that the

Winsorizing We follow the same protocol for winsorizing as in administrative data.

Appendix G Google Search

G.1 Obtaining data from Google Search results

We scraped data that would be seen on Google using a script that mimics users searching from different comunas with various search terms. For each search, we randomly selected a comuna-search term pair. The set of comunas comes from our baseline survey and is weighted by the number of RCT participants from each comuna. The set of search terms comes from the Google ads campaign that we used to recruit people to our experiment, weighted by their frequency. During each search, we changed the geolocation parameter in Google to match the selected comuna and searched Google using the selected search term. We scraped all available information from each result on the first page, which includes both ads and regular search results. The information we scraped includes the content provider, link, title, text snippet, and the position of the result on the page. The scraper ran from October 19th, 2023, to February 5th, 2024, resulting in 6,677,889 Google search results from 101,852 comuna-search term pairs.

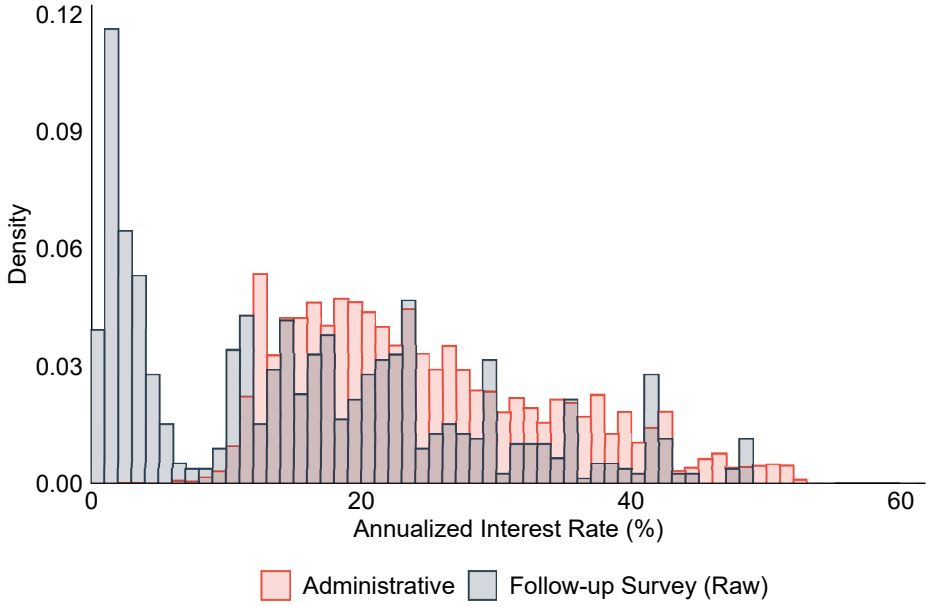
We scraped these results using a desktop emulator, mimicking what a user would see if they opened Google on a desktop computer. Ideally, we would have also scraped the same results using mobile emulation to check if people see anything different when searching Google on mobile phones. However, we were unable to manipulate location information with mobile emulation. This is because Google adopts different functions to determine the user's location, with the geolocation parameters of the browser used on desktops, and the user's IP address used on mobile phones. It was not operationally feasible to fake the IP addresses of each Chilean comuna, so we were unable to scrape mobile results.

We used OpenAI's Assistant API to extract and classify variables from raw text scraped from Google Search results pages. We also tested traditional text processing techniques to extract these variables from the Google search results, but found that employing an advanced large language model (LLM) like GPT-4 yielded higher accuracy for this complicated natural language processing task. To extract interest rate numbers from raw scraped text with rule-based text processing code, we would have needed to exhaust every possible pattern in which an interest rate could occur in a sentence, as well as exclude all possible false positive cases. This task becomes increasingly more challenging as the number of observations increases. On the other hand, a well-trained LLM will be able to comprehend the whole sentence and correctly identify whether it contains a consumer

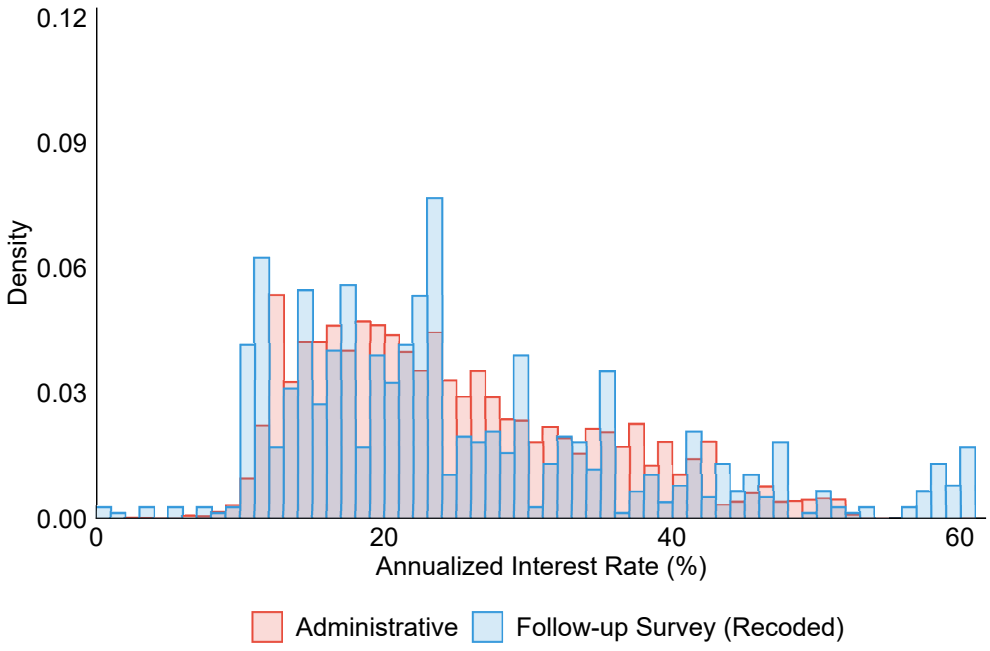
responses collected by the new surveyors did not pass its quality checks, and thus the survey firm stopped collecting additional surveys, dropped these observations, and provided us with a data set of survey responses collected through the initially-planned end date of March 2024.

Figure F.1: Interest Rate Distributions

(a) Raw Annualized Interest Rates in Survey vs. Administrative Data



(b) Recoded Survey Rates vs. Administrative Data



This figure shows the distributions of reported annualized interest rates from our follow-up phone survey in panel (a) and the recoded rates in panel (b), both compared to the interest rates in administrative data. Participants could choose to respond to our follow-up survey question with interest rates expressed as a monthly or annual rate. To correct for loan rates that were reported as annual rates but were implausibly low and likely intended to be reported as monthly rates (i.e. less than 10% annual), we instead assume that these rates are monthly rates and multiply them by 12 to create annual rates. We do not show rates above the 95th percentile of the distribution for legibility.

loan interest rate. We used OpenAI's Assistant API, which allows users to create and tailor an "assistant" for a specific task and use it repeatedly.

We used the "gpt-4-turbo-preview" model of the Assistant API, the state-of-art text processing model at the time, along with its built in tool "retrieval." The key variables to extract were interest rates and the corresponding banks that offered the rates. We also configured the assistant to identify the language, country, and loan type, so that we could filter only results that were Spanish-language consumer loan-related results from Chile. We also had the assistant identify whether the interest rate is a monthly or annual rate, a consumer loan (rather than a credit card, mortgage, or other product), and whether the interest rate excluded fees or was an APR including fees. Given that it is a closed-source LLM, the results generated by our assistant may not be fully reproducible in the future due to the stochastic nature of the model and model updates.

G.2 Comparison of Rates Seen on Google and Received Rates

To compare rates that participants in our RCT would have seen on Google to rates that they actually received, we must match our scraped Google Search data with the administrative loan data. We restrict our sample to the 30,979 individuals in the administrative data who took out a loan. For each of these participants, we keep one unique consumer loan from the administrative data, which is the first loan taken after treatment based on the date that the individual participated in the RCT and the date that the loan was taken out. In the rare case that the individual took two loans on the same day (0.55% of the sample), we take the largest loan taken on the first day after treatment that any loans were taken out by that individual. We further restrict to loans obtained during the months in which we ran the script to scrape Google search results, i.e., October 2023 to February 2024. While this limits the sample size, it mitigates the concern that any observed differences could be due to changes in interest rates over time.

For each individual who took out a consumer loan, we matched the interest rates they would have seen on Google—based on comuna and bank—with the interest rate they received in the administrative data. We then annualized any monthly interest rates by multiplying them by 12 and excluded any scraped results that included only CAE but not interest rates.

Appendix H Predictors of Biased Beliefs

We first assess whether baseline characteristics predict how biased participants' beliefs are. To do this in a data-driven way that guards against overfitting with many covariates, we estimate a series of regularized regression models using elastic net (Friedman, Hastie, and Tibshirani, 2010). The outcome variables measure how biased an individual's beliefs are, compared to the distribution of

interest rates that they face conditional on borrower and loan characteristics. For the interest rate, unlike in Figure 3, we take the difference between their belief about the rate they expect to obtain and the median of the conditional distribution they face rather than the rate they actually obtained to avoid losing sample size (as the majority of participants did not obtain a loan within one year after participating in the RCT). In other words, we compute

$$r_i^{bias} \equiv r_i^{belief} - r_i^{admin} \tag{7}$$

For dispersion, we take the range between the highest and lowest rates they thought a bank would offer them as their belief about dispersion and subtract the corresponding range from the conditional distribution, as in Figure 4:

$$dispersion_i^{bias} \equiv \left(\overline{r}_i^{belief} - \underline{r}_i^{belief}\right) - \left(\overline{r}_i^{admin} - \underline{r}_i^{admin}\right) \tag{8}$$

For each measure r_i^{bias} and $dispersion_i^{bias}$, we then create four measures: (i) the measure as defined in equations (7) and (8), which can take positive or negative values; (ii) its absolute value; (iii) a binary indicator for underestimating, $r_i^{bias} < -1\,\text{pp}$ and $dispersion_i^{bias} < -1\,\text{pp}$; and (iv) a binary indicator for overestimating, $r_i^{bias} > 1\,\text{pp}$ and $dispersion_i^{bias} > 1\,\text{pp}$.[22]

For each model, we use 20-fold cross-validation to identify the best combination of the regularization strength (λ) and mixing parameters (α), to minimize the mean squared error (MSE), where α interpolates between ridge and lasso penalties. We then extract the estimated coefficients from the model corresponding to the optimal (λ, α) pair for each outcome. This approach allows us to isolate the most predictive features while reducing overfitting.

Results. Table H.1 presents the coefficient estimates from the selected elastic net models that render the minimum MSE. Individuals with more-biased beliefs tend to be younger and have lower incomes, less education, and less experience with financial products (i.e., bank accounts and loans). Furthermore, those looking for a smaller loan amount also tend to have more-biased beliefs.

Appendix I Testing for Heterogeneous Treatment Effects

To test for heterogeneous treatment effects of the price comparison tool in a disciplined manner, we implement the machine-learning methodology proposed by Chernozhukov, Demirer, Duflo, and Fernández-Val (2025). We focus on whether there is predictable heterogeneity in the treatment effect of the price comparison tool on belief updating and on search, negotiation, and loan

[22] We allow for slight biases within ± 1 pp without considering these under- or overestimating.

Table H.1: What Predicts Biased Beliefs?

	Prior minus administrative data							
	Interest Rate				Dispersion			
	Level (1)	Abs. value (2)	Under (3)	Over (4)	Level (5)	Abs. value (6)	Under (7)	Over (8)
Personal characteristics								
Age	-0.304	-0.217	0.002	-0.001	-0.485	-0.295	0.003	-0.003
log(Income)	-0.966	-4.852	-0.031	0.014	-3.024	-2.577	0.028	-0.027
Incomplete high-school	11.663	11.119	-0.072	0.061	9.453	8.718	-0.042	0.048
Complete high-school	3.593	4.286	-0.025	0.034	3.963	4.226	-0.021	0.028
Complete 2-year program	1.012	1.907	0.001		0.925	1.259		0.011
Financial products								
Bank account	-0.606	-1.513			-2.585	-2.559	0.021	-0.027
Any loan	-1.315	-2.303		-0.000	-2.137	-1.963	0.004	-0.011
Desired loan characteristics								
log(Loan Amount)	-0.867	-5.866	-0.021	0.006	-0.697	-6.241		-0.007
log(Maturity (years))	0.551		0.004		-2.541	-1.564	0.026	-0.024
Observations	14,401	14,401	14,401	14,401	13,677	13,677	13,677	13,677

This table shows coefficient estimates from regularized regression models using elastic net. For each outcome variable—capturing biased beliefs about the interest rate a participant will obtain,defined in equation (7), and interest rate dispersion, defined in equation (8), we regress the outcome on a common set of covariates. The model includes baseline characteristics from Table 2 and missingness indicators, with missing values imputed by sample means (the coefficients for the missingness indicators are omitted from the table for legibility). Among the education categories in the baseline characteristics, we omit the indicator for whether a participant completed a 5-year degree program or higher to avoid perfect collinearity. For each outcome, we use 20-fold cross-validation to select the best combination of the regularization strength (λ) and mixing parameters (α), with the objective of minimizing mean squared error. The displayed coefficients correspond to the estimates at the optimal (λ, α) pair for each outcome. Coefficients shrunk to zero are left blank.

terms as a function of baseline characteristics.[23] We measure belief updating the same way as in equation (1), i.e., $Posterior_i - Prior_i$ for the rate the participant expects to obtain, the lowest and highest rates a bank could offer them, and the dispersion in rates.

Following the description of Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) in Gertler, Higgins, Scott, and Seira (2025), we consider the conditional average treatment effect (CATE):

$$s_0(Z) := E[Y(1) \mid Z] - E[Y(0) \mid Z],$$

where $Y(1)$ is the potential outcome if treated with the price comparison tool and $Y(0)$ is the potential outcome if assigned to the control group. Z is a vector of baseline characteristics.[24]

[23] We use the same set of baseline characteristics as in Table 2: age, monthly income, education categories, binary variables for whether the participant had a bank account and had previously taken out a loan, and desired loan terms (loan amount and maturity). We omit the indicator for whether a participant completed a 5-year degree program or higher to avoid perfect collinearity.

[24] Missing baseline characteristics are imputed with sample means and missingness indicators are included in Z.

Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) use machine learning (ML) proxies $S(Z)$ for the CATE, $s_0(Z)$, and provide a method to formally test whether there are heterogeneous treatment effects without data mining.

The best linear predictor (BLP) of $s_0(Z)$ given $S(Z)$ is the solution to

$$\min_{b_1,b_2} \mathrm{E}\left[s_0(Z) - b_1 - b_2 S(Z)\right]^2,$$

which, if it exists, is

$$\mathrm{BLP}\left[s_0(Z) \mid S(Z)\right] = \beta_1 + \beta_2 \left(S(Z) - \mathbb{E}[S(Z)]\right),$$

where $\beta_1 = \mathbb{E}[s_0(Z)]$ and $\beta_2 = \mathrm{Cov}\left[s_0(Z), S(Z)\right] / \mathrm{Var}[S(Z)]$. If $S(Z)$ is a perfect proxy for $s_0(Z)$, then $\mathrm{Cov}\left[s_0(Z), S(Z)\right] = \mathrm{Var}[S(Z)]$ and thus $\beta_2 = 1$. In general, $\beta_2 \neq 1$ due to noisy predictions $S(Z)$. If $S(Z)$ is purely noise and uncorrelated with $s_0(Z)$, then $\beta_2 = 0$. Furthermore, if there is no heterogeneity, i.e., $s_0(Z) = s$, then $\beta_2 = 0$ because $\mathrm{Cov}\left[s_0(Z), S(Z)\right] = \mathrm{Cov}\left[s, S(Z)\right] = 0$. Rejecting the hypothesis that $\beta_2 = 0$ therefore indicates that both (i) $S(Z)$ is a relevant predictor and, more importantly for our purposes, (ii) there is heterogeneity in $s_0(Z)$. Therefore, by estimating the β_2 coefficient of $\mathrm{BLP}\left[s_0(Z) \mid S(Z)\right]$ and testing whether the coefficient is significantly different from zero, we can empirically test for heterogeneity in the treatment effect across individuals.[25]

In the ML step, we consider support vector machine, LASSO, and random forest, and choose the model that maximizes $\Lambda \equiv |\beta_2|^2 \mathrm{Var}(S(Z)) \propto \mathrm{Corr}^2(s_0(Z), S(Z))$.[26] Maximizing Λ is equivalent to maximizing the correlation between the ML proxy $S(Z)$ and the CATE $s_0(Z)$, or equivalent to maximizing the R^2 in the regression of $s_0(Z)$ on $S(Z)$. The models that maximize Λ for our analysis on $Posterior_i - Prior_i$ are as follows: LASSO for expected rate, lowest rate, and the highest rate; and random forest for dispersion. The models that maximize Λ for our analysis on search and negotiation behavior and loan outcomes are as follows: LASSO for number of offers and Pr(take loan) measured in survey data; support vector machine for number of institutions searched, Pr(negotiate), and log interest rate offered (all measured in survey data), as well as Pr(take loan) measured in administrative data; and random forest for number of institutions applied and log interest rate taken (all measured in survey data), as well as log interest taken measured in administrative data.

[25] The coefficient estimates will depend on the data split that takes place before training the machine learning models. The coefficient estimates, p-values, and lower and upper bounds of their confidence intervals that we report in Table I.1 and Table I.4 are the medians of the estimates and confidence intervals from 1,000 such distinct splits, as recommended by Chernozhukov, Demirer, Duflo, and Fernández-Val (2025).

[26] More specifically, since each of the 1,000 data splits leads to a different Λ, we maximize the median Λ from the 1,000 data splits.

I.1 Belief Updating

First, we reject the null hypothesis of no heterogeneity in treatment effects of the tool for belief updating about the rate participants expect to obtain and the highest rate a bank would offer them, but we fail to reject the null hypothesis of no heterogeneity for the lowest rate a bank would offer them and dispersion (Table I.1).

Table I.1: Test for Heterogeneous Treatment Effects on Belief Updating

	Expected rate (1)	Lowest rate (2)	Highest rate (3)	Dispersion (4)
ATE	16.05***	10.86***	30.16***	15.86***
	[0.000]	[0.000]	[0.000]	[0.000]
	(13.18, 18.96)	(8.56, 13.15)	(24.61, 35.70)	(12.24, 19.46)
HTE	0.57***	0.34	0.56***	0.15
	[0.001]	[0.145]	[0.001]	[0.200]
	(0.23, 0.91)	(-0.12, 0.83)	(0.22, 0.91)	(-0.08, 0.38)
Observations	4,699	4,651	4,578	4,302
Control Mean Posterior	29.22	22.65	47.45	23.18
Control Median Posterior	18	12	25	10.68
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows average treatment effect (ATE) and heterogeneous treatment effect (HTE) coefficient estimates from the Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) machine learning procedure described in Appendix I. We conduct this test for the outcomes in Table 3. p-values testing the null hypothesis that the parameter is equal to zero are in square brackets, while 95% confidence intervals are in parentheses. Coefficients, p-values, and lower and upper bounds of confidence intervals are medians over 1,000 data splits. The bin density fixed effects (FEs) are deciles of how many loans are in the price comparison tool histogram that would have been shown to the participant based on their borrower and loan characteristics, had they been assigned to the price comparison tool arm. If any of the participant's loan amount, maturity, or income is outside of the support of the administrative data, the bin density variable would be a missing value; for these individuals, we create an additional fixed effect category that we include in the regression, i.e., we include an additional dummy equal to 1 if the bin density is a missing value. Among the education categories in the baseline characteristics, we omit the indicator for whether a participant completed a 5-year degree program or higher to avoid perfect collinearity. The smaller sample size compared to Table 3 is due to the exclusion of participants treated with the simple tool. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Given that we detected heterogeneous treatment effects for updating about the expected and highest interest rates, we follow Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) to determine what characteristics predict treatment effects of the tool on belief updating. To do this, the data is split into training and estimation samples, and after training the ML model on the training sample, the model is used to estimate predicted individual treatment effects for each individual in the estimation sample. Individuals in the estimation sample are then grouped into quintiles based on predicted individual treatment effects, and "group average treatment effects" (GATES) for each quintile are estimated. Next, differences in the average characteristics Z across the groups with the

highest and lowest predicted treatment effects (i.e., the most-affected and least-affected groups) are estimated in a classification analysis (CLAN).

Results Table I.2 presents a GATES analysis, estimating treatment effects for the quintiles with the lowest and highest predicted treatment effects. Consistent with our findings in Table I.1, we find that the difference in treatment effects of the tool for the most- and least-affected quintiles is statistically significant (column 3). Interestingly, although there are statistically significant heterogeneous treatment effects, the tool leads participants to increase their expectations about the rate they will obtain and dispersion even in the least-effected quintile (by 10.2 pp, compared to 24.7 pp in the most-affected quintile).

In contrast, and again consistent with the results in Table I.1, we find that the difference in the treatment effect between the most affected and the least affected quintiles for belief updating about the lowest rate and dispersion is not statistically significant at the 5% level.

In Table I.3, we present a CLAN analysis to understand which individual traits are associated with the largest impacts on belief updating for the two outcomes for which we observe a statistically significant heterogeneous treatment effect of the tool. We find that the same characteristics that predict more-biased beliefs also predict a larger treatment effect of the tool on beliefs about both the expected rate and the highest rate. In particular, the treatment effect of the tool is larger for participants who are younger and have lower incomes, less education, and less experience with financial products (i.e., bank accounts and loans), as well as those looking for smaller loans.

I.2 Search, Negotiation, and Loan Terms

Across all of the tests on search and negotiation behavior and loan terms, we find no detectable heterogeneity in treatment effects (Table I.4). This suggests that the treatment effects of the tool on negotiation, the probability of taking out a loan, and interest rates do not differ by the characteristics of the participants, but rather tend to be spread evenly across participants.

Table I.2: Heterogeneous Treatment Effects: GATES Analysis

	Most affected quintile (1)	Least affected quintile (2)	Difference (3)	Observations (4)
Expected rate	24.67***	10.26***	14.38***	4,699
	[0.000]	[0.002]	[0.002]	
	(18.28, 30.97)	(3.80, 16.68)	(5.27, 23.47)	
Lowest rate	15.09***	9.05***	6.22*	4,651
	[0.000]	[0.001]	[0.090]	
	(10.05, 20.17)	(3.94, 14.15)	(-0.96, 13.40)	
Highest rate	45.08***	19.29***	25.65***	4,578
	[0.000]	[0.002]	[0.004]	
	(32.94, 57.33)	(6.87, 31.65)	(8.39, 43.06)	
Dispersion	20.10***	13.25***	6.46	4,302
	[0.000]	[0.001]	[0.263]	
	(12.04, 28.18)	(5.19, 21.39)	(-4.91, 17.91)	

The table shows sorted group average treatment effect coefficients for the outcomes in Table 3. The sample is divided into 5 groups, based on the quintiles of the machine learning proxy predictor $S(Z)$. Coefficients, p-values, and lower and upper bounds of confidence intervals are medians over 1,000 data splits. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table I.3: Heterogeneous Treatment Effects: Classification Analysis (CLAN)

	Expected rate (1)	Highest rate (2)
Personal characteristics		
Age	-11.29***	-8.94***
	[0.000]	[0.000]
	(-12.39, -10.17)	(-10.10, -7.72)
log(Income)	-1.68***	-1.68***
	[0.000]	[0.000]
	(-1.88, -1.49)	(-1.88, -1.49)
Incomplete high-school	0.12***	0.11***
	[0.000]	[0.000]
	(0.09, 0.15)	(0.08, 0.14)
Complete high-school	0.54***	0.47***
	[0.000]	[0.000]
	(0.49, 0.59)	(0.41, 0.52)
Complete 2-year program	0.03	0.06**
	[0.173]	[0.013]
	(-0.01, 0.07)	(0.01, 0.10)
Financial products		
Bank account	-0.47***	-0.42***
	[0.000]	[0.000]
	(-0.52, -0.42)	(-0.47, -0.36)
Any loan	-0.32***	-0.34***
	[0.000]	[0.000]
	(-0.37, -0.26)	(-0.39, -0.29)
Desired loan characteristics		
log(Loan Amount)	-2.54***	-2.65***
	[0.000]	[0.000]
	(-2.69, -2.39)	(-2.81, -2.51)
log(Maturity (years))	-0.60***	-0.55***
	[0.000]	[0.000]
	(-0.67, -0.52)	(-0.63, -0.48)
Observations	4,699	4,578

Classification analysis for the outcomes in Table 3 that exhibit significant HTE coefficients in Table I.1 – expected rate and highest rate. The sample is divided into 5 groups, based on the quintiles of the CATE proxy predictor $S(Z)$. We report the mean difference in the baseline characteristics from Table 2 between the highest and lowest $S(Z)$ quintiles. For example, for age, we subtract the mean age of the lowest $S(Z)$ quintile from the mean age of the highest $S(Z)$ quintile. Among the education categories in the baseline characteristics, we omit the indicator for whether a participant completed a 5-year degree program or higher to avoid perfect collinearity. Coefficients, p-values, and lower and upper bounds of confidence intervals are medians over 1,000 data splits. $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table I.4: Test for Heterogeneous Treatment Effects on Search, Negotiation, and Loan Terms

	Survey Data							Administrative Data	
	N of inst. searched (1)	N of inst. applied (2)	N of offers (3)	Pr(negotiate) (4)	Log interest rate offered (5)	Pr(take loan) (6)	Log interest rate taken (7)	Pr(take loan) (8)	Log interest rate taken (9)
ATE	0.012	0.024	0.072	0.038*	-0.119	0.036	-0.102	0.008	0.013
	[0.904]	[0.733]	[0.119]	[0.057]	[0.120]	[0.232]	[0.241]	[0.183]	[0.298]
	(-0.185, 0.208)	(-0.117, 0.166)	(-0.019, 0.163)	(-0.001, 0.076)	(-0.270, 0.031)	(-0.023, 0.095)	(-0.274, 0.068)	(-0.004, 0.021)	(-0.011, 0.037)
HTE	0.172	0.311*	-0.007	0.254	0.230	-0.001	-0.374	0.289	0.207
	[0.417]	[0.087]	[0.979]	[0.474]	[0.355]	[0.998]	[0.441]	[0.121]	[0.117]
	(-0.248, 0.596)	(-0.046, 0.661)	(-0.554, 0.540)	(-0.454, 0.960)	(-0.258, 0.711)	(-0.542, 0.542)	(-1.351, 0.599)	(-0.073, 0.680)	(-0.052, 0.466)
Observations	2,172	2,095	2,084	2,065	354	2,081	232	30,718	5,982

This table shows average treatment effect (ATE) coefficient estimates and heterogeneous treatment effect (HTE) coefficient estimates from the Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) machine learning procedure described in Appendix I. We conduct this test for the outcomes in Table 4. p-values testing the null hypothesis that the parameter is equal to zero are in square brackets, while 95% confidence intervals are in parentheses. Coefficients, p-values, and lower and upper bounds of confidence intervals are medians over 1,000 data splits. Missing baseline characteristics are imputed with sample means and missingness indicators are included as additional regressors. Among the education categories in the baseline characteristics, we omit the indicator for whether a participant completed a 5-year degree program or higher to avoid perfect collinearity. The smaller sample size compared to Table 4 is due to the exclusion of participants treated with the simple tool. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix J Differential Treatment Effects for "Compliers"

Given that the price comparison tool had an effect on the probability of taking out a loan, in this section we attempt to isolate "always-takers" (those who would have taken out a loan regardless of treatment) and "compliers" (those who only take out a loan if treated). Our goal is to understand whether the lack of a treatment effect of the price comparison tool on interest rates in the admin data—in contrast with the effect in survey data—is due to a negative treatment effect for always-takers being offset by a selection effect where compliers are less-creditworthy borrowers receiving higher interest rates. If this were the case, the selection effect leading compliers to be included in the regression only in the treatment group would pull the treatment effect towards zero.

We attempt to isolate always-takers and compliers by again implementing Chernozhukov, Demirer, Duflo, and Fernández-Val (2025). First, we divide the sample into two equally sized subsamples. Using one half of the sample, we train a machine learning model that predicts both $E[Y(1) \mid Z]$ and $E[Y(0) \mid Z]$, where $Y(\cdot)$ is a binary variable measuring loan take-up. We select the machine learning model using the same method as in Appendix I, which in this case is support vector machine. Using the trained model, we predict $E[Y(1) \mid Z]$ and $E[Y(0) \mid Z]$ for each observation in the other half of the data. This follows the "honest" approach (Athey and Imbens, 2016) where the model is trained on one subsample while the predictions are obtained using a different subsample.

Always-takers are defined as participants with $E[Y(1) \mid Z] = E[Y(0) \mid Z] = 1$, while compliers are defined by $E[Y(1) \mid Z] = 1, E[Y(0) \mid Z] = 0$. Thus, an estimand that we find performs well for isolating compliers is

$$\frac{E[Y(1) \mid Z]}{E[Y(0) \mid Z]} = \frac{P(Y = 1 \mid T = 1, Z)}{P(Y = 1 \mid T = 0, Z)}, \tag{9}$$

as this measure increases as $E[Y(1) \mid Z] \to 1$ and as $E[Y(0) \mid Z] \to 0$. Furthermore, we find it performs better than a more standard estimand of the treatment effect $E[Y(1) \mid Z] - E[Y(0) \mid Z]$, because this difference can be large even for groups in which $E[Y(0) \mid Z]$ is not close to 0.

We thus estimate (9) with the individual-level potential outcome estimates $\hat{E}[Y(1) \mid Z]$ and $\hat{E}[Y(0) \mid Z]$ that we obtain from the Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) method:

$$\frac{\hat{E}[Y(1) \mid Z]}{\hat{E}[Y(0) \mid Z]}. \tag{10}$$

To estimate (10) for each individual, we iterate this process 1,000 times and then take the median of (10) within individual across iterations.[27]

[27] Since the initial sample split divides the data in half, on average 500 estimates of (10) are estimated for each individual across the 1,000 iterations, but there is variation in this across individuals due to the random nature of the sample splits.

We then sort observations by their median values of (10), and define the top decile sorted by this measure as compliers and the bottom decile as always-takers. Restricting our focus to observations in these two groups, we test two things. First, we check whether our out-of-sample compliance prediction predicts take-up well: the treatment effect on take-up should be 0 for the always-takers and positive for the compliers. Second, we examine whether the treatment effect on the interest rate differs across these two groups. To achieve these two goals, we run the following regression:

$$y_i = \alpha + \beta_1 \mathbb{1}\{\text{Price Comparison Tool}\}_i + \beta_2 \mathbb{1}\{\text{Complier}\}_i \tag{11}$$

$$+ \beta_3 \mathbb{1}\{\text{Price Comparison Tool}\}_i \times \mathbb{1}\{\text{Complier}\}_i + \varepsilon_i, \tag{12}$$

where y_i is loan take-up or log interest and $\mathbb{1}\{\text{Compliers}\}_i$ is an indicator variable equal to 1 when the observation is labeled as a complier. The estimate β_1 can be interpreted as the treatment effect for the always-takers, while the estimate β_3 can be interpreted as the differential treatment effect for the compliers relative to the always-takers.

Table J.1 presents the treatment effects for the always-takers and the differential treatment effect for the compliers, which are discussed in the main text.

Table J.1: Differential Treatment Effects for Compliers vs. Always-Takers

	Pr(take loan)	Log interest rate taken
	(1)	(2)
Treatment Effect (Always-Takers)	−0.001	−0.058
	(0.014)	(0.040)
Differential Treatment Effect for Compliers	0.062***	0.092**
	(0.020)	(0.046)
Observations	6,143	1,374

The table shows the effect of the price comparison tool on search, negotiation, and loan terms using administrative data. It shows the β_1 and β_3 estimates from specification (11). Column (1) is a dummy variable equal to 1 if the participant obtained a consumer loan within 1 year after participating in the RCT, according to administrative data from the CMF. Column (2) is the natural logarithm of the interest rate on the loan that the participant took out, also according to administrative data from the CMF. Both are measured using administrative data The sample is restricted to those that are in the bottom 10% and top 10% of the individual-level median of (10) across iterations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.