Search and Negotiation with Biased Beliefs in Consumer Credit Markets*

Erik Berwart Sean Higgins Sheisha Kulkarni Santiago Truffa

February 12, 2025
Latest version available here

Abstract

How do inaccurate beliefs about the distribution of interest rates affect search and outcomes in consumer credit markets? In collaboration with Chile's financial regulator, we conducted a randomized controlled trial with 112,063 loan seekers where we 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 consumers thought interest rates were lower than they actually were, and the price comparison tool caused them to increase their expectations about the interest rate they would obtain by 56%. Consumers also underestimated price dispersion, and our price comparison tool caused them to increase their estimates of dispersion by 69%. 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. We also cross-randomized whether we asked participants their beliefs about the distribution of interest rates, and find that merely asking these questions led them to search at 4% more institutions and obtain 10% lower interest rates.

*Berwart: Comisión para el Mercado Financiero, eberwart@cmfchile.cl. Higgins: Department of Finance, Kellogg School of Management, Northwestern University, 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 the Lab for Inclusive Fintech (LIFT), UC Berkeley. We are very grateful to Deniz Aydın, Claire Célérier, Isis Durmeyer, John R. Grigsby, and Chris Hansman for discussing our paper. We thank conference and seminar participants at AIEA seminar, BI Norwegian Business School, Boulder Summer Conference on Consumer Financial Decision Making, Canadian Economics Association, CEPR WEFIDEV Workshop on Finance and Development, CFPB, Conference on Digital Experimentation, CSIO/TSE Joint Workshop on Industrial Organization, EFA, Federal Reserve Bank of Philadelphia, HEC Montréal, IPA–GPRL Researcher Gathering, Junior Household Finance Seminar, Lab for Inclusive FinTech at UC Berkeley, NBER Summer Institute Household Finance, New York University, SITE Financial Regulation, and University of Washington Summer Finance Conference for helpful comments and discussions. We thank Arturo Charleston, Vicente Corral, Andrés Cruz, Matías Fuentes, Jora Li, Erick Molina, Limin Peng, Carlos Restituyo, 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). In our setting of Chilean consumer loans, the *same* consumer receives substantially different interest rate offers from different banks: based on survey data we collected, the average within-consumer range in interest rate offers for those receiving more than one offer was 8 percentage points (pp), compared to an average annual interest rate of 29%. This dispersion in interest rates can have substantial real costs for borrowers, as they must choose between not searching much and incurring higher loan costs, or incurring potentially high search costs to obtain a lower interest rate.[1] The option of searching to obtain a lower interest rate may be costly for several reasons, including time and travel costs, high rejection rates, and the cognitive effort required to compare financial products.[2]

While there has been substantial work on search costs in consumer credit markets, little is known about whether consumers have accurate beliefs about the distribution of interest rates that banks would offer them. In a standard sequential search model, inaccurate beliefs would affect the perceived benefits of search, and could have important implications for both search behavior and equilibrium pricing. If consumers underestimate the first moment of the distribution of interest rates banks would offer them (as we find most consumers do), they will search more than is optimal because they will overestimate the expected value of 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 value of another draw. In addition, inaccurate beliefs could decrease consumers' ability to successfully negotiate with the bank once they receive an offer.

We ask how inaccurate beliefs about the distribution of interest rates affect search behavior and loan outcomes in consumer credit markets. We conducted a randomized controlled trial (RCT) with 112,063 Chileans searching for loans where we both measure priors about the interest rate

[1] Several papers have shown that not searching much in consumer credit markets has a high cost due to interest rate dispersion. In the US auto loan market, the average borrower pays $488 more in present value for a $17,000 car due to not searching (Argyle, Nadauld, and Palmer, 2023). In mortgage markets, borrowers also pay substantially more due to not searching: 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. While these estimates are for the US, other countries' mortgage markets have similarly high levels of dispersion, including Canada (Allen, Clark, and Houde, 2019), Italy (Guiso, Pozzi, Tsoy, Gambacorta, and Mistrulli, 2022), and the UK (Coen, Kashyap, and Rostom, 2024).

[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). These cognitive costs may be particularly important in consumer credit markets because financial products can be cognitively costly to understand (Célérier and Vallée, 2017; Kulkarni, Iberti, and Truffa, 2023) and may include shrouded costs (Campbell, Jackson, Madrian, and Tufano, 2011; Ferman, 2016; Alan, Cemalcilar, Karlan, and Zinman, 2018).

distribution and show loan seekers in the treatment group a price comparison tool designed to correct inaccurate priors. 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 the administrative data.[3] We then had them fill out a short baseline survey that randomized whether we asked them their expectations about the distribution of interest rates and how much they will search (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 built the price comparison tool using administrative data from Chile's financial regulator, the Comisión para el Mercado Financiero (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.

We first document that the majority of participants have inaccurate beliefs about both the first and second moments of the interest rate distribution. While there is significant heterogeneity in expectations, 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. Furthermore, borrowers who underestimated the rate they would obtain did so by on average 14.9 pp. Their estimates of dispersion in the interest rates banks would offer them are also much lower than suggested by administrative data: specifically, 75% of participants underestimate dispersion in the interest rates bank would offer them compared to administrative data. This finding is robust to various measures of dispersion; 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.[4]

We then cross-randomized all participants (i.e., including both those who were and were not asked questions on their beliefs about interest rates) into one of two treatment arms designed to

[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.

[4]We borrow this measure from the macroeconomic uncertainty literature (Coibion, Georgarakos, Gorodnichenko, Kenny, and Weber, 2024) and prefer it as it 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, and use the implied standard deviation of the distribution under certain functional form assumptions as an alternative measure of dispersion.

correct inaccurate beliefs about the distribution of interest rates, or a control group. We measure effects of these treatments on (i) expectations about interest rates, (ii) search behavior, (iii) whether they take out a loan, and (iv) the terms of the loan they take out. Both treatments were built using administrative data on loan and borrower characteristics for the universe of consumer loans in Chile, i.e., data from over 2.1 million loans to approximately 1.3 million borrowers. The first is a price comparison tool that shows participants the conditional distribution of interest rates that similar borrowers obtained for similar loans over the previous six months. The second is a simple tool where we estimated the cost savings from searching at additional banks, by running simulations using the same conditional distribution. Finally, the control group was shown a video that was designed to take the same amount of time as using the tools but that did not contain useful information for search.

Immediately after seeing the price comparison tool, simple tool, or control video, participants assigned to the elicit beliefs treatment were asked again about their expectations about the distribution of interest rates and how much they will search. When treated with the price comparison tool, participants update and report expecting to receive a 16 pp *higher* interest rate on the loan they obtain, or a 55.6% increase compared to the control mean posterior belief of a 29.5% expected annual interest rate.[5] The price comparison tool also led participants to increase their expectation of how much price dispersion they face in the market by 16 pp, or 68.6% relative to the control mean posterior of 23.4 pp dispersion in annual interest rates.[6] 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 (statistically significant at the 10% level), while the coefficient on dispersion is very close to 0 at –0.04 pp, and is not statistically significant.

To measure effects on search behavior, loan take up, 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. The lack of a treatment effect on

[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—prior to 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."

[6] 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%, while the control median posterior of dispersion is 10.8 pp.

search could be due to offsetting effects of updating inaccurate beliefs about the first and second moments of the distribution, since people increased their beliefs about the rate they would obtain (which should reduce search) but also increased their beliefs about dispersion (which should increase search).

Despite not affecting the number of institutions at which participants searched, the price comparison tool led them to obtain 13% more offers, to receive 11.9% lower interest rate offers (both measured in the follow-up survey), and to be 5% more likely to take out a loan (measured in the administrative data).[7] We find evidence that increased negotiation explains how borrowers obtained more loan offers and more favorable terms: the price comparison tool increases the probability of negotiating by 39%. This is consistent with prior evidence that negotiation can play an important role in consumer credit markets (Allen, Clark, and Houde, 2019; Allen and Li, forthcoming).

Randomizing whether we elicited priors about the interest rate distribution and the number of institutions at which participants intended to search was motivated by evidence from Zwane, Zinman, Van Dusen, Pariente, Null, Miguel, Kremer, Karlan, Hornbeck, Giné, Duflo, Devoto, Crepon, and Banerjee (2011) that survey questions can have treatment effects on real-world behavior and outcomes. In particular, being asked the question itself may have made people realize there was more dispersion than they thought, since we would not ask them questions about the distribution if there were very little dispersion.We find that the elicit priors treatment led participants to search at 0.13 (4%) more institutions and to obtain 10% lower interest rates than participants who were not asked their priors about interest rates or expectations about search. Unlike the price comparison tool treatment, eliciting priors does not have an effect on the number of offers received.

The effect of eliciting priors, which also leads people to obtain lower interest rate offers but requires them to search more to do so, is consistent with the negotiating mechanism that we find drives the effect of the price comparison tool on loan outcomes. (This is not something we can test directly because we do not observe priors for those to whom we did not ask the questions.) Thus, eliciting priors would have increased their priors about dispersion *without* resolving their uncertainty about the distribution. As a result, those in the elicit priors arm are not able to better negotiate without searching more—which requires knowing concrete information about the distribution—but do update about the benefits of search and hence search more. The simple tool did not affect the number of institutions at which participants searched for information or applied for a loan, which is not surprising given that the simple tool did not lead people to update their priors about the interest rate distribution.

Taken together our results suggest that correcting biased beliefs about the interest rate distri-

[7]Consistent with this, participants are 9.7% more likely to take out a loan in the survey data, but unlike in the administrative data, the effect on the probability of taking out a loan is not statistically significant in the survey data, where we have a smaller sample size.

bution can improve borrower loan terms: providing information to borrowers about feasible rates so they can negotiate allows them to benefit from lender competition without paying high search costs. However, our price comparison tool intervention was data-intensive: it requires loan-level data with interest rates and other loan terms merged with relevant borrower characteristics used in banks' pricing decisions; not many regulators collect these detailed data. On the other hand, merely asking participants their beliefs about the interest rate distribution caused them to receive similar reductions in interest rates. This is a much more scalable intervention, but it required loan seekers to search at more lenders (and thus incur search costs) to receive these lower rates.

Our paper makes two main contributions. First, unlike other papers on search in consumer credit markets, we are able to distinguish between search costs and biased beliefs because we collect data on beliefs and conduct an RCT that we show shifts priors. Other papers on search in financial markets typically assume that people are drawing from a known distribution of prices. For example, this assumption is made by Hortaçsu and Syverson (2004) for S&P 500 index funds, Honka, Hortaçsu, and Vitorino (2017) and Yankov (forthcoming) for deposit accounts, Argyle, Nadauld, and Palmer (2023) for auto loans, Allen, Clark, and Houde (2013) and Agarwal, Grigsby, Hortaçsu, Matvos, Seru, and Yao (2024) for mortgages, and Cuesta and Sepúlveda (2021) for consumer loans. However, if this is not true—as we document in the context of consumer loans in Chile—this assumption leads to biased estimates of search costs (Koulayev, 2013). Whether high search costs or inaccurate priors explains the lack of search in consumer credit markets matters because the optimal policy responses to each of these interventions are quite different.

Second, we show how experimentally-induced changes in priors about the distribution of interest rates affect search and decision-making in the market for consumer loans. Prior work has studied how priors affect financial decision-making on both the assets and liabilities sides of the household balance sheet.[8] On the assets side, individuals who have experienced low stock market returns are more pessimistic about future stock market returns and are less likely to participate in the stock market (Malmendier and Nagel, 2011). Experimentally-induced increases in expectations about house price growth cause increases in real estate investments (Armona, Fuster, and Zafar, 2019), and residual variation in past returns also positively correlates with investments after controlling for priors about future house price changes (Liu and Palmer, 2023).[9] On the 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). The majority of first-time payday

[8] In other contexts beyond financial markets, papers have studied how priors affect search in the education market (Kapor, Neilson, and Zimmerman, 2020; Arteaga, Kapor, Neilson, and Zimmerman, 2022; Agte, Allende, Kapor, Neilson, and Ochoa, 2024) and labor market (Jäger, Roth, Roussille, and Schoefer, forthcoming; Bandiera, Bassi, Burgess, Sulaiman, Vitali, and Rasul, forthcoming).

[9] The real estate investments investigated in these papers are incentivized but experimental, i.e., they are stylized decisions made by participants as part of an incentivized survey experiment. In contrast, we study decisions made by loan seekers in the real world after receiving our treatment.

borrowers underestimate the dollar amount of the fees they will pay (Bertrand and Morse, 2011), while more experienced payday borrowers have more accurate priors (Allcott, Kim, Taubinsky, and Zinman, 2022); an intervention correcting priors by providing information about the dollar cost of paying off a payday loan over a given time period reduced loan demand (Bertrand and Morse, 2011; Wang and Burke, 2022).

2 Institutional Context

2.1 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 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 ($N = 2{,}197{,}716$ 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.

2.2 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.[10]

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, Iberti, and Truffa (2023) 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.

2.3 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 survey (when participants were looking for a loan), the three most important features of the loan were all functions of the interest rate: 26% of participants reported that the total loan cost was the most important feature, 22% reported monthly payment, and 20% reported the interest rate or annualized percentage rate (APR, which is known as the *carga anual equivalente*, or CAE, in Chile). Similarly, in our follow-up survey, the most common reason for choosing a particular lender was a lower interest rate, with 41% 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"

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

(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 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, 16% reported that they chose a particular lender because that was the only offer they received. Furthermore, 25% of borrowers chose a particular lender because they were quickly approved by that institution; being quickly approved and only receiving one offer were the second- and third-most common reasons for choosing a particular lender.

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 10% of participants) and whether it is a bank in which they already had an account (8%). In the follow-up survey, when taking out a loan the less important features included whether the loan payment could be automatically deducted from payroll (5%), whether they were a client of that bank (3%), trust in the institution (3%), and getting approved for a higher loan amount (3%).

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 A.2): 60% 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% of participants reporting that they planned to stop searching after they were approved by one institution.[11]

2.4 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 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 5 and

[11]These survey questions were not mutually exclusive, and based on the responses it appears that loan seekers implement a combination of strategies.

describe the details of our procedure in Appendix B. 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 A.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 B 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 a number of third-party comparison websites for consumer loans, and 12% of participants reported using such a tool when searching for a loan. Table A.1, panel B, describes the inputs required by each third-party comparison website.

The first and most popular third-party comparison tool, ComparaOnline, operates as a quote aggregator and is run by a private-sector company. Consumers input their desired loan size and maturity and receive quotes for loans from different institutions (see Figure A.3 for screenshots from their website). However, ComparaOnline does not ask for any borrower characteristics, and thus the interest rate quotes it provides are not conditional on borrower characteristics.

Destacame is another comparison website run by a private-sector company. Prospective borrowers input information related to their loan search, as well as borrower characteristics including current employment status, tenure at their current job, monthly income, and current financial products they have. After entering this information, they do not immediately receive interest rate quotes, but instead financial institutions can submit products for the borrower to consider. When we tested Destacame's website, the menu of products submitted by financial institutions took 1–2 weeks to appear, and did not include loan terms such as the interest rate or loan amount (see Figures A.4 and A.5 for screenshots). Thus, a consumer using Destacame to search for loans would still need to formally apply for a loan in order to receive an interest rate quote. Destacame's business model is to sell services such as a credit counseling service to improve the consumer's probability of being approved for a loan.

Rankia is a website that also advertises itself as a simulator to "help people make better financial decisions." For consumer loan seekers, they provide a simulator that asks consumers the use of the loan and the size of the loan they would like (Figure A.6). However, regardless of the responses

to these questions, the output is always the same. The output includes an article titled "Best Consumer Loans for 2024," a link to the website of Banco Internacional, and a table of interest rates for a sample loan across eight banks, where Banco Internacional's loan has the lowest interest rate (Figure A.7). The terms in this table do not change regardless of the inputs entered by the user in Rankia's simulator.

Chile's consumer protection agency, SERNAC, also runs a comparison website. Like ComparaOnline and Rankia, SERNAC does not ask for any borrower characteristics (Figure A.8), and thus the interest rate quotes it provides are not conditional on borrower characteristics. SERNAC only displays loan sizes in increments of one million pesos (up to 10 million pesos) and maturities in increments of one year. According to SERNAC (2015), they collect interest rate data from bank websites, but unlike our exercise described in Appendix B, do not condition on any borrower characteristics.

3 Experimental Design

3.1 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. 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 A.9 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 the CMF logo and a description of our study. This page also included the informed consent to participate in the study. 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, 36,439 people (23.9% of clicks) consented to participate; we blocked people from participating more than once using their national ID. Of those who consented to participate, 112,063 (308%) 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; 46,051 consumer loan seekers (126.4% of those who consented) con-

tinued 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. These are the two sample sizes (112,063 and 46,051) for our respective research questions on the impact of eliciting priors and the impact of the price comparison tool.[12]

3.2 Elicit Priors 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 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 selectd 25% of the sample in order to test whether these survey questions have a treatment effect (Zwane, Zinman, Van Dusen, Pariente, Null, Miguel, Kremer, Karlan, Hornbeck, Giné, Duflo, Devoto, Crepon, and Banerjee, 2011), 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 priors treatment the same interest rate expectation and search questions to test whether their expectations were affected by treatment.

[12]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 priors 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 priors treatment and the subsample who did not receive the elicit priors treatment, relative to a control group in which the same proportions did and did not receive the elicit priors treatment; thus, if there is an interaction effect between the tool treatments and the elicit priors treatment, it would be reflected in the effect of the tool treatments that we estimate (Muralidharan, Romero, and Wüthrich, 2023).

3.3 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 2.1 million loans to 1.3 million borrowers. 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 C 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.

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 2–6 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 D 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 A.10 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.

4 Data

4.1 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.[13] Importantly, credit bureaus in Chile do not report continuous credit scores, but rather 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 are also able to follow repayment of the loan in monthly intervals after its issuance to evaluate outcomes such as delinquency and default. 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 eventual loans they obtain. 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. 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 A.12 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.[14] 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: 51.1% 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 37.42% of the RCT sample compared to 40.44% 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 A.13). 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

[13] 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.

[14] 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.

to 34 months in the overall population.

4.2 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 priors for those assigned to the elicit priors treatment (described in Section 3.2), we asked participants about their sociodemographic characteristics and detailed questions regarding their existing banking relationships and other financial products they have.

We also asked participants questions to determine how they form priors 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 priors 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 (denoted "control" in the table, although this is different than the control group in the assignment to the price comparison tool, simple tool, or control). 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 priors questions is balanced: the p-value of the omnibus F-test regressing the elicit priors 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 bal-

anced, as could be expected by chance: those assigned to the elicit priors 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.[15] The 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 priors 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).

4.3 Endline Phone Survey

We surveyed participants via phone at least six months after they participated in the RCT. We attempted to contact 42,250 participants (38% of our 112,063 sample), and ultimately collected 6,441 completed surveys, for a 15.5% response rate. Table A.2 shows that response rates are balanced across both the elicit priors 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, and the loan terms they were offered after this negotiation. We also include questions to understand the mechanism behind the potential effect on search, as well as other measures of financial well-being, such as total debt, total savings, and ability to cope with shocks.

[15] 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 priors treatment, and the elicit priors treatment caused some participants to stop participating in the survey. (Note that participants who abandoned the survey during the elicit priors module are still tracked in the administrative data and included in the sample to estimate the effect of eliciting priors.)

5 Results

5.1 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 priors 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 14.9 pp lower than the rate they ultimately receive.

Figure 4 shows the distribution of the difference between a participant's priors on dispersion—measured as the highest rate they think a bank could offer them minus the lowest rate they think a bank could offer them—and the difference in the highest and lowest rates we observe in the administrative data 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 priors 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 (75%) *underestimate* price dispersion, but also that there is a long right tail of participants who substantially *overestimate* dispersion.

Given that priors are inaccurate, how are consumers forming their priors?

We find that 41% of participants report having seen advertisements by banks, 44% used bank websites in the past (prior to the current search), 12% have used third-party comparison websites in the past, and 23% have asked friends and family about interest rates. To assess whether bank advertisements, bank websites, comparison websites, or family and friends might cause people to have inaccurate priors, we compare rates our participants *would* have seen in each of these contexts with the rate they ultimately received in our bank administrative data (Figure 5).

For advertisements, 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 Google searches using that keyword and geolocation pair, and scrape the resulting first page of Google results—including both Google ads and regular Google search results. More details are provided in Appendix E. The difference between interest rates that are shown in Google ads

or search results and the rates people actually obtain are heavily negatively skewed with 73.53% of ads advertising rates that were lower than what participants ultimately received from the same bank in our administrative data (Figure 5a).

For bank websites and 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 each comparison website and scrape the resulting loan terms. More details are provided in Appendix B. Bank simulators tend to show inaccurate rates, as the difference between the rate a bank website showed and the rate a participant obtained can be as much as 27 pp lower or 20 pp higher than the rate they ultimately receive. 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 third-party comparison websites, the difference between rates shown on comparison websites and the rates obtained 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 eleven 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. They also can be significantly inaccurate with some rates being more than 20 percentage points lower than the rate borrowers received (Figure 5d).

5.2 Price Comparison Tool Leads to Large Updates in Rate Expectations

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 levels (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.[16]

Table 3 shows the results. On average, participants increased their expected rate expectation by 16 pp, or 55.6% relative to the control mean posterior of 29.5%.[17] 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 11 pp and their expectation about the highest interest rate a bank would offer them by 30 pp. Their expectations about dispersion also increase by 16 pp compared to a control mean posterior of 23.4 pp of dispersion, an increase of 68.6%. Tables A.3 and A.4 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 A.5 and A.6 show the same pattern when we log-transform expectations 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 between the highest and lowest rate, which is a scale-invariant measure. Table A.7 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.

The net effect of updating priors in this way about the first and second moments of the distribution is ambiguous. Since participants increased their estimates about the first moment, this could lead them to search less. In the absence of the price comparison tool, after receiving a draw these participants would think it is a bad offer and continue searching, whereas after seeing the price comparison tool they would know it is a reasonable offer and stop searching. However, since

[16] 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.

[17] 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%.

participants also increased their estimates about the second moment, this could lead them to search more because their estimates of the benefits of search have increased.[18]

5.3 Price Comparison Tool Leads to More Offers and Lower Rates

Table 4 shows the effect of our treatments on search behavior and loan outcomes. We run the following regression:

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

Neither the price comparison tool nor the simple tool led people to search at more institutions or to formally apply for loans at more institutions. Furthermore, Figure A.14 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. Nevertheless, we find that the price comparison tool led participants to receive 0.069 (13.1%) more offers in the survey data and to be 1 pp (5%) more likely to take out a loan in the administrative data. According to our survey data, borrowers who did not take out a loan overwhelmingly did not make the purchase or investment for which the loan was earmarked.

Participants also receive 11.9% lower interest rate offers and 10.5% lower interest rates on the loans they take out according to the survey data. Figures A.15 and A.16 show the distribution of interest rates offered and interest rates taken, respectively. The distribution of interest rates offered for the price comparison tool is a uniform leftward shift from the control group's interest rate distribution. The p-value of the Kolmogorov-Smirnov (KS) test of whether the CDFs in the treatment arms and control arm differ is 0.126 for the price comparison tool and 0.08609 when we pool the tool arms for increased power. Taken together, these results suggest that the price comparison tool shifted the distribution of offers to the left of that of the control group. The CDFs for interest rates taken look similar, but the KS test fails to reject that the treatment and control arms are drawn from the same distribution ($p = 0.187$).

In administrative data, we do not find lower interest rates on the loan participants took out. The discrepancy between the estimates in administrative and survey data is likely due to some borrowers obtaining loan offers and loans from institutions that are not regulated by the CMF and hence not included in our data; these include credit cooperatives such as Coopeuch, Caja Los Andes, and Caja Los Heroes, as well as FinTech lenders. Indeed, the discrepancies in the fraction of borrowers who take out a loan in the control group between our survey data (36.9%) and administrative data (19%) suggests that the administrative data are missing a substantial number of

[18] Merely learning the distribution could also lead participants to search less, as they may have been searching to learn about the distribution in the absence of the tool (De Los Santos, Hortaçsu, and Wildenbeest, 2017).

loans that are being picked up in the survey data and for which treated borrowers are able to obtain lower interest rates after seeing the price comparison tool.[19]

We test two potential mechanisms behind the tool enabling people to search better in Table A.18. First, people may be searching at different institutions; in other words, a participant may have already decided to search at three banks, but after seeing the tool the participant thought harder about which three banks to search. We do not find evidence of this: we find a null treatment effect on searching at a different institution than the two institutions that people listed as the institutions where they planned to search in the baseline survey (Table A.18 column 1).

Second, resolving uncertainty about the distribution may lead people to have a better sense of their outside option when negotiating with the bank, which allows them to extract more surplus in the form of lower interest rates or a higher probability of approval. We do find evidence that the tool leads people to convey information from the tool before applying (Table A.18, column 2) and to be 3.7 pp more likely to negotiate after applying, relative to a control mean of 9.7% of people negotiating, or a 39% increase in the relative probability of negotiating (column 3), statistically significant at the 1% level. Furthermore, the number of institutions at which participants negotiate increases by 43.5% in the price comparison tool arm (column 4), statistically significant at the 1% level.

5.4 Eliciting Priors Leads to More Search and Lower Rates

Table 5 shows the effect of eliciting participants' priors on search behavior and loan outcomes. We run the following regression:

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

We find that merely asking these questions led consumers to search at 0.13 more institutions, or a 4% increase compared to the control mean of 3.357. 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), but lower in magnitude, suggesting 1.1% lower interest rates.[20]

[19] After sixteen years of attempting to regulate these non-bank institutions, the Chilean government passed a law in July 2024 that will lead to these other institutions reporting to the CMF and being included in the administrative data going forward.

[20] Again, comparing the percent in the control group who took out loans in the survey and administrative data suggests that this discrepancy is due to not all loans appearing in the administrative data because some loans are from institutions not regulated by the CMF.

6 Conclusion

We document that consumers have inaccurate priors about both the first and second moment of the distribution of interest rates. Almost 73% percent of participants thought they would obtain an interest rate lower than what they actually received and almost 75% percent underestimated the dispersion of interest rate offers they could get. There are several sources that could have led borrowers to have inaccurate priors, including Google results and comparison websites, which show loan seekers lower rates than what they will actually obtain 85% and 73% of the time, respectively. The presence of inaccurate priors suggests that consumers are likely to have different search behavior than that predicted by models where agents are assumed to perfectly know the distribution of rates from which they are drawing. This can lead to biased estimates of search costs in structural models, as well as a misunderstanding of how banks price loans in equilibrium.

We designed two tools to correct inaccurate priors and test their effects on search behavior and loan outcomes using an RCT. The first was a a price comparison tool that showed participants a histogram of interest rates that borrowers with similar characteristics obtained on similar loans in the last six months. The second is a simple tool that showed participants the expected benefits of searching at more banks. Both tools were built using on 2.1 million loans from 1.3 million borrowers using administrativee data from the CMF, the financial institutions and market regulator of Chile. We recruited participants online through Google ads and both surveyed and treated them online, and followed up with a subset of participants with a phone survey at least six months after they participated, in order to collect rich data on their search behavior and loan outcomes.

We find that the price comparison tool caused participants to update their priors about the interest rate they would obtain upwards by 16 pp or 55.6%. It also caused them to update their priors about dispersion in the rates banks could offer them by 16 pp or 68.6%. The price comparison tool led participants to receive 13% more offers and 11% lower interest rates. It also made them 5% more likely to take out a loan. Participants appear to have achieved lower rates by negotiating 39% more. In contrast, participants who saw the simple tool did not change their priors about the rate they would obtain or dispersion. Unsurprisingly, then, the simple tool had no impact on participants' search behavior or loan terms.

We also cross-randomized whether we asked participants their priors about the distribution of interest rates and how much they would search. Merely eliciting priors led participants to search at 0.13 more institutions and receive 9.6% lower interest rate offers on average.

These findings show that, on the one hand, there are cost-effective ways to help people obtain lower interest rates without incurring additional search costs (by showing the price comparison tool, which can resolve their uncertainty about the distribution of interest rates and lead them to negotiate better). On the other hand, the price comparison tool requires substantial data that

many regulators do not have. Our results also show that a less data intensive and more scalable intervention—merely asking questions about priors—also leads people to obtain lower interest rates. However, because eliciting priors does not resolve their uncertainty about the interest rate distribution, obtaining lower interest rates with this less data intensive and more scalable intervention does require people to search more (rather than to negotiate better without searching more) in order to obtain those lower rates.

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 (2019). "Search Frictions and Market Power in Negotiated-Price Markets." *Journal of Political Economy* 127(4), 1550–1598.

Allen, Jason and Shaoteng Li (forthcoming). "Dynamic Competition in Negotiated Price Markets." *Journal of Finance*.

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 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.

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, Munshi Sulaiman, Anna Vitali, and Imran Rasul (forthcoming). "The Search for Good Jobs: Evidence from a Six-year Field Experiment in Uganda¤." *Journal of Labor Economics*.

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*.

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.

CMF (2024). "SISTEMA CONTABLE (Instrucciones Generales)." https://www.sbif.cl/sbifweb3/internet/archivos

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.

De Los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest (2017). "Search With Learning for Differentiated Products: Evidence from E-Commerce." *Journal of Business & Economic Statistics* 35(4), 626–641.

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

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.

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.

Honka, Elisabeth, Ali Hortaçsu, and Maria Ana Vitorino (2017). "Advertising, Consumer Awareness, and Choice: Evidence from the U.S. Banking Industry." *The RAND Journal of Economics* 48(3), 611–646.

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 (forthcoming). "Worker Beliefs About Outside Options." *The Quarterly Journal of Economics*, qjae001.

Kapor, Adam J., Christopher A. Neilson, and Seth D. Zimmerman (2020). "Heterogeneous Beliefs and School Choice Mechanisms." *American Economic Review* 110(5), 1274–1315.

Koulayev, Sergei (2013). "Search With Dirichlet Priors: Estimation and Implications for Consumer Demand." *Journal of Business & Economic Statistics* 31(2), 226–239.

Kulkarni, Sheisha, Gonzalo Iberti, and Santiago Truffa (2023). "Removing the Fine Print: Standardization, Disclosure, and Consumer Loan 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.

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.

SERNAC (2015). "Estudio de simulación de créditos en línea: SERNAC detectó hasta casi 3 millones de pesos de diferencia en crédito de $4 millones en 48 meses. - Portal SERNAC." https://www.sernac.cl/portal/619/w3-article-7302.html.

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.

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.

Wang, Jialan and Kathleen Burke (2022). "The Effects of Disclosure and Enforcement on Payday Lending in Texas." *Journal of Financial Economics* 145(2, Part B), 489–507.

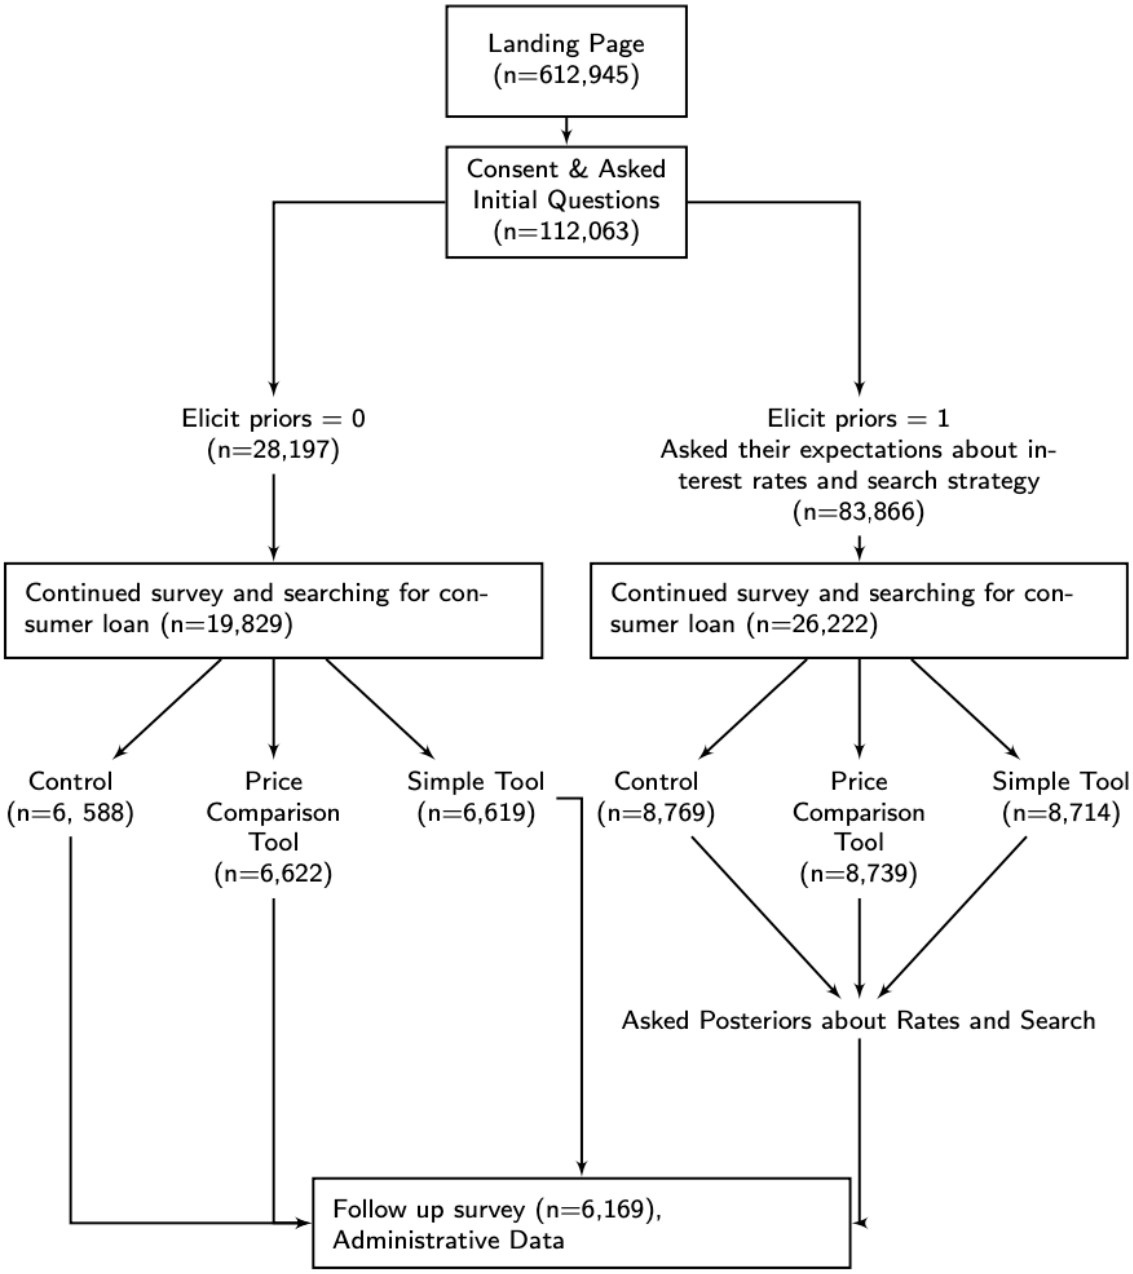
Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion (2022). "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications." *Journal of Economic Perspectives* 36(3), 157–184.

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.

Yankov, Vladimir (forthcoming). "In Search of a Risk-Free Asset: Search Costs and Sticky Deposit Rates." *Journal of Money, Credit and Banking*.

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.

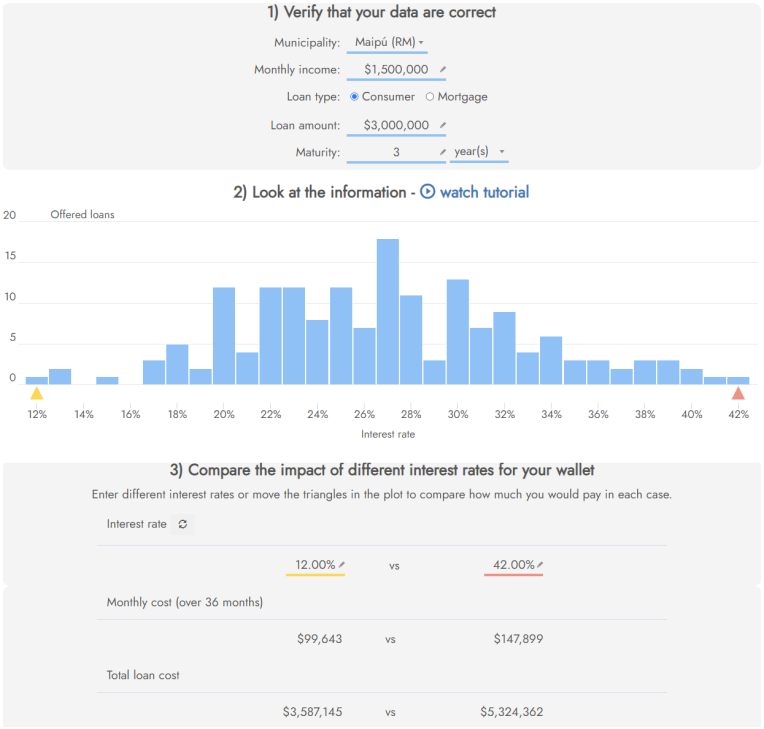
Figure 1: RCT Flowchart



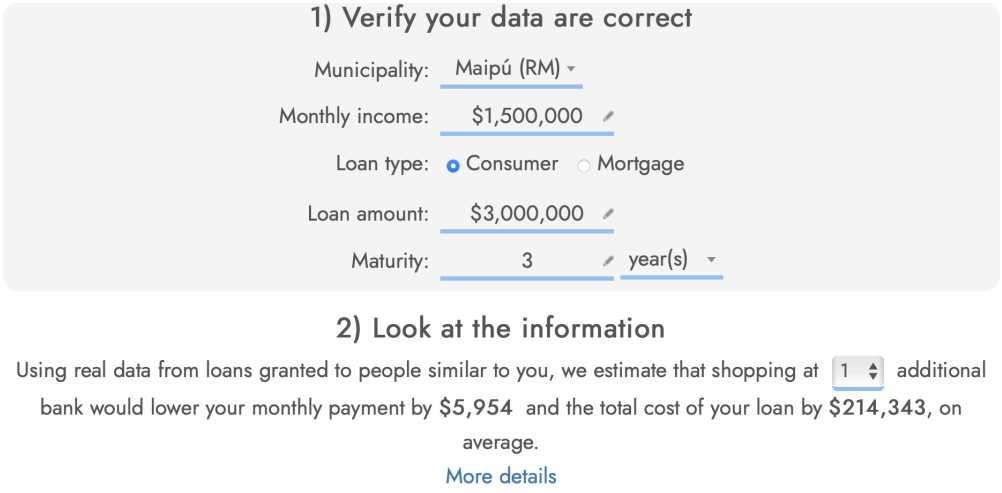
This figure shows the progression of the participants through our study after they reached our landing page from our Google advertisements. They are randomized at two key points: when they are assigned either "Elicit priors = 0" or "Elicit priors = 1" and subsequently when they are cross-randomized to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

Figure 2: Screenshot of Comparison Tool Treatments

(a) Interest Rate Price Comparison Tool

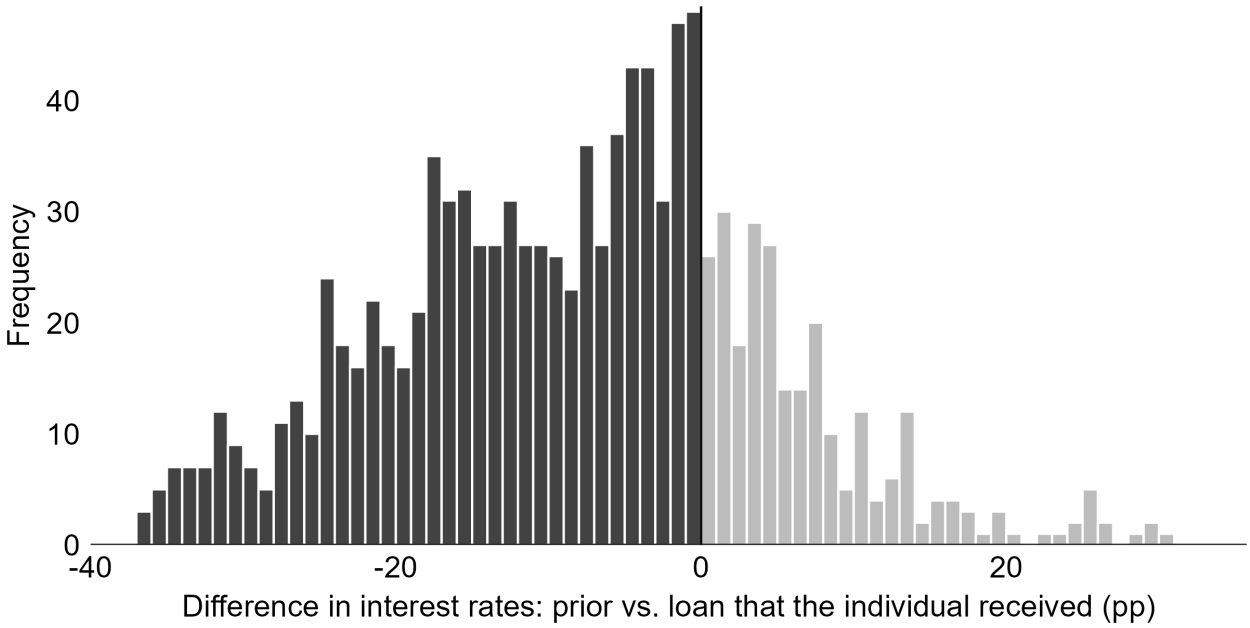


(b) Simple Tool



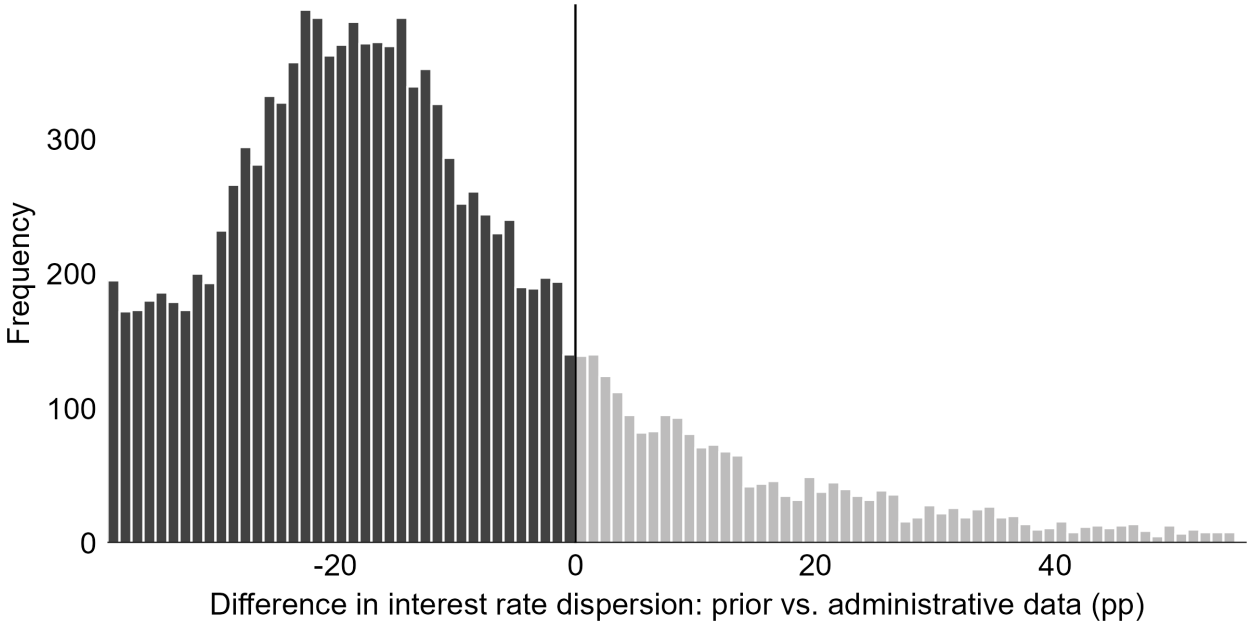
This figure shows a screenshot of English translations of our price comparison tool (panel a) and simple tool highlighting the benefits of search (panel b). For both tools, prospective borrowers already entered the borrower and loan characteristics in the top panel of the tool in our baseline survey; 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. Appendix D documents the construction of the histograms and the data behind them. 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.

Figure 3: Difference in Interest Rates Between Prior and Rate the Individual Received (pp)



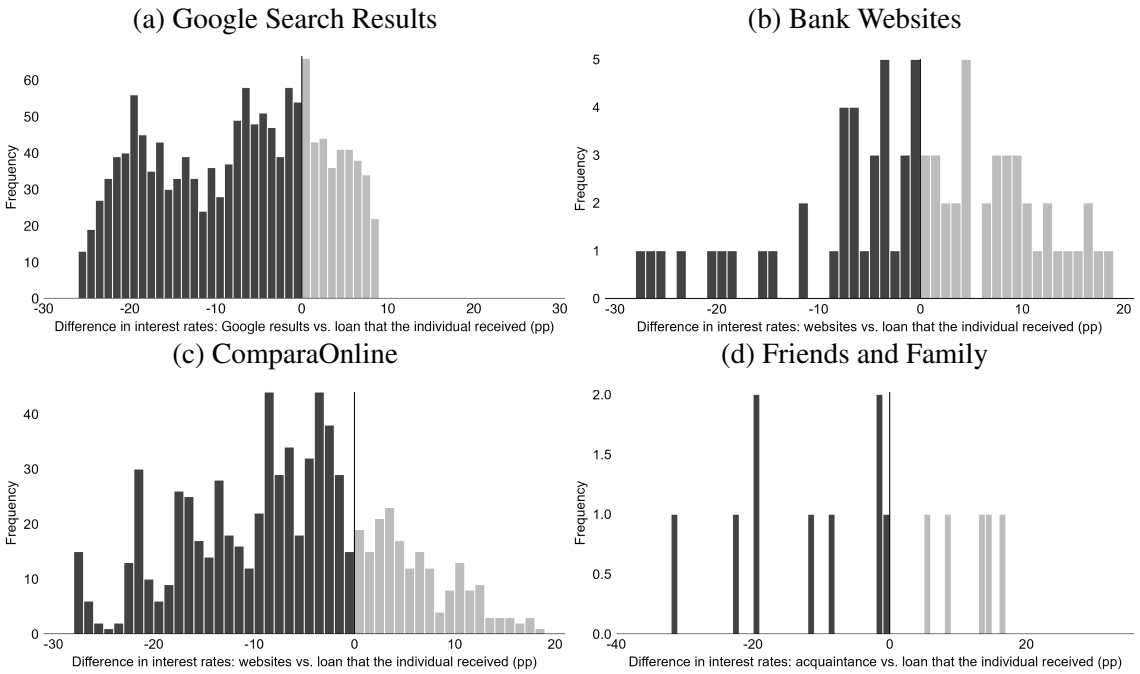
This figure shows that participants tend to underestimate the interest rate they ultimately obtain. The figure is a histogram of the difference between a participant's prior expectations about the interest rate they will get on the loan they take out (as reported in our baseline survey) and the interest rate they ended up receiving on the loan they took out in our administrative data. We construct this figure by restricting to the subset of participants in the control group who took out a loan after participating and comparing the interest rate they obtained on the loan in the administrative data to the prior they had reported in the baseline survey. 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 1,091. 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 Prior and Administrative Data (pp)



This figure shows that participants tend to 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 14,149. The percentage of people who underestimated the dispersion, i.e., the percentage of the sample in the negative portion of the histogram, is 75%.

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



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 (see Appendix E for more detail). There are 1,379 observations, of which 73.53% 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 (see Appendix B for more detail). 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 online comparison tool ComparaOnline for the bank where they obtained a loan and the rate they received from that bank in the administrative data (see Appendix B for more detail). There are 749 observations, of which 74.1% are negative. Panel (d) shows a histogram of the difference between the interest rates that a participant's friends and family told them they received in the baseline survey and the rate they received in the administrative data. There are 16 observations, 68.8% of which are negative.

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

	Elicit Priors = 0 Mean (1)	Elicit Priors (2)	N (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 priors treatment for the full sample. We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta \mathbb{1}(\text{Elicit Priors})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy indicating whether participant i was assigned to the elicit priors treatment. Column (1) shows α which is the mean for the $\mathbb{1}(\text{Elicit Priors})_i = 0$ group. Column (2) shows β which is the difference in means between the $\mathbb{1}(\text{Elicit Priors})_i = 0$ and $\mathbb{1}(\text{Elicit Priors})_i = 1$ 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: $\mathbb{1}(\text{Elicit Priors})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$, where $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy equal to 1 if participant i was assigned to the elicit priors treatment. 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. "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, and 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 priors treatment (rather than a prior module), and the elicit priors treatment affected whether participants continued in the survey. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Balance of Pre-Treatment Characteristics Across Tool Treatment Arms

		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]	
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. "Loan characteristics" refer to characteristics of the loan they are 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, and cash advance. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Interest Rate Expectations

	Expected rate (1)	Lowest rate (2)	Highest rate (3)	Dispersion (4)
Simple Tool	0.79*	0.82**	−0.29	−0.04
	(0.44)	(0.36)	(0.82)	(0.68)
Price Comparison Tool	16.38***	10.90***	30.20***	16.06***
	(1.23)	(0.96)	(2.31)	(1.50)
Observations	6,409	6,364	6,269	5,907
Control Mean Posterior	29.46	22.82	47.88	23.42
Control Median Posterior	18	12	25.2	10.8
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows the effect of the simple tool and price comparison tool on priors. 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 had they been assigned to the price comparison tool arm, based on their borrower and loan characteristics. 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. There are 1,129 observations with dispersion = 0 (highest rate = lowest rate). Continuous variables are winsorized at the 5th and 95th percentiles by treatment arm. Tables A.3-A.7 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 (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Effect of Price Comparison Tool and Simple Tool on Search and Loan Terms

	Survey Data							Administrative Data	
	N of inst. searched	N of inst. applied	N of offers	Pr(negotiate)	Log interest rate offered	Pr(take loan)	Log interest rate taken	Pr(take loan)	Log interest rate taken
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(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 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 (but who were nevertheless included in the survey because they did reach the module in which the elicit priors treatment was 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. 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. If the participant does not remember all the rates they were offered, we use additional questions on the lowest and highest rates they recalled receiving (unless they did not know or refused to answer these additional questions). 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 from a regulated institution 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. Continuous variables are winsorized at the 5th and 95th percentiles by treatment arm. Tables A.8–A.12 show balance tests for the subsamples used in the various columns of this table. 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: Effect of Eliciting Priors on Search and Loan Terms

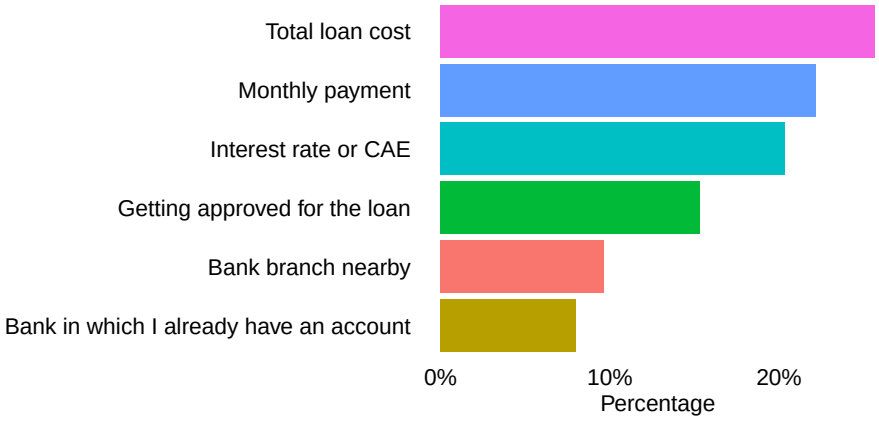
	Survey Data							Administrative Data	
	N of inst. searched (1)	N of inst. applied (2)	N of offers (3)	Pr(negotiate) (4)	Log interest rate offered (5)	Pr(take loan) (6)	Log interest rate taken (7)	Pr(take loan) (8)	Log interest rate taken (9)
(Intercept)	3.357***	1.192***	0.579***	0.111***	3.553***	0.360***	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 Priors	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 priors treatment on search behavior 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. 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. If the participant does not remember all the rates they were offered, we use additional questions on the lowest and highest rates they recalled receiving (unless they did not know or refused to answer these additional questions). 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 from a regulated institution 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. Continuous variables are winsorized at the 5th and 95th percentiles by treatment arm. Tables A.13–A.17 show balance tests for the subsamples used in the various columns of this table. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

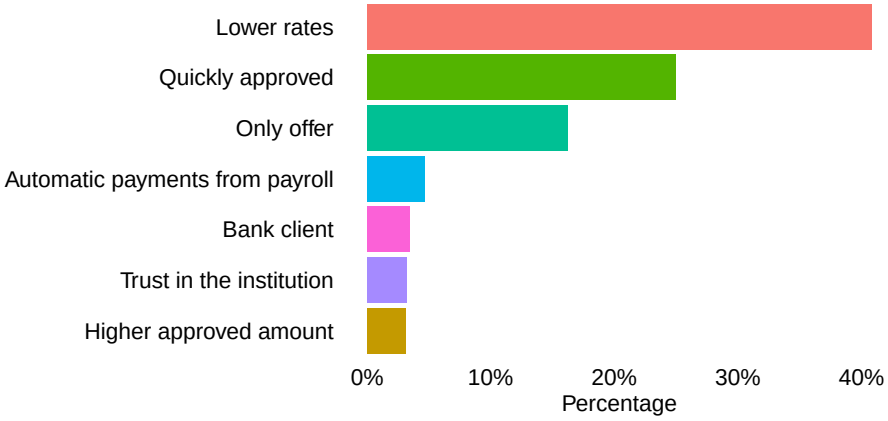
Appendix A Additional Figures and Tables

Figure A.1: Stated Importance of Loan Features

(a) Most Important Stated Loan Feature (Baseline Survey)



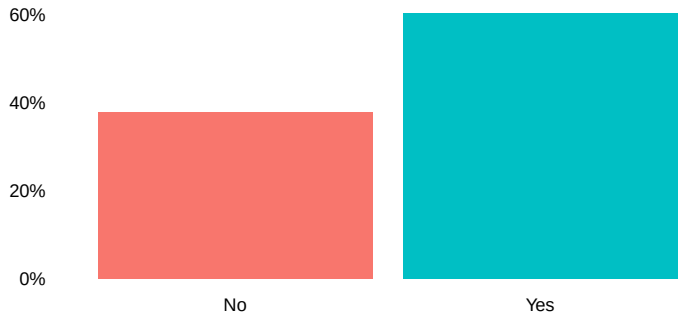
(b) Reason for Choosing Lender (Follow-up Survey)



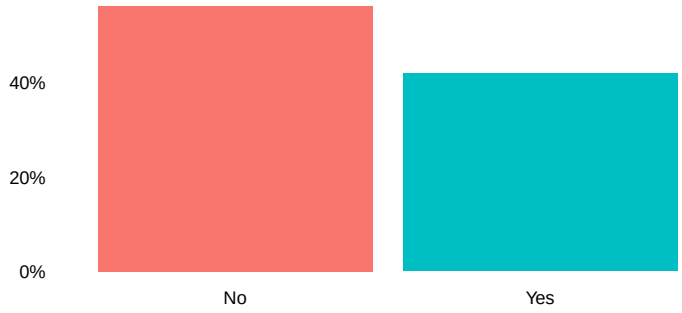
This figure shows the most important features of a loan reported in the baseline survey and the reason they chose to take out a particular offer in the follow-up survey. Panel (a) shows results from the baseline survey, conducted when participants were searching. It shows the reasons that participants 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 A.2: Sequential Search, Simultaneous Search, and Searching for Approval

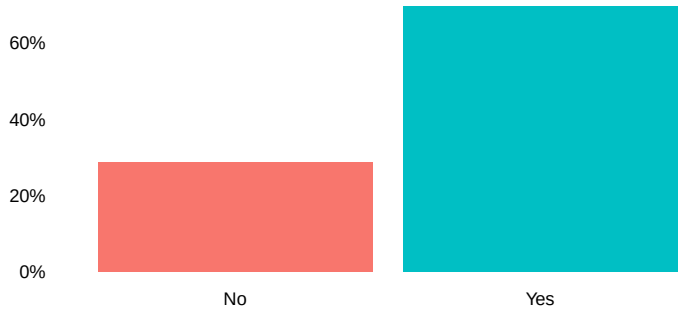
(a) Sequential Search: Target Interest Rate



(b) Simultaneous Search: Target N Offers or N Banks



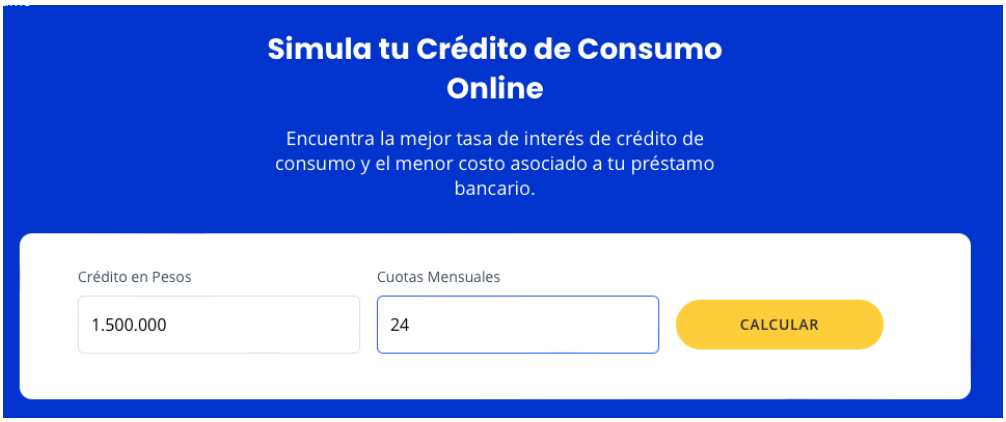
(c) Searching for Approval: Search Until 1 Offer



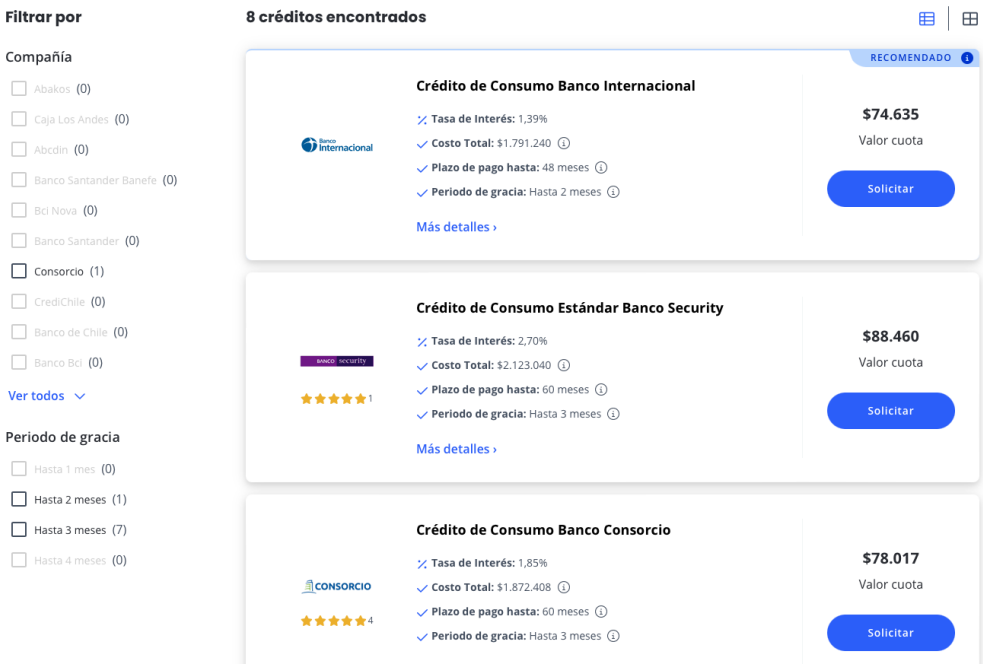
This figure shows the results of asking participants in our follow-up phone survey questions about their search strategy. We asked four yes/no questions to make the three plots: (a) "Did you plan to search until you reached a target interest rate and then stop searching?"; (b) "Did you have a target number of offers you would like to receive from financial institutions to stop looking?' or "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 counted the number of answers to the questions, and specifically, for panel (b), we reported the number of participants who answered "yes" to either of the two questions.

Figure A.3: ComparaOnline

(a) Input



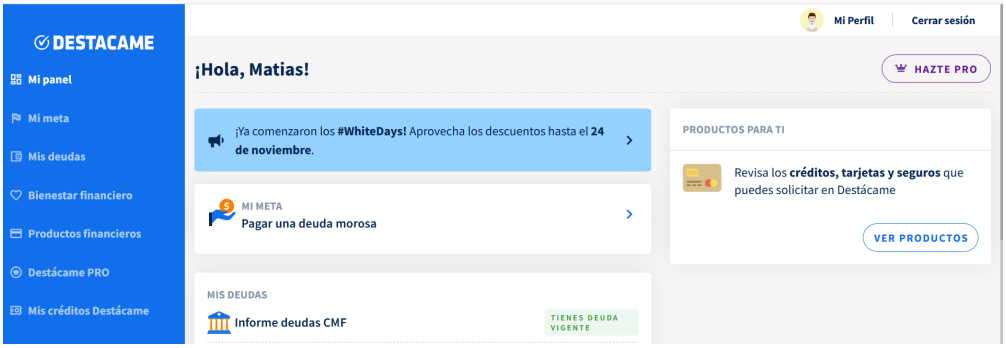
(b) Output



This figure shows the user interface of ComparaOnline. Their website provides rate quotes to prospective customers and direct customers to financial institutions. It functions as a quote aggregator that displays the interest rates that banks report they would (but are not required to) offer. Last accessed on May 15, 2024.

Figure A.4: Destacame

(a) Beginning Page



(b) Input Page 1



(c) Input Page 2



(d) Input Page 3



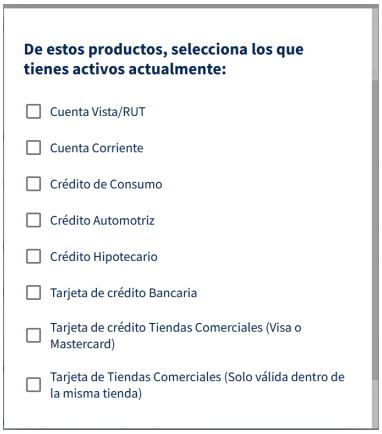
This figure shows the user interface of Destacame and the input pages of its consumer loan simulator. Last accessed on May 15, 2024.

Figure A.5: Destacame

(a) Input Page 4



(b) Input Page 5



(c) Input Page 6



(d) Output Page



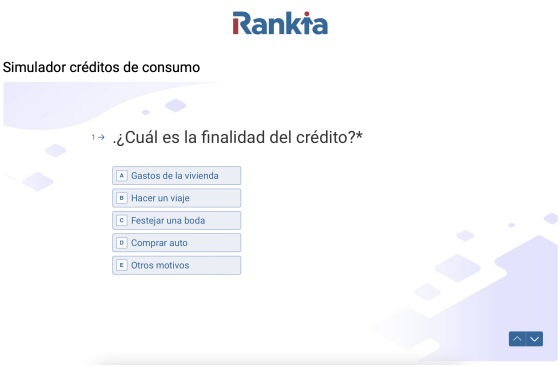
This figure shows the user interface of the input webpages and output webpages of Destacame's consumer loan simulator. Last accessed on May 15, 2024.

Figure A.6: Rankia Input Pages

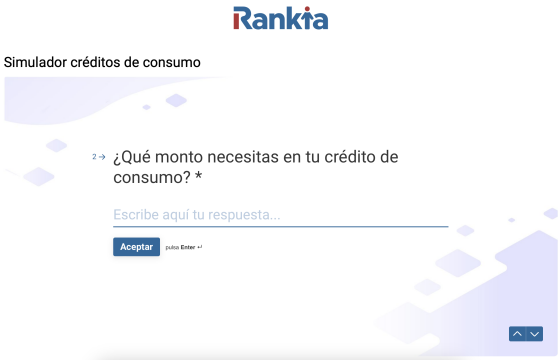
(a) Input Page 1



(b) Input Page 2



(c) Input Page 3



(d) Input Page 4



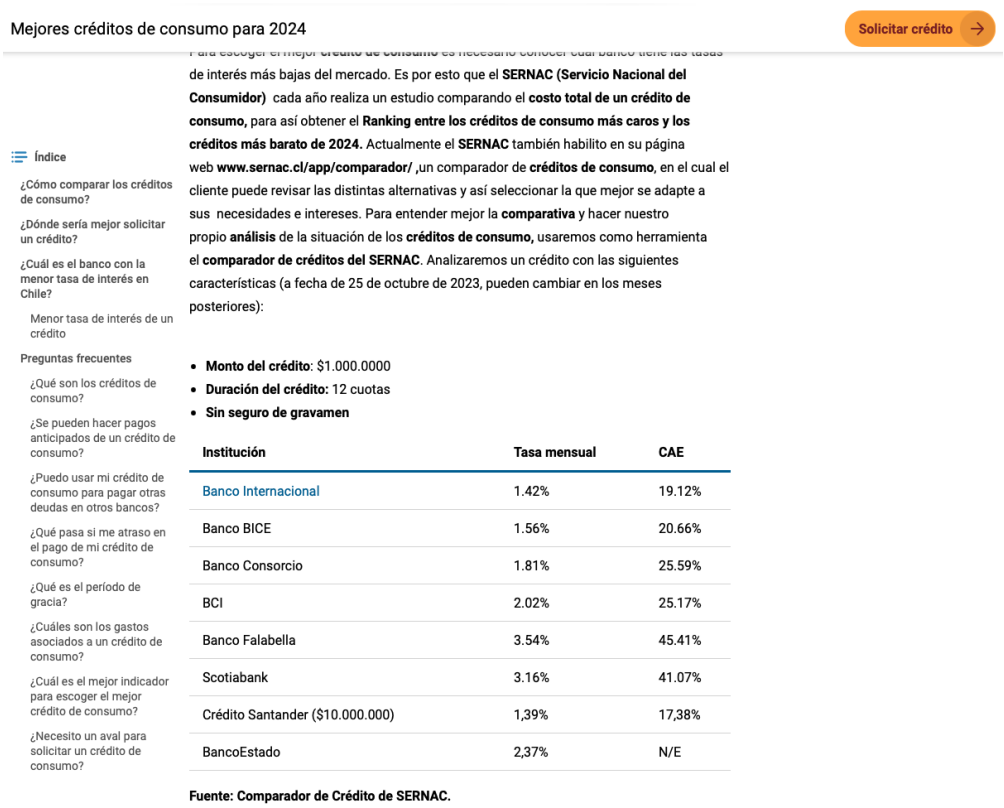
This figure shows the input pages of Rankia's consumer loan simulator. Last accessed April 3, 2024.

Figure A.7: Rankia Output Pages

(a) Top of the Output Page



(b) Middle of the Output Page



This figure shows the output page of Rankia's consumer loan simulator. Last accessed April 3, 2024.

Figure A.8: SERNAC



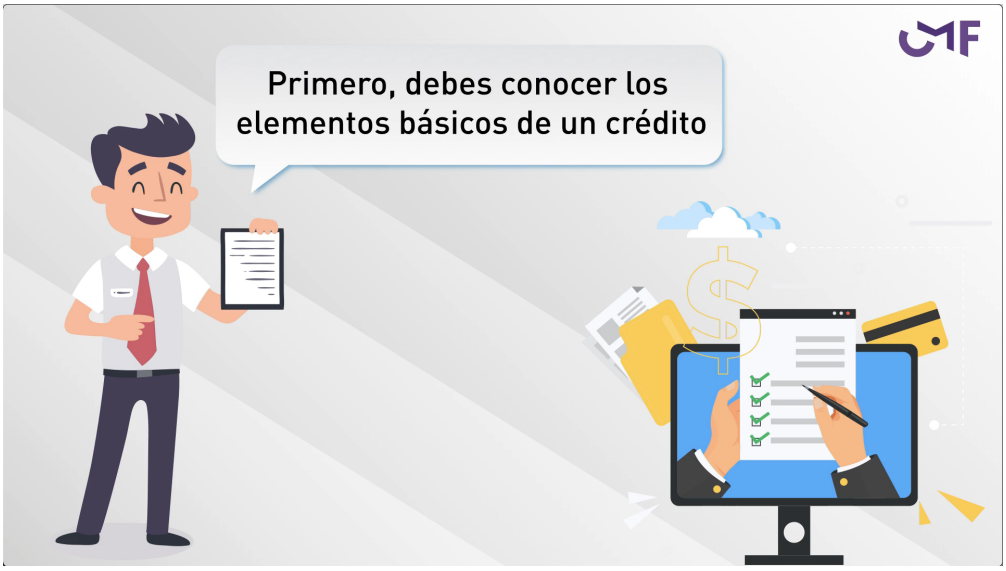
This figure shows the user interface of SERNAC for entering input information. It was last accessed on October 29, 2023. SERNAC is the Chilean equivalent of a consumer financial protection bureau. Prospective borrowers enter their inputs and in tern receive interest rate quotes for the desired loan searched. The underlying data are derived from online bank loan simulators. We only show a screenshot of the inputs of the SERNAC simulator because the website of their simulator, `https://www.sernac.cl/app/comparador/`, is currently down (as of April–July 2024) and thus we are unable to obtain a screenshot of the outputs. We scraped data from SERNAC and captured a screenshot of the inputs screen prior to the website being down.

Figure A.9: Sample Google Advertisement for Participants Recruitment



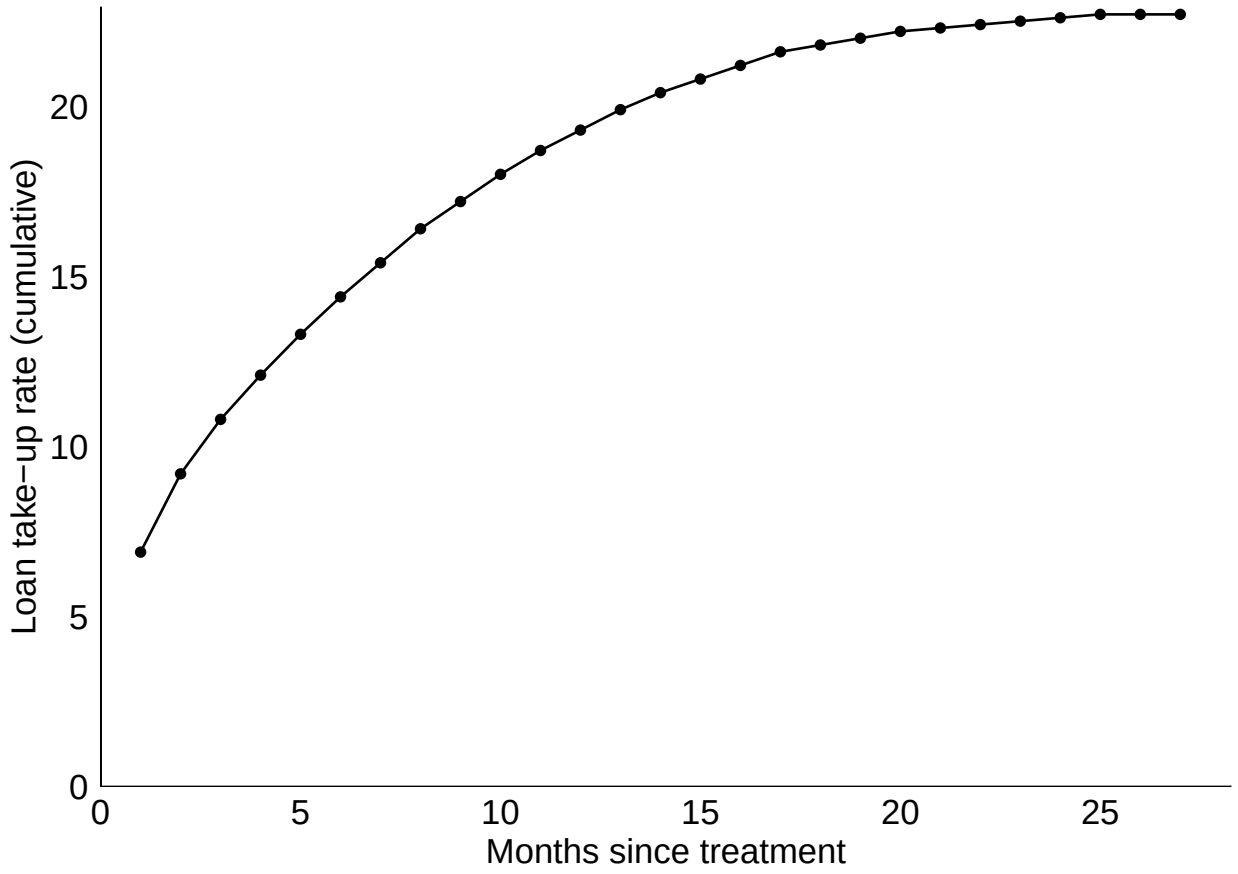
This figure shows an English translation of one of our Google advertisements that we targeted to people searching for keywords related to consumer loans in Chile to recruit them as participants in the RCT.

Figure A.10: Screenshot of Control Video



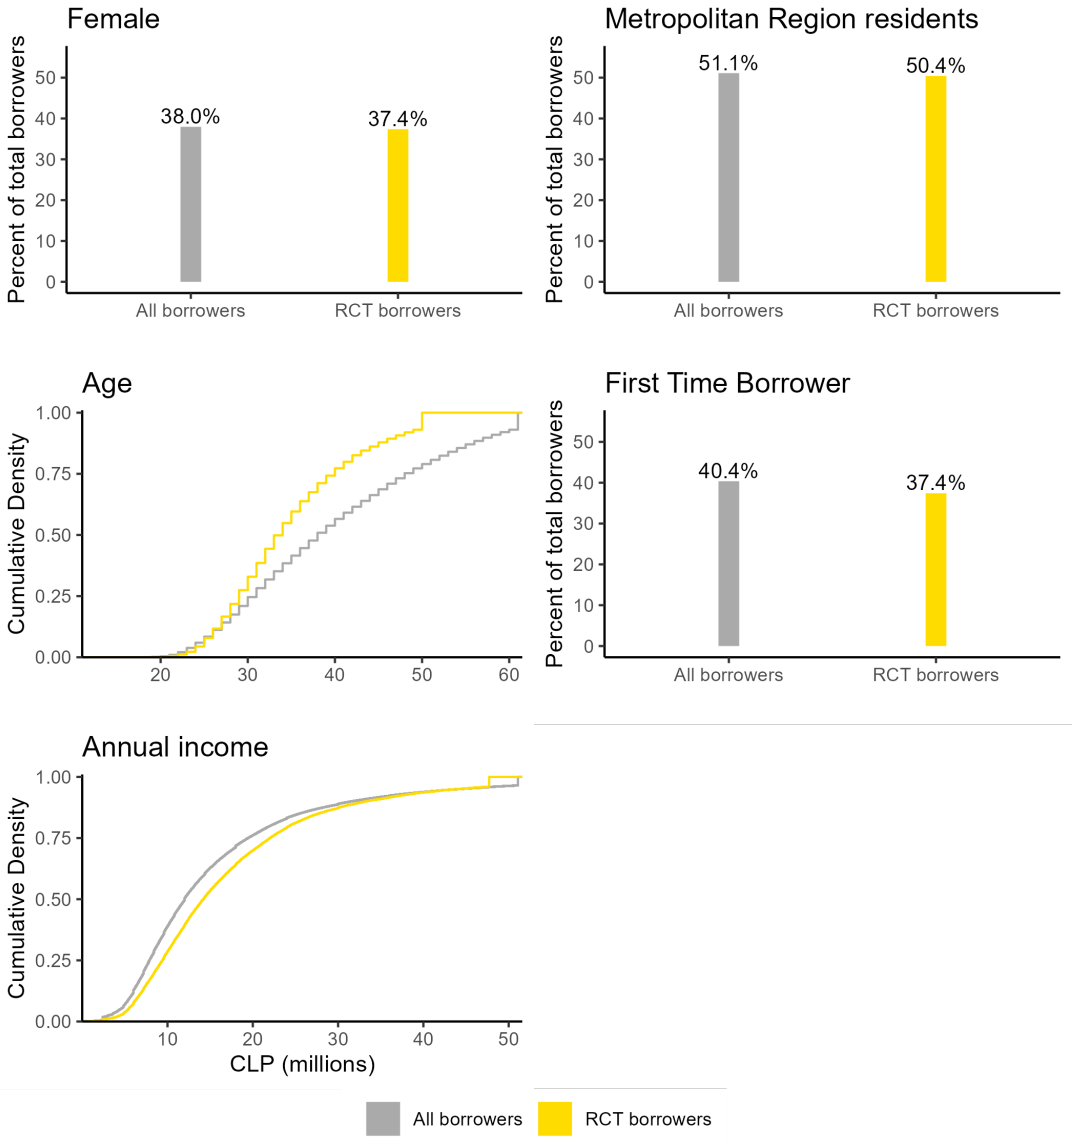
This figure shows a screenshot of the animated video shown to the control group. The video lasts 1 minute and 35 seconds and was developed by the Comisión Mercado Financiero (CMF) to provide basic loan terminology, but not provide information that would affect search.

Figure A.11: Participant Loan Take Up Rate Since Treatment



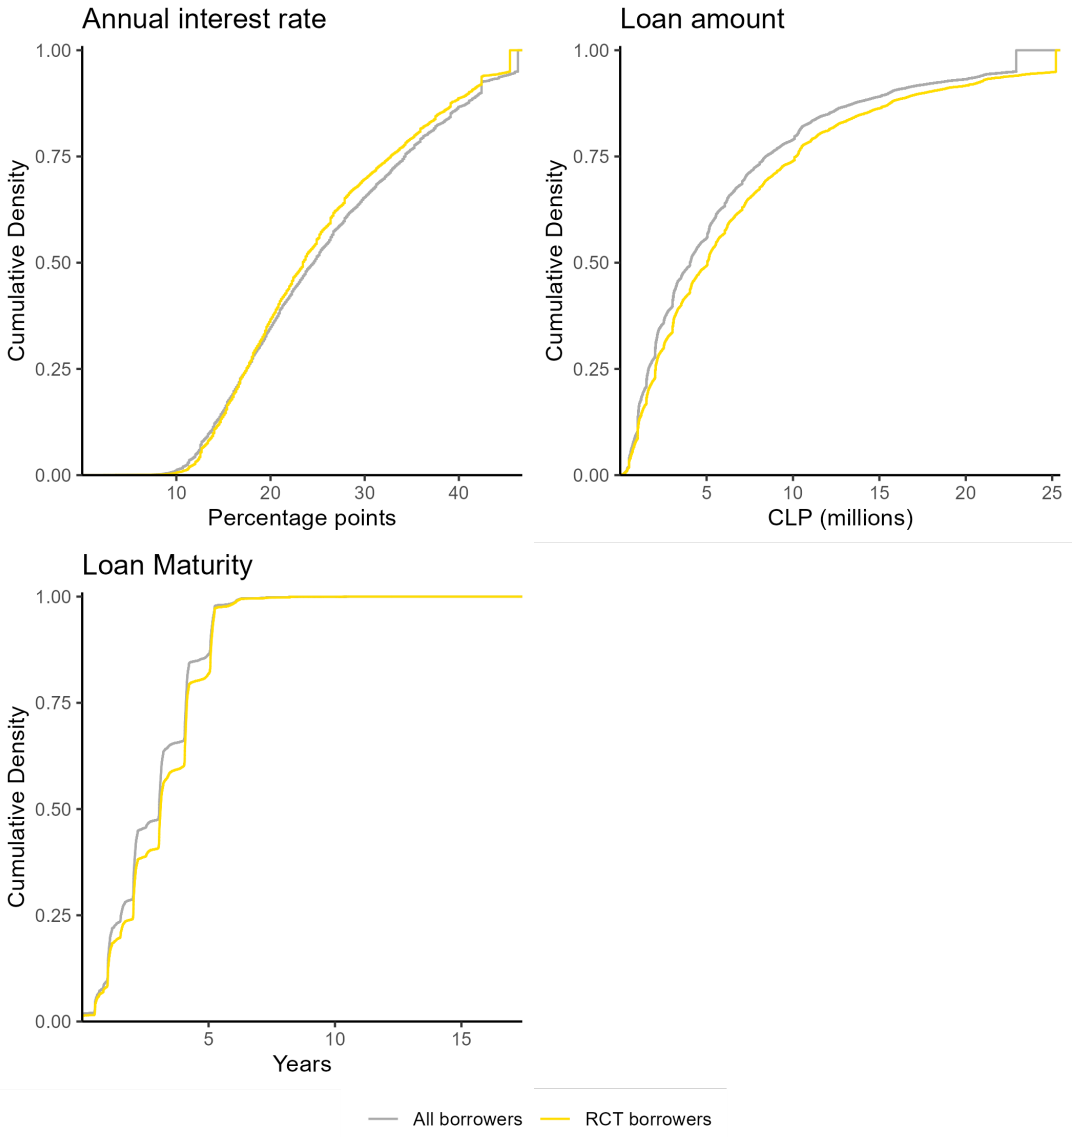
The figure shows the cumulative loan take-up rates of consumer loan borrowers of the 46,051 participants who were assigned to either our control video, price comparison tool, and simple tool. Overall, 10,448 of our RCT participants ended up taking out a consumer loan. We define loan take up as the participants having a loan in our administrative data on bank consumer loans.

Figure A.12: External Validity: Personal Characteristics



This figure shows the comparison of borrower attributes for all bank consumer borrowers taking out bank consumer loans in the sample period with borrowers who received a bank consumer loan and participated in our RCT. We have 27,130 loans taken out by RCT borrowers and 1,348,637 loans taken out by all consumer loan takers from November 2021 to February 2024.

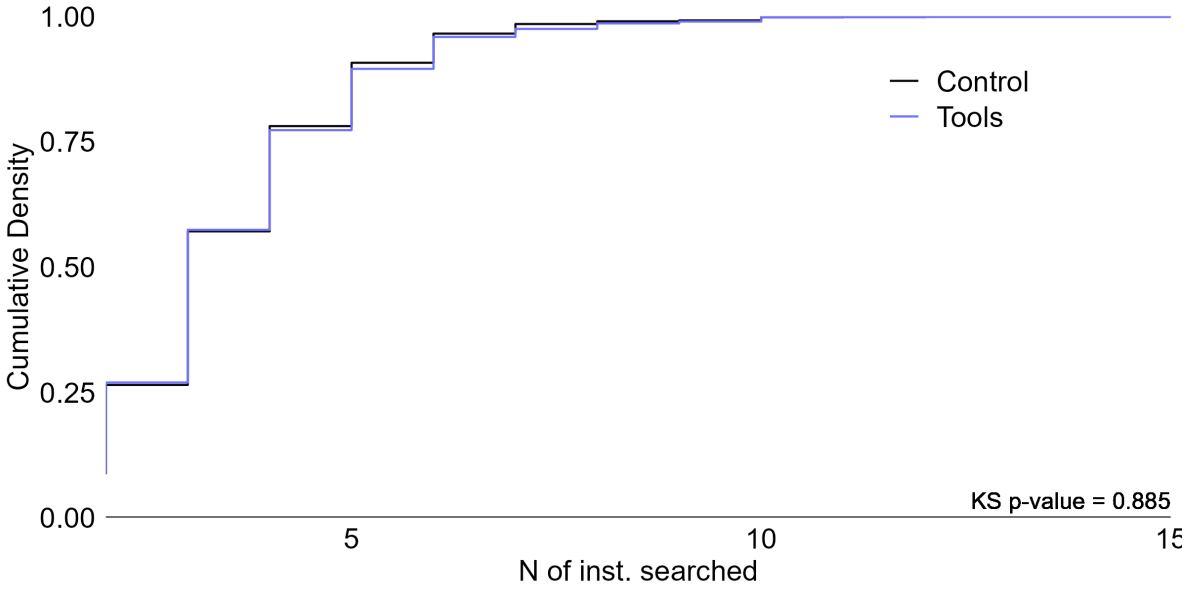
Figure A.13: External Validity: Loan Terms



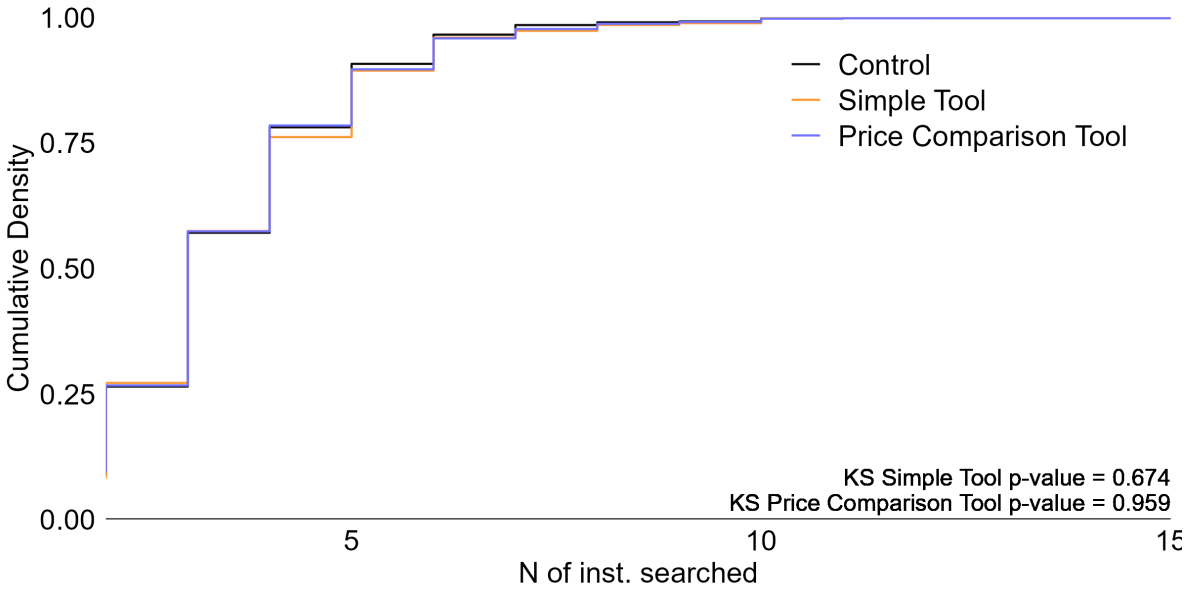
This figure shows the comparison of borrower loan terms for all bank consumer borrowers taking out bank consumer loans in the sample period with borrowers who received a bank consumer loan and participated in our RCT. We have 27,130 loans taken out by RCT borrowers and 1,348,637 loans taken out by all consumer loan takers from November 2021 to February 2024.

Figure A.14: Cumulative Distribution Function of Number of Institutions Searched

(a) Pooled Tool Treatments



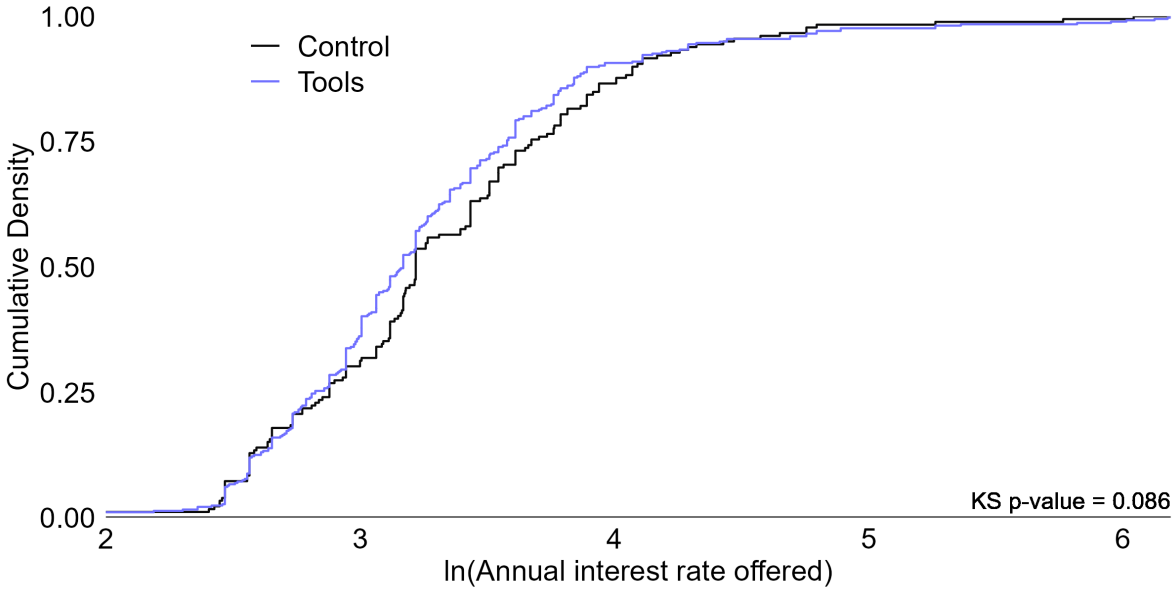
(b) By Treatment Arm



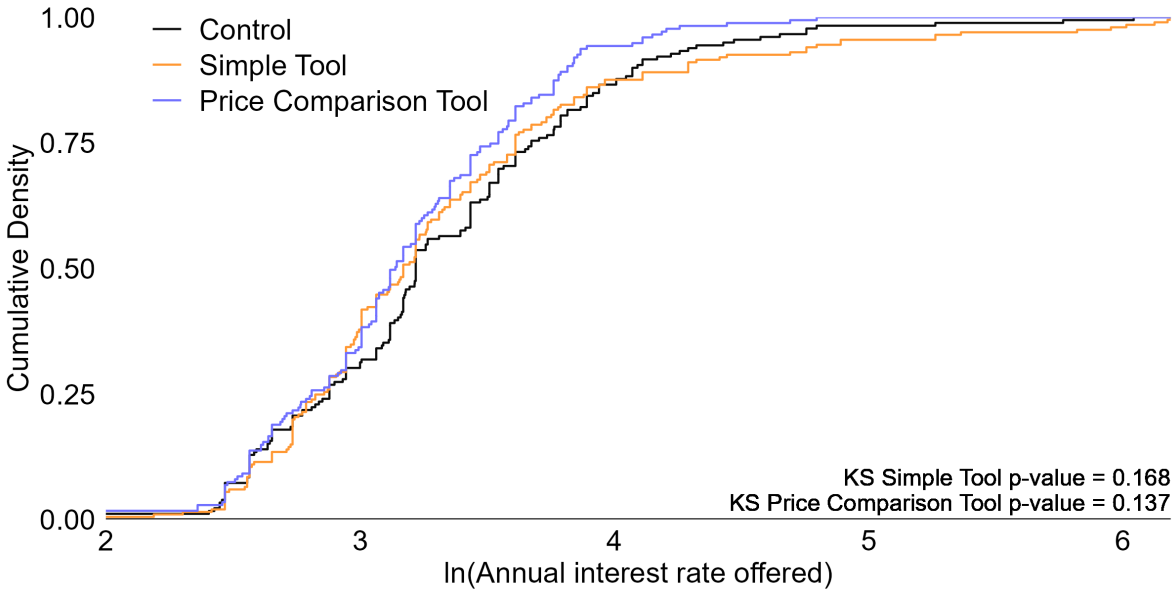
This figure shows the cumulative distribution function (CDF) of 3,283 participants that reported the number of institutions they searched at in our follow-up survey. Panel (a) shows the CDFs estimated by two groups: the tools treatment that pool participants who were assigned to the price comparison tool or the simple tool treatment, and the control group. The Kolmogórov-Smirnov test p-value is 0.88471 and is calculated using Monte Carlo simulations with 10,000 replications. The number of observations is 3,283. Panel (b) shows the CDFs estimated separated by treatment arm. The Kolmogórov-Smirnov test is estimated comparing each tool treatment to the control group and is calculated by Monte Carlo simulations with 10,000 replications. The p-value for the price comparison tool is 0.9588 and for the simple tool is 0.67373.

Figure A.15: Cumulative Distribution Function of Log Interest Rates Offered

(a) Pooled Tool Treatments



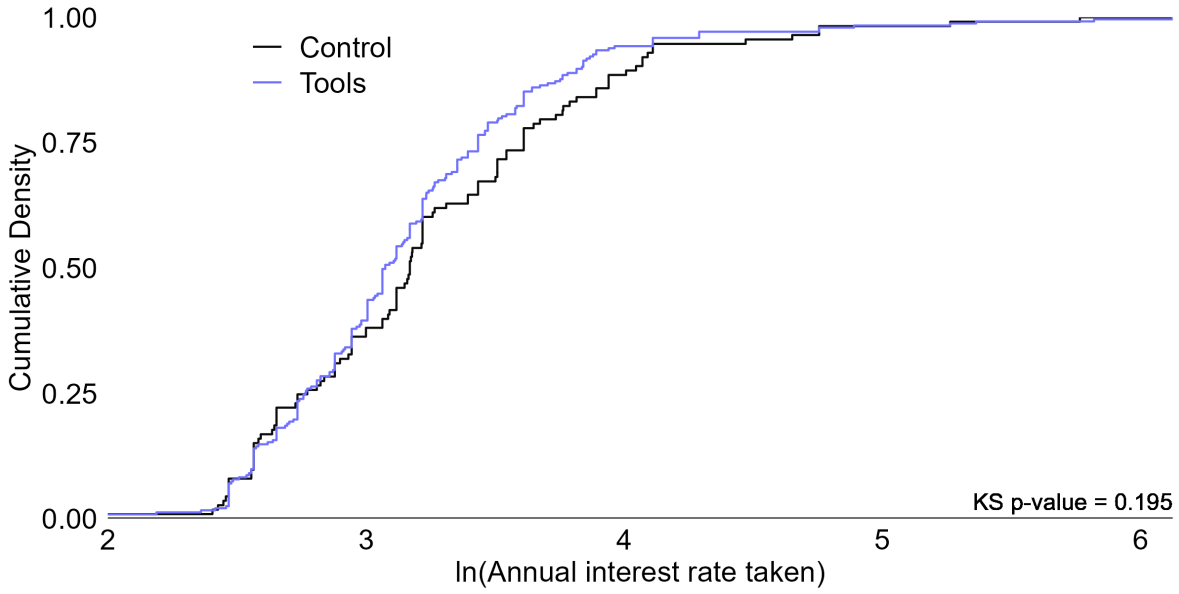
(b) By Treatment Arm



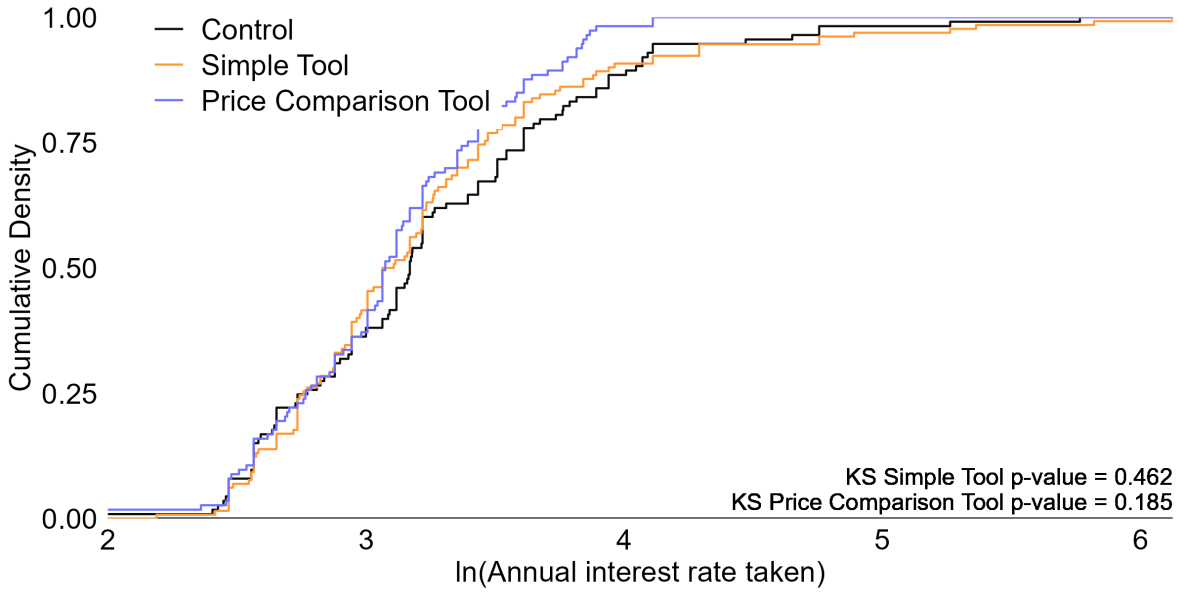
This figure shows the cumulative distribution function (CDF) of the natural logarithm of 555 offered interest rates that participants reported receiving in our follow-up survey. Panel (a) shows the CDFs estimated by two groups: The tools treatment that pool participants who were assigned to the price comparison tool or the simple tool treatment, and the control group. The Kolmogórov-Smirnov test p-value is 0.08609 and is calculated using Monte Carlo simulations with 10,000 replications. The number of observations is 555. Panel (b) shows the CDFs estimated separated by treatment arm. The Kolmogórov-Smirnov test is estimated comparing each tool treatment to the control group and is calculated by Monte Carlo simulations with 10,000 replications. The p-value for the price comparison tool is 0.12639 and for the Simple Tool is 0.16478.

Figure A.16: Cumulative Distribution Function of Log Interest Rates Taken

(a) Pooled Tool Treatments



(b) By Treatment Arm



This figure shows the cumulative distribution function (CDF) of the log interest rates that the participants who took out a loan reported receiving from banks in our follow-up phone survey. Panel (a) shows the CDFs estimated by two groups, the tools treatment that pools participants who were assigned to the price comparison tool or the simple tool, and the control group. The Kolmogórov-Smirnov test p-value is 0.18678 and is calculated by Monte Carlo simulations with 10,000 replications. The number of observations is 356. Panel (b) shows the CDFs estimated separated by treatment arm. The Kolmogórov-Smirnov test is estimated comparing each tool treatment to the control group. The p-value for the price comparison tool is 0.18538 and for the simple tool is 0.46575. The Kolmogórov-Smirnov test is calculated by Monte Carlo simulations with 10,000 replications. The number of observations is 356.

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

		Borrower characteristics								Loan characteristics			
	Scraped (1)	Name (2)	RUT (3)	Document number (4)	Income (5)	Phone number/email (6)	Comuna (7)	Employment condition (8)	Other active loans (9)	Loan amount (10)	Maturity (11)	First payment date (12)	Insurance options (13)
Panel A: Bank websites													
Banco Santander			Y		Y					Y	Y	Y	Y
Banco Estado			Y							Y	Y	Y	Y
Banco de Chile			Y		Y					Y	Y	Y	Y
Banco BCI		Y	Y		Y	Y				Y	Y	Y	Y
Scotiabank	Y		Y		Y					Y	Y	Y	Y
Banco BICE	Y									Y	Y		Y
Banco Falabella	Y		Y							Y	Y	Y	Y
Banco Internacional	Y		Y	Y		Y				Y	Y	Y	Y
Banco Ripley			Y										
Banco Security	Y		Y							Y	Y	Y	Y
Consorchio	Y	Y	Y		Y	Y				Y	Y		Y
Coopeuch	Y		Y			Y				Y	Y		
Panel B: Comparison Tools													
ComparaOnline	Y									Y	Y		
Destacame		Y	Y	Y	Y	Y		Y	Y				
SERNAC	Y									Y	Y		
Rankia						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. Column (1) shows whether we were able to scrape data from each bank website or third-party comparison tool. We were not able to scrape data from some bank websites for various reasons. The websites of Banco Santander and Banco Estado had robust anti-bot firewalls in place that prevented scraping. Banco de Chile's simulator consistently returned errors when attempting to initiate the simulation process. Banco BCI's simulator required the user to have a BCI digital account in order to obtain interest rate information. Finally, access to Banco Ripley's consumer loan simulator was only available as a paid service. Column (2), "Name", refers to the name of the person searching for information. Column (3), "RUT", refers to the *rol único tributario*, the national ID number in Chile. Column (4), "Document number", refers to the serial number on the national identity card which is distinct from the national ID number or RUT. Column (7), "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 (12), "First payment date", can be either any specific day chosen by the customer, or the date the simulator is used plus one or more complete months, depending on the simulator. Screenshots providing more details about each bank website and third-party comparison website, as well as the process we used to scrape data from these sites, are provided in Appendix B.

Table A.2: Follow-Up Survey Response Rate

	Pr(answer the survey)	
	(1)	(2)
(Intercept)	0.157***	0.153***
	(0.004)	(0.004)
Simple Tool	−0.004	
	(0.006)	
Price Comparison Tool	−0.006	
	(0.006)	
Elicit Priors		0.004
		(0.004)
Observations	20,831	37,286

This table tests for differential response rates to the follow-up survey by tool treatment status and by elicit priors 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. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Interest Rate Expectations 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	−0.23	0.38	0.95	−1.98**
	(0.90)	(0.73)	(1.51)	(0.83)
Price Comparison Tool	19.16***	14.44***	39.11***	20.93***
	(1.62)	(1.27)	(3.02)	(1.84)
Observations	6,409	6,364	6,269	5,907
Control Mean Posterior	29.46	22.82	47.88	23.42
Control Median Posterior	18	12	25.2	10.8
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows the effect of the simple tool and price comparison tool on priors. 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 had they been assigned to the price comparison tool arm, based on their borrower and loan characteristics. 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. There are 1,129 observations with dispersion = 0 (highest rate = lowest rate). Continuous variables are winsorized at the 5th and 95th percentiles by treatment arm. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4: Interest Rate Expectations Without Controlling for Priors

	Expected rate (1)	Lowest rate (2)	Highest rate (3)	Dispersion (4)
Simple Tool	−0.98	0.12	−0.67	−2.95***
	(1.24)	(0.99)	(2.09)	(1.02)
Price Comparison Tool	22.37***	17.51***	44.35***	23.49***
	(1.90)	(1.53)	(3.49)	(1.99)
Observations	7,330	7,190	7,084	6,888
Control Mean Posterior	30.285	23.189	48.624	23.968
Control Median Posterior	18	12	25	12.2
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows the effect of the simple tool and price comparison tool on priors. We estimate the following specification: $Posterior_i = \beta_1 \mathbb{1}(\text{Simple Tool})_i + \beta_2 \mathbb{1}(\text{Price Comparison Tool})_i + \lambda_{b(i)} + \varepsilon_i$, $Prior_i$ 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 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 had they been assigned to the price comparison tool arm, based on their borrower and loan characteristics. 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. There are 1,129 observations with dispersion = 0 (highest rate = lowest rate). Continuous variables are winsorized at the 5th and 95th percentiles by treatment arm. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Log Interest Rate Expectations Controlling for Priors

	ln(Expected rate) (1)	ln(Lowest rate) (2)	ln(Highest rate) (3)	ln(Dispersion) (4)
ln(Prior)	0.689***	0.698***	0.682***	0.574***
	(0.011)	(0.011)	(0.011)	(0.013)
Simple Tool	−0.037	−0.003	−0.038	−0.092***
	(0.024)	(0.023)	(0.025)	(0.034)
Price Comparison Tool	0.310***	0.271***	0.367***	0.343***
	(0.029)	(0.028)	(0.030)	(0.040)
Observations	6,409	6,364	6,269	5,907
Control Mean Posterior	2.736	2.505	3.163	2.317
Control Median Posterior	2.944	2.565	3.266	2.468
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows the effect of treatment on priors. 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 had they been assigned to the price comparison tool arm, based on their borrower and loan characteristics. 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. There are 1,129 observations with dispersion = 0 (highest rate = lowest rate). Continuous variables are winsorized at the 5th and 95th percentiles by treatment arm. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: Log Interest Rate Expectations Without Controlling for Priors

	ln(Expected rate) (1)	ln(Lowest rate) (2)	ln(Highest rate) (3)	ln(Dispersion) (4)
Simple Tool	−0.050	−0.020	−0.055	−0.124***
	(0.034)	(0.034)	(0.036)	(0.041)
Price Comparison Tool	0.413***	0.385***	0.470***	0.403***
	(0.035)	(0.035)	(0.038)	(0.045)
Observations	7,330	7,190	7,084	6,888
Control Mean Posterior	2.73	2.491	3.148	2.299
Control Median Posterior	2.944	2.565	3.258	2.416
Bin Density FEs	Yes	Yes	Yes	Yes

This table shows the effect of treatment on priors. 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 $Prior_i$ 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 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 had they been assigned to the price comparison tool arm, based on their borrower and loan characteristics. 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. There are 1,129 observations with dispersion = 0 (highest rate = lowest rate). Continuous variables are winsorized at the 5th and 95th percentiles by treatment arm. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Interest Rate Expectations Normalized Dispersion

	Normalized Dispersion
	(1)
Simple Tool	−0.01
	(0.01)
Price Comparison Tool	0.04***
	(0.01)
Observations	5,907
Control Mean Posterior	0.672
Control Median Posterior	0.667
Bin Density FEs	Yes

This table shows the effect of the simple tool and price comparison tool on priors. 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 had they been assigned to the price comparison tool arm, based on their borrower and loan characteristics. 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. There are 1,129 observations with dispersion = 0 (highest rate = lowest rate). Continuous variables are winsorized at the 5th and 95th percentiles by treatment arm. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.8: Balance of Tool Treatment Arms for Sample in Table 4, Column (1)

		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]	
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 sample in Table 4, column (1). 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. "Loan characteristics" refer to characteristics of the loan they are 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, and cash advance. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.9: Balance of Tool Treatment Arms for Sample in Table 4, Column (2)

		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	36.452***	-0.347	0.073	0.544	3,167
	(0.304)	(0.427)	(0.437)	[0.58]	
log(Income)	13.565***	-0.003	0.000	0.002	3,117
	(0.033)	(0.051)	(0.049)	[0.998]	
Incomplete high-school	0.021***	0.002	0.001	0.066	3,094
	(0.004)	(0.007)	(0.006)	[0.936]	
Complete high-school	0.377***	-0.031	-0.025	1.174	3,094
	(0.015)	(0.021)	(0.021)	[0.309]	
Complete 2-year program	0.206***	0.038**	0.011	2.214	3,094
	(0.013)	(0.018)	(0.018)	[0.109]	
Complete 5-year program or higher	0.396***	-0.010	0.012	0.534	3,094
	(0.015)	(0.022)	(0.022)	[0.586]	
Financial products					
Bank account	0.654***	0.018	0.028	0.921	3,050
	(0.015)	(0.021)	(0.021)	[0.398]	
Any loan	0.699***	0.029	0.002	1.379	3,074
	(0.014)	(0.020)	(0.020)	[0.252]	
Loan characteristics					
log(Loan Amount)	15.002***	0.037	0.031	0.229	3,002
	(0.041)	(0.060)	(0.059)	[0.795]	
log(Maturity (years))	1.361***	0.015	0.007	0.139	2,877
	(0.019)	(0.028)	(0.027)	[0.87]	
Omnibus F-statistic					
Price Comparison Tool		1.113			2,094
		[0.338]			
Simple Tool			1.599*		2,131
			[0.066]		
Number of participants by arm	1,058	1,036	1,073		3,167

This table tests the balance of pre-treatment characteristics across treatment arms for the sample in Table 4, column (2). 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. "Loan characteristics" refer to characteristics of the loan they are 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, and cash advance. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.10: Balance of Tool Treatment Arms for Sample in Table 4, Column (3)

		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	36.447***	-0.319	0.052	0.435	3,147
	(0.305)	(0.428)	(0.438)	[0.647]	
log(Income)	13.564***	-0.002	0.002	0.003	3,097
	(0.034)	(0.051)	(0.050)	[0.997]	
Incomplete high-school	0.021***	0.002	0.001	0.064	3,074
	(0.005)	(0.007)	(0.006)	[0.938]	
Complete high-school	0.376***	-0.030	-0.025	1.147	3,074
	(0.015)	(0.021)	(0.021)	[0.318]	
Complete 2-year program	0.206***	0.038**	0.011	2.206	3,074
	(0.013)	(0.018)	(0.018)	[0.11]	
Complete 5-year program or higher	0.396***	-0.010	0.013	0.564	3,074
	(0.015)	(0.022)	(0.022)	[0.569]	
Financial products					
Bank account	0.655***	0.016	0.027	0.853	3,032
	(0.015)	(0.021)	(0.021)	[0.426]	
Any loan	0.701***	0.027	0.001	1.149	3,056
	(0.014)	(0.020)	(0.020)	[0.317]	
Loan characteristics					
log(Loan Amount)	15.000***	0.042	0.033	0.283	2,983
	(0.041)	(0.060)	(0.059)	[0.754]	
log(Maturity (years))	1.361***	0.015	0.004	0.152	2,861
	(0.020)	(0.028)	(0.027)	[0.859]	
Omnibus F-statistic					
Price Comparison Tool		1.063			2,083
		[0.386]			
Simple Tool			1.621*		2,116
			[0.061]		
Number of participants by arm	1,052	1,031	1,064		3,147

This table tests the balance of pre-treatment characteristics across treatment arms for the sample in Table 4, column (3). 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. "Loan characteristics" refer to characteristics of the loan they are 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, and cash advance. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.11: Balance of Tool Treatment Arms for Sample in Table 4, Column (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	36.455***	-0.318	0.051	0.43	3,143
	(0.305)	(0.429)	(0.438)	[0.65]	
log(Income)	13.564***	-0.004	0.009	0.03	3,093
	(0.034)	(0.051)	(0.049)	[0.971]	
Incomplete high-school	0.021***	0.002	0.001	0.066	3,070
	(0.005)	(0.007)	(0.006)	[0.936]	
Complete high-school	0.376***	-0.030	-0.025	1.131	3,070
	(0.015)	(0.021)	(0.021)	[0.323]	
Complete 2-year program	0.206***	0.037**	0.011	2.117	3,070
	(0.013)	(0.019)	(0.018)	[0.121]	
Complete 5-year program or higher	0.396***	-0.010	0.013	0.546	3,070
	(0.015)	(0.022)	(0.022)	[0.579]	
Financial products					
Bank account	0.655***	0.016	0.028	0.913	3,028
	(0.015)	(0.021)	(0.021)	[0.402]	
Any loan	0.701***	0.025	0.000	1.082	3,052
	(0.014)	(0.020)	(0.020)	[0.339]	
Loan characteristics					
log(Loan Amount)	15.003***	0.039	0.032	0.249	2,979
	(0.041)	(0.060)	(0.059)	[0.78]	
log(Maturity (years))	1.362***	0.014	0.003	0.149	2,857
	(0.020)	(0.028)	(0.027)	[0.861]	
Omnibus F-statistic					
Price Comparison Tool		1.041			2,080
		[0.409]			
Simple Tool			1.62*		2,114
			[0.061]		
Number of participants by arm	1,051	1,029	1,063		3,143

This table tests the balance of pre-treatment characteristics across treatment arms for the sample in Table 4, column (3). 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. "Loan characteristics" refer to characteristics of the loan they are 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, and cash advance. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.12: Balance of Tool Treatment Arms for Sample in Table 4, Column (8)

	Control Mean	Difference relative to control mean: Price Comparison Tool	Difference relative to control mean: Simple Tool	Joint test F-stat	N
	(1)	(2)	(3)	(4)	(5)
Personal characteristics					
Age	35.112***	0.054	0.233	0.61	8,868
	(0.157)	(0.220)	(0.222)	[0.544]	
log(Income)	13.905***	0.028	0.023	1.263	8,746
	(0.014)	(0.019)	(0.019)	[0.283]	
Incomplete high-school	0.007***	0.002	0.000	0.267	8,715
	(0.002)	(0.002)	(0.002)	[0.765]	
Complete high-school	0.244***	-0.012	-0.015	0.997	8,715
	(0.008)	(0.011)	(0.011)	[0.369]	
Complete 2-year program	0.205***	0.011	0.015	1.066	8,715
	(0.008)	(0.011)	(0.011)	[0.345]	
Complete 5-year program or higher	0.544***	-0.001	-0.001	0.002	8,715
	(0.009)	(0.013)	(0.013)	[0.998]	
Financial products					
Bank account	0.863***	0.017*	0.005	2.005	8,731
	(0.006)	(0.009)	(0.009)	[0.135]	
Any loan	0.882***	0.001	0.003	0.067	8,761
	(0.006)	(0.008)	(0.008)	[0.936]	
Loan characteristics					
log(Loan Amount)	15.429***	0.059**	0.043	2.063	8,491
	(0.021)	(0.030)	(0.030)	[0.127]	
log(Maturity (years))	1.426***	0.040***	0.022	3.361**	8,266
	(0.011)	(0.015)	(0.015)	[0.035]	
Omnibus F-statistic					
Price Comparison Tool		1.367			5,905
		[0.154]			
Simple Tool			0.69		5,847
			[0.797]		
Number of participants by arm	2,884	3,021	2,963		8,868

This table tests the balance of pre-treatment characteristics across treatment arms for the sample in Table 4, column (3). 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. "Loan characteristics" refer to characteristics of the loan they are 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, and cash advance. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.13: Balance by Elicit Priors Treatment for Sample in Table 5, Column (1)

	Elicit Priors = 0 Mean	Elicit Priors	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 priors treatment for the sample in Table 5, column (1). We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta \mathbb{1}(\text{Elicit Priors})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy indicating whether participant i was assigned to the elicit priors treatment. Column (1) shows α which is the mean for the $\mathbb{1}(\text{Elicit Priors})_i = 0$ group. Column (2) shows β which is the difference in means between the $\mathbb{1}(\text{Elicit Priors})_i = 0$ and $\mathbb{1}(\text{Elicit Priors})_i = 1$ 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: $\mathbb{1}(\text{Elicit Priors})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$, where $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy equal to 1 if participant i was assigned to the elicit priors treatment. 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. "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, and 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 priors treatment (rather than a prior module), and the elicit priors treatment affected whether participants continued in the survey. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.14: Balance by Elicit Priors Treatment for Sample in Table 5, Column (2)

	Elicit Priors = 0 Mean (1)	Elicit Priors (2)	N (3)
Personal characteristics			
Age	36.700***	-0.325	5,565
	(0.252)	(0.295)	
log(Income)	13.610***	0.031	5,466
	(0.032)	(0.037)	
Incomplete high-school	0.025***	0.000	5,435
	(0.004)	(0.005)	
Complete high-school	0.339***	-0.012	5,435
	(0.012)	(0.014)	
Complete 2-year program	0.211***	-0.001	5,435
	(0.011)	(0.012)	
Complete 5-year program or higher	0.424***	0.013	5,435
	(0.013)	(0.015)	
Financial products			
Bank account	0.689***	0.010	5,352
	(0.012)	(0.014)	
Any loan	0.738***	-0.009	5,395
	(0.011)	(0.013)	
Omnibus F-statistic		0.719	5,565
		[0.721]	
Number of participants by arm	1,518	4,047	5,565

This table tests the balance of pre-treatment characteristics by elicit priors treatment for the sample in Table 5, column (2). We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta \mathbb{1}(\text{Elicit Priors})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy indicating whether participant i was assigned to the elicit priors treatment. Column (1) shows α which is the mean for the $\mathbb{1}(\text{Elicit Priors})_i = 0$ group. Column (2) shows β which is the difference in means between the $\mathbb{1}(\text{Elicit Priors})_i = 0$ and $\mathbb{1}(\text{Elicit Priors})_i = 1$ 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: $\mathbb{1}(\text{Elicit Priors})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$, where $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy equal to 1 if participant i was assigned to the elicit priors treatment. 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. "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, and 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 priors treatment (rather than a prior module), and the elicit priors treatment affected whether participants continued in the survey. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.15: Balance by Elicit Priors Treatment for Sample in Table 5, Column (3)

	Elicit Priors = 0 Mean (1)	Elicit Priors (2)	N (3)
Personal characteristics			
Age	36.719***	-0.371	5,525
	(0.254)	(0.297)	
log(Income)	13.606***	0.035	5,427
	(0.032)	(0.037)	
Incomplete high-school	0.025***	0.000	5,396
	(0.004)	(0.005)	
Complete high-school	0.340***	-0.013	5,396
	(0.012)	(0.014)	
Complete 2-year program	0.212***	0.000	5,396
	(0.011)	(0.013)	
Complete 5-year program or higher	0.423***	0.014	5,396
	(0.013)	(0.015)	
Financial products			
Bank account	0.690***	0.009	5,315
	(0.012)	(0.014)	
Any loan	0.739***	-0.010	5,358
	(0.012)	(0.014)	
Omnibus F-statistic		0.783	5,525
		[0.657]	
Number of participants by arm	1,504	4,021	5,525

This table tests the balance of pre-treatment characteristics by elicit priors treatment for the sample in Table 5, column (3). We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta \mathbb{1}(\text{Elicit Priors})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy indicating whether participant i was assigned to the elicit priors treatment. Column (1) shows α which is the mean for the $\mathbb{1}(\text{Elicit Priors})_i = 0$ group. Column (2) shows β which is the difference in means between the $\mathbb{1}(\text{Elicit Priors})_i = 0$ and $\mathbb{1}(\text{Elicit Priors})_i = 1$ 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: $\mathbb{1}(\text{Elicit Priors})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$, where $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy equal to 1 if participant i was assigned to the elicit priors treatment. 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. "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, and 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 priors treatment (rather than a prior module), and the elicit priors treatment affected whether participants continued in the survey. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.16: Balance by Elicit Priors Treatment for Sample in Table 5, Column (4)

	Elicit Priors = 0 Mean (1)	Elicit Priors (2)	N (3)
Personal characteristics			
Age	36.729***	-0.385	5,516
	(0.254)	(0.297)	
log(Income)	13.611***	0.029	5,418
	(0.032)	(0.037)	
Incomplete high-school	0.025***	0.000	5,387
	(0.004)	(0.005)	
Complete high-school	0.340***	-0.013	5,387
	(0.012)	(0.014)	
Complete 2-year program	0.211***	0.000	5,387
	(0.011)	(0.013)	
Complete 5-year program or higher	0.423***	0.013	5,387
	(0.013)	(0.015)	
Financial products			
Bank account	0.691***	0.008	5,306
	(0.012)	(0.014)	
Any loan	0.738***	-0.010	5,349
	(0.012)	(0.014)	
Omnibus F-statistic		0.767	5,516
		[0.673]	
Number of participants by arm	1,502	4,014	5,516

This table tests the balance of pre-treatment characteristics by elicit priors treatment for the sample in Table 5, column (4). We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta \mathbb{1}(\text{Elicit Priors})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy indicating whether participant i was assigned to the elicit priors treatment. Column (1) shows α which is the mean for the $\mathbb{1}(\text{Elicit Priors})_i = 0$ group. Column (2) shows β which is the difference in means between the $\mathbb{1}(\text{Elicit Priors})_i = 0$ and $\mathbb{1}(\text{Elicit Priors})_i = 1$ 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: $\mathbb{1}(\text{Elicit Priors})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$, where $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy equal to 1 if participant i was assigned to the elicit priors treatment. 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. "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, and 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 priors treatment (rather than a prior module), and the elicit priors treatment affected whether participants continued in the survey. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.17: Balance by Elicit Priors Treatment for Sample in Table 5, Column (8)

	Elicit Priors = 0 Mean	Elicit Priors	N
	(1)	(2)	(3)
Personal characteristics			
Age	35.217***	-0.024	21,102
	(0.111)	(0.128)	
log(Income)	14.042***	0.003	20,852
	(0.009)	(0.011)	
Incomplete high-school	0.007***	0.000	20,802
	(0.001)	(0.001)	
Complete high-school	0.207***	0.002	20,802
	(0.006)	(0.006)	
Complete 2-year program	0.199***	0.000	20,802
	(0.005)	(0.006)	
Complete 5-year program or higher	0.586***	-0.002	20,802
	(0.007)	(0.008)	
Financial products			
Bank account	0.887***	0.008	20,828
	(0.004)	(0.005)	
Any loan	0.889***	-0.002	20,892
	(0.004)	(0.005)	
Omnibus F-statistic		0.456	21,102
		[0.93]	
Number of participants by arm	5,409	15,693	21,102

This table tests the balance of pre-treatment characteristics by elicit priors treatment for the sample in Table 5, column (8). We run the following regression separately for each baseline covariate k: $X_i^k = \alpha + \beta \mathbb{1}(\text{Elicit Priors})_i + \varepsilon_i$, where X_i^k is a baseline covariate for participant i and $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy indicating whether participant i was assigned to the elicit priors treatment. Column (1) shows α which is the mean for the $\mathbb{1}(\text{Elicit Priors})_i = 0$ group. Column (2) shows β which is the difference in means between the $\mathbb{1}(\text{Elicit Priors})_i = 0$ and $\mathbb{1}(\text{Elicit Priors})_i = 1$ 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: $\mathbb{1}(\text{Elicit Priors})_i = \delta + \sum_{k=1}^{K} \gamma_k X_i^k + \varepsilon_i$, where $\mathbb{1}(\text{Elicit Priors})_i$ is a dummy equal to 1 if participant i was assigned to the elicit priors treatment. 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. "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, and 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 priors treatment (rather than a prior module), and the elicit priors treatment affected whether participants continued in the survey. Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.18: Mechanisms Behind Effects of Price Comparison Tool

	Search at different inst. than planned (1)	Pr(convey info from tool before applying) (2)	N of inst. negotiated (3)	Pr(successfully negotiate) (4)	N of inst. successfully negotiate (5)
(Intercept)	0.730***	0.000	0.114***	0.054***	0.059***
	(0.019)		(0.011)	(0.007)	(0.008)
Simple Tool	0.011	0.138***	0.016	0.004	0.009
	(0.027)	(0.039)	(0.018)	(0.010)	(0.012)
Price Comparison Tool	0.011	0.145***	0.050***	0.019*	0.022*
	(0.027)	(0.043)	(0.018)	(0.011)	(0.012)
Observations	1,614	190	3,114	3,103	3,103

This table shows the effect of the simple tool and price comparison tool on search and negotiating behavior. 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 (but who were nevertheless included in the survey because they did reach the module in which the elicit priors treatment was assigned). Column (1) is a dummy variable equal to 1 if the participant reported searching at at least one different institution in the follow-up survey than the institutions they reported planning to search first and second in the baseline survey. This sample excludes participants assigned to $\mathbb{1}(\text{Elicit priors})_i = 0$ as they did not receive the questions about which institutions they planned to search, as well as those assigned $\mathbb{1}(\text{Elicit priors})_i = 1$ who did not report which institutions they planned to search first and second in the baseline survey or who did not know or refused to answer how many institutions searched or the names of those institutions in the follow-up survey. Column (2) is a dummy variable equal to 1 if the participant indicated they conveyed information from one of our tools to at least one institution where they searched and 0 if they did not or were assigned to the control group. Compared to the other columns, this sample has fewer observations because the question was added to the survey only in January 2024. Column (3) is a count variable that includes the number of institutions with which participants tried to negotiate. Compared to column (3) of Table 4, the sample also excludes participants who did not know or refused to answer whether they negotiated for all their loan offers. Column (4) is a dummy variable equal to 1 if the participant received negotiated loan terms with at least one institution and 0 if they did not, or if the participant did not negotiate. Compared to column (3), the sample also excludes participants who did not know or refused to answer whether the institution changed anything in the offer that the participant tried to negotiate. Column (5) is a count variable of the number of institutions at which participants were able to successfully negotiate (i.e., received negotiated loan terms). It uses the same sample as column (4). Heteroskedasticity-robust standard errors (not clustered since the unit of randomization is the individual) are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix B Bank Websites and Comparison Websites

In our follow-up survey data, 44% of participants report using bank websites during their search and 12% report using third-party comparison websites (aggregators). Thus, these channels are likely a way that some consumers form their prior beliefs about loan interest rates. To investigate whether the tools on banks' own websites, known in Chile as "simulators," provide accurate information. If not, these simulators potentially contribute to participants' holding inaccurate priors. We scraped data from seven banks' consumer loan simulators and two aggregators, ComparaOnline and SERNAC. ComparaOnline includes 15 institutions: Banco de Chile, Banco Estado, Banco Falabella, Banco Internacional, Banco Itaú/Corpbanca, Banco Ripley, Banco Santander, Banco Security, Scotiabank, Banca.me, Coopeuch, Caja 18, Caja Los Andes, Caja Los Héroes, and Oriencoop. SERNAC includes 12 institutions: Banco BCI, Banco BICE, Banco Consorcio, Banco de Chile, Banco Estado, Banco Falabella, Banco Internacional, Banco Ripley, Banco Santander, Scotiabank, Banco Security, and Caja Los Andes. 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 feeds these inputs into each website (including bank simulators and comparison websites) and scrapes the output. Next, we compared the rates participants would have seen on these websites with the rates they actually received in the administrative data.

B.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. Prospective borrowers input their personal information along with desired loan amount, terms, and other details. The simulator then generates loan terms including the interest rate, the "carga anual equivalente (CAE)"—comparable to the annual percentage rate (APR) in the U.S.—the "costo total del crédito (CTC)," which represents the total loan cost, the monthly cost, and the details of application costs and insurance costs. We have included screenshots of the input and output pages from these simulators for 11 major Chilean banks in Section B.5. The input variables required by each simulator are also tabulated in Table A.1, panel A. All bank websites require information on loan amount and maturity. All but Banco BICE also require the consumer's RUT (national ID number), but in tests that we show below, we find that the interest rates and other loan terms 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.

At the time of our data collection, we were not able to get data from some of the existing bank simulators. The simulators of Banco Santander and Banco Estado had robust anti-bot firewalls

in place which prevented us from batch-processing queries to these simulators. Banco de Chile's simulator consistently returned errors when attempting to initiate the simulation process. Banco BCI's simulator required the user to have a BCI digital account prior to using the simulator. Finally, access to Banco Ripley's consumer loan simulator was only available as a paid service.

B.2 Description of Comparison Websites

There are three main third-party or government-run comparison websites providing estimated loan terms from multiple banks, also known as aggregators: ComparaOnline, SERNAC, and Rankia. In addition, there is a website called Destacame that provides assistance to loan seekers and shows them potential products, but does not include estimated loan terms.

ComparaOnline ComparaOnline, operates as a quote aggregator and is run by a private-sector company. Consumers input their desired loan size and maturity and receive quotes for loans from different institutions (see Figure A.3). However, ComparaOnline does not ask for any borrower characteristics, and thus the interest rate quotes it provides are not conditional on borrower characteristics. Banks may have an incentive 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 their websites over other banks.

Destacame Destacame is another comparison website run by a private-sector company. Prospective borrowers input information related to their loan search, current employment status, tenure at the current job, monthly income, and current financial products, and institutions can submit products for the borrower to consider. Additionally, Destacame sells services such as credit counseling to improve the consumer's probability of being approved for a loan (see Figures A.4 and A.5). However, when we tested Destacame's website, the products submitted by financial institutions for the borrower to consider took 1–2 weeks to appear, and did not include loan terms such as the interest rate or loan amount. Thus, a consumer using Destacame to search for loans would still need to formally apply for a loan in order to receive an interest rate quote.

SERNAC SERNAC is the Chilean government's consumer protection bureau. SERNAC hosts a comparison website that requests only two inputs: the loan amount and maturity; it does not condition on any borrower characteristics. Furthermore, it only provides limited options for these two inputs. The available options for loan amount are 1 million to 10 million Chilean pesos in increments of 1 million, while for maturity, the available options include 12 to 60 months in increments of 12 months. According to SERNAC (2015), they collected banks' simulation data using

a method similar to ours (but not conditioning on any borrower characteristics such as income), using the consumer loan simulators available on the websites of financial institutions.

Rankia Rankia does not offer its own comparison tool. Instead, it hosts a portal (see Figure A.6) that directs users to an article on their site titled "Mejores créditos de consumo para 2024 (Best Consumer Loans for 2024)" (Figure A.7). At the top of the article, there is a button to Banco Internacional's consumer loan simulator. Further down, the article introduces SERNAC's annual study, which compares the total costs of consumer loans across a subset of institutions supported on their platform. Data from SERNAC's annual study is laid out on a table displaying the monthly interest rates and CAE (APR) for a 12-month, $10,000,000-peso loan with no insurance from eight different banks. These values remain consistent regardless of any user inputs in the initial simulation portal.

B.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 are set to be the median values and remain constant. We set the test size to 100 observations and tested the five bank websites where randomly generated ID information was used. As shown by Figure B.1 , despite occasional variations in interest rates for different RUTs from Banco Falabella and Scotiabank, 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. B.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

Table B.1: Simulators' Min/Max Configuration

	Min	Max
Scotiabank	No grace period, no insurance	Maximum grace period (6 months), Seguro Desgravamen (life insurance), and Seguro Cesantía (severance insurance)
Banco BICE	No insurance	Desgravamen con ITP (disability insurance) and Protección Laboral (labor protection)
Banco Falabella	No grace period, no insurance	Maximum grace period (2 months), Desgravamen Hospitalizacion (life and hospitalization insurance)
Banco Internacional	No grace period, no insurance	Maximum grace period (the end of the next month), Seguro Desgravamen
Banco Security	No grace period, no insurance	Maximum grace period (3 months), Seguro de desgravamen
Consorchio	No insurance	Seguro de Desgravamen

Note: This table shows the min/max inputs we used to get simulation results from banks that allow users to select their preferred grace period and insurance options.

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. Table B.1 shows the details of the minimum and maximum input configurations. 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 consumer-loan seeker: monthly interest rates, 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).

B.4 Comparison of Websites' Rates and Received Rates

To compare rates to the rate an individual in our RCT would have seen on bank and comparison websites to rates 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 the administrative data who had taken 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 ComparaOnline and the other from the bank's website, we keep all quotes. The matched sample has 14,354 observations. The main results are shown in Figure B.5.

The interest rate quotes shown by banks and comparison websites are highly inaccurate (Figure B.5). 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, as seen from Table E.3(1). Thus, they do not provide quotes conditional on all the relevant borrower characteristics that influence the interest rate. Secondly, 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. Beyond institutional features, there may be other factors related to the loan search process that can explain discrepancies between rate quotes seen on websites and actual loan rates. First, our scraped simulator data could be different from the loans participants ultimately

took out, either due to the change of loan requirements or because our tool endogenously changed their search strategy on desired terms. Second, the discrepancy could be due to banks offering the same loan on different days when the bank might have changed their pricing model. We assess each of these potential explanations.

First, we consider the possibility that consumers changed their loan characteristics from their baseline requirements to originated loan terms. For example, a participant might go on a bank website, enter their baseline characteristics, and get an initial rate estimate. Participants may then change the characteristics of their loan in response to this estimate. If the estimate is more than the participant can afford, they may reduce the amount they borrow or extend the maturity of their loan to reduce their monthly payment. Consequently, the discrepancy between simulated and actual rates may be entirely explained by borrowers changing their loan terms. 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 baseline survey and their actual loan terms. The scatter plots are presented in Figure B.3 for banks and Figure B.4 for ComparaOnline. Loan size differences are presented in panel a and maturity differences in panel b. The majority of difference points between baseline loan size and maturity are clustered along the vertical line at zero. However, there is still substantial variation in interest rates received despite these main loan terms not changing. We regress the rate difference on the differences in loan terms and find that the R^2 of these regressions are 0.173 for loan amount and 0.265 for maturity, respectively. This suggests that the differences between the observed and simulated rates can be explained by participants changing their loan terms throughout their search.

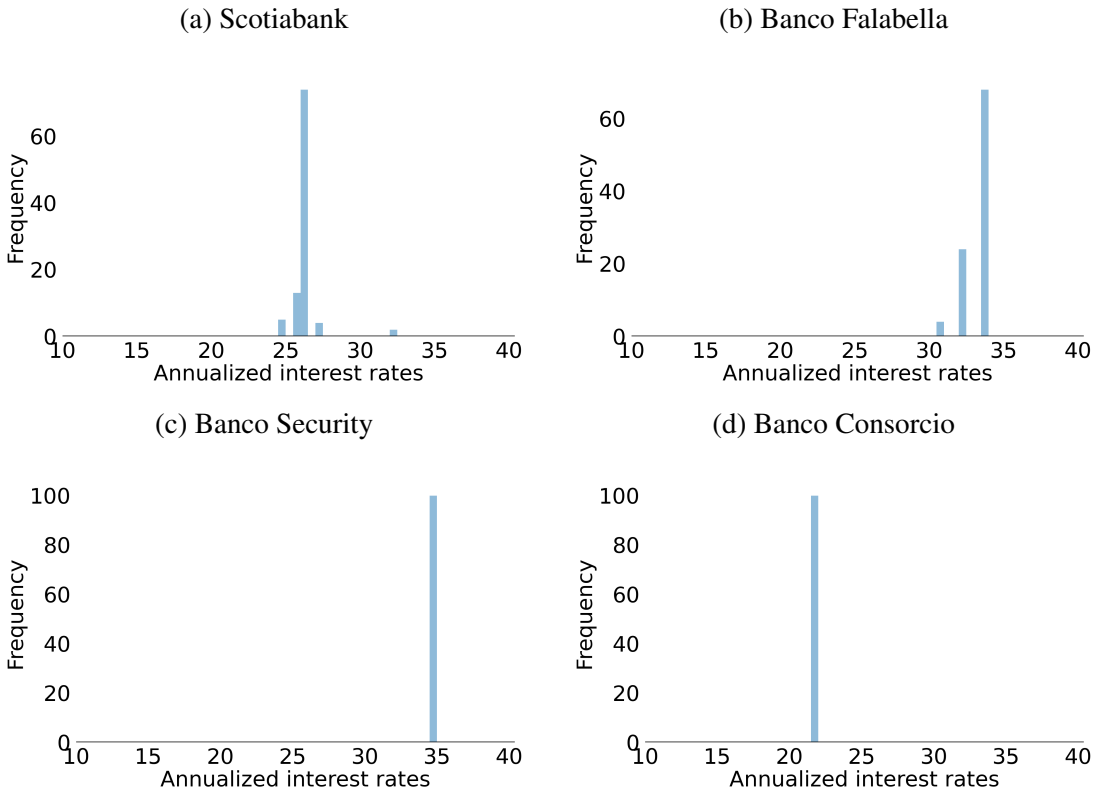
We also consider whether the final loan terms may have changed as a result of interacting with either our price comparison tool. To eliminate this effect, we compare the interest rate differences only for borrowers in our control group. The resulting histograms are presented in Figure B.5. 40.77% of rate differences are negative within bank simulator websites, and 61.94% of rate differences are negative using rates from comparison websites, as compared to 38.58% and 60.95% respectively in our full sample. The similarity in percentages suggests that this discrepancy in rates is not driven by our treatments.

Lastly, these differences may be explained by a difference in timing between when quotes were scraped and when borrowers took out loans. To address this concern, we use a restricted sample of borrowers that took out loans between September 28th and October 9th, 2023 when we scraped our data. The difference between the simulator rates and actual loan rates are plotted in Figure B.6. For these 21 borrowers, 66.67 % of rate differences are negative for bank simulators and 72.65% are negative for the ComparaOnline simulator. Given that these quotes are contemporaneous with loan issuance, we should expect the rate differences to be *smaller* here than in the full sample if timing of quotes is a factor in the rate differences. As compared to our full sample negative differences

of 38.58% and 60.95%, these negative differences are actually larger, though there could be an element of small sample bias.

Finally, we consider the comparison website SERNAC. Since SERNAC collected its data from bank simulators, we exclude its data from our main results to avoid repetition. We replicate the rate comparison exercises for our full sample using only SERNAC data. Specifically, we merged SERNAC's simulation results with the administrative data by matching each consumer's received loan amount and maturity to the closest available options in SERNAC, as well as by bank. Additionally, we also attempted to match each consumer's received loan amount with the closest *higher* available SERNAC loan amount option. Figure B.7 shows that for around 70% of the observations, SERNAC suggested *higher* interest rates than consumers would actually receive. This is a higher positive percentage than the bank simulator websites (Figure B.5), possibly due to the size and maturity increments available in the SERNAC simulator.

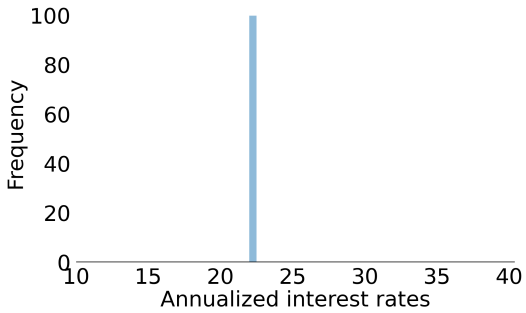
Figure B.1: Tests for Randomly Generated RUT



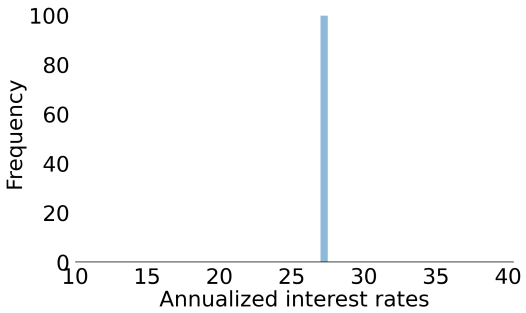
This figure shows the annualized interest rates from bank websites, derived from a test where we varied only the RUT while maintaining other inputs constant. Specifically, we standardized inputs such as income, loan amount, and loan maturity by using the median values from the baseline sample: a monthly income of $740,000 pesos, a loan amount of $2,500,000 pesos, a maturity period of 3 years, and a fake phone number of 2-181352 (where "2" is the area code). The x-axis represents the annualized interest rates calculated by each bank's website simulation. Each simulation consists of 100 observations. Notably, Banco Internacional required both the simulator user's RUT and document number for identification purposes. Since using a fake RUT was not feasible, we used the authentic RUT and document number of one of our Chilean research assistants to gather the simulation data used in the main results.

Figure B.2: Tests for Randomly Generated Area Codes

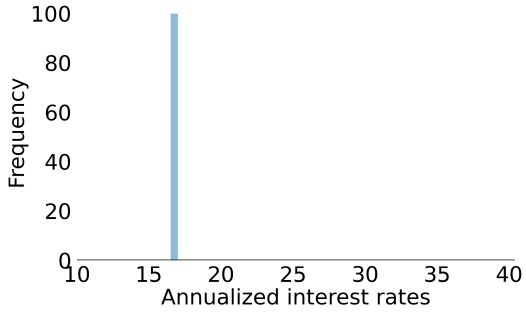
(a) Banco Consorcio



(b) Coopeuch

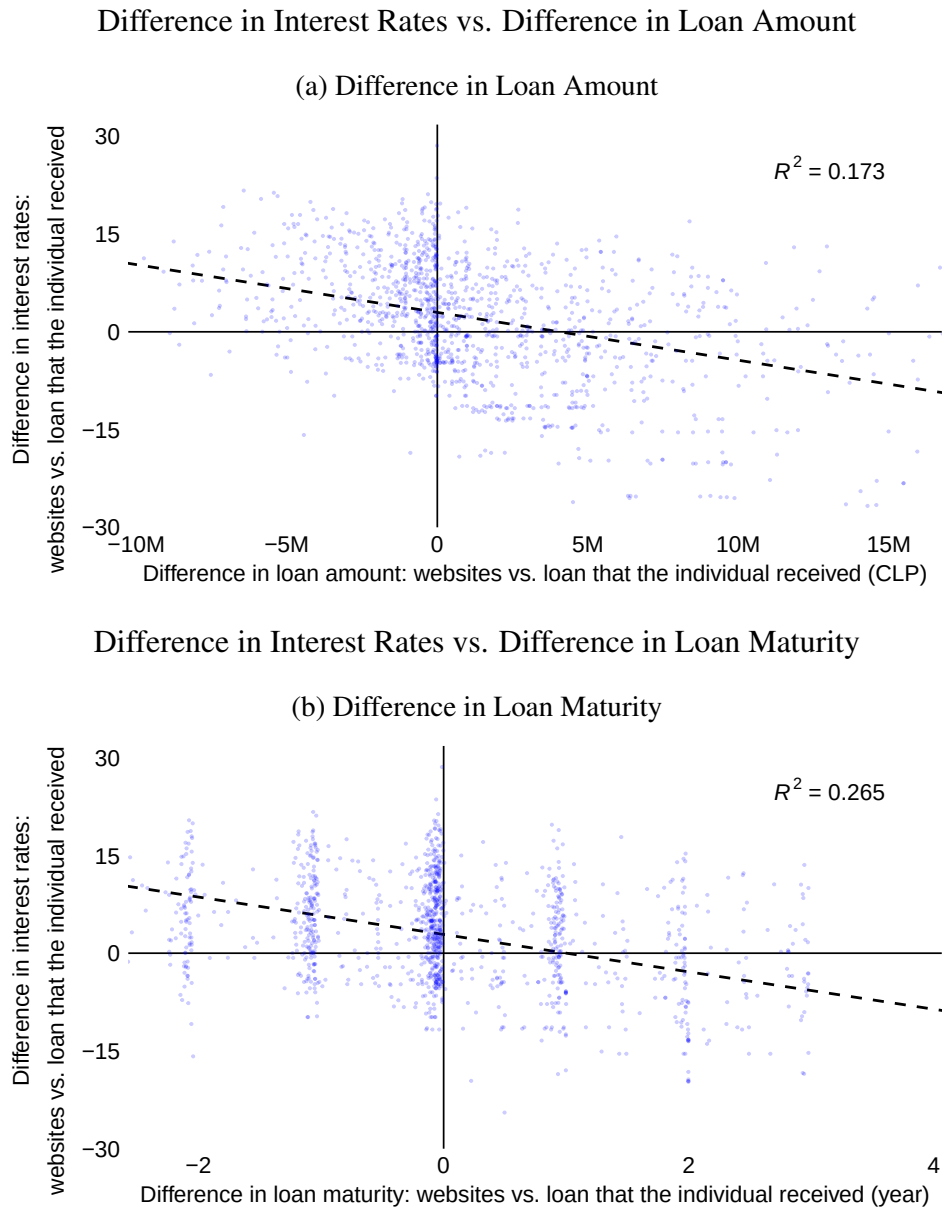


(c) Banco Internacional



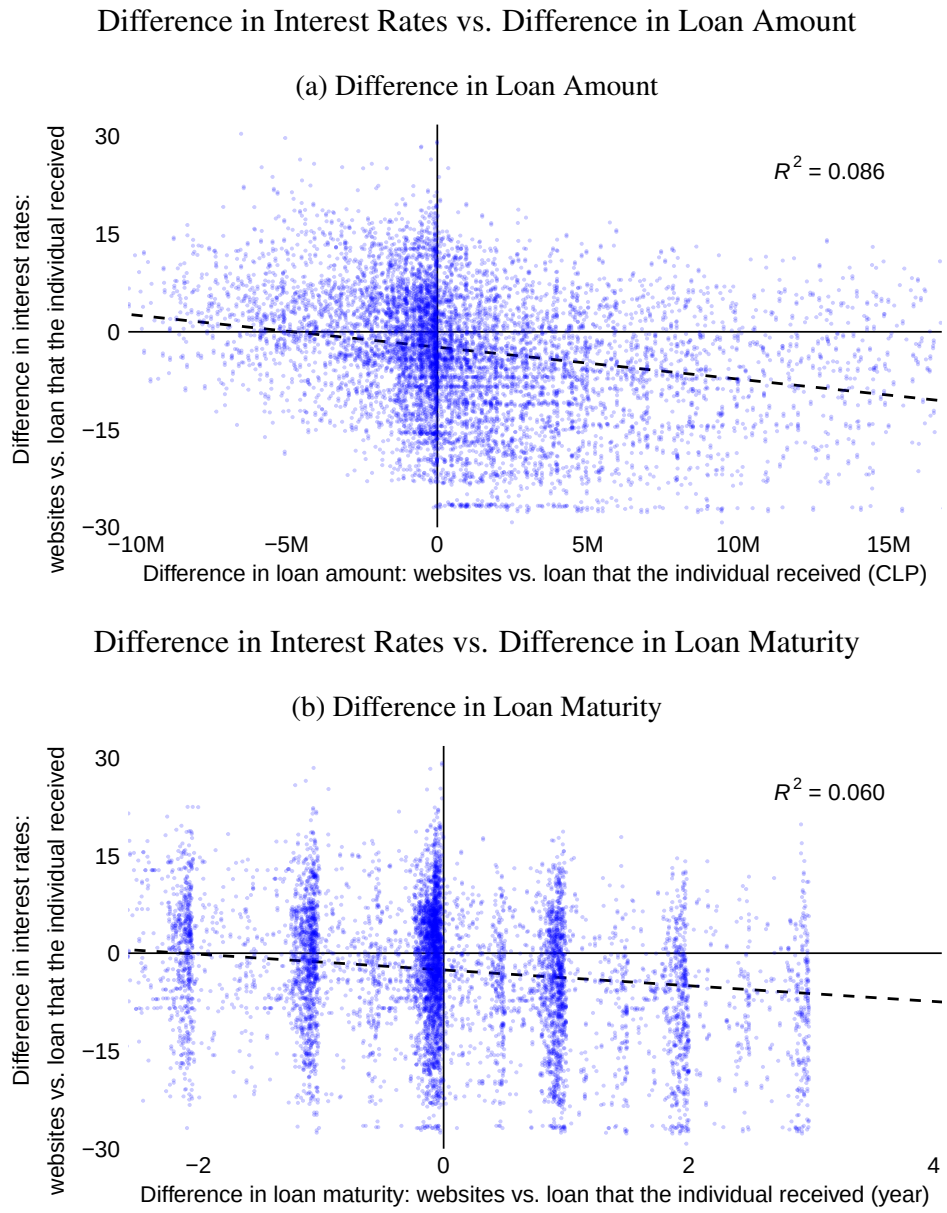
This figure shows the annualized interest rates from bank websites, derived from a test where we varied only the RUT while maintaining other inputs constant. Specifically, we standardized inputs such as income, loan amount, and loan maturity by using the median values from the baseline sample: a monthly income of $740,000 pesos, a loan amount of $2,500,000 pesos, a maturity period of 3 years, and a fake RUT of 32954440-0. The x-axis represents the annualized interest rates calculated by each bank's website simulation. Each simulation consists of 100 observations.

Figure B.3: Difference in Interest Rates Between Bank Websites and Loan That the Individual Received



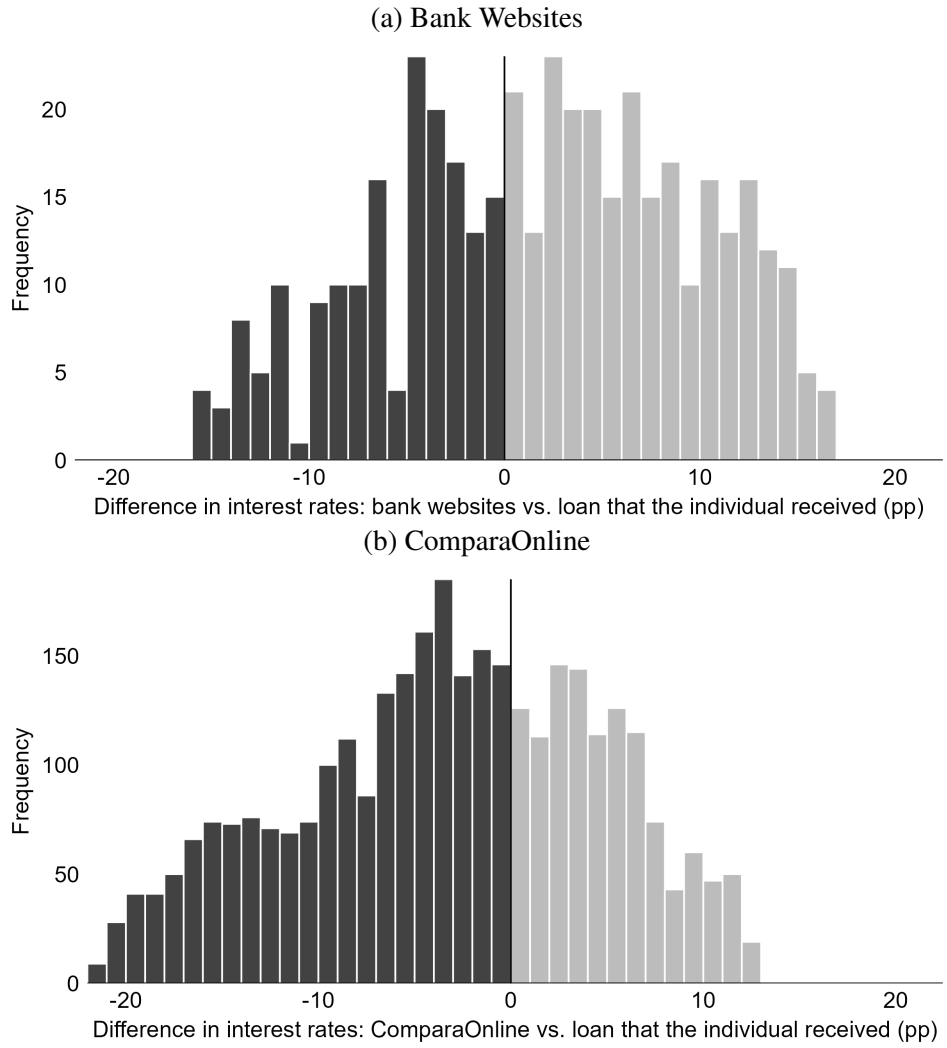
This figure shows the correlation of the difference in interest rates with the difference in loan amount (panel a) or the difference in loan maturity (panel b). Interest rates on bank websites are the rates displayed on consumer loan simulation websites, given an individual's baseline survey inputs. 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 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 B.4: Difference in Interest Rates Between ComparaOnline and Loan That the Individual Received



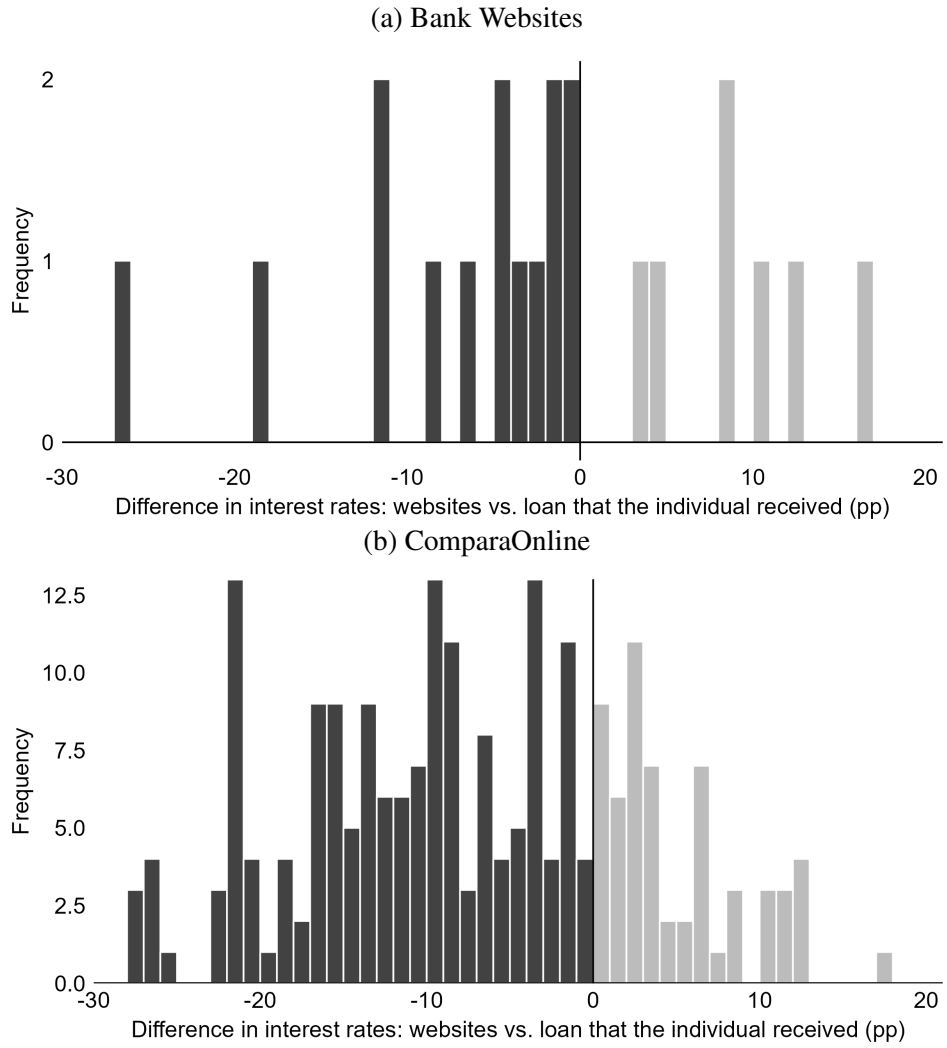
This figure shows the correlation of the difference in interest rates with the difference in loan amount (panel a) or the difference in loan maturity (panel b). Interest rates on bank websites are the rates displayed on consumer loan simulation websites, given an individual's baseline survey inputs. 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.

Figure B.5: Difference in Interest Rates Between Simulation Results and Loan That the Individual Received: Control Group Only



This figure shows the distribution of differences in interest rates for the control group. Interest rates on websites are the rates displayed on consumer loan simulation websites given an individual's baseline survey inputs. Difference in interest rates is calculated as: interest rates on websites - interest rates of the loan that the individual received. Panel a) includes all simulated rates from bank websites. There are 466 observations, 40.77% of which are negative. Panel b) plots the interest rates shown by ComparaOnline. There are 3,560 observations, 61.94% of which are negative. The top and bottom 5 percentile of differences in interest rates are excluded from each histogram for legibility.

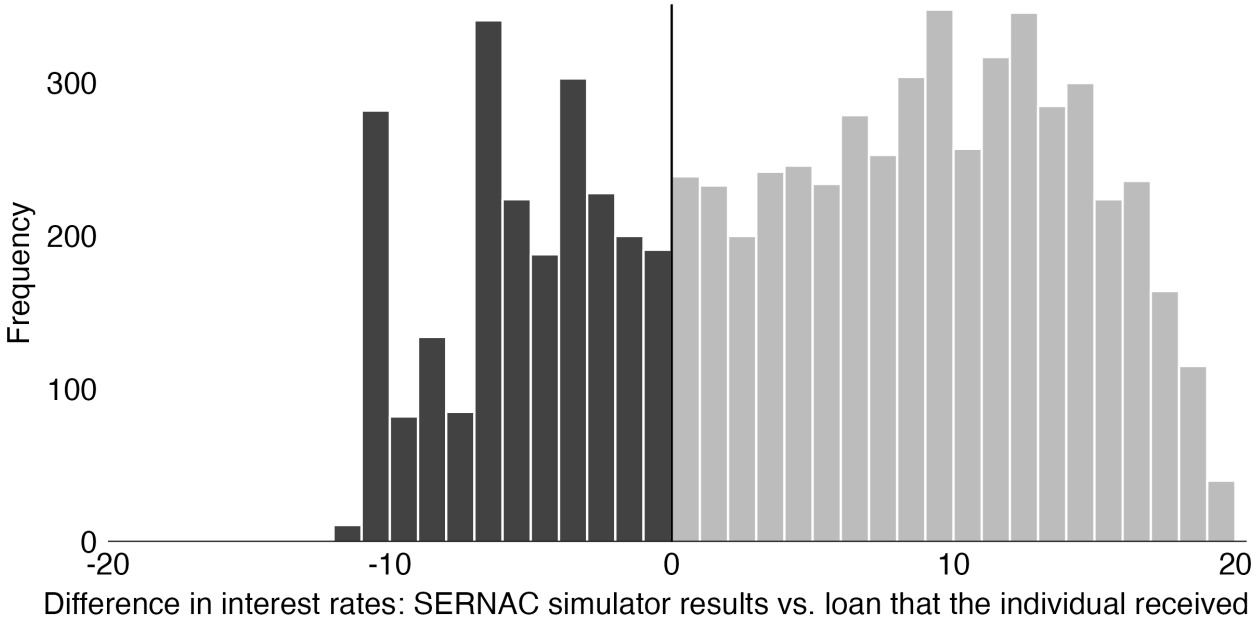
Figure B.6: Difference in Interest Rates Between Simulation Results and Loan That the Individual Received: Control Group and Restricted Sample Only



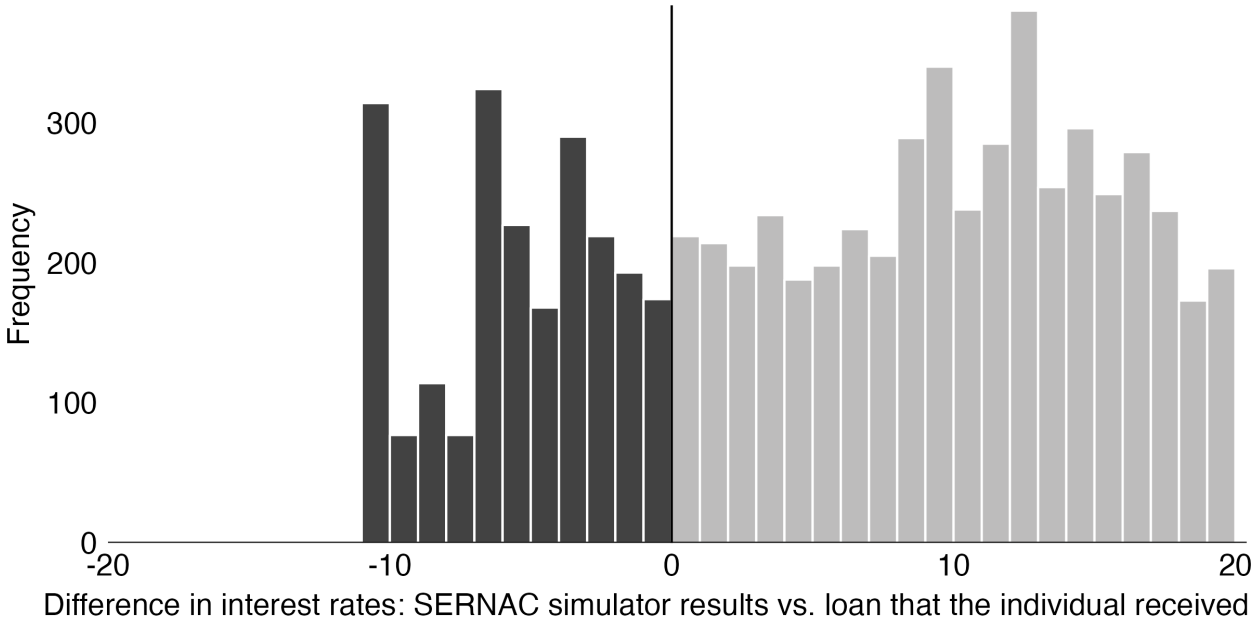
This figure shows the distribution of differences in interest rates for the control group. Interest rates on websites are the rates displayed on consumer loan simulation websites given an individual's baseline survey inputs. Difference in interest rates is calculated as: interest rates on websites - interest rates of the loan that the individual received. Panel a) includes all simulated rates from bank websites. There are 21 observations, 66.67% of which are negative. Panel b) plots the interest rates shown by ComparaOnline. There are 223 observations, 72.65% of which are negative. The top and bottom 5 percentile of differences in interest rates are excluded from each histogram for legibility.

Figure B.7: Difference in Interest Rates Between SERNAC Simulation Results and Loan That the Individual Received

(a) Matched by the Closest Loan Amount Option



(b) Matched by the Closest *Higher* Loan Amount Option

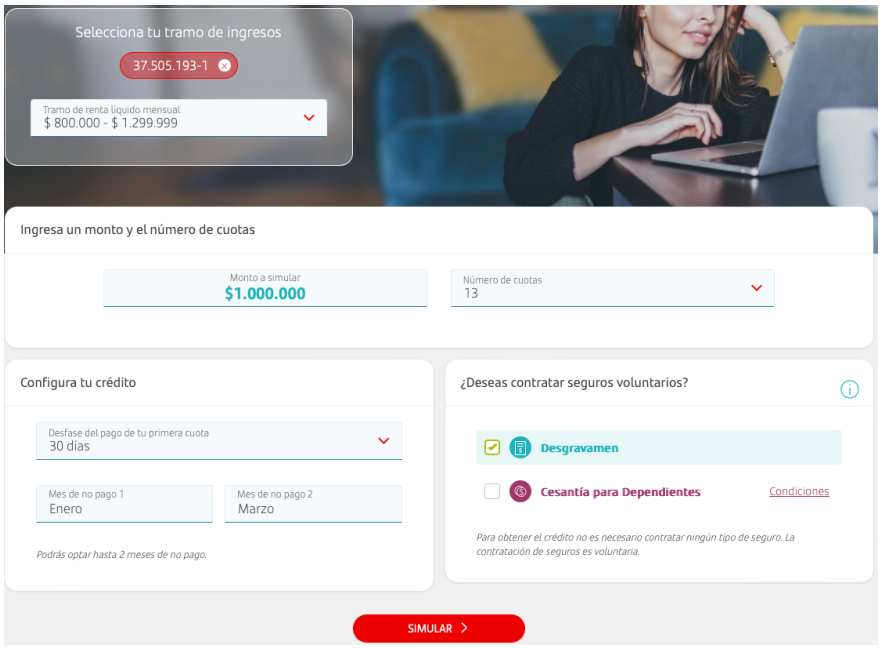


These figures display the distribution of the average interest rate differences between the SERNAC simulation data and administrative data (Interest rates from SERNAC minus those from administrative data). The histogram is truncated at the 5th and 95th percentiles. The numbers of observations are 7,131 and 7,155, and the percentages of negative differences are 31.82% and 30.43%, respectively.

B.5 Interface of Banks' Consumer Loan Simulators

Figure B.8: Banco Stantander

(a) Input Page



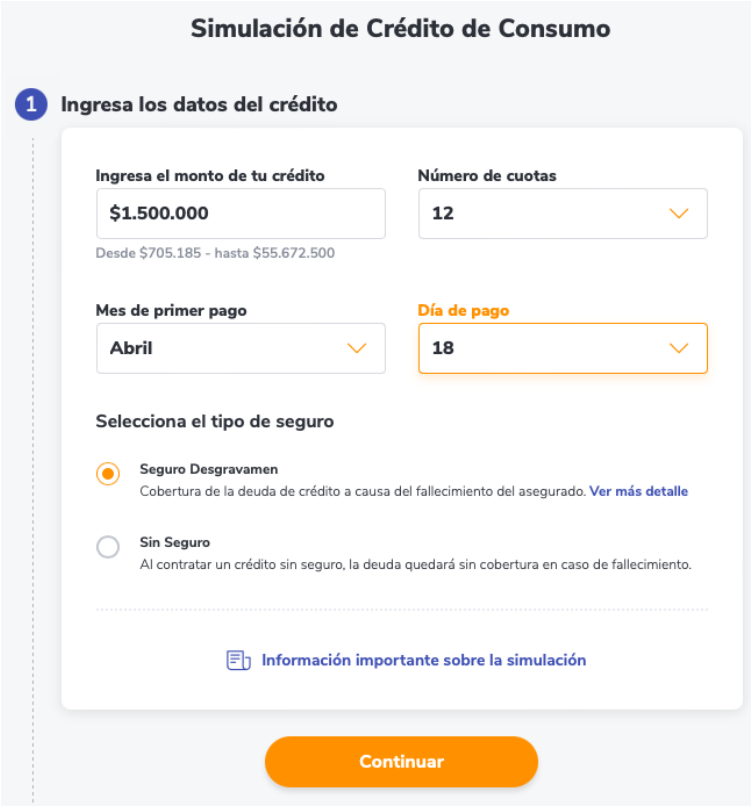
(b) Output Page



This figure shows the input and output interface of Banco Santander's consumer loan simulator as of April 3, 2024.

Figure B.9: Banco Estado

(a) Input Page



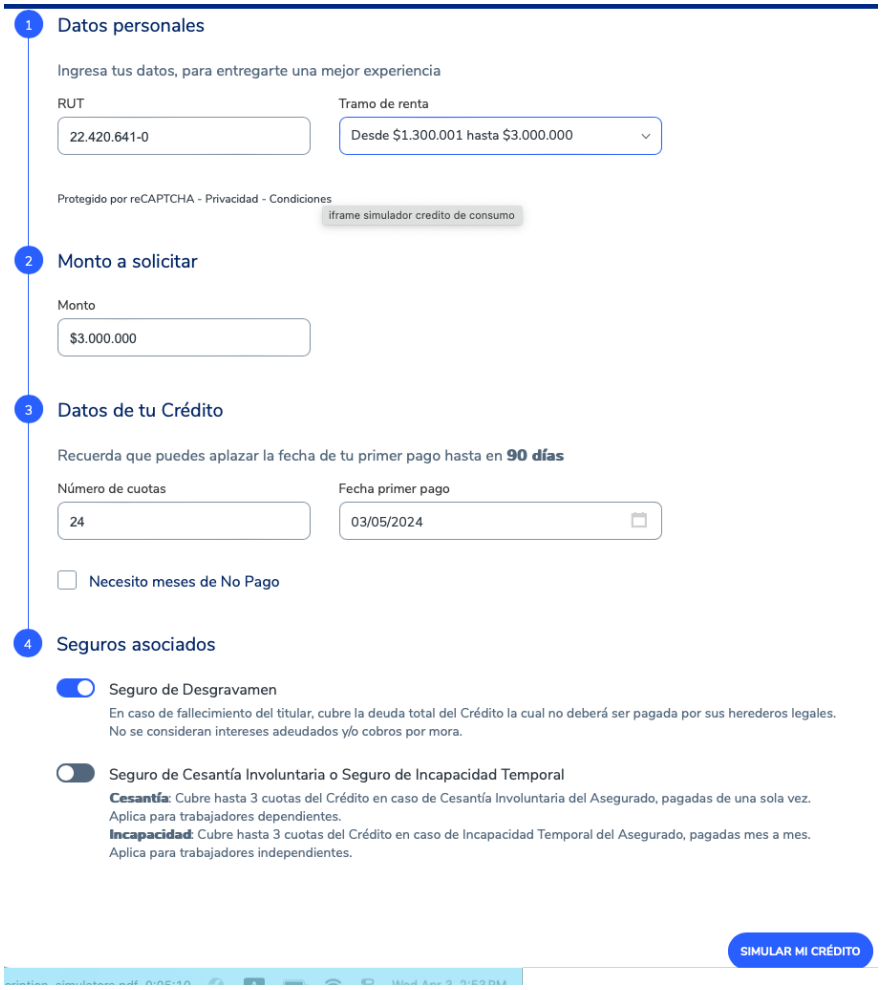
(b) Output Page



This figure shows the input and output interface of Banco Estado's consumer loan simulator as of April 3, 2024.

Figure B.10: Banco de Chile

(a) Input Page



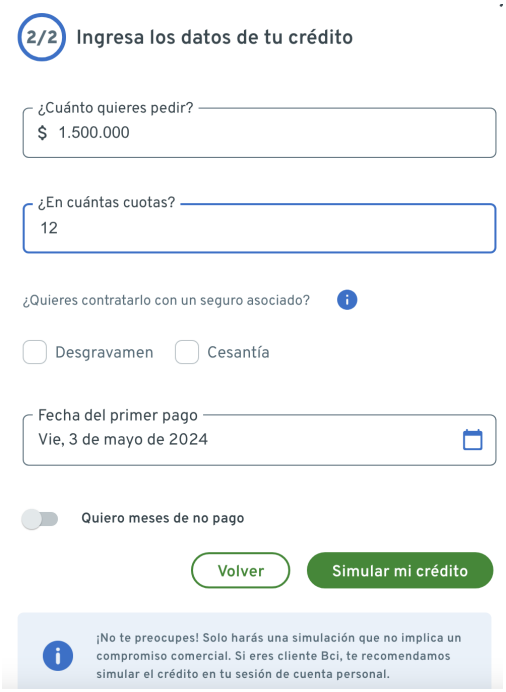
This figure shows the input interface of Banco de Chile's consumer loan simulator, captured on April 3, 2024. Due to persistent errors during the initiation of the simulation process, the output page could not be displayed.

Figure B.11: Banco BCI

(a) First Input Page



(b) Second Input Page



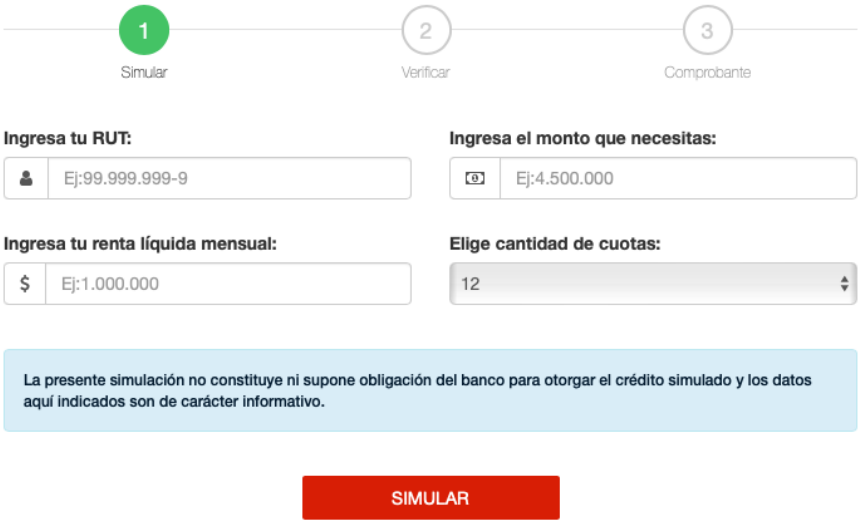
(c) Output Page



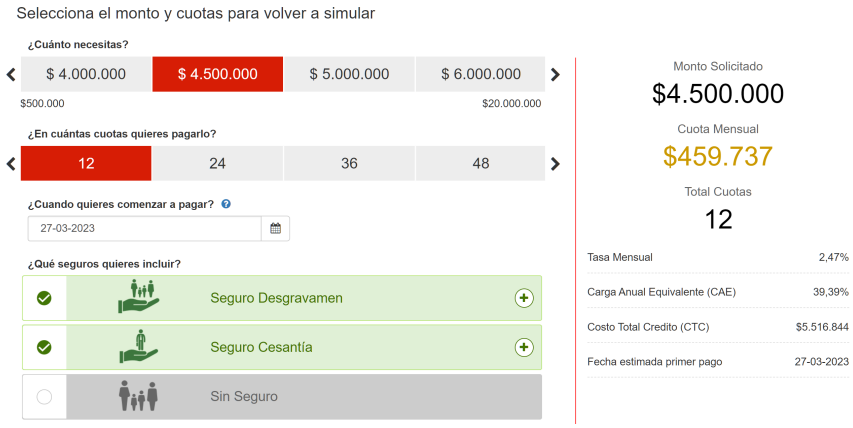
This figure shows the input and output interface of Banco BCI's consumer loan simulator as of April 3, 2024.

Figure B.12: Scotiabank

(a) Input Page

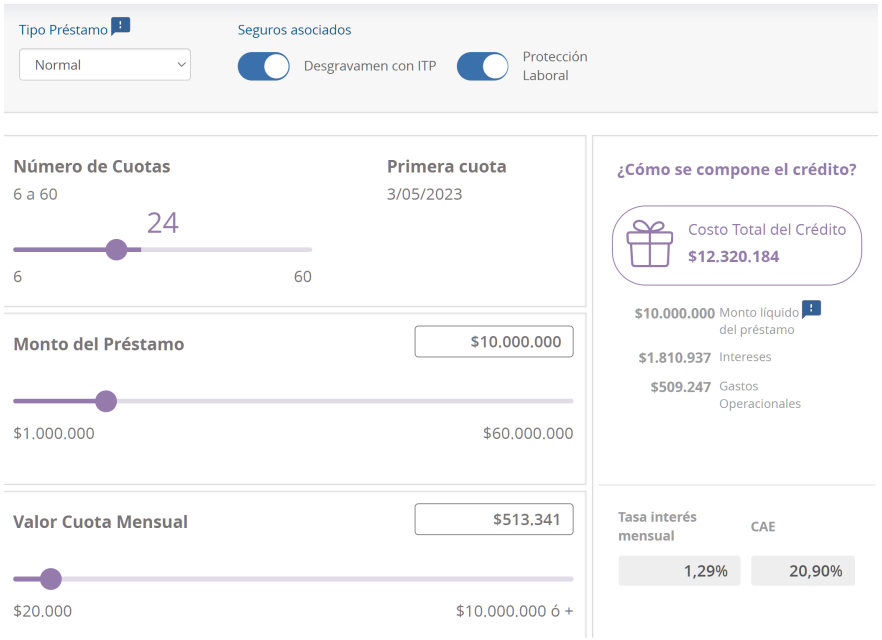


(b) Output Page



This figure shows the input and output interface of Scotiabank's consumer loan simulator as of April 3, 2024.

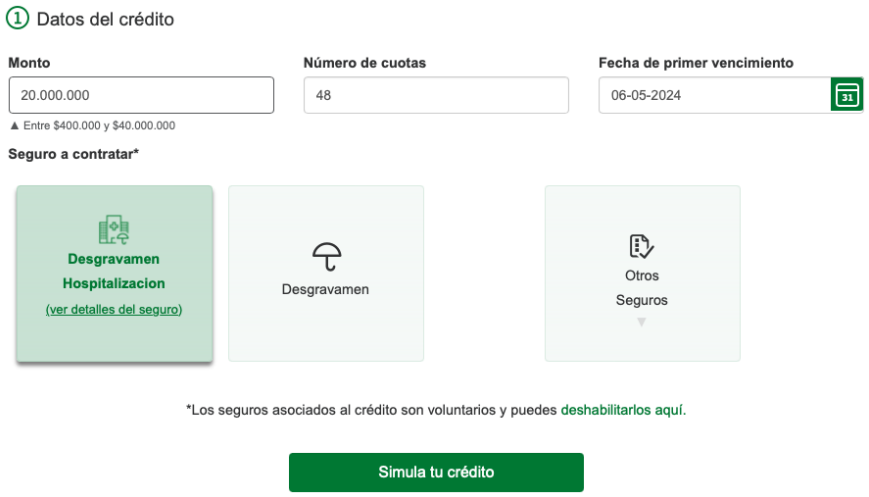
Figure B.13: Banco BICE



This figure shows the input and output interface of Banco BICE's consumer loan simulator as of April 3, 2024.

Figure B.14: Banco Falabella

(a) Input Page



(b) Output Page



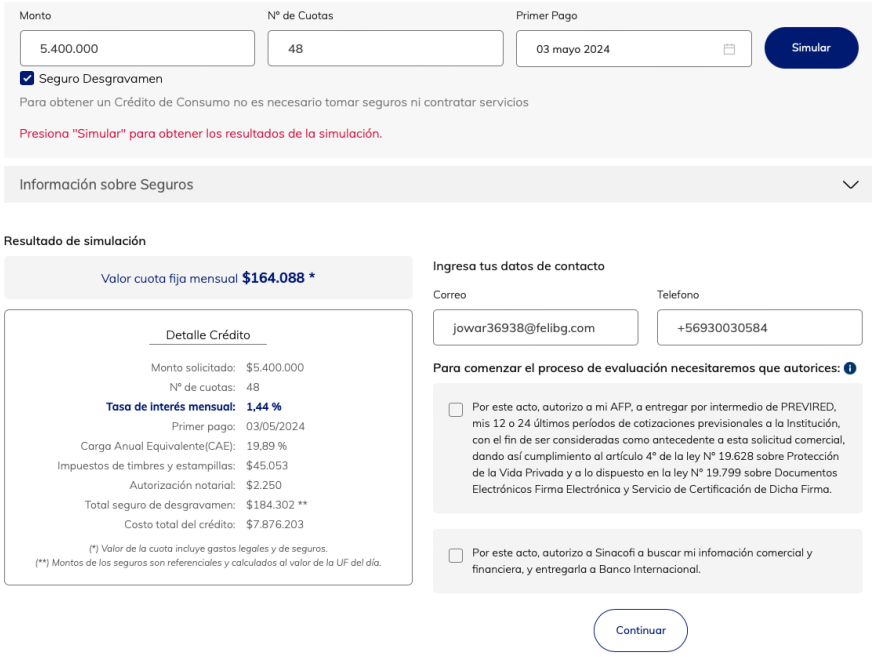
This figure shows the input and output interface of Banco Falabella's consumer loan simulator as of April 3, 2024.

Figure B.15: Banco Internacional

(a) Input Page



(b) Output Page



This figure shows the input and output interface of Banco Internacional's consumer loan simulator as of April 3, 2024.

Figure B.16: Banco Security

(a) Input Page



(b) Output Page



This figure shows the input and output interface of Banco Security's consumer loan simulator as of April 3, 2024.

Figure B.17: Consorcio

(a) First Input Page



(b) Second Input Page



(c) Output Page



This figure shows the input and output interface of Banco Consorcio's consumer loan simulator as of April 3, 2024.

Figure B.18: Coopeuch

(a) Input Page



(b) Output Page



This figure shows the input and output interface of Coopeuch's consumer loan simulator as of April 3, 2024.

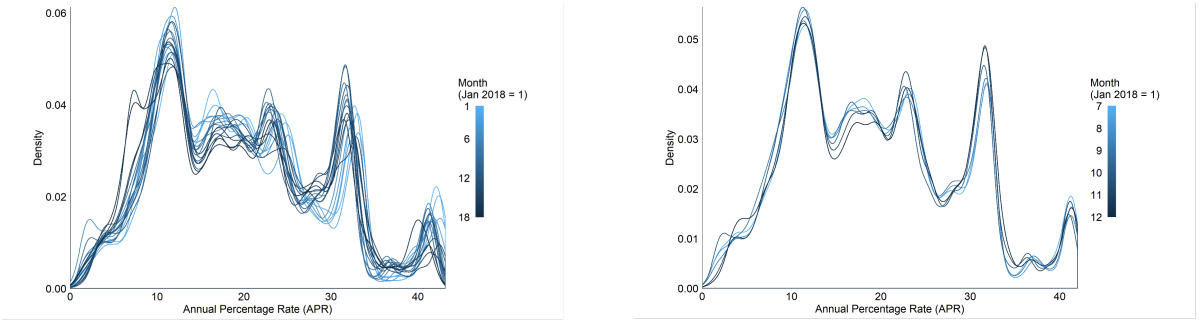
Appendix C Length of Past Data Shown

Loans differ from other products in that consumers cannot merely compare the current prices of loans at different banks and decide which to buy; instead, they must apply for a loan at each bank that they want to include in their comparison and see whether they are approved. Thus, we view a tool based on actual loans that were obtained by similar consumers in the market as more relevant than a lot of existing price comparison tools that instead collate current information on rates that banks report that they would offer to consumers of different types. The information banks report on current rates that could be collated in this way is (i) not sufficiently dis-aggregated by consumer type and (ii) an inaccurate measure of the rates consumers actually receive, because the information banks are required to report is what loan and interest rate a consumer would be technically eligible for (but this does not reflect the probability that a consumer of that type is approved for the loan in practice). By showing the distribution of interest rates that were actually obtained by similar consumers (based on income and neighbourhood of residence), we provide consumers with a sense of what loans they could actually obtain in the marketplace. Since they see the entire distribution of APRs, they will also have a sense of the probability of banks offering various rates and of being approved for those rates.

In order to provide consumers with this 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 rather than the past 18 months, 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 much information we have to 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 C.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 C.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. Furthermore, we will refresh the data underlying the tool each month so that the tool always shows the most recent 18 months of data.

Figure C.1: Distribution of Interest Rates by Time



(a) 18 Months

(b) 6 Months

This figure shows the distribution of interest rates obtained by consumers for consumer loans, with data sourced from the CMF (Financial Market Commission). Panel (a) displays the distribution of interest rates over the 18 months from January 2018 to June 2019. It illustrates how interest rates have varied over time, providing insight into trends in the loan market during this period. Panel (b) shows the distribution of interest rates for the most recent six months from July 2018 to December 2019. It demonstrates that despite monthly variability, 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.

Appendix D Search Benefit Calculation

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 municipality, income quartile, loan size, and maturity for 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 bin maturity and loan size to calculate the monthly interest payments of loan l_0 and l_1. We then consider two cases:

If $l_1 < l_0$, we calculate the present value of the difference in monthly payments over the life of the loan.

If $l_1 \geq l_0$, we set the "benefit of searching at one additional bank" to zero as the participant could take out the first loan.

We then repeat this drawing of two quotes 1,000,000 times. We then set the "benefit of searching at one additional bank" to be the mean of all the present value differences in monthly payments. 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 $n \in \{3,4,5\}$ additional banks, we simulate the process for drawing l_n loans from $J-n$ banks 1,000,000 times and take the mean. As before, all benefits are calculated in relation to the first draw l_0, and any differences in monthly payments that are greater than or equal to l_0 are coded as zeros.

We repeat this procedure for all constructed bins. If there are less than 5 loans issued in a bin and less 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 5 loans and two unique rates, we expand the bin again to include comunas that border the bordering comunas to the reference comuna.

Appendix E Google Search

E.1 Obtaining data from Google Search results

The data were scraped by mimicking users searching from different comunas with various search terms. For each search, we randomly selected a comuna-search term pair. The comuna population data are derived from our baseline survey and weighted by the number of participants from each comuna. The search term population is sourced from our Google Ad campaign. We collected the search terms that led people to our price comparison tool and weighted them by their frequency in

searches. During each search, we changed the geolocation parameter in Google to match the selected comuna and searched Google using the selected term. We scraped all available information from each result on the first page, including the content provider, link, title, text snippet, and the position of the results on the page. The scraper ran from November 8th, 2023, to February 4th, 2024, resulting in 6,677,889 Google search results from 101,852 comuna-search term pairs.

We scraped these results using desktop emulator, mimicking what a user would see if they opened Google on a desktop computer. Ideally, we would have also scrape 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 the mobile results.

We used OpenAI's Assistant API to extract 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 tasks. To extract interest rate numbers from raw scraped text with rule-based text processing code, we would have 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. For instance, in each of the examples presented in Table E.1, the percentage number in the sentence carries a distinct meaning.

On the other hand, a well-trained LLM will be able to comprehend the whole sentence and correctly identify whether it contains a consumer loan interest rate. At the time of our data processing, OpenAI provides two APIs, Chat API and Assistant API. The Assistant API allows users to create and tailor an "assistant" for a specific task and use it repeatedly. It also contains built-in tools tuned for particular tasks, including "code_interpreter", "retrieval", and "function".

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 and whether the interest rate excluded fees or was an APR including fees. The prompt we sent to the API can be found in Section E.3.

As 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. However, we argue

that the high level of accuracy achieved by our approach will ensure consistent results for future replicators.

Table E.1: Examples of Non-Interest Rate Percentage Number

Original text (in Spanish)	English translation	Meaning of the percentage number
May 23, 2017 — 776 (Banco Condell). Es decir, una diferencia de $856.692 (33,9%) entre el monto más barato y el más caro. SIMULADOR DE CRÉDITO. El mismo Sernac ...	May 23, 2017 - 776 (Condell Bank). That is, a difference of $856,692 (33.9%) between the cheapest and the most expensive amount. CREDIT SIMULATOR. The same Sernac ...	difference
Requisitos · Impuesto al Crédito: 0.066 % por fracción de mes, aplica sobre el monto total del crédito. · Impuesto al Crédito: Tope máximo 0.8 % equivalente a ...	Requirements - Credit Tax: 0.066 % per fraction of month, applied on the total amount of the credit. - Credit Tax: Maximum cap of 0.8 % equivalent to ...	tax
Sistemáticamente, el Banco de Chile (CHILE) ha sido el banco más rentable de Chile a lo largo de los años, con un ROA medio del 1,8% en los últimos 10 años, superando a toda su competencia local. Como comparación, su ROA es 30 puntos básicos más que Banco Santander Chile (BSANTANDER).	Systematically, Banco de Chile (CHILE) has been the most profitable bank in Chile over the years, with an average ROA of 1.8% over the last 10 years, outperforming all of its local competition. As a comparison, its ROA is 30 basis points higher than Banco Santander Chile (BSANTANDER).	ROA
Ahorra hasta un 15% en tasas de interés en tu crédito para compra vehicular en Compara Online ... El equipo de RadarCupón te aconseja: Compara Online te regala ...	Save up to 15% on interest rates on your vehicle purchase credit at Compara Online ... The RadarCoupon team advises you: Compara Online gives you free ...	discount
El límite que se suele establecer es de entre un 25% y 35% de tus ingresos, es decir, si la cantidad que has pedido supera en este porcentaje a tus ingresos lo más normal es que el banco deniegue tu solicitud y te quedes sin el préstamo o crédito que habías pedido.	The limit that is usually established is between 25% and 35% of your income, that is to say, if the amount you have requested exceeds your income by this percentage, the bank will normally deny your application and you will not receive the loan or credit you had requested.	percentage

This table shows examples of sentences that contain percentage numbers which are not interest rates.

E.2 Comparison of Google Search Displayed 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. Initially, we restrict our sample to the 30,979 individuals in the administrative data who had taken 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.

Next, for each individual who took out a consumer loan, we matched the interest rates they would have seen on Google—based on comuna—with the interest rate they received in the administrative data. We then annualized the monthly interest rates by multiplying them by 12 and excluded any scraped results that included only CAE but not interest rates. This matching process resulted in 15,579 observations (162 unique comuna).

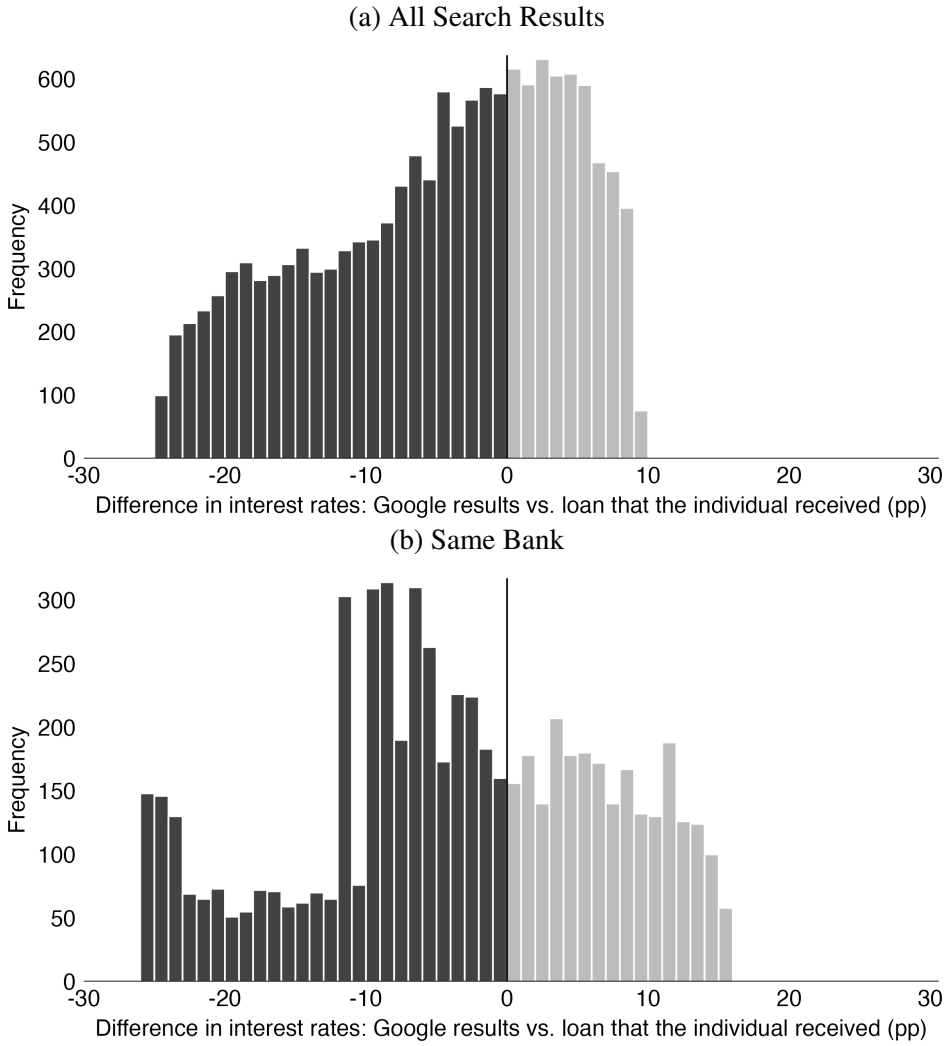
One concern is that interest rate differences across comunas may stem from banks advertising with varying intensity or offering different loan terms in different areas. To address this, we also matched Google Search results with administrative data at the comuna and bank level using encrypted bank identifiers in the administrative data. For instance, if an individual took a loan from Bank A, we merged the interest rate they received with the rates from Bank A displayed on Google Search. This yielded 6,920 observations (282 unique comuna-bank pairs). Figure E.1 shows the results: 64.11% of the rates shown on Google were lower than what people actually received for the sample matched by comuna, and 61.94% were lower for the sample matched by comuna-bank. These results indicate that accounting for bank-by-comuna differences did not significantly reduce the number of disparities between advertised rates and the rates people actually received.

Another concern is the potential inaccuracy of Google Search interest rates due to the discrepancy between the data scraping date and the actual loan date. To address this, we created a restricted sample containing individuals who received loans from October 24, 2023, to February 4, 2024, when interest rates were collected from consumer loan simulator websites. This sample included 1,533 observations (from 141 unique comunas). Figure E.2 illustrates the distribution of the differences between Google Search results and the actual loan rates received in the restricted window. 73.53% of the rates shown on Google were lower than what people actually received for the sample matched by comuna, and 77.37% were lower for the sample matched by comuna-bank. Similar to the results with the unrestricted sample, it indicates that people generally see lower interest rates on Google compared to the rates they actually obtain.

Lastly, our price comparison tool might have affected people's loan searching and terms negotiation behaviors, which might have affected the loan terms that the person eventually received. To

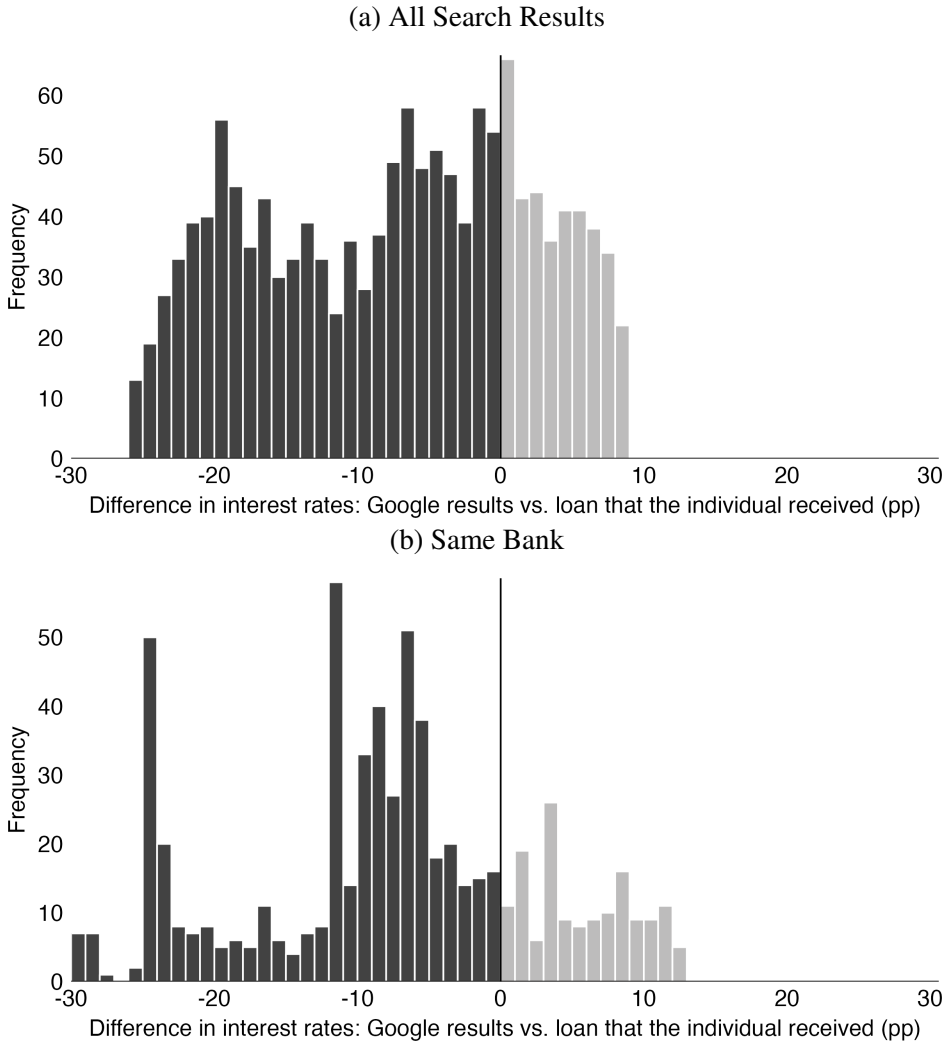
eliminate this effects, we also compared the interest rates that came from Google Search results and those of the loans that people actually received with only the control group. The results are shown in Figures E.4 and E.3, from which we can see that restricting the sample to only the control group does not change the results much. Notably, 64.18% of the interest rates shown to the control group in Google Search results were lower than the rates they ultimately obtained in their actual loan agreements.

Figure E.1: Difference in Interest Rates Between Google Search Results and Loan That the Individual Received (pp), Unrestricted Sample



This figure shows the distribution of the average interest rate differences between the scraped Google Search data and administrative data (Interest rates from Google Search minus those from administrative data) for the full sample. Panel A merges the Google Search data and administrative data solely by comuna (equivalent to 5a), while Panel B merges by comuna-bank. The histogram is truncated at the 5th and 95th percentiles. In Panel A, the number of observations is 14,029. The percentage of negative differences is 64.11%. In Panel B, the number of observations is 6,243. The percentage of negative differences is 61.94%.

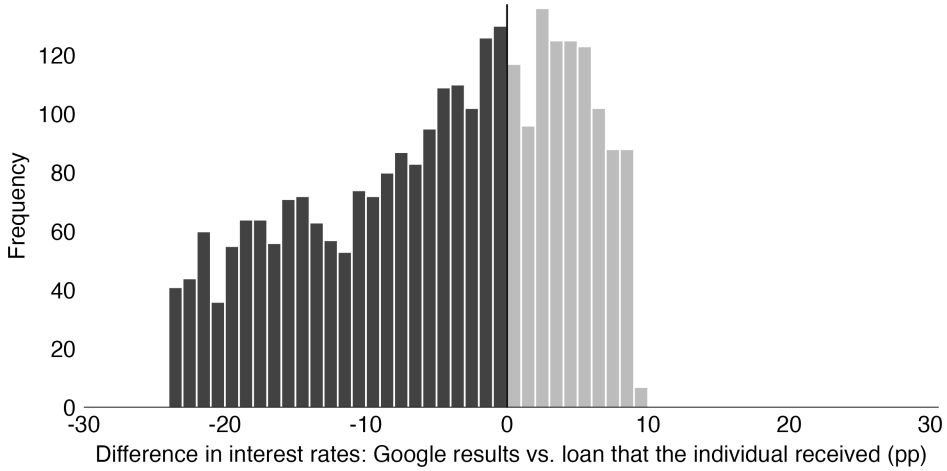
Figure E.2: Difference in Interest Rates Between Google Search Results and Loan That the Individual Received (pp), Restricted Sample



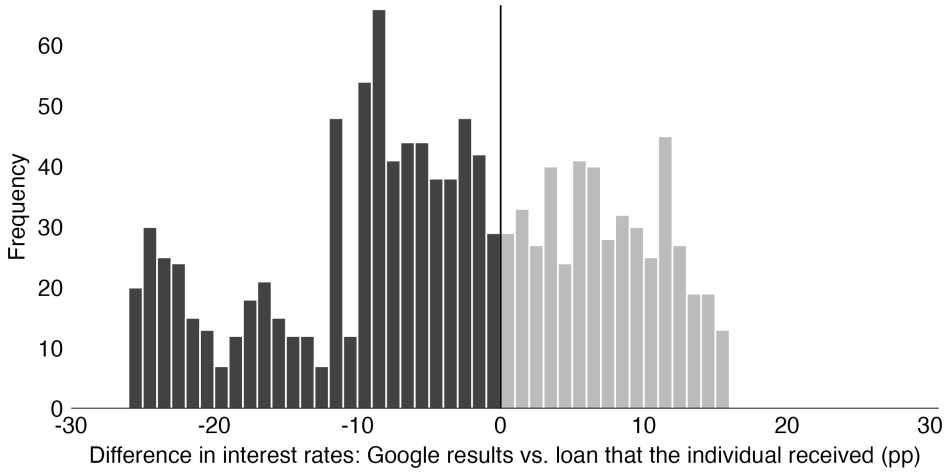
This figure shows the distribution of the average interest rate differences between the scraped Google Search data and administrative data (Interest rates from Google Search minus those from administrative data) at the comuna level. Panel A merges the Google Search data and administrative data solely by comuna, while Panel B merges by comuna-bank. The sample contains only individuals who have received loans from October 24, 2023 to Feburary 4, 2024 when interest rates were collected from Google Search. The histogram is truncated at the 5th and 95th percentiles. In Panel A, the number of observations is 1,379. The percentage of negative differences is 73.53%. In Panel B, the number of observations is 654. The percentage of negative differences is 77.37%.

Figure E.3: Difference in Interest Rates Between Google Search Results and Loan That the Individual Received (pp), Control Group, Unrestricted Sample

(a) All Search Results Sample

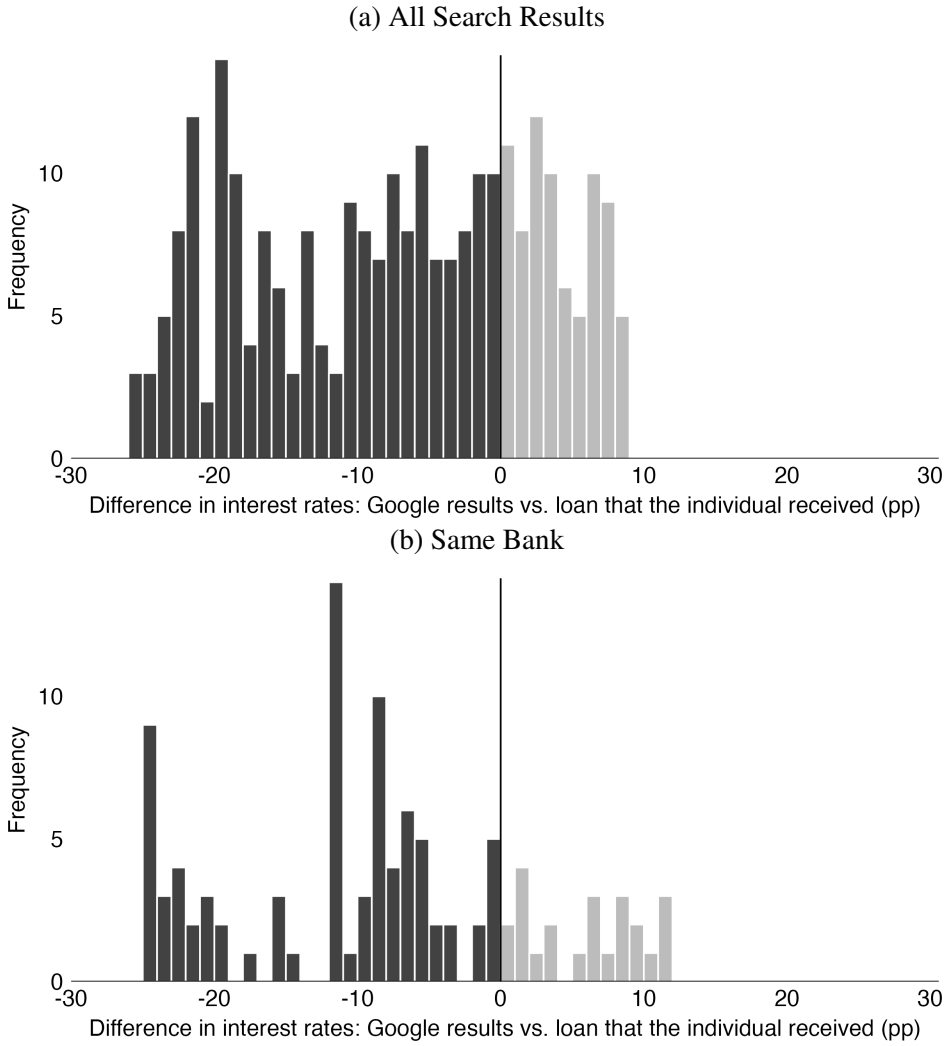


(b) Same Bank, Unrestricted Sample



This figure shows the distribution of the average interest rate differences between the scraped Google Search data and administrative data (Interest rates from Google Search minus those from administrative data) for the full sample. Panel A merges the Google Search data and administrative data solely by comuna, while Panel B merges by comuna-bank. The histogram is truncated at the 5th and 95th percentiles. In Panel A, the number of observations is 2,811. The percentage of negative differences is 64.18%. In Panel B, the number of observations is 1,207. The percentage of negative differences is 60.89%.

Figure E.4: Difference in Interest Rates Between Google Search Results and Loan That the Individual Received (pp), Control Group, Restricted Sample



This figure shows the distribution of the average interest rate differences between the scraped Google Search data and administrative data (Interest rates from Google Search minus those from administrative data) for the full sample. Panel A merges the Google Search data and administrative data solely by comuna, while Panel B merges by comuna-bank. The sample contains only individuals who have received loans from October 24, 2023 to Feburary 4, 2024 when interest rates were collected from Google Search.The histogram is truncated at the 5th and 95th percentiles. In Panel A, the number of observations is 264. The percentage of negative differences is 71.21%. In Panel B, the number of observations is 105. The percentage of negative differences is 78.1%.

E.3 Prompt for the Assistant API

Task: Analyze text scraped from Google Search results to extract and organize loan-related data, with a focus on interest rates, Equivalent Annual Cost (CAE), and other pertinent details. Given that our research is centered on a sample of Chilean loan takers, avoid relying on "common sense" assumptions typical of English-speaking countries when making inferences in your analysis.

Interest rate data should only be included if explicitly referred to in the context of borrowing or lending. It's ok if no clear period for the rate is provided. Do not extract percentages or fees that refer to one-time charges, service fees, transaction costs, etc. When unsure, provide a descriptive note regarding the ambiguity rather than extracting incorrect data.

Input Format: JSON-formatted strings sent directly as text snippets.

Output Format: Format your findings in JSON with the following variables:

1. tasa_anual: A dictionary mapping bank names to their respective annual interest rates (include both nominal and effective rates). If specific bank names are not mentioned, mark the bank name as "unknown". Note that if the returned variable is a dictionary, the key should always be the name of a bank. Do not use nested dictionaries. If you try to infer the time frame of the interest rate when it's not given, write your reasoning under the "note" variable described below.

2. tasa_mensual: A dictionary mapping bank names to their monthly interest rates. Similar rules as tasa_anual for unnamed banks.

3. tasa_unidentified: A dictionary mapping bank names to unclear time frame interest rates. Similar rules as tasa_anual for unnamed banks. Avoid categorizing percentages that are not related to interest rates under this variable.

4. cae: A dictionary mapping bank names to their 'Carga Anual Equivalente' (CAE), excluding any values included in apr_number. CAE, a term commonly used in Spanish-speaking countries, is analogous to the APR (Annual Percentage Rate). It denotes the effective annual cost of a loan, encompassing both interest and additional charges. Apply the same rules as for tasa_anual when bank names are not specified. Populate this variable only when the term 'CAE' is explicitly mentioned. If 'CAE' is not directly referred to, use apr_var and apr_number, as outlined below.

5. apr_var: A string indicating a non-CAE APR term (like APR, TEA, CAT).

6. apr_number: A dictionary mapping bank names to their APR value. If it's a "CAE" in the direct term, list the values in "cae" instead. If this variable is not NaN, the previous variable "apr_var" must not be NaN. Similar rules as tasa_anual for unnamed banks.

7. loan_type: Classify the type of loan or financial product as “consumer_loan”, “mortgage_loan”, “credit_card”, “deposit”, “policy”, or “unknown” (where “policy” refers to the central bank’s policy rate). Based solely on the given text. If you try to infer the loan type based on the text, write your reasoning under the “note” variable described below.
8. language: Language abbreviation (e.g., "es" for Spanish).
9. country: Country where the loan is offered, or "unknown" if uncertain.
10. note: If uncertain about the context of a percentage figure, provide a descriptive note to explain the ambiguity. Additionally, explicitly state any implicit assumptions made while interpreting the text.

Special Instructions:

- In cases where multiple banks or entities are mentioned with specific rates, organize this data in a dictionary format under the relevant variable (e.g., tasa_anual).
- Do not perform rate calculations.
- If a search result lacks financial data but is relevant to banks or loans, still return a JSON object with variables 8 and 9. Ensure no variable contains an empty list. Exclude variables from the JSON output if there are no values to report.
- Do not assume the time frame of an interest rate unless it is explicitly mentioned.
- Use your best judgment for determining “language” and “country”. For consumer loans, do not assume that the time frame of the interest rate is annual; in some countries, monthly rates are more commonly used.

Examples:
Example 1:
Input:
{’content_provider’: ’condusef.gob.mx’, ’link’: ’https://www.condusef.gob.mx > ...’, ’title’: ’¿Sabes cuál es la tasa de interés y el CAT que te cobran por tu crédito de ...’, ’question’: ’¿Cuál es el banco que ofrece la mejor tasa de interés?’, ’text’: ’respecto a este producto. La tasa de interés anual más alta para este tipo de productos la cobra Banorte con 44%, seguida de Banregio con 43%, y posteriormente se ubica HSBC con 39.9%; en tanto que el CAT más alto es igualmente de Banorte con 63.1%.’}

Output:

{‘‘tasa_anual’’: {’’Banorte’’: 44, ’’Banregio’’: 43, ’’HSBC’’: 39.9}, ’’ apr_number’’: {’’Banorte’’ : 63.1}, ’’apr_var’’: ’’CAT’’, ’’loan_type’’: ’’unknown’’, ’’language’’: ’’es’’, ’’country’’: ’’Mexico’’}

Example 2:

Input:

{’content_provider’: ’mrfinan.com’, ’link’: ’https://mrfinan.com/mx/prestamos/prestamo -hasta-100000-pesos’, ’title’: ’Préstamo hasta 100 mil Pesos’, ’text’: ’Préstamos hasta 100 mil pesos. 2 MINUTOS | GRATIS | SIN COMPROMISO. 3- 36 Meses. En Buró de Crédito. CAT mínimo 1.58%’}

Output:

{ "apr_var": ‘‘CAT’’, "apr_number": {‘‘unknown’’: 1.58}, "loan_type": "consumer_loan", "language": "es", "country": "Mexico" }

Example 3:

Input:

{’content_provider’: ’didiglobal’, ’link’: ’https://web.didiglobal.com/mx/prestamos/’, ’title’: ’DiDi Préstamos - Rápido, Fácil y Seguro. | DiDi México’, ’text’: ’Deja tú lo fácil que es solicitar un préstamo, la tasa de interés ordinaria va desde el 5% hasta el 12%. (*Tasa ordinaria mensual estimada).’}

Output:

{ "tasa_mensual":{ "DiDi México": [5, 12] }, "loan_type": "consumer_loan", "language": "es", "country": "Mexico" }