Credit Failures and Entrepreneurial Risk Aversion*

William Fuchs † Ganesh Iyer‡ Przemyslaw Jeziorski§

October 5, 2024

Abstract

This paper examines the influence of adverse experiences resulting from the adoption of a new credit technology on entrepreneurial risk aversion. In a randomized controlled trial that manipulated access to credit for small retail entrepreneurs in Kenya, we show that experience of credit failure significantly amplifies entrepreneurs' risk aversion. Our design separates the causal effect of credit from selection effects and reveals the critical role of selection into credit. The more risk-loving entrepreneurs endogenously adopt the credit treatment; however, they also become substantially more risk averse as a result of the failed credit experience. We find that these two effects are comparable in magnitude. The analysis of heterogeneous treatment effects identifies key demographic factors, indicating that younger male entrepreneurs managing smaller businesses are more likely to adopt the new credit product. However, these same entrepreneurs exhibit particularly large treatment effects upon a credit failure, leading to a disproportionate increase in their risk aversion. Furthermore, we show that this increased risk aversion has a significant impact on business decisions, leading to a lower counterfactual credit adoption rate of 19%, compared to the actual adoption rate of 26%.

*We would like to thank Jori Barash, Andrey Ordin, Paul Gertler, Prateek Mahajan, and Francesco Trebbi, and participants of the seminars at UC Berkeley, University of Toronto, and UT Austin. We are also grateful to Mastercard, Unilever and Technoserve for their research support and funding.

†University of Texas Austin, UC3M and FTG

‡giyer@berkeley.edu, UC Berkeley, Haas School of Business

§przemekj@berkeley.edu, UC Berkeley, Haas School of Business

1 Introduction

Economic models commonly assume that agents are endowed with stable risk preferences that are unchanged by past outcomes and experiences. However, recent evidence from neuroscience indicates that our response to a set of risky choices is influenced by how our brain processes the outcomes of these choices, suggesting that the wiring of our brain influences risk preferences.[1] Importantly, as documented in the literature on brain plasticity, the brain wiring is not fixed and is affected by past experiences.[2] This provides the rationale for possible *experience effects* whereby our past experiences affect our current economic choices (see Malmendier and Nagel, 2011; Guiso, Sapienza, and Zingales, 2018). Using a randomized control trial (RCT) among small retail entrepreneurs, we show that these experience effects can indeed alter risk preferences and can significantly change future decisions.

Concretely, we study experience effects in the context of entrepreneurship and ask some fundamental yet unexplored questions in this area: Do experiences from past business decisions alter entrepreneurs' risk preferences? Furthermore, how does a shift in risk attitudes affect their inclination towards subsequent entrepreneurial activities? To answer these questions we introduce exogenous variation in entrepreneurial experience by randomly manipulating an opportunity to make a significant, albeit potentially risky, business decision, namely the adoption of a new credit technology. Following their decision, we then measure the entrepreneurs' risk preferences and assess the treatment effect of this experience. Finally, we utilize a structural model to compute counterfactual business decisions under the modified preferences. This methodology effectively navigates the challenge of estimating the causal influence of past experiences on risk preferences, as these experiences may also originate from entrepreneurs' endogenous propensities to pursue risky endeavors. By preventing entrepreneurs in the control group from undertaking the risky action, our approach allows for a clear distinction between the treatment effect and the selection into risky activities.

The study manipulated access to Jaza Duka[3], which is a first-of-its-kind modern and

[1]See Christopoulos, Tobler, Bossaerts, Dolan, and Schultz (2009). The wiring may depend on age and gender and play an important role in whether an individual would choose a risky gamble. See for example Paulsen, Carter, Platt, Huettel, and Brannon (2012) for an fMRI study on the role of age on risk preferences and Sapienza, Zingales, and Maestripieri (2009) for evidence on gender differences in financial risk aversion.

[2]See Knutson, Wimmer, Kuhnen, and Winkielman (2008).

[3]"Jaza Duka" means "fill up your store" in Swahili.

massively accessible retail credit line launched by Mastercard in Kenya. The goal of the credit product was to alleviate the financial constraints of small shop owners by providing working capital to grow their sales. The experiment was carried out in the Malindi region of Kenya on a population of 999 credit-eligible small retail shops. In September 2019 we offered credit to half of the population, while the remaining half did not receive the credit offer. Between November and December 2019, we measured risk preferences using a methodology similar to Holt and Laury (2002) by offering the subjects a series of gambles. The measure of intent-to-treat (ITT) reveals that the opportunity to adopt credit results in more risk averse preferences ex-post. In particular, 7.8% more subjects reject a fair gamble in the treatment group, compared to the control group, which is significant at the 1% level. The gap is larger than the ex-ante differences in risk aversion across gender and age in the control group. The negative average impact of the credit experience on future risk taking can be explained by the adverse experience with the credit product upon adoption and a large default rate. In particular, a staggering 69% of adopters eventually end up having their credit cards restricted. Thus, our main result uncovers that a negative experience from the failed use of a new credit technology leads to more risk averse preferences.

Moving beyond ITT, we apply the potential outcomes framework to identify the local average treatment effect of adopting credit and the extent of selection into credit. The instrumental variables estimator of the treatment on the treated shows that 29.6% of adopters switch from risk-loving to risk-averse as a result of adopting credit. We also measure selection into and out of credit based on ex-ante (untreated) preferences. We show that the selection into credit is substantially larger (over twice the size) than the selection out of credit. Those that adopt credit have substantially lower ex-ante risk aversion than the population. Non-adopters have 8.5% higher incidence of risk aversion compared to the control group. In contrast, adopters have 23.5% lower ex-ante incidence of risk aversion.

Note that the impact of selection into credit is of similar magnitude as the treatment effect of credit adoption. Since treatment increases risk aversion, these two forces cancel each other. In other words, the more ex-ante risk-loving entrepreneurs in the population are the ones that adopt the credit line and are exposed to potential failures. But as a result of this negative experience these entrepreneurs end up becoming ex-post risk-averse.

We identify a more substantial impact of credit on younger entrepreneurs, aligning with the notion that individuals with a less extensive experience base are more significantly influenced. We also find that males are ex-ante less risk averse, are more likely to adopt credit, and modify their preferences to a greater extent compared to females. In particular, female entrepreneurs show a treatment effect on risk aversion that is about half the size of the one observed for males. Moreover, we find that smaller stores experience a relatively larger treatment effect when compared to larger stores. These differences in the treatment effects can be explained by the heterogeneity in the extent of failures in credit usage. In particular, women and larger stores have smaller default rates, thus their credit experience should likely lead to muted or even reversed effects on risk preferences. Indeed, we find that a small segment of females operating larger stores experience positive treatment effects.

To have a clearer understanding of the heterogeneity of the treatment effects, we analyze individual-level wholesale purchase ledgers that involve credit with an aim to identifying risky actions involving credit usage, that can potentially lead to credit failures. We focus on new product adoption, defined as buying a new Stock Keeping Unit (SKU), because it correlates with a larger default rate. To separate the impact of SKU adoption from credit adoption, we augment our experimental variation with supply-side data regarding the effectiveness of sales representatives to cross-sell new products. Our estimates show that the less risk averse retailers among the credit adopters use the credit line to adopt new products that they have never purchased before. We further show that adopting a new SKU increases risk aversion above and beyond adopting credit. In particular, in comparison with the control group, adopting credit without the adoption of a new SKU increases the risk premium by 31%, while adopting both credit and the new SKU increases the risk premium by 44%.

In order to quantify the impact of the change in risk preferences on future credit decisions we construct a structural model that endogenizes credit and SKU adoption. Using the model we compute counterfactual scenarios in which we recalculate credit adoption given the effect of negative experiences on risk aversion. The counterfactual adoption is significantly lower at 19% as compared to the baseline of 27%, underscoring the idea that experiences of past failures have the potential to significantly stunt future entrepreneurial risk taking. In addition, we show that adoption of new SKUs falls from 16% to 10% under the ex-post, more risk-averse, preferences.[4]

Our findings point to some important policy implications for fostering entrepreneurial innovation and technology adoption that optimally balance risk taking and innovation incentives. Early failures can be especially costly in terms of dissuading the adoption of future innovations. This points to the need for more gradual technology introduction policies coupled with learning programs which make the adoption of innovations safer and mitigate the risk of early failures. These considerations may be particularly important for targeting younger entrepreneurs running smaller enterprises who may be more vulnerable to failures that affect their risk aversion.

1.1 Related literature

The role of "experience effects" in the economics literature was first proposed by Malmendier and Nagel (2011). Our paper makes three distinct contributions to this literature. First, we study experience effects in the context of a randomized controlled trial that manipulates exposure to risky business decisions.[5] This helps us to rule out the possible endogeneity of risk preferences to the shock which generates the experience effects in the first place. Crucially, we are able to separate out the causal effect of the treatment from selection effects and show that accounting for selection matters as it is of similar magnitude (and opposite in sign) as the treatment effect of credit. Nevertheless, the adoption of credit conditional on selection still has a substantial causal effect on increasing risk aversion.

Second, we bring to bear the role of past experiences on the important context of entrepreneurship and risky decision making by small subsistence level entrepreneurs. Thus, our findings that adverse first-time experiences with credit can increase risk aversion of entrepreneurs, provide new evidence about how past failures might stunt efficient entrepreneurial risk taking and innovation in the future.

[4]We also use the structural model and information on daily profits to rule out wealth effects (Paravisini, Rappoport, and Ravina, 2017) or mental accounting effects (Thaler and Johnson, 1990) as possible explanations of our results. See Section 6 for a detailed discussion.

[5]With the exception of Cohn, Engelmann, Fehr, and Maréchal (2015), who use a priming manipulation, the existing literature typically relies on observational data for identification (see Callen, Isaqzadeh, Long, and Sprenger, 2014; Cameron and Shah, 2015; Hanaoka, Shigeoka, and Watanabe, 2018; Brown, Montalva, Thomas, and Velásquez, 2019; Jakiela and Ozier, 2019; Shum and Xin, 2022).

Lastly, most of the previous work has focused on the belief formation process. Basically, the main point made by this literature is that lived experiences carry a stronger weight than objective data that is available but which the individual in question did not directly experience or suffer from. In this context, Malmendier and Nagel (2016), Malmendier and Wellsjo (2023) and Kuchler, Piazzesi, and Stroebel (2023) show how past inflation experiences and home ownership status affect an entrepreneur's future beliefs about inflation and home prices influencing their future home ownership and financing decisions. In this paper since we are presenting our retailers with a series of objective gambles, we argue that the effect we document is related to changes in risk aversion. We can show strong validity for the elicited risk aversion measures: e.g., subjects in the treatment group who end up not adopting credit are significantly more risk averse compared to the control. Further, more risk averse individuals also report lower loan uptakes in the past year and higher savings. Related to the literature, we also find that the effects are stronger for the young compared to the old.[6].

In an important paper Sapienza, Zingales, and Maestripieri (2009) describe a related analysis based on the panel data of Italian investors. Like us, they have a similar question eliciting choices between a safe and risky option. Importantly, they asked the same question in 2007 and 2009, that is, before and after the financial crises.[7] They show that this measure of risk aversion increases significantly between both waves. They argue that changes in wealth or future expectations are not consistent with the data, and thus these changes must be attributed to a change in preferences or to the salience of negative outcomes. Like them, we show that neither wealth nor future expectations can explain our findings. We differ in that they cannot rule out that for some exogenous reason there was change in risk-aversion which in turn helped cause the crisis.[8] Indeed, the empirical asset pricing literature basically assumes that the causality goes the other way. Researchers in asset pricing argue that "Discount rates vary a lot more than we thought. Most of the puzzles and anomalies that we face amount to discount rate variation we don't understand." (see Cochrane, 2011). As

[6]This could be because with fewer experiences the brain reacts more to a new stimulus. This is consistent with the evidence that brain plasticity decreases with age (see Burke and Barnes, 2006)

[7]See Sahm (2012)for another study using panel data from the Survey of Consumer Finances albeit relying on a more qualitative question that does not allow one to identify whether the changes are due to beliefs about future prospects from preferences.

[8]Asriyan, Fuchs, and Green (2019) for a rational model that could provide such type of rationalization based on beliefs rather than preferences.

mentioned before the randomized controlled trial in our study allows us to clearly establish causality from the negative experience to the change in risk-aversion. In addition, it allows us to separate out the selection effects (both selection into credit and selection out of credit) from the causal effect and to examine their implications.

2 Setting and Experimental Design

The study involves manipulating the availability of Jaza Duka, a credit program that is a collaborative effort between Mastercard, Unilever, and the Kenya Commercial Bank (KCB). Jaza Duka was introduced in early 2017 to address the financial constraints faced by micro-retailers by providing them with a credit line to access working capital. The core idea behind Jaza Duka is to offer liquidity to small retailers to mitigate stock shortages, enable them to purchase larger pack sizes, and facilitate the opportunity to experiment with new products. Jaza Duka enables retailers to take on additional risks by buying more inventory or trying new products to grow their businesses.

Retailers were required to have a tenure of at least 12 months with Unilever and to demonstrate a sustainable stream of wholesale purchases. Unlike traditional credit programs, Jaza Duka did not require the retailer to have a prior credit history, bring collateral, or be part of a lending group. These features made Jaza Duka a genuinely accessible credit product and the first experience with formal credit for many small retail entrepreneurs who would otherwise not have the credit qualifications required by conventional loans. The low exposure to credit also makes this an ideal setting to document the experience effects.

Jaza Duka functions akin to a modern credit card. However, the credit provided through Jaza Duka could only be utilized to purchase Unilever products. This arrangement had the aim of creating a mutually beneficial situation where Mastercard could use the Unilever purchase history as a substitute for credit scoring. At the same time, Unilever could potentially benefit from increased sales if retailers utilized the credit facility.

Participating stores are provided with a 17-day grace period for credit repayment, during which no interest is charged on the outstanding balance. This interest-free option was especially appealing to the country's Muslim population, many of whom do not approve of credit interest. Within each repayment cycle, stores are required to pay at least 50% of their balance to prevent their credit line from being restricted. If a store fails to meet this payment requirement, their card swiping ability is restricted, and they may eventually be cut off from transacting in the cash channel with Unilever.

Our surveys help us to understand the financial sophistication in the market. We found that many retailers do not maintain written books for their businesses, indicating a lack of formal accounting practices. Additionally, a considerable percentage of retailers do not have a clear understanding of what an interest rate is. This low level of financial sophistication, coupled with the low level of prior experience with formal credit and the complexity of the Jaza Duka program's rules, implies that the adoption of this credit product can be a risky business strategy for these retailers.

Our study population consists of all retailers associated with a single distributor, Banjara, that met the individual credit qualification criteria. The study consisted of 999 retailers who were randomly assigned to either the control or treatment groups. The randomization process was stratified by the size of the stores, which was determined by their pre-experimental volume of purchases from Unilever. The control group consisted of stores that were not provided the credit offer in September 2019. However, they were assured that they would receive credit in the future, with the plan to open credit for them simultaneously with the originally planned roll-out date for Malindi (approximately one year later). Conversely, all retailers in the treatment group were granted access to credit following an accelerated schedule of September 2019. By strategically manipulating the market-level roll-out date, we could create a group that received credit and a control group that did not, without any credit being withheld. This design allowed us to compare the effects of immediate credit access (treatment group) with no access to credit access (control group) on the retailers' risk preferences.[9]

We assume that the offer of credit, when not executed, does not change the risk preferences of the store owners. The decision to adopt credit was optional and there were no negative consequences for not adopting it other than not having the additional liquidity

[9]Control group was originally slated to receive credit 6 months after treatment group, but because of overall delays, and later impact of COVID-19 restrictions, the control did not receive credit until after the pandemic.

that taking the credit would have enabled. In addition, the implementation of our control was strict; that is, none of the retailers in the control arm could adopt credit. Given these aspects of the design, the conditions of Unconfoundedness, Monotonicity, and Ignorability of Noncompliance, as described by Angrist, Imbens, and Rubin (1996), are satisfied in our study.

Our setting also conforms to the Stable Unit Treatment Value Assumption (SUTVA), which requires no peer effects or spillover effects among the stores. Specifically, the offer of credit to one store and the decision to adopt it should not impact other stores. In our setting, there are several reasons why SUTVA is plausible. The Malindi market we chose for our study is predominantly rural, with many stores being the only ones in their respective villages. These villages are also not generally well-connected by formal roads. Further, the brief interval between the treatment initiation and its measurement would necessitate an unusually swift spread of information to materially infringe upon SUTVA.

In the next section, we discuss the descriptive statistics of the data, including the variation in ex-ante risk aversion among the participants.

3 Data

The research uses data collected as part of a large RCT program in Malindi that evaluated the Jaza Duka credit product and accompanying business training offered by Mastercard. The data was collected using three surveys: a baseline survey conducted in person between March and April 2019, a post-treatment survey conducted in person between November and December 2019, and finally a follow-up telephone survey conducted in December 2020. The baseline survey contained questions about demographic, store assortment, competition, and business practices. The post-treatment survey was completed before the onset of the COVID-19 pandemic and was the same as the baseline with added questions about risk aversion, time value of money, and psycho-metric measures. The follow-up end-line telephone survey was conducted after the pandemic and focused on responses on the impact of the economic shock. We also obtained individual-level wholesale ledgers from Unilever that contain aggregate pre-study data and product-level transaction data for the period of the study. The relevant data for this study are: cross-sectional measures of risk aversion (post-treatment survey), demographics (all three surveys), store profits and revenue (post-treatment survey), business decision making, such as loan taking behavior (all three surveys and wholesale ledgers), patterns of leveraged purchases (wholesale ledgers).

We focus our analysis on the sample of 582 stores (labeled as the rational sample) which are run by individuals who respond to the post-treatment survey,[10], identify themselves as sole owners, and who pass the comprehension test (see 3.2 for the description of the test). Table 1 contains summary statistics of our sample, entrepreneurs whose stores were operating, credit eligible, completed mid-line survey and reported as sole owners, across both control and treatment arms.[11] The sample consists of 65% males, as recorded in a post-experiment survey. Age was measured in a follow-up telephone survey and averaged 39 years. The next panel of the table shows education levels. The majority, or 76%, of respondents posses 6-12 grade education level and 10% of the sample holds a college degree. The shop owners population is more educated than the average Kenyan, for example, Statista reported that only 3.5% of Kenyan residents have college degrees.[12]

The third panel of Table 1 contains the distribution of the store size according to the volume reported in the Unilever wholesale ledger. The ledger was assessed using a proprietary score by Unilever which was reported to us prior to the experiment. The score is used by the Unilever sales force to optimize their effort. According to this measure 59% of the stores are

[10]In the Appendix A contains a detailed attrition analysis. The attrition had two primary reasons: stores closed or became credit ineligible between recruitment and roll-out (69 stores); or failed to provide the post-treatment survey (150 stores) leaving us with a sample of 780 stores. We compared all groups using the Unilever pre-experiment ledger data (available for all 999 stores) and demographic information from the baseline survey (available for 891 stores). We find no observable drivers of attrition present in the post-treatment survey. Additionally, Table 7 details balance assessments between sole owners (607 stores) and the rest of the sample population (173 stores) and finds no discernible differences in store size, educational attainment, or experience with financial products. However, it is worth noting that owners tended to be slightly older and there were a greater proportion of males among owners as compared to non-owners.

[11]In terms of demographics, whenever feasible and including gender, we rely on measurements from the post-treatment survey, which is available for all 582 subjects. Education data was collected in the baseline survey, so we define the dummy variables as "responded to the baseline survey and reported a specific education level." Age was recorded in the telephone survey conducted 5 months after post-treatment survey. Age is available for 421 stores. For missing age data, we use the average age when estimating heterogeneous treatment effects. To avoid endogeneity, we use revenue from the baseline, which is available for 500 stores. The attrition decribed in this footnote does not affect most of our conclusions since these variables are utilized only sparingly.

[12]https://www.statista.com/statistics/1237796/distribution-of-population-in-kenya-by-highest-level-of-education-completed/

categorized as small, 29% are medium and 12% are large. We also measured self-reported pre-treatment revenue and find that an average store generates 11,543 Kenyan Shillings per day which amounts to circa $100. The median store revenue is approximately 7,000 Shillings.

The Table confirms that 50% of the subjects were offered credit. Of those, 39% signed up and made at least one purchase on credit. Males had a 42% credit adoption rate, while the female adoption rate was 33%. These numbers are useful in our analysis to assess overall default rate. A more useful measure is the credit adoption rate before risk preferences were elicited in the post-treatment survey. This number amounts to 27%, 30% and 20% for the entire population, males and females, respectively.

The Appendix B contains randomization checks using observable characteristics. We find that out of 16 coefficients none are significant. The smallest obtained p-value is for revenue and amounts to 0.22. Additionally, an F-test, which considers all the variables in the regression together, delivers a p-value of 0.86, and r-squared of 2%, which confirm that the two groups are well-balanced.

3.1 Risk-Taking and Default Behavior

In this section we provide descriptive and model-free evidence of risk taking and credit default behavior. This sets the stage for understanding how prior adverse experiences of credit usage can influence individual risk preferences. According to Table 1, a substantial 69% of those who took out a loan failed to repay on time and had their credit cards restricted. Additionally, 12% experienced hard default, defined as a 180-day delinquency leading to the permanent closure of the account. Such a high default rate may indicate that Jaza Duka was, ex-post, not a beneficial experience for most participants. In part this could be because of the lack of information about loan terms. For instance, according to our post-treatment survey, 30% of credit users report that they were unaware of the interest rate they needed to pay, and 18% were unaware of the length of the repayment period. When asked about their experience with Jaza Duka 25% report it is "fair" or "bad."

For those that adopt credit, the initial credit transaction stands out, being 33% larger than an average purchase (p-value = 0.036). While this fact alone should not immediately provoke concern, given that Jaza Duka is primarily designed to facilitate the expansion of store inventory, a potentially more troubling trend emerges when considering users who eventually face hard credit restrictions. These individuals purchase nearly 50% more in their first credit transaction, indicating that an early over-reliance on credit might be a precursor to a future default.

Furthermore, 56% of those adopting credit introduced new SKUs with their first swipe—items that were previously not part of their inventory. To put this into context, in 2019, 20% of regular purchase events conducted by credit adopters involved the purchase of new SKUs. Again, this development need not immediately raise concerns. After all, these retailers might have wanted to purchase these new SKUs but were formerly precluded due to liquidity constraints. And the stated aim of the Jaza Duka credit program was to enhance the variety and volume of a store's inventory. However, a cause for concern, aside from a large overall default rate, is the disparity in default rates between stores that incorporated new SKUs and those that did not. Specifically, among the credit adopters who defaulted, over 60% introduced a new SKU with their first swipe. In contrast, only 45% of non-defaulters introduced a new SKU (p-value of the difference being 0.107). Additionally, nearly 80% of hard defaulters purchased new SKUs on credit, which is 30 percentage points higher than for non-hard defaulters (p-value of the difference being 0.065). This difference indicates that incorporating new SKUs might represent an additional layer of risk-taking that credit adopters may engage in. More critically, the data suggest that this risk generated adverse financial outcomes, culminating in card restriction and possible hard default.

It is noteworthy that more males than females chose to experiment with new SKUs. This gender difference may contribute to the gender gap in default rates, although the evidence is only correlational. For example, variation in risk aversion may drive both the adoption of new SKUs and default rates. We conduct an analysis accounting for this endogeneity in Section 5.2.

3.2 Risk aversion

Our main outcome variable is an elicited measure of risk aversion using the standard low-value gambles methodology of Holt and Laury (2002). In particular, during an in-person survey, the subjects received a series of hypothetical gambles. Each subject is presented with

	Average	Count	Standard deviation	Total sample
Male	0.65	381		582
Age	38.79		9.08	421
No education, can not read	0.02	9		582
No education, can read	0.04	22		582
Class 1-5	0.08	48		582
Class 6-12	0.76	443		582
Vocational Training	0	2		582
College	0.10	58		582
Small Unilever Segment	0.59	344		582
Medium Unilever Segment	0.29	168		582
Large Unilever Segment	0.12	70		582
Offered credit	0.50	290		582
Pre-treatment revenue	11,543		14,998	482
Used credit, if offered	0.39	112		290
Used credit, Male	0.42	79		190
Used credit, Female	0.33	33		100
Used credit before the survey, if offered	0.27	77		290
Used credit before the survey, Male	0.30	57		190
Used credit before the survey, Female	0.20	20		100
Eventually restricted, if used credit	0.69	77		112
Eventually restricted, Male	0.73	58		79
Eventually restricted, Female	0.58	19		33
Hard default, if used credit	0.12	14		112
Hard default, Male	0.11	9		79
Hard default, Female	0.15	5		33
New SKU on first credit purchase, if used credit	0.56	63		112
New SKU on first credit purchase, Male	0.58	46		79
New SKU on first credit purchase, Female	0.52	17		33

Table 1: Descriptive statistics, final sample.

a choice between two options: a sure payoff of 100 Kenyan Shillings (approximately $1), and a gamble paying π Kenyan Shillings with probability of 50% or 0, otherwise.[13] Since we anticipated that the concept of probability could be difficult for participants to understand we used simple 50/50 gambles instead of a collection of gambles with varying probabilities.

The first gamble in the sequence sets π to 100 Shillings. A rational decision-maker should always reject this gamble, as it is Pareto dominated. Therefore, this choice serves as

[13] The exact wording was: "which one of the following two options will you choose to receive: 100 vs π - 50% of the time."

a means to detect either a lack of comprehension by the subject or potential mis-coding by an enumerator. Some of our analyses that rely on the assumption of rationality cannot be repeated with subjects who fail the comprehension test, unless additional assumptions are made. To facilitate the analysis we drop 25 subjects that accept all the gambles, including the Pareto-dominated gamble. We do not expect significant selection, as there are no statistical differences between subjects who passed and those who failed the test in terms of credit adoption rates, SKU adoption, or default rates (see Appendix A, Table 8). However, results that can be derived from the full sample, such as regressions using the percentage of risk-averse subjects as the dependent variable, were successfully replicated.

A particularly useful measure of risk preferences is the rejection rate of an actuarially fair gamble: 200 with 50% probability and otherwise 0. Individuals rejecting this gamble are risk-averse. Otherwise, they are either risk neutral or risk loving. In the next sections we introduce other more complex measures that aggregate information from many gambles at the individual level; however, in this section, we use fraction of risk-averse individuals to describe the variation in risk aversion in our population.

An advantage of our measure of risk aversion is that the probabilities pertaining to risky choices are objective and easy to understand for the participants. The questionnaire explicitly outlines the likelihood of winning and losing, as well as the stakes involved. This stands in contrast to risky choices involving business decisions, which may often entail some ambiguity regarding the stakes and the odds. A common caveat when utilizing more ambiguous risky choices to assess changes in risk aversion is that the treatment may influence beliefs about the underlying economic primitives in addition to altering preferences for them. Conflating beliefs and preferences may limit transferability of the effect across domains, if the beliefs about stakes and odds are domain specific. For example, those that suffered through a large market crash might be less willing to invest in the stock market, but how would they change their attitudes towards a new medical treatment? Employing transparent and objective gambles in both treatment and control conditions allow us to attribute the change to risk attitudes rather than beliefs alleviating these concerns.

3.2.1 Validity of the Risk Aversion Measure

While our design is consistent with the classic Holt and Laury (2002) procedure, it does prompt a question regarding whether the recovered risk preferences can predict actual risky business decisions, such as adopting loans or new products. The existing literature speaks to the translation of risk preferences from such gambles to both hypothetical and actual higher-stakes situations. For instance, Holt and Laury (2002) themselves note that "behavior is slightly more erratic under the high-hypothetical treatments," while also observing that "...behavior is largely unaffected when hypothetical payoffs are scaled up..." It has been shown that agents may exhibit slightly increased risk aversion when the gambles are scaled up. If this holds true in our case, our estimates of changes in risk aversion may be on the conservative side.

To further validate our measures of risk aversion, we correlate them with individual covariates, reported willingness to take risky business actions, and actual risky business decisions. An immediate measure pertaining to a real-world risky action involves comparing the control group to individuals in the treatment group who reject the credit offer. Both of these groups do not have credit and should exhibit the same measure of risk aversion if the mere fact of receiving a credit offer does not alter preferences, and there was no selection into accepting credit. Notably, we observe that 85% of individuals in the control group reject the fair gamble, compared to 92% among those without credit in the treatment groups (p-value for the difference is 0.019). This leads us to two conclusions: i. there is significant selection on risk preferences when adopting credit, and ii. our measure of risk aversion using the gambles successfully captures this cross-sectional difference in risk aversion.

We also explore the measures that correlate with our ex-ante risk aversion measure in the control group by using the pre-treatment (baseline) survey which was conducted over eight months prior to the elicitation of risk preferences.[14] The time gap allows us to assess the validity and time consistency of our measure of risk aversion. We find that more risk-averse entrepreneurs have lower demand for credit (i.e., number of loan uptakes in the past 12 months). Risk-averse entrepreneurs save more money, and they also tend to have more cash savings than savings at a bank. We also find that older individuals and Muslims are more

[14]This analysis is reported in Appendix C.

risk averse.

4 Results

We conduct our analysis in three distinct steps. Firstly, we examine the impact of the credit offer on risk aversion. These are intent-to-treat (ITT) effects since some owners did not accept the credit offer. We explore the heterogeneity of the ITT effects based on various covariates such as age, gender, business size, and big5 personality metrics. Additionally, we provide both model-free and structural estimates of the ITT effects.

In the the second step, we estimate the effect of the average treatment of the treated (ATT). To do so, we assume that the mere offer of credit has no impact on risk aversion if the recipient does not utilize the credit line. Furthermore, we propose a mechanism that influences risk aversion by analyzing credit default data and SKU-level purchase data. Through the use of a structural model, we demonstrate that the ATT effect is more significant for individuals who purchased new SKUs during their initial credit purchase. This finding is consistent with the experience effect hypothesis since buying new SKUs correlates with the subsequent default.

In the third step, we study the role of wealth as a potential driver of our results. We use observed proxies for wealth and show that changes in wealth are not the primary driving force behind our findings.

4.1 Causal Effect of the Credit Offer

Since the credit offer was randomized, we can estimate its causal effect by comparing averages across treatment and control arms. Our control group is strict because it had no credit available. Strict control is convenient because all adoption of credit before the survey was conducted (that is, 20% of the treatment group) can be attributed to the experimental manipulation. To obtain the impact of the manipulation on risk aversion we start by analysing the raw data pertaining to gamble choices. Figure 1 illustrates the rejection rates for each gamble, accompanied by 95% confidence intervals (CI). The first gamble that is not Pareto-dominated gives the chance of winning KSh 150. This gamble is rejected with a large,

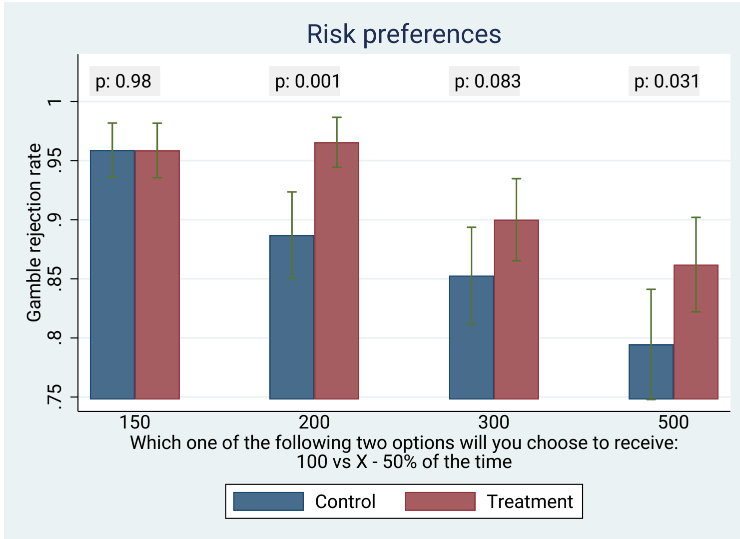

Figure 1: Gamble rejection rate (582 stores, rational sample). Average of the raw data of responding to gamble questions. Treatment is the offer of credit. Brackets are 95% CIs for each bar.

95.8%, likelihood in both arms. Acceptance of this gamble would imply a considerable level of risk-loving behavior.

The next gamble in our study was designed to be actuarially fair, to use it as a metric for gauging the proportion of risk-averse individuals. Our analysis revealed a notable difference in the rates at which the gamble was declined by the two groups under scrutiny: the control group exhibited an 88.7% rejection rate, whereas the treatment group showed a 96.5% rejection rate, resulting in a statistically significant 7.8% increase among those offered credit (p-value less than 0.001). In other words, we find that 7.8% of individuals in the treatment group converted from risk-loving or -neutral to risk-averse preferences. The remaining two gambles, that give a chance of winning 300 and 500, reveal similar patterns with p-values of 0.083 and 0.031, respectively.

Moving to the regression analysis, we estimate models using data on the individual gambles as well as data aggregated across all gambles at the individual level. Column (1) of Table 1 presents the regression results in which the dependent variable is the rejection rate of the fair gamble, revealing a treatment effect of 7.85% with a p-value below 1%. This result is a replication in the rational sample of the earlier finding illustrated in Figure 1. In the Appendix E, Table 12 we report results for various super-samples. For instance, we re-estimated the model for the sample of owners, co-owners and employees (total of 755 stores). We only find significant treatment effects for owners. We also included irrational stores (total of 607 stores), and find a slightly decrease of ITT to 6.7%, yet it remains statistically significant at 1% level.

	(1) Risk aversion	(2) Risk premium LB	(3) Risk aversion	(4) Risk aversion	(5) Risk premium LB	(6) Risk premium LB
Credit	0.0785***	11.51**				
	(0.0215)	(4.799)				
Not adopted credit			0.0848***		14.53***	
			(0.0234)		(5.210)	
Adopted credit			0.0611*	0.296***	3.157	43.35**
			(0.0332)	(0.0865)	(7.408)	(18.78)
IV/ATT	no	no	no	yes	no	yes
N	582	582	582	582	582	582

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Dependent variable choosing a safe amount is "Which of the following two options do you choose to receive, 100 vs 200 - 50% of the time?." Credit variable is a dummy for the offer of credit (ITT). LB indicates lower bound on the risk premium at the threshold.

In our scenario where all gambles have two equally probable outcomes, namely a payout of zero or a payout of π, each gamble can be defined by its payout value, denoted as π_t. For each store i, define rejection threshold, $\bar{\pi}_i$, so that they reject the gamble if and only if $\pi_t \leq \bar{\pi}_i$. For simplicity, we assume that store reject the gamble if they are indifferent. Also define a *risk premium at the threshold* (RPAT) as the excess expected value of the gamble for which the user is indifferent between accepting or rejecting. In other words, it is the maximum amount of money a user is willing to lose in expectation in order to avoid the risk of the gamble. Formally, $\text{RPAT}_i = 0.5\bar{\pi}_i - 100$.

We specify that the risk premium is "at the threshold" to distinguish it from the *risk premium for the fair gamble* – the amount a user needs to be compensated to be indifferent between accepting and rejecting a fair gamble. Denote each gamble by t. We solicit responses to 6 gambles with different values of π_t. Thus, our data partially identifies the threshold $\bar{\pi}_i$. For instance, if the consumer rejects gamble $\pi = 500$, their rejection threshold is greater or equal to 500, but the data does not identify the upper bound on the threshold. Similarly, if the user accepts the gamble $\pi = 500$ and rejects gamble $\pi = 300$, we know that $\bar{\pi}_i \in [300, 500)$. By induction, we can partially identify the CDF of $\bar{\pi}_i$.

The above example applies induction from above by ordering the gambles in descending order and scanning until the first rejected gamble. All subsequent decisions are disregarded, assuming that rational subjects would reject all gambles of lesser value. Alternatively, the gambles can be arranged in ascending order and data can be retained until the first acceptance occurs. These two approaches yield essentially identical empirical conclusions. In what follows we use the induction from above because it provides a more elegant exposition delivering a right-continuous CDF.

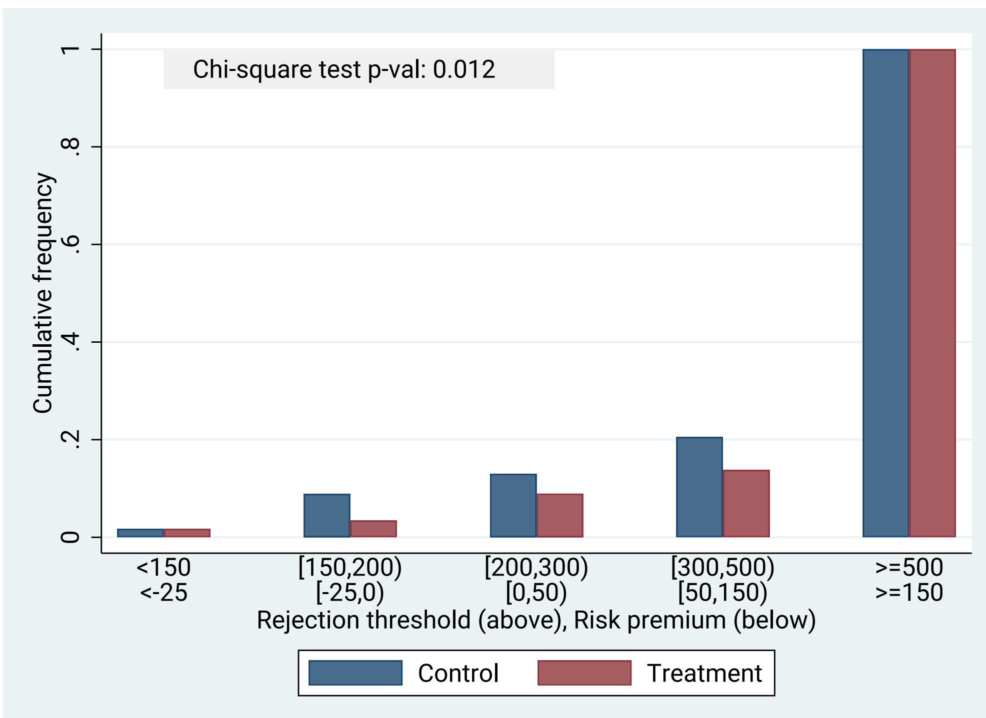

Figure 2: CDF of the gamble rejection threshold. We performed standard chi-square test for correlation of threshold and treatment arm.

Figure 2 depicts the empirical CDF of the threshold in the control and treatment group obtained using the empirical distribution of $\bar{\pi}_i$. Each point on the X-axis represents possible rejection thresholds, and corresponding risk premiums "at the threshold", that are implied by the gamble choices. The empirical mass of the rejection thresholds in the treatment group is shifted towards higher thresholds, which indicates larger risk aversion. This is with the exception of the the first two bars, which indicate an insignificant shift towards lower thresholds in the treatment group. We also performed a χ^2 test for the correlation between thresholds and manipulation arm. We obtained a p-value of 0.012, which indicates that the offer of credit affects the CDF of rejection thresholds; thus, altering the risk preferences.

Column (2) in Table 2 contains regressions of the lower bound of risk premium on the credit offer dummy. We use the lower bound because the upper bound is sometimes equal to infinity. This issue prevents us from using the upper bounds in the regression, as well any functions of it, such as the middle of the interval. We find that on average the lower bound on the risk premium increases by 11.5 Shillings, which is 11.5% of the value of the certain payoff.

4.2 Selection into and Causal Effect of Credit

In this section, we define and estimate causal impact of credit and selection in and out credit. For this purpose, we use a simplified notation for potential outcomes that is similar to Duflo, Glennerster, and Kremer (2007). We observe a single measure of risk aversion per individual, solicited after the credit was utilized. However, thanks to the randomization, we observe 3 groups of subjects: stores that were offered the credit and adopted it (IN), stores that were offered the credit but did not adopt it (OUT), and control stores that did not receive any credit offer.

We assume that the mere offer of credit does not impact risk-aversion. Rather the risk aversion outcome is conditional on the adoption decision.[15] Thus, we consider two potential outcomes of risk aversion for store i: when adopted credit, Y_i^A, and when not adopted credit, Y_i^{NA}, regardless of the experimental assignment.

The treatment effect (ATT) is given by $E[Y_i^A - Y_i^{NA}|\text{IN}]$. It the can be shown that under

[15]Note also that subjects could not obtain the Jaza Duka credit without having the credit offer. Therefore, all subjects who adopted credit are *compliers*, and there are no *defiers*.

standard assumptions of Angrist, Imbens, and Rubin (1996) (see Appendix D):

$$
\begin{aligned}
&\underbrace{E[Y_i^A|\mathrm{IN}] - E[Y_i^{NA}|\mathrm{OUT}]}_{\text{observed difference between adopters and non-adopters}} = \\
&\underbrace{E[Y_i^A - Y_i^{NA}|IN]}_{\text{ATT}} + \underbrace{E[Y_i^{NA}|IN] - E[Y_i^{NA}|OUT]}_{\text{ex-ante difference between adopters and non-adopters}} = \\
&\underbrace{E[Y_i^A - Y_i^{NA}|IN]}_{\text{ATT}} + \underbrace{E[Y_i^{NA}|IN] - E[Y_i^{NA}]}_{\text{selection into credit}} - \underbrace{(E[Y_i^{NA}|OUT] - E[Y_i^{NA}])}_{\text{selection out of credit}}
\end{aligned}
\tag{1}
$$

If the impact of credit on risk aversion were assessed using observational data, or without the inclusion of a control group, the observed differences between adopters and non-adopters would encompass both causal (treatment) effects and selection biases. In our context, an observational data approach would be akin to disregarding the control group and solely comparing the credit adopters to non-adopters within the treatment group.

The bias in the observational estimate, attributable to selection, is articulated through Equation (1). Specifically, the observational estimate (which is the observed difference in risk aversion between the adopters and non-adopters of credit) represents the sum of the causal effect of credit adoption and the difference in the ex-ante risk aversion between adopters and non-adopters. Note that the ex-ante risk aversion of adopters is not directly observable. The difference in this ex-ante risk aversion can be further dissected into two components: *selection into credit* and *selection out of credit*. These are defined as the average differences in risk aversion between the respective group (i.e., adopters and non-adopters) and the overall population mean.

The benefit of our randomized study design lies in its ability to not only estimate the causal effect but to also dissect both selection effects. Studies in the experience effects literature which rely only on observational panel data, in general, cannot separate out causal and selection effects and analyze their implications. To achieve this, we employ both reduced form and structural approaches. In the remainder of this section, we focus on presenting the results derived from the reduced form approach.

Column (3) of Table 2 contains an OLS regression of the indicator function for risk aversion on dummies for credit adopters and non-adopters in the treatment group. Compared to the control group, both credit adopters and non-adopters display a higher proportion of risk-averse individuals, by 6.1% and 8.5% respectively. Thus adopters have a 2.4% lower proportion of risk-averse subjects compared to non-adopters.

Column (4) presents the IV estimate for ATT. It indicates that 29.6% of credit adopters shifted from being risk-loving or neutral to risk-averse as a result of the treatment. The difference between ex-ante percentage of risk-loving and risk-averse populations is calculated at 32.0%, signifying that credit adopters were, ex-ante, significantly more risk-loving.

To decompose the estimated difference in risk aversion between credit adopters and non-adopters into selection effects, we begin by estimating the ex-ante risk aversion of non-adopters. This is achieved by analyzing the average risk aversion among non-adopters in the treatment group. We estimate that 97.2% of non-adopters in the treatment group are risk-averse, compared to an 88.7% average across the general population, as determined from the control group. Therefore, the selection out of credit – i.e., the increased likelihood of being ex-ante risk-averse among those who did not adopt credit – is approximately 8.5%. Consequently, the remaining disparity can be attributed to selection into credit, which we calculate to be -23.5%.

The fact that selection into credit is larger than selection out of credit could be due to the smaller size of the credit-adopting population. However, it might also suggest a heavy-tailed distribution among those opting into credit, indicating that individuals with a significantly lower aversion to risk are more inclined to engage with credit opportunities. The implication of this finding is that in similar settings with low adoption of the risky option, one can expect downward bias in the observational estimates, mostly due to the selection into the risky option of less risk-averse individuals.

This substantial selection into credit supports the hypothesis that engaging with credit is akin to embarking on a risky venture, with individuals who are less risk-averse being more inclined to take up credit. While this selection into credit effect is notable (i.e., -23.5%), it is still somewhat less pronounced in magnitude than the treatment effect (29.6%). This implies that the experience of adopting credit transforms the preferences of adopters, aligning them just above the typical level of risk aversion seen in the general population, namely, 6.1% above as reported in column (3) in Table 2.

We conducted the same analysis using a pooled sample of all gambles and using the risk premium as the dependent variable. The results are presented in Columns (5) and (6) of Table 2. Our findings reveal that credit adopters have an *ex-post* risk premium that is 11.37 Shilling lower that of the non-adopters. The IV estimate suggests that the experience of credit raised the risk premium of adopters by 43.35 Shillings. These two numbers imply that ex-ante difference between adopters and non-adopters was 54.72 Shillings, or 54.72% of the risk-free payoff. The risk premium of non-adopters is equal to 133.45 Shillings, and the population average is 118.92 Shillings. This implies that the magnitude of the selection out of credit equals 14.53 Shillings and that of selection into credit -40.19 Shillings.

Overall, the findings of the risk premium analysis are similar to the those with the risk aversion measure. The selection into credit is substantial, but the treatment effect is even more pronounced. One difference worth noting is that the risk premium method yields higher selection estimates than the ATT. This difference can be attributed to the fact that the risk premium approach incorporates data from the tails of the risk distribution, which is informed by more skewed gambles. For example, the data reveal that individuals who were risk-averse before the introduction of credit became more so following their experience with credit. If the selection for such individuals was more extensive than the ATT, it might not be fully apparent in an analysis that relies exclusively on the fair gamble.

4.3 Heterogeneous treatment effects

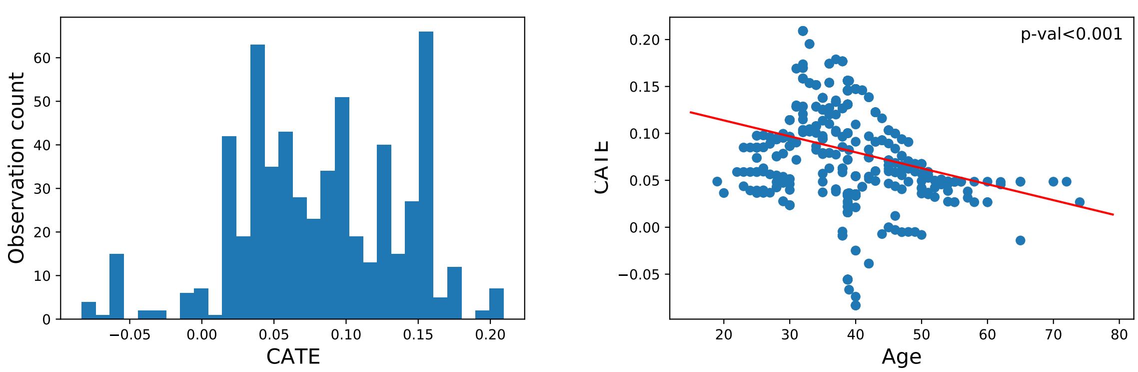

Figure 3: Results from Double Machine Learning Causal Forest. Distribution of Conditional Average Treatment Effects in the population (left panel). Histogram of CATEs against age with the linear regression line (right panel).

Our exploration of heterogeneous treatment effects begins with an examination of covariates using a causal forest methodology (see Battocchi, Dillon, Hei, Lewis, Oka, Oprescu, and Syrgkanis, 2019). We focus primarily on demographic variables, including store size, age, gender, and religion, and employ acceptance of a fair gamble as the outcome variable. It's important to note that the results obtained from the machine learning analysis should be interpreted as intent-to-treat (ITT).

In the left panel of Figure 3, we present a histogram of estimated heterogeneous conditional average treatment effects (CATE). The average effect is approximately 0.082, closely resembling our ITT estimates. The standard deviation of the treatment effects is around 0.057, indicating moderate dispersion. Notably, approximately 6% of the estimates are positive, suggesting that for a small segment of entrepreneurs, the offer of credit decreased their risk aversion.

Our specific interest lies in understanding the impact of age on the strength of the treatment effect. As mentioned earlier, age may be an important moderator because as we previously argue, younger individuals' who likely have lower accumulated stock of past experiences would be more susceptible to being influenced by the current experience. In the right panel of Figure 3, we regress CATE on age, which indicates a negative and obtain statistically significant relationship, confirming our hypothesis.

To further rank the demographic moderators in terms of their influence on the treatment effect, we employ a decision tree. This tree is constructed by sequentially selecting a variable and its split that has the most predictive power in explaining the variation in the treatment effects. The resulting tree is depicted in Figure 4. The three most important variables are gender, age and store size, with gender emerging as the factor with the most predictive power. Notably, women display a treatment effect on risk aversion that is almost half as pronounced as that observed in men. Several mechanisms could be contributing to this difference. Primarily, a greater number of men than women took up credit; consequently, the intent-to-treat effect should be more substantial for men than for women. Furthermore, we demonstrate in the latter part of this section, even after accounting for the gap in adoption rates, men still show a more significant effect than women. This disparity can be ascribed to differing experiences with credit, evidenced by men having a default rate of 73% compared to

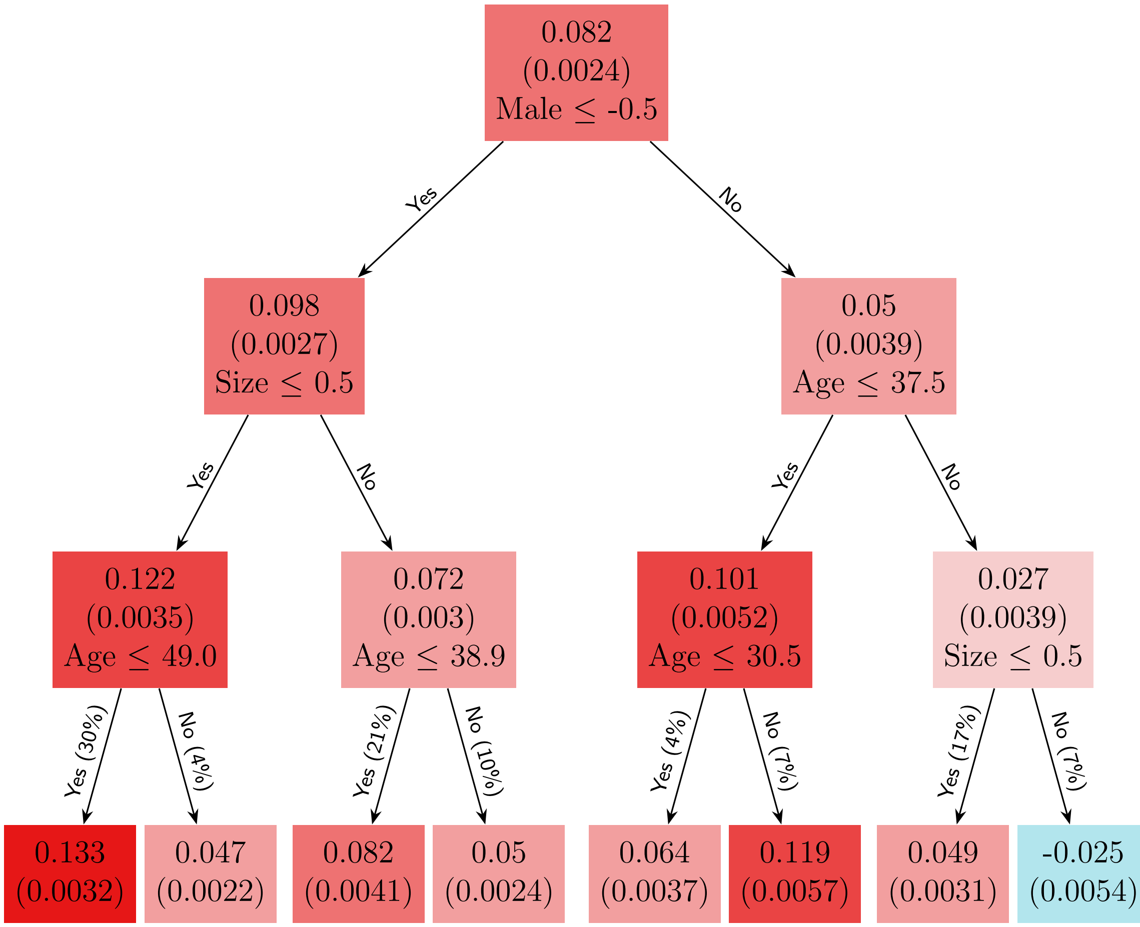

Figure 4: Decision tree depicting decomposition of Conditional Average Treatment Effects.

a 58% default rate for women (with a p-value of 0.089). The differential experience can stem from distinct patterns of credit usage across the genders, ex., 58% of men who took credit adopted new SKUs on credit compared to 52% of women (with a p-value little over 0.1). Appendix E.4 confirms that cross-gender differences occur mostly on the intensive margin, when using credit. In particular, ATT for males is more than twice that of females.

Store size is the most important predictive variable for males and second most after age for females. Smaller stores experience a much larger treatment effect on risk aversion compared to larger stores. In the data, we observe that larger stores that adopted credit generally have lower default rates, supported by a χ^2 test p-value of 0.052. This difference may be explained by better credit experience and more effective credit usage by larger stores.

Age is the most important variable for females and second most important for males. The magnitudes are similar across genders but the relevant cutoff age is higher for males than females. The younger group shows a larger effect, with a possible exception of very young women (the sample size of very young women is too small to draw statistically robust conclusions). Additionally, for subjects over 39 and those under 39 years old, the default rates are 71% and 67% respectively (difference not statistically significant); thus, younger entrepreneurs updated more despite experiencing the same or smaller level of default.

We also repeated the same analysis, including psychometric variables on top of demographics. Most of the results are closely aligned with those that use only demographic variables. One additional insight is that among all the psychometric variables, "finding fault with others" emerges as the strongest moderator of the treatment effect. This suggests that individuals who attribute their failures to external factors rather than to themselves may be less likely to internalize the failure and to adjust their risk preferences.

5 Utility Model

This section evaluates the determined effects of credit adoption through various structural models that assume utility-maximizing behavior. This allows us to specify structural parameters that encapsulate intrinsic risk preferences, and distinguish them from other possible driving forces, in particular from wealth effects. Further, the models endogenize the observed real business decisions, such as credit adoption and usage patterns, enabling us to gauge the economic impact of the preference shifts. Notably, the analysis provides insight into how the entrepreneurs might modify their future business decisions in light of their altered preferences. By adopting this approach, we can measure the impact of the credit adoption experience on preferences and use it to assess its influence on future entrepreneurship.

We examine two models. In Section 5.1, we analyze a model that endogenizes the gamble rejection and credit uptake decisions. This permits us to predict what credit adoption rates might have been under the ex-post preferences shaped by the experience with credit. Importantly, in this section we contrast the estimates for Constant Absolute Risk Aversion (CARA) and Constant Relative Risk Aversion (CRRA) utility models and show that accounting for wealth effects does not change our main findings. In Section 5.2, we expand the model to endogenize the entrepreneurs' decision to adopt new SKUs. This addition highlights which credit use patterns lead to the most significant preference shift. Further, we can estimate counterfactual credit usage patterns, in addition to adoption rates, under the preferences shaped by the credit experience.

5.1 Credit Adoption Model

In this subsection, we develop a structural model, which relates the change in preferences to credit adoption. The strength of this model lies in its close connection to our descriptive analysis, as it leans on the experimental variation for identification. This model yields two vital outputs: firstly, it offers a distribution of both ex-ante and ex-post risk preferences, accounting for wealth effects. Partialling out wealth, or controlling for wealth effects, allows us to investigate whether changes in revealed risk attitudes are primarily driven by a fundamental shift in preferences, or whether they are a result of wealth effects. Secondly, the structural model facilitates the calculation of credit uptake under updated preferences or, more specifically, the hypothetical decision to adopt credit under circumstances similar to the original decision but driven by post-experience preferences.

Consider a utility function $u(\pi_t; \gamma_i, w_i)$, where γ_i is a structural preference parameter embodying inherent risk preferences, and w_i is the current level of wealth. We postulate the following simultaneous equations model of risk aversion and credit adoption:

$$\gamma_i = \bar{\gamma} + \epsilon_i + \Delta D_i \tag{2}$$

$$D_i = \begin{cases} 1 & \text{if } V + \nu_i > 0 \text{ and } Z_i = 1 \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where D_i indicates credit adoption, $Z_i = 1$ indicates the treatment arm, and $Z_i = 0$ indicates the control arm.

The risk preference parameter is composed of three terms. The first term $\bar{\gamma}$ term represents population average ex-ante risk aversion without the credit offer, i.e., in the control group. The second term, ϵ_i embodies individual level differences in ex-ante risk aversion and it explains the variation in gamble take up in the control group. The third term ΔD_i

represents the treatment effect of credit adoption D_i on the level of risk aversion.[16]

The second equation details the decision to adopt credit, D_i. The term V denotes a population average surplus, while ν_i signifies the idiosyncratic surplus. The sum $V + \nu_i$ represents the certainty equivalent of the net present value of adopting credit after subtracting adoption costs. For example, this could include payoffs from purchasing additional inventory with credit, encompassing both new products and more of the existing stock. It also takes into account returns from extra cash available after leveraging some existing purchases and the costs of credit, such as monitoring, transaction costs, interest rates, and potential defaults.

Given that $V + \nu_i$ is a certainty equivalent its value depends on risk preferences; thus, the model must allow for correlation of $V + \nu_i$ and γ_i. Because of this correlation, as previously mentioned, we generally anticipate that $E[\epsilon_i|D_i] \neq 0$. This endogeneity issue was the primary reason for conducting the field experiment. To address endogeneity, we allow for an arbitrary relationship between V and $\bar{\gamma}$, and we examine the joint distribution of ϵ_i and ν_i, acknowledging their potential correlation. The identification of the model depends on the exclusion of Z_i from Equation (2) (unconfoundedness).

We parameterize the model by considering two utility functions: CARA, and CRRA. For the CARA utility, the level of risk aversion is determined by a single parameter, and the corresponding utility function is given by:

$$u(\pi_t; \gamma_i) = \frac{1 - \exp(-\gamma_i \pi_t)}{\gamma_i}. \tag{4}$$

This framework is convenient because of the absence of wealth effects. It allows us to establish a benchmark for considering the importance of wealth effects in driving our results.

Since Jaza Duka credit resulted in a significant amount of default, it is likely that treated entrepreneurs end up with different wealth levels than their untreated counterparts. We allow the data to indicate if survey responses are effected by various ex-post measures of wealth. In other words, we would like to determine if the increase in risk aversion in the

[16]We assume no unobserved variation in the impact of credit on risk aversion, Δ. While it is possible to account for heterogeneous Δ, we have chosen not to pursue that direction. Instead, we incorporate observable heterogeneity in our analysis wherever the available data permits. For further discussion of unobserved heterogeneity in treatment effects, see page 74 of Heckman and Robb (1986).

treatment arm is driven by a decrease in wealth related to high rates of default. To answer this question it is helpful to consider the CRRA utility function, i.e.,

$$u(\pi_t;\gamma_i,w_i)=\frac{(w_i+\pi_t)^{1-\gamma_i}}{1-\gamma_i}. \tag{5}$$

If wealth was observable for each individual when making the decision between a gamble and risk-free outcome, we could condition on the level of wealth. Unfortunately, we do not have wealth information. We adopt two approaches to proxy for wealth. First, we use daily profits which is likely the most relevant driver of differences in wealth induced by credit. Second, we estimate the wealth effects directly, using the data on multiple gambles.

To close the parametric specification of the model, we postulate that the joint distribution $F(\epsilon,\nu)$ is Gaussian with mean 0 and with the variance-covariance matrix applying standard Probit normalization for the adoption equation as:

$$\begin{bmatrix}\sigma_\epsilon^2 & \rho\sigma_\epsilon \\ \rho\sigma_\epsilon & 1\end{bmatrix}.$$

We estimate the model using Simulated Maximum Likelihood Estimation (SMLE) using implied gamble and credit adoption choices. The unit of observation is a single entrepreneur (the unit of randomization). We obtain standard errors by using a non-parametric bootstrap which samples entrepreneurs with replacement from the empirical distribution. Standard errors are clustered the entrepreneur level.

5.1.1 CARA Utility Model

Table 3 presents the results of the estimation. Column (1) displays the outcomes of the CARA model that does not allow for wealth effects. This approach closely mirrors the model-free analysis, attributing all experimental variations in gamble acceptance rates to inferred differences in risk preferences.

Accounting for the findings from the previous section, we allow for heterogeneity in the risk aversion distributions and treatment effects between males and females. We again observe slight but statistically insignificant disparities in the base risk aversion between men

		(1) CARA No wealth effect	(2) CRRA Constant wealth	(3) CRRA Proportional to daily profits	(4) CRRA Estimated wealth effect
RA intercept, $\bar{\gamma}_{\text{FEMALE}}$		1.76***	3.17***	2.93***	3.20***
		(0.17)	(0.25)	(0.46)	(0.25)
RA intercept, $\bar{\gamma}_{\text{MALE}}$		1.73***	3.07***	2.87***	3.06***
		(0.16)	(0.28)	(0.45)	(0.29)
RA Dispersion, σ_γ		1.21***	2.13***	2.03***	2.14***
		(0.13)	(0.16)	(0.29)	(0.17)
RA-adoption correlation, ρ		−0.39***	−0.45***	−0.52***	−0.44***
		(0.13)	(0.12)	(0.13)	(0.12)
Impact of adoption, Δ_{FEMALE}		0.16*	0.10***	0.31	−0.09***
		(0.09)	(0.02)	(0.36)	(0.02)
Impact of adoption, Δ_{MALE}		0.67**	1.36***	1.35***	1.27***
		(0.29)	(0.40)	(0.39)	(0.42)
Wealth intercept, $\bar{w}$		-	59.31***	51.59***	68.47***
			(10.89)	(16.94)	(16.57)
Wealth slope, ϕ		-	-	0.00***	-
				(0.00)	
Wealth effect, Δw		-	-	-	17.62
					(25.78)
Risk premium for the fair gamble	Control group	52.3	67.8	64.7	62.8
		(2.09)	(2.32)	(2.32)	(3.29)
	ITT	2.99*	4.13***	5.24**	3.88**
		(1.73)	(1.52)	(2.28)	(1.52)
	SEL out of	5.09**	6.48***	7.76***	6.33***
		(2.26)	(2.21)	(2.69)	(2.28)
	SEL into	−13.80**	−17.67***	−21.18***	−17.38***
		(5.83)	(5.45)	(7.48)	(5.59)
	ATT	11.10*	15.41***	19.53**	14.52***
		(6.41)	(5.49)	(8.86)	(5.38)
Credit adoption	Baseline	27%	27%	27%	27%
	Counterfactual under ex-post preferences	20%	19%	16%	18%
	Hypothetical Good experience	32%	37%	40%	36%
N		582	582	582	582

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Results from estimation of the structural model. Model (3) uses individual-level daily profits as a proxy for wealth. Models (4) and (5) use average daily profits – averages are taken separately for the control group, credit non-adopters, and credit-adopters.

and women. It is further validated that men show a much greater impact of credit adoption on risk aversion compared to women. Nevertheless, we discern a weakly significant effect for females, roughly equivalent to the cross-gender difference in ex-ante preferences.

Next, we calculate the implied treatment and selection effects. Since we estimate struc-

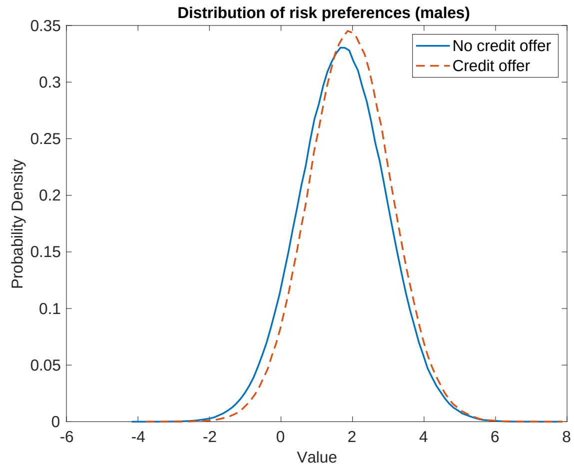

Figure 5: The figure displays the distribution of γ across the male population, encompassing both adopters (treated) and non-adopters (untreated). The solid blue line represents the baseline (ex-ante) risk aversion. The dashed red line shows the ex-post risk aversion following adoption of Jaza Duka.

tural preferences, we can determine the exact implied changes in risk premiums for the fair gamble, instead of having bounds at threshold of acceptance.[17] The ITT, in terms of risk premium, amounts to 3 Shillings (or 6% of the control group average risk premium). This means that stores in the treatment group are willing to pay up to 3 additional Shillings to avoid a fair gamble, over and above 52.3 Shillings in the control group. We also confirm the conclusions from the reduced form analysis, especially the observation that selection into credit (i.e., individuals who opt for credit initially have lower risk aversion) is significantly larger than than the selection out of credit. Furthermore, we find that the increase in the risk aversion after adoption (i.e., the ATT) is of approximately of the same magnitude as the selection into credit. This aligns with the reduced form results, particularly regarding the risk premium at the threshold.[18]

Beyond the above results the structural analysis allows us to compute credit adoption counterfactuals based on our estimated credit adoption model. In the bottom section of the table, we present counterfactual scenarios in which we recalculate the adoption of credit, considering the effect of experiences of individuals on their risk aversion. Specifically, we

[17]Formally, risk premium is defined as $u^{-1}\left(\frac{1}{2}u(200;\gamma_i,w_i)+\frac{1}{2}u(0;\gamma_i,w_i)-u(100;\gamma_i,w_i);\gamma_i,w_i\right)$.

[18]Similar to the reduced form analysis, we observe that the comparison between selection and ATT depends on the dependent variable. When we recalculate selection and ATT using percentage of risk-averse individuals as the dependent variable, we again find that the treatment effect is larger in magnitude than the selection into credit.

conduct a new simulation of credit adoption decisions using the preferences that individuals would have if they had already undergone the credit experience. For each individual, we generate credit adoption shocks, denoted as ν_i, from a conditional distribution that adjusts their ex-ante preferences, represented by γ_i, to $\gamma_i = \bar{\gamma} + \epsilon_i + \Delta$. Since ν_i is inversely related to γ_i (greater risk aversion leads to lower credit adoption), the counterfactual adoption rates for credit are reduced. To be precise, the model's initial adoption prediction stands at 27%, whereas the counterfactual adoption rate decreases to 20%.

This exercise serves as an additional means of measuring the impact of our findings, this time utilizing real-world decisions rather than hypothetical scenarios. This phenomenon may provide one explanation for the "adoption puzzle" observed in developing countries (see de Janvry, Sadoulet, Dar, and Emerick, 2016), wherein entrepreneurs tend to under-adopt new practices that are theoretically advantageous. The suggestion of our analysis is that such under-adoption may result from past setbacks in analogous circumstances and the subsequent increase in risk aversion. In this sense our findings indicate a quantifiable implication of credit adoption on entrepreneurship. For instance, if Mastercard were to introduce another round of Jaza Duka, addressing the issues identified in the initial rollout, they should anticipate 30% lower adoption rates compared to the original wave.

Section 4.3 provides some indicative evidence that positive credit experiences could potentially have the opposite effect on preferences.[19] This opens the possibility of exploring the our model's counterfactuals to hypothesize symmetrically opposite effects of positive experiences on preferences. We simulate a hypothetical scenario with a treatment effect of $-\Delta$, and find that counterfactual credit adoption rises from 27% to 32%. Compared to negative experiences, the opposite shift in preferences has a slightly muted effect on adoption suggesting a smaller impact on the marginal adopter.

Beyond the first-order effects, our model suggests that the combination of selection and treatment influences the dispersion of risk preferences. Selection implies that more risk-tolerant individuals are more likely to be treated, but may revert to being more risk-averse afterward. In this way, negative experiences could lead to the homogenization of risk pref-

[19]For instance, in a small proportion of individuals (6%), the credit offer led to a reduction in risk aversion; specifically. According to Figure 4, these individuals were mostly older females with larger stores.

erences, as illustrated in Figure 5.[20] Conversely, positive experiences (i.e., $-\Delta$) could have the opposite effect of making the ex-post risk preferences more heterogeneous.[21]

5.1.2 CRRA Utility Model

Columns (2) to (4) of Table 3 encompass the estimates of the CRRA model, each with a distinct econometric specification for wealth. First, as a benchmark, in Column (2) we estimate the model assuming wealth is the same for all subjects. Both wealth and the coefficient of relative risk aversion are estimated. This can be done since we offer several different gambles to each individual. While this model overlooks potential wealth effects from credit adoption, it facilitates a direct comparison of the CARA and CRRA functional forms. The initial six rows outline the primitives of the utility function. We yet again discern a modest gender-based difference in risk aversion and a faint effect of credit adoption on women. We find that the wealth estimate, $\bar{w}$, that most accurately reflects the gamble choices is approximately 60 Shillings, indicating that participants do not factor in their total wealth when deciding between gambles. This is consistent with previous studies with low stake gambles.[22] It is also worth noting, that our estimates for γ are close to the median value of 3.77 from a meta-analysis of 92 studies by Elminejad, Havranek, and Irsova (2022).

Looking at the second panel we see that the results in terms of risk premium are within 10% of those obtained with the CARA model. The effect on the counterfactual exercise is larger with adoption falling 1 additional percentage point to 19%.

The next step in our analysis is to enhance the specification to consider the potential influence of credit on wealth. We start by using the data on daily profits (collected in the same survey as the risk aversion measure) to create a proxy for wealth. We assume that an entrepreneur's daily profits can serve as a direct representation of their wealth, as for many

[20] This reasoning presumes that the treatment effect is homogeneous or that its variability is not too large.

[21] We should emphasize that these results on the dispersion of risk preferences are based on one-shot experiences with credit adoption. Our experiment is not designed to directly examine the long-term effects of how past experiences can shape the dynamics of this dispersion.

[22] For instance, Palacios-Huerta and Serrano (2006) demonstrates that to attain reasonable levels of the coefficient of relative risk aversion in relation to small gambles, the implied level of wealth must be correspondingly modest. Otherwise, such calibration may result in preferences characterized by extreme levels of relative risk aversion, leading to the anomalies highlighted by Rabin and Thaler (2001).

entrepreneurs, the store is their only source of income.[23] We consider the following equation for wealth: $w_i = \bar{w} + \phi \times (\text{DAILY PROFITS})_i$. Column (3) offers results from this model. All the values are very close to those in Column (2), indicating that even after accounting for wealth, there remains a large unexplained variation in risk-taking across the three relevant groups.

The estimates of selection and treatment effects reflect only change in preference after controlling for changes in wealth. To isolate the influence of intrinsic risk preferences, we (i) use updated structural parameters estimated when keeping wealth fixed, and (ii) fix the level of wealth to the mean in the control group when computing all risk premiums. The minimal difference between treatment effects from Columns (1), (2), and (3) suggests that the wealth repercussions of credit have minimal explanatory power concerning the influence of credit adoption on gamble decisions.

While using individual-level daily profits provides a granular view of financial standings, it may also capture inherent variability the survey, potentially leading to attenuation of the estimate of ϕ and underestimation of wealth effects. To address this, instead of using individual responses, we use average reported daily profits and compute the averages separately for the three crucial groups: the control group, the treatment group without credit, and the treatment group with credit. This specification adjusts for average cross-sectional wealth differences. The results of this exercise are identical to those obtained using disaggregated wealth data. Additionally, for robustness, we considered a quadratic specification in profits and also obtained identical conclusions (exact numbers for both robustness checks are not reported here for brevity).

Next, we directly estimate the wealth effect of credit in addition to estimating changes in γ. Specifically, we introduce an extra parameter Δw that measures the change in wealth due to adopting credit, such that: $\bar{w}_i = \bar{w} + D_i \Delta w$

The extra parameter is identified since we observe multiple gambles for every individual in the treatment and control group; thus, wealth and risk aversion parameters in both groups are recoverable from our data, under the functional assumption of CARA preferences.

The results are presented in Column (5). The wealth effect amounts to approximately

[23] According to our survey, only 18% of entrepreneurs report having other income besides the store.

25% of the initial wealth. This wealth effect is plausible because it corresponds to losing most of the Unilever sales, as our subjects reported that Unilever constitutes approximately 30% of their total revenue. Since most of the credit takers have not completely lost access to Unilever products, we believe that the point estimate is on the larger side, making our estimates of treatment effects conservative. The treatment effects at the bottom of the table are purged of wealth effects by fixing wealth at $\bar{w}$ and varying only γ. Despite this adjustment, we observe negligible impact of directly accounting for wealth effects on our results. In particular, ITT decreases from 4.13 to 3.88. Thus, we conclude that wealth changes cannot explain the shift in risk-taking behavior.[24]

5.1.3 Wealth Expectations and Discounting

Beyond the impact of current wealth levels on risk preferences, one can anticipate that expectations of future wealth might also play a significant role. For instance, more optimistic expectations about future returns may lead to a different risk preference than more pessimistic expectations. If adopting credit affected these expectations, one might detect it as a change in γ, even while keeping current wealth constant. To eliminate this potential confounding factor, we gauged expectations about future wealth by asking participants the question, "After 12 months from now, what do you think will be your daily revenue?" This question was posed in both the baseline and midline surveys, yielding panel data. Utilizing both cross-sectional and panel data variation, we conducted a series of regressions in an attempt to discern the impact of the treatment arm on future expectations. No significant differences in expectations were detected as detailed in the Appendix E, Table 13.

In addition to measuring risk preferences, we also assessed time preferences by posing a series of questions such as "Which of the following two options would you prefer: 300 in 1 week or X in a month?" where X equaled 310, 350, 400, 500, and 600. If either current or

[24] Our estimate of the wealth effect is not statistically significant, indicating that estimating two-dimensional treatment effects is pushing the boundary of our data. Nevertheless, treatment effects, after purging the wealth effects, remain precisely estimated. To offer another interpretation of this result, we recomputed the standard errors by calibrating the value of the wealth effect to the point estimate of 25%. As expected, this resulted in smaller standard errors across the board, strengthening our conclusions. To push this argument to the extreme, we re-estimated the model by calibrating the wealth effect to 50% of the initial wealth. The ITT decreased from 4.13 to 3.30, indicating that the majority of the impact of credit on preferences remains substantial even with under assumption of significant decrease in wealth.

future wealth changed as a result of the treatment, it is plausible that such changes would be reflected in time preferences. For example, if the treatment negatively affected current cash flows, one might expect subjects to exhibit more impatience. Conversely, if credit negatively impacted future cash flows, subjects might demonstrate a greater propensity to save. When we analyzed the acceptance rates in the time-preference questions, we found no significant differences at the 5% level between the treatment and control groups, as detailed in the Appendix E, Figure 7. This provides further evidence that wealth effects and future wealth expectations are not significant, as they should lead to change in time preferences. Moreover, this finding supports the hypothesis that risk and time preferences are distinct entities, as formalized in the theory of dynamic choice and temporal uncertainty resolution following Kreps and Porteus (1978) and Epstein and Zin (1991).

5.2 Impact of New SKU Adoption

To this point, we have established that individuals who take on the risk of leveraging their purchases of Unilever products become more risk-averse. In this subsection, we explore the mechanisms underlying this phenomenon. We have already noted substantial default rates that contribute to the shift in preferences. To gain a better understanding of this impact, it is useful to zoom into the wholesale ledgers at the level of SKU transactions. As mentioned in Section 3.1, a particular action, that is, purchasing new SKUs on credit, is significantly correlated with default. Thus, this decision, over and above adopting credit, leads to a bad experience. In the analysis that follows, we investigate whether the decision to adopt new SKUs generates a larger treatment effect on preferences than simply adopting credit to buy familiar SKUs. To study this, we extend our model to endogenize SKU adoption on the first credit transaction in addition to modeling the credit adoption. Our goal is to differentiate the selection into adopting a new SKU from the causal impact of SKU adoption on preferences.

In econometric language, the introduction of new SKUs is an additional kind of non-random selection (see Angrist and Imbens, 1995), beyond the decision to adopt credit. Thus, our analytical structure must now account for three types of compliance: not adopting credit, adopting credit without new SKU adoption on the first credit transaction, and adopting credit with new SKU adoption on the first credit transaction. We adopt a structural method to estimate the effects of credit and SKU adoption on credit, modifying a model to endogenize the adoption of SKUs.

As our study featured a binary treatment, discerning three-way compliance requires extra instruments (exclusion restrictions) – variables that affect the uptake of new SKUs but do not correlate with the risk aversion of individuals prior to the study. By examining our SKU-specific data and initial survey, we identified two promising variables. First, our dataset contains details on a fleet of vehicles distributing Unilever products. Out of the 7 vans, 3 have markedly better records in promoting new SKUs than the other 4. Second, stores which report a market share of Unilever products exceeding 30% are more likely to adopt new SKUs on credit. This suggests that stores dominated by Unilever sales exhibit a greater propensity to experiment with new SKUs when credit is available.

Regrettably, our dataset becomes sparse when we segment it by van ID and Unilever market shares within the sample of credit adopters. Therefore, we decided to formulate an interaction instrument. We define a "low type" as a store that is served by a less successful van ID *and* has a smaller share of Unilever products. "Low types" have a 40% probability of adopting new SKUs on the initial credit swipe, compared to a probability exceeding 60% for the remaining "high types." This difference is statistically significant at the 5% level, which resembles a first-stage test for weak instruments. We have attempted other combinations to define a "low type" by adjusting the market share thresholds and incorporating more marginal van IDs. However, other combinations result in an imbalanced distribution of types and a reduction in statistical power.

Fortuitously, given that we observe the Unilever market share and van ID for all subjects, we can correlate the low-type dummy with the risk aversion observed in the control group. We regress the dummy variable for accepting a risk-neutral gamble, as well as the acceptance threshold, on the 'low-type' dummy. We find p-values of 0.9 and 0.34, respectively, which suggests no significant correlation between the low-type dummy and observed risk aversion in the control group. This serves as a test for the necessary exclusion restriction.

To examine the impact of new SKUs we augment the credit adoption equation with a

SKU adoption equation:

$$\gamma_i = \bar{\gamma} + \Delta^1 D_i + \Delta^2 \text{SKU}_i + \epsilon_i \tag{6}$$

$$D_i = \begin{cases} 1 & \text{if } D^0 + \beta^1 \text{LOW}_i + \nu_i^1 > 0 \text{ and } Z_i = 1 \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

$$\text{SKU}_i = \begin{cases} 1 & \text{if } \text{SKU}^0 + \beta^2 \text{LOW}_i + \nu_i^2 > 0 \text{ and } D_i = 1 \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

where LOW_i denotes a dummy variable for the low type. This dummy could impact both credit adoption and SKU adoption. However, it is omitted from the risk aversion equation, thereby enabling identification. The goal is to estimate two treatment effects Δ^1 and Δ^2 that indicate the degree of risky activity.

To control for selection into risky activity we further posit the following distribution for the unobservable factors:

$$\begin{bmatrix} \sigma_\epsilon^2 & \rho_1 \sigma_{\nu^1} & \rho_2 \sigma_{\nu^2} \\ \rho_1 \sigma_{\nu^1} & 1 & 0 \\ \rho_2 \sigma_{\nu^2} & 0 & 1 \end{bmatrix}$$

Importantly, we permit selection based on risk aversion for taking the risky actions, such as adopting credit and adopting SKUs on credit. Nevertheless, due to sparsity of the data, we rule out direct correlation between unobservable factors that might jointly drive credit adoption and SKU adoption. Admittedly, this does present a constraint in our analysis. However, it should be noted that we allow for correlation between these actions via risk aversion. Specifically, individuals who exhibit greater risk aversion might be less likely both to adopt credit and to incorporate SKUs through credit.

Table 4 displays the results of our estimation. The average values, dispersion, and correlations are nearly identical to those observed in the baseline CARA model. Notably, the new parameter introduced, which represents the correlation between the SKU adoption shock and risk aversion, is determined to be -0.08. This value signifies a negative relationship between risk aversion and the adoption of new SKUs.[25]

[25] It is worth noting that the correlation between risk aversion and SKU adoption is smaller than the correlation with credit adoption. This observation may appear counter-intuitive; however, it can be explained

		CARA
RA intercept, $\bar{\gamma}_{\text{FEMALE}}$		1.74^{***}
		(0.18)
RA intercept, $\bar{\gamma}_{\text{MALE}}$		1.73^{***}
		(0.16)
RA Dispersion, σ_γ		1.26^{***}
		(0.15)
RA-adoption correlation, $\rho_!$		-0.52^{***}
		(0.03)
RA-newSKU correlation, ρ_2		-0.08^{**}
		(0.03)
Impact of adoption, Δ^1_{FEMALE}		0.07^{***}
		(0.02)
Impact of adoption, Δ^1_{MALE}		0.59^{**}
		(0.26)
Impact of new sku, Δ^2		0.31^{*}
		(0.16)
Risk premium for the fair gamble	Control group	51.4
		(2.04)
	ITT	4.80^{***}
		(1.47)
	SEL out of credit	7.07^{***}
		(2.05)
	SEL into credit, out of new SKU	-16.55^{***}
		(5.12)
	SEL into credit and new SKU	-28.44^{***}
		(3.72)
	ATT adoption and no new SKU	16.27^{***}
		(5.76)
	ATT adoption and new SKU	23.02^{***}
		(6.27)
Credit adoption	Baseline	27%
	Counterfactual preferences: credit, no new SKU	21%
	Counterfactual preferences: credit, new SKU	17%
SKU adoption	Baseline	16%
	Counterfactual preferences: credit, no new SKU	13%
	Counterfactual preferences: credit, new SKU	10%
N		582

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Structural Estimates with Credit and SKU Adoption

by the fact that SKUs can only be adopted on credit when credit itself is available. Consequently, SKU adoption necessitates both a substantial credit shock and a substantial SKU shock. Therefore, on average, individuals who adopt SKUs are likely to be less risk-averse than those who adopt credit alone. We confirm this later by computing the relevant selection measures.

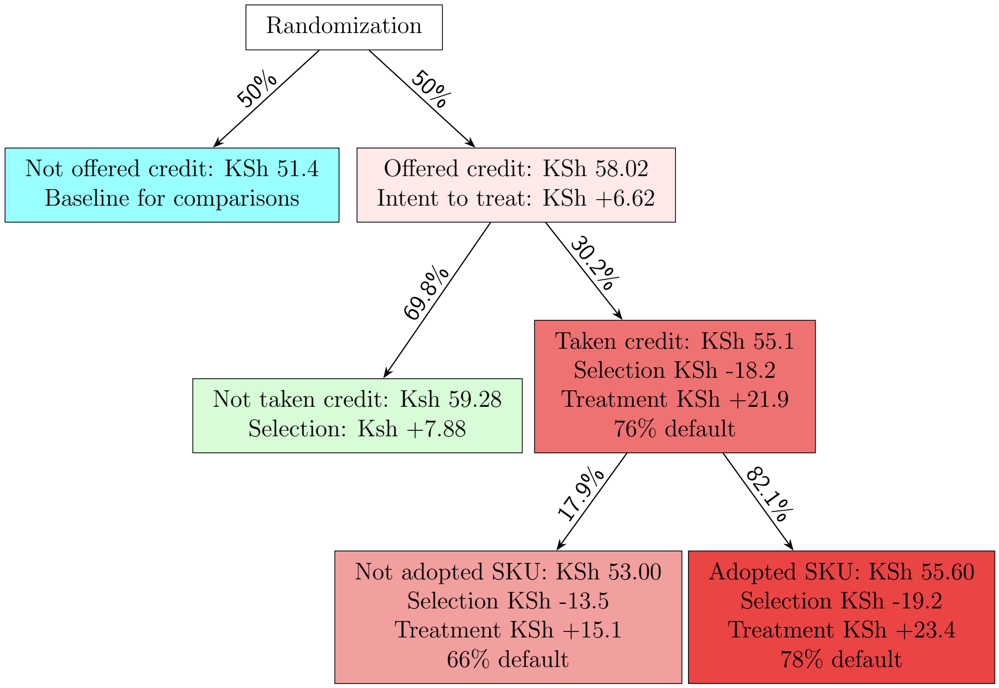

Figure 6: Decision tree depicting the average risk premium for each group and selection (relative to the control) and treatment effects for the high-type males.

The treatment effect of credit, denoted as Δ^1, appears to be smaller than in the baseline model. However, its interpretation is that this effect applies to those who did not adopt new SKUs. Essentially, credit adoption by itself modifies risk preferences, albeit to a lesser extent than the use of that credit for riskier endeavors. This additional modification is captured by the parameter Δ^2, which is positive. We find that stores self-selecting to adopt new SKU are less risk averse than stores that adopted credit but did not adopt a new SKU. We also find that the treatment effect for the stores that adopted a new SKU is over 40% larger than for stores that did not adopt new SKU but adopted credit.

The process can be depicted using the decision tree. Figure 6 contains such tree for the high-type males. In the initial section of the tree, we show a division into the control group and the treatment group. In the control group, we have an average risk premium of 51.4 Shillings, which serves as our baseline without any credit offer. The treatment group shows an average intent-to-treat value of 6.62 Shillings, suggesting a approximately 13% increase in the risk premium over the baseline.

Moving forward, the treatment group is further divided into credit adopters and non-adopters. Among the high-type males, 30.2% adopted credit, and 69.8% did not adopt credit before our survey. As result of the selection out of credit, credit non-adopters exhibited, ex-ante, a 7.88 Shillings larger risk premium than the baseline. Conversely, credit adopters had 18.2 Shillings smaller ex-ante risk premiums. Replicating our earlier results, we show that the treatment effect amounts to 21.9 Shillings and is larger than the selection into credit.

The model decomposes the above effects depending on the SKU adoption decision. For instance, new SKU non-adopters are 13.5 Shillings, while SKU adopters are 19.2 Shillings less ex-ante risk averse, constituting respectively 26% and 37% of the baseline risk premium. The gap shows that adoption of the new SKUs is undertaken by significantly less risk-averse individuals. Additionally, our analysis reveals that the impact of credit on preferences depends on the nature of its usage. The shift towards risk aversion is larger for those users who expand to purchasing new SKUs. This phenomenon can be partly attributed to the fact that new SKU adopters are inherently less risk-averse, facilitating more pronounced shifts in their preferences post-treatment. However, given the association between SKU adoption and higher default rates, one might anticipate such adoption to culminate in more adverse experiences, thus increasing risk aversion levels.

We detect considerable incremental impact of the SKU-adoption experience on preferences, over and above the experience of adopting credit (the acceptance rate for credit drops from 27% to 17%, and SKU adoption from 16% to 10%). This suggests that a poor first time roll out of new products may lead to significant frictions in the propensity to adopt products in the future. Importantly, this analysis also isolates the impact of changing risk preferences as distinct from other channels that do not rely on preference changes. It demonstrates that considering changing risk preferences based on experience is an important channel when examining the dynamics of entrepreneurial decisions.

6 Conclusion

For millions of small entrepreneurs around the world the experiences from past decisions are an important channel that govern their risk taking and the adoption of future innovations. This paper examines whether past experiences, in particular, the experience of failures arising from a business decision to adopt a new credit technology, can increase risk aversion and influence future risk taking. Using a randomized controlled trial that deployed a new credit technology under Mastercard's Jaza Duka program to small retailers, we show that those who adopted the new credit line and then experienced failure and default become significantly more risk averse.

Our analysis is able to separate out the causal effect of credit adoption from the selection effects. The selection into credit is more pronounced than the selection out of credit – i.e., those who adopt credit have substantially lower ex-ante risk aversion. But the post-adoption treatment effect which increases risk aversion is even more pronounced and swamps the effect of the selection into credit. Thus, the effect of experienced credit failures has the potential to homogenize the risk preferences in the population. Taken together these results have material implications for entrepreneurial risk taking and innovation. The more risk-loving entrepreneurs in the population are the ones driving the adoption of the new credit technology. But it is precisely these entrepreneurs who may end up becoming overly risk averse thereby foregoing valuable credit opportunities and dampening future entrepreneurial performance.

We highlight key demographic factors that moderate both selection and treatment effects, bearing significant implications for decision-making and policy formulation. Prior to our intervention, males exhibit markedly lower levels of risk aversion compared to females. However, the impact of adopting credit on their risk aversion is nearly double that observed in females. In a similar vein, younger entrepreneurs and those managing smaller enterprises display a greater propensity for risk and are more inclined to adopt credit. Yet, following a setback post-adoption, their preferences shift dramatically, leading to a significant increase in risk aversion. From a policy standpoint, it emerges that younger male entrepreneurs, especially those running smaller enterprises, are typically the first to adopt new technologies. However, it is this demographic that may be more vulnerable to failures and the consequent impacts on risk aversion and their future entrepreneurial decisions. Depending on the prominence of these groups, our insights could play an important role in shaping policies aimed at optimally balancing risk-taking and innovation.

References

Angrist, J. D., and G. W. Imbens (1995): "Two-stage least squares estimation of average causal effects in models with variable treatment intensity," *Journal of the American statistical Association*, 90(430), 431–442.

Angrist, J. D., G. W. Imbens, and D. B. Rubin (1996): "Identification of causal effects using instrumental variables," *Journal of the American statistical Association*, 91(434), 444–455.

Asriyan, V., W. Fuchs, and B. Green (2019): "Liquidity sentiments," *American Economic Review*, 109(11), 3813–3848.

Battocchi, K., E. Dillon, M. Hei, G. Lewis, P. Oka, M. Oprescu, and V. Syrgkanis (2019): "EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation," https://github.com/py-why/EconML, Version 0.x.

Brown, R., V. Montalva, D. Thomas, and A. Velásquez (2019): "Impact of violent crime on risk aversion: Evidence from the Mexican drug war," *Review of Economics and Statistics*, 101(5), 892–904.

Burke, S. N., and C. A. Barnes (2006): "Neural plasticity in the ageing brain," *Nature reviews neuroscience*, 7(1), 30–40.

Callen, M., M. Isaqzadeh, J. D. Long, and C. Sprenger (2014): "Violence and risk preference: Experimental evidence from Afghanistan," *American Economic Review*, 104(1), 123–148.

Cameron, L., and M. Shah (2015): "Risk-taking behavior in the wake of natural disasters," *Journal of human resources*, 50(2), 484–515.

Christopoulos, G. I., P. N. Tobler, P. Bossaerts, R. J. Dolan, and W. Schultz (2009): "Neural correlates of value, risk, and risk aversion contributing to decision making under risk," *Journal of Neuroscience*, 29(40), 12574–12583.

Cochrane, J. H. (2011): "Presidential address: Discount rates," *The Journal of finance*, 66(4), 1047–1108.

Cohn, A., J. Engelmann, E. Fehr, and M. A. Maréchal (2015): "Evidence for countercyclical risk aversion: An experiment with financial professionals," *American Economic Review*, 105(2), 860–885.

de Janvry, A., E. Sadoulet, M. Dar, and K. Emerick (2016): "The agricultural technology adoption puzzle: what can we learn from field experiments?," Discussion paper, FERDI working paper.

Duflo, E., R. Glennerster, and M. Kremer (2007): "Using randomization in development economics research: A toolkit," *Handbook of development economics*, 4, 3895–3962.

Elminejad, A., T. Havranek, and Z. Irsova (2022): "Relative risk aversion: a meta-analysis," .

Epstein, L. G., and S. E. Zin (1991): "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis," *Journal of political Economy*, 99(2), 263–286.

Guiso, L., P. Sapienza, and L. Zingales (2018): "Time varying risk aversion," *Journal of Financial Economics*, 128(3), 403–421.

Hanaoka, C., H. Shigeoka, and Y. Watanabe (2018): "Do risk preferences change? evidence from the great east japan earthquake," *American Economic Journal: Applied Economics*, 10(2), 298–330.

Heckman, J. J., and R. Robb (1986): "Alternative methods for solving the problem of selection bias in evaluating the impact of treatments on outcomes," *Drawing inferences from self-selected samples*, pp. 63–107.

Holt, C. A., and S. K. Laury (2002): "Risk aversion and incentive effects," *American economic review*, 92(5), 1644–1655.

Jakiela, P., and O. Ozier (2019): "The impact of violence on individual risk preferences: Evidence from a natural experiment," *Review of Economics and Statistics*, 101(3), 547–559.

Knutson, B., G. E. Wimmer, C. M. Kuhnen, and P. Winkielman (2008): "Nucleus accumbens activation mediates the influence of reward cues on financial risk taking," *NeuroReport*, 19(5), 509–513.

Kreps, D. M., and E. L. Porteus (1978): "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica: journal of the Econometric Society*, pp. 185–200.

Kuchler, T., M. Piazzesi, and J. Stroebel (2023): "Housing market expectations," in *Handbook of Economic Expectations*, pp. 163–191. Elsevier.

Malmendier, U., and S. Nagel (2011): "Depression babies: Do macroeconomic experiences affect risk taking?," *The quarterly journal of economics*, 126(1), 373–416.

——— (2016): "Learning from inflation experiences," *The Quarterly Journal of Economics*, 131(1), 53–87.

Malmendier, U., and A. S. Wellsjo (2023): "Rent or buy? inflation experiences and homeownership within and across countries," *Journal of Finance, forthcoming.*

Palacios-Huerta, I., and R. Serrano (2006): "Rejecting small gambles under expected utility," *Economics Letters*, 91(2), 250–259.

Paravisini, D., V. Rappoport, and E. Ravina (2017): "Risk aversion and wealth: Evidence from person-to-person lending portfolios," *Management science*, 63(2), 279–297.

Paulsen, D. J., R. M. Carter, M. L. Platt, S. A. Huettel, and E. M. Brannon (2012): "Neurocognitive development of risk aversion from early childhood to adulthood," *Frontiers in human neuroscience*, 5, 178.

Rabin, M., and R. H. Thaler (2001): "Anomalies: risk aversion," *Journal of Economic perspectives*, 15(1), 219–232.

Sahm, C. R. (2012): "How much does risk tolerance change?," *The quarterly journal of finance*, 2(04), 1250020.

Sapienza, P., L. Zingales, and D. Maestripieri (2009): "Gender differences in financial risk aversion and career choices are affected by testosterone," *Proceedings of the National Academy of Sciences*, 106(36), 15268–15273.

Shum, M., and Y. Xin (2022): "Time-Varying Risk Aversion? Evidence from Near-Miss Accidents," *The Review of Economics and Statistics*, 104(6), 1317–1328.

Thaler, R. H., and E. J. Johnson (1990): "Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice," *Management science*, 36(6), 643–660.

Online Appendix

A Attrition analysis

Our initial study cohort comprised 999 stores deemed eligible for credit as of November 2018, coinciding with the commencement of data acquisition from Mastercard. Subsequent to this period, spanning from November 2018 to April 2019, 69 stores ceased operations or were rendered ineligible for credit. Creditworthiness adjudication fell under the purview of KCB Bank Kenya, the entity responsible for loan underwriting, predominantly relying on an analysis of purchasing patterns pertaining to Unilever products. A store's forfeiture of eligibility stemmed from a decline in Unilever purchase volume below a reference threshold established by the bank. No baseline or post-treatment surveys were conducted for these 69 stores. Nonetheless, comprehensive pre-treatment Unilever ledger data encompassing parameters such as annual Unilever turnover, minimum quality threshold (representing a measure of anticipated sales performance at the store level), year-over-year alterations in annual turnover, and a proprietary store size metric categorized as a discrete variable (0, 1, 2) were available for all 999 stores. Detailed comparison between the 69 discontinued stores and the remaining 930 establishments is provided in Table 5. Our analysis reveals a consistent pattern indicating that the discontinued stores exhibit markedly lower metrics across all aforementioned parameters. Such observations align with expectations, given that creditworthiness determinants were intrinsically tied to store characteristics. In light of these findings, it is pertinent to conceptualize our study population as comprising of ongoing establishments capable of sustaining credit eligibility over a foreseeable temporal horizon. Out of the 930 stores initially included in our analysis, 29 stores lacked both baseline and post-treatment survey data. This resulted either due to their non-consent for participation in the study, or the inaccessibility of the closed stores to enumerators during the data collection process. Additionally, 121 stores completed the baseline survey but failed to provide data for the post-treatment assessment. In aggregate, a total of 150 stores did not furnish post-treatment survey responses despite their eligibility for credit at the onset of the program. Table 6 presents an attrition analysis comparing the 150 stores to the remaining 780 establishments. Initially, we scrutinize the

	(1) Closed or ineligible	(2) Open and eligible	(3) (1) vs. (2)	(4) p-value
Annual turnover	76,501.42	123,231.10	-46,729.67	0.02
	(18,110.20)	(5,262.00)	(20,067.67)	
Minimum quantity threshold	10,452.15	27,831.34	-17,379.19	0.00
	(2,894.33)	(1,154.36)	(4,341.22)	
Year-over-year annual turnover	12,271.63	40,053.26	-27,781.63	0.00
	(10,711.62)	(2,308.58)	(9,013.31)	
Store size	0.12	0.52	-0.41	0.00
	(0.05)	(0.02)	(0.09)	
N	69	930	999	

Standard errors in parentheses.

Table 5: Characteristics of stores that closed or became credit ineligible during the study.

balance of treatment arms across both samples, observing no discernible disparities in arm distribution. Subsequently, leveraging pre-study Unilever ledger data, we evaluate the size characteristics of the 150 dropout stores in comparison to the 780 continuing stores, detecting no disparities across various size metrics. Further examination entails an analysis utilizing demographic data collected during the baseline survey. It is imperative to note the exclusion of 139 stores that completed the post-treatment survey but failed to furnish baseline data, thereby precluding their inclusion in the ensuing t-tests. Despite thorough scrutiny, we failed to discern any dimensions correlating with attrition. Specifically, no disparities were discerned concerning gender, age, education level, revenue, profits, or prior experience with loan products among the participating stores. For the 780 stores that completed the post-treatment survey, we proceeded to evaluate the extent of store ownership. Our hypothesis is that the impact on risk aversion would manifest among those who are store owners, because it is they who have direct decision making authority and responsibility for credit decisions and repayment. Conversely, we expect this effect to be negligible among employees who do not have the responsibility for credit, and who may lack awareness regarding credit. Within the cohort, 48 stores were operated by employees, 100 stores identified themselves as co-owners, and 25 responses to the ownership question were missing. We direct our focus towards on the remaining 607 stores self-identified as sole owners who unlike the co-owners would have the exclusive responsibility for any credit decisions. Table 7 details balance assessments between owner and non-owner demographics. Notably, no discernible differences were observed in store size, educational attainment, or experience with financial products.

	(1) No survey	(2) Completed survey	(3) (1) vs. (2)	(4) p-value
Credit Arm	0.52	0.49	0.03	0.53
	(0.04)	(0.02)	(0.04)	
Annual turnover	132881.17	121375.31	11505.86	0.42
	(16676.34)	(5395.69)	(14309.47)	
Minimum quantity threshold	27453.40	27904.02	-450.62	0.89
	(2931.19)	(1256.45)	(3140.22)	
Year-over-year annual turnover	35537.24	40921.72	-5384.48	0.39
	(5032.77)	(2576.89)	(6277.64)	
Store size	0.54	0.52	0.02	0.75
	(0.06)	(0.02)	(0.06)	
Baseline survey: Male	0.67	0.67	-0.00	0.98
	(0.04)	(0.02)	(0.05)	
Baseline survey: Age	35.68	38.32	-2.64	0.18
	(1.83)	(0.40)	(1.99)	
Baseline survey: Class 1-5	0.05	0.06	-0.01	0.70
	(0.02)	(0.01)	(0.02)	
Baseline survey: Class 6-12	0.66	0.62	0.04	0.35
	(0.04)	(0.02)	(0.04)	
Baseline survey: College	0.06	0.09	-0.03	0.23
	(0.02)	(0.01)	(0.02)	
Baseline survey: Monthly revenue	10118.58	11320.97	-1202.39	0.40
	(1342.60)	(563.10)	(1421.22)	
Baseline survey: Monthly profits	1878.68	2234.39	-355.71	0.37
	(237.69)	(167.63)	(399.52)	
Baseline survey: Did you take a loan?	0.14	0.16	-0.02	0.57
	(0.03)	(0.01)	(0.03)	
Baseline survey: Owner	0.57	0.61	-0.04	0.38
	(0.04)	(0.02)	(0.04)	
Baseline survey: Co-owner	0.12	0.11	0.01	0.69
	(0.03)	(0.01)	(0.03)	
Baseline survey: Employee	0.12	0.11	0.01	0.66
	(0.03)	(0.01)	(0.03)	
N	150	780	930	

Standard errors in parentheses.

Table 6: Characteristics of stores that did not complete the post-treatment survey. Demographics are furnished in the baseline survey; 139 stores that completed the post-treatment survey but failed to furnish baseline data were excluded.

However, it is worth noting that owners tended to be slightly older compared to non-owners. There were also a higher proportion of owners who were males as compared to non-owners. It is relevant to acknowledge that while some women may be identified as co-owners, their involvement and awareness of credit-related matters might be limited, potentially leading to a muted or absent effect among this demographic. As we proceed to present our primary findings for stores that identified as owners, we revisit our analysis concerning the co-owner and employee sub-populations in order to empirically test the aforementioned hypothesis. Finally, we provide table of descriptive statistics for the entire sample of owners, including

	(1) Non-owner	(2) Owner	(3) (1) vs. (2)	(4) p-value
Credit Arm	0.46	0.50	-0.04	0.37
	(0.04)	(0.02)	(0.04)	
Annual turnover	124,709.26	120,425.11	4,284.15	0.74
	(12,149.85)	(6,011.91)	(12,994.91)	
Minimum quantity threshold	29,589.04	27,423.77	2,165.26	0.47
	(3,125.51)	(1,347.50)	(3,025.22)	
Year-over-year annual turnover	46,733.81	39,265.23	7,468.58	0.23
	(6,304.28)	(2,780.90)	(6,200.79)	
Store size	0.47	0.53	-0.06	0.31
	(0.05)	(0.03)	(0.06)	
Baseline survey: Male	0.60	0.69	-0.09	0.05
	(0.04)	(0.02)	(0.04)	
Baseline survey: Age	35.73	38.90	-3.17	0.00
	(0.88)	(0.44)	(1.02)	
Baseline survey: Class 1-5	0.07	0.06	0.01	0.63
	(0.02)	(0.01)	(0.02)	
Baseline survey: Class 6-12	0.57	0.63	-0.06	0.15
	(0.04)	(0.02)	(0.04)	
Baseline survey: College	0.09	0.09	0.00	0.89
	(0.02)	(0.01)	(0.02)	
Baseline survey: Monthly revenue	10,097.79	11,665.91	-1,568.12	0.25
	(908.97)	(674.46)	(1,359.04)	
Baseline survey: Monthly profits	1942.89	2316.60	-373.71	0.36
	(204.62)	(206.95)	(404.72)	
Baseline survey: Did you take a loan?	0.19	0.16	0.03	0.37
	(0.03)	(0.02)	(0.03)	
N	173	607	780	

Standard errors in parentheses.

Table 7: Characteristics of owners vs. non-owners.

the owners that pass or fail the rationality test. Please refer to the Table 8 for details. This Table could be compared directly with Table 1; however, because only 25 subjects failed comprehension test statistical comparisons are under-powered. A simple comparison reveals nearly identical credit, SKU adoption and default rates across rational and irrational subjects. Also, we notice a near uniform distribution of irrational subjects amongst education levels.

	Average	Count	Standard deviation	Total sample
Male	0.65	397		607
Age	38.90		9.17	439
No education, can not read	0.01	9		607
No education, can read	0.04	22		607
Class 1-5	0.08	49		607
Class 6-12	0.76	463		607
Vocational Training	0	2		607
College	0.10	62		607
Small Unilever Segment	0.59	357		607
Medium Unilever Segment	0.29	176		607
Large Unilever Segment	0.12	74		607
Offered credit	0.50	304		607
Pre-treatment revenue	11,666		15,081	500
Used credit, if offered	0.38	117		304
Used credit, Male	0.41	81		198
Used credit, Female	0.34	36		106
Used credit before the survey, if offered	0.26	79		304
Used credit before the survey, Male	0.29	58		198
Used credit before the survey, Female	0.20	21		106
Eventually restricted, if used credit	0.69	81		117
Eventually restricted, Male	0.74	60		81
Eventually restricted, Female	0.58	21		36
Hard default, if used credit	0.12	14		117
Hard default, Male	0.14	9		81
Hard default, Female	0.11	5		36
New SKU on first credit purchase, if used credit	0.56	65		117
New SKU on first credit purchase, Male	0.58	47		81
New SKU on first credit purchase, Female	0.50	18		36

Table 8: Descriptive statistics, sample containing rational and irrational subjects.

B Randomization checks

	(1)
	Credit Arm
Male	0.0201
	(0.0462)
Age	-0.000877
	(0.00111)
children	-0.0000657
	(0.000174)
Class 1-5	-0.0755
	(0.116)
Class 6-12	-0.0332
	(0.0928)
College	-0.0647
	(0.112)
Store size	0.0155
	(0.0330)
Have you NOT taken a loan in the last 12 months	0.0627
	(0.0640)
Pre-treatment revenue	-0.00000212
	(0.00000171)
Pre-treatment revenue missing	0.0298
	(0.102)
Do you have cash savings?	-0.0657
	(0.0916)
Save at the bank	-0.00452
	(0.0527)
Do you have a written business plan?	-0.0429
	(0.0513)
Do you keep financial records?	-0.0121
	(0.0509)
Would you consider mobile money for loan payments	-0.00126
	(0.0498)
Gives customer credit	0.0510
	(0.0492)
Credit arm	582

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Randomization checks

C Correlates of risk aversion in the control group

	(1)
	Risk aversion
Age	0.00494*
	(0.00256)
Age missing	0.219*
	(0.113)
Female	0.109**
	(0.0432)
Muslim	0.113**
	(0.0497)
Class 6-12	0.133**
	(0.0607)
College	0.118
	(0.0853)
I see myself as someone who is relaxed and handles stress well	0.0575**
	(0.0230)
I see myself as someone who tends to find fault with others	-0.0362**
	(0.0170)
I see myself as someone is generally trusting	-0.00779
	(0.0154)
I see myself as someone is reserved	0.00534
	(0.0134)
I see myself as someone who tends to be lazy	-0.0124
	(0.0227)
I see myself as someone who has few artistic interests	0.00433
	(0.0154)
I see myself as someone who is outgoing and sociable	-0.0110
	(0.0164)
I see myself as someone who does a thorough job	-0.00288
	(0.0231)
I see myself as someone who gets nervous easily	0.0121
	(0.0155)
I see myself as someone who has an active imagination	-0.0132
	(0.0200)
N	253

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Correlates of risk aversion in the control group. We employ the data from and baseline, endline (psychometrics), and telephone surveys data or the control group, and we eliminate the absent observations. This culling procedure yields 253 observations, while the total count of the control group stands at 303.

	(1) Risk aversion	(2) Risk aversion
Have you taken a loan in the last 12 months	-0.101*	-0.105*
	(0.0574)	(0.0594)
Do you have cash savings?		0.124*
		(0.0666)
Save at the bank		-0.111**
		(0.0500)
Would you consider mobile money for loan payments		-0.0560
		(0.0459)
Customer credit		-0.0430
		(0.0461)
Revenue		0.00000188
		(0.00000146)
N	277	277

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 11: Correlates of risk aversion in the control group. We employ the data from and baseline and telephone surveys data or the control group, and we eliminate the absent observations. This culling procedure yields 277 observations, while the total count of the control group stands at 303.

D Potential outcomes

In this appendix, we provide additional discussion on treatment effects and selection using a notation compatible with the original potential outcomes framework by Angrist, Imbens, and Rubin (1996), henceforth AIR.

Denote the credit adoption by D_i, where $D_i = 1$, if the individual adopts credit, and 0 otherwise. Further, denote assignment to the treatment and control arms as Z_i, where 1 and 0 are treatment (credit offer) and control (no credit offer), respectively.

The potential risk-aversion outcomes for individual i are given by $Y_i(Z_i, D_i)$. In a general setting both credit adoption and arm assignment could affect risk-aversion. We postulate an exclusion restriction that simply receiving the credit offer does not impact risk aversion for non-adopters. Formally, $Y_i(1,0) - Y_i(0,0) = 0$. To close the argument also assume that arm assignment would not have mattered for adopters, if they could use the credit in the control group, $Y_i(1,1) - Y_i(0,1) = 0$; although, $Y_i(0,1)$ does not occur in our sample because one cannot obtain credit without the credit offer. Following the convention, we simplify the notation such that $Y_i(D_i) = Y_i(Z_i, D_i)$.

We define the causal effect as in the AIR framework

$$Y_i(1, D_i(1)) - Y_i(0, D_i(0)) = Y_i(1) - Y_i(0).$$

Subjects could not obtain credit without the credit offer, i.e., $D_i(0) = 0$. For this reason, all individuals that adopted credit (treated subjects) are *compliers*, and there are no *defiers*, so the average treatment effect (LATE) and average treatment effect on the treated (ATT) are the same. Henceforth, we use ATT as our measure of causal effects, i.e.,

$$\text{ATT} = E[Y_i(1) - Y_i(0) | D_i(1) = 1].$$

After some manipulations, we obtain a version of Equation (9) in the AIR notation:

$$
\begin{aligned}
&\underbrace{E[Y_i(1)|D_i(1)=1]-E[Y_i(0)|D_i(1)=0]}_{\text{observed difference between adopters and non-adopters}} = \\
&\underbrace{E[Y_i(1)-Y_i(0)|D_i(1)=1]}_{\text{ATT}} + \underbrace{E[Y_i(0)|D_i(1)=1]-E[Y_i(0)|D_i(1)=0]}_{\text{ex-ante difference between adopters and non-adopters}} = \\
&\underbrace{E[Y_i(1)-Y_i(0)|D_i(1)=1]}_{\text{ATT}} + \underbrace{E[Y_i(0)|D_i(1)=1]-E[Y_i(0)]}_{\text{selection into credit}} - \underbrace{(E[Y_i(0)|D_i(1)=0]-E[Y_i(0)])}_{\text{selection out of credit}}
\end{aligned}
\tag{9}
$$

E Additional results

E.1 Treatment effect and ownership

	(1) Risk aversion	(2) Risk aversion	(3) Risk aversion	(4) Risk aversion	(5) Risk aversion
Credit Arm	0.0663***	0.0663***	-0.0735	-0.0429	0.0465**
	(0.0256)	(0.0255)	(0.0629)	(0.0987)	(0.0237)
Credit Arm × Co-owner	-0.140**				
	(0.0682)				
Credit Arm × Employee	-0.109				
	(0.0959)				
Co-owner	0.0697				
	(0.0470)				
Employee	0.0381				
	(0.0623)				
Sample	ALL	OWNERS	CO-OWNERS	EMPLOYEES	OWNERS & CO-OWNERS
N	755	607	100	48	707

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 12: Effect of credit offer on risk aversion for co-owners and employees including rational and irrational subjects.

E.2 Effect of credit on income expectations

	(1) Post-treatment income expectations	(2) Post-treatment income expectations	(3) Difference in income expectations
Offered credit	-12970.2	-22795.6	-27432.2
	(12329.4)	(17215.7)	(17433.7)
Pre-treatment income expectations		0.566^{***}	
		(0.135)	
Post-treatment income expectations	414	287	287

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 13: Impact of credit offer on the expectations of future income. Expectations are measured using the question, "After 12 months from now, what do you think will be your " daily revenue?" which was collected before and after treatment. Model (1) uses a cross-section of post-treatment expectations. Model (2) uses panel data. Model (3) uses the difference in expectations as a dependent variable.

E.3 Effect of credit on time preferences

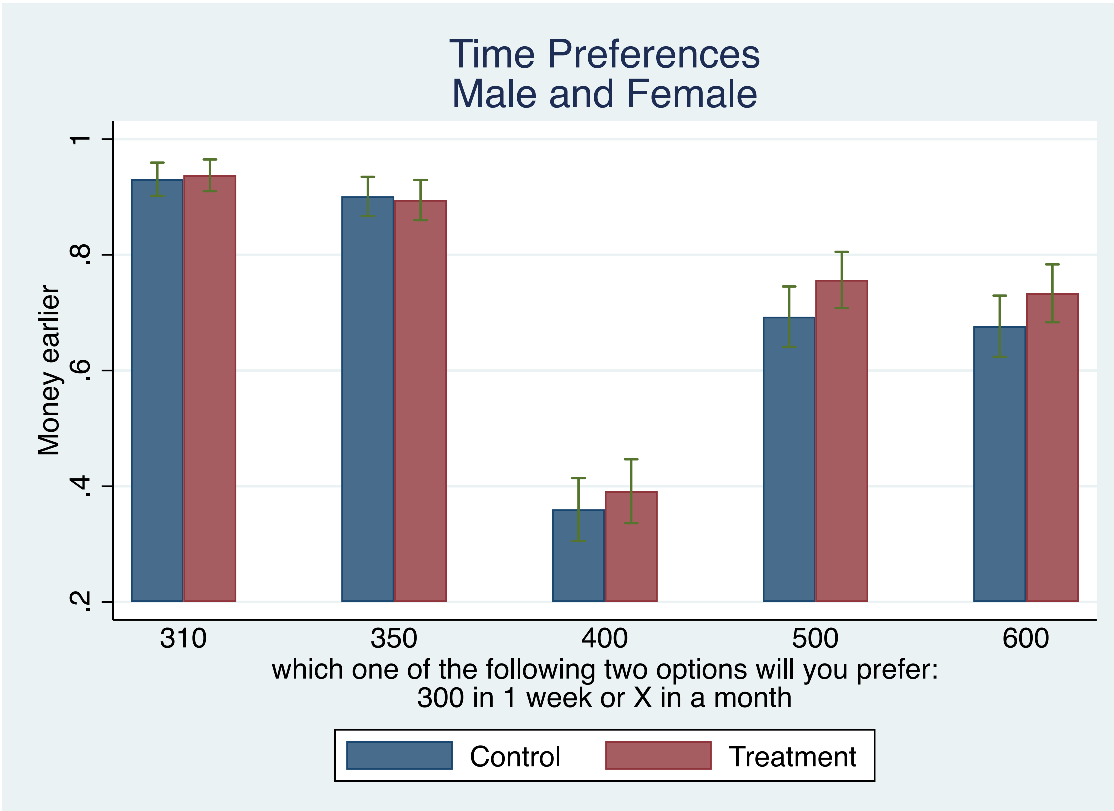

Figure 7: Measure of time preferences. None of the differences are significant at 5% level.

E.4 Heterogeneity in ATT by gender

	(1) Risk aversion	(2) Risk aversion	(3) Risk aversion	(4) Risk aversion	(5) Risk aversion	(6) Risk aversion	(7) Risk aversion
Credit	0.0292	0.105***	0.0292				
	(0.0365)	(0.0271)	(0.0348)				
Male	-0.0517						
	(0.0318)						
Male/Credit	0.0754*						
	(0.0451)						
Not adopted credit				0.101***		0.0542	
				(0.0300)		(0.0366)	
Adopted credit				0.113***	0.349***	-0.0708	0.146
				(0.0400)	(0.0976)	(0.0599)	(0.180)
Sample	ALL	MALE	FEMALE	MALE	MALE	FEMALE	FEMALE
IV/ATT	no	no	no	no	yes	no	yes
N	582	381	201	381	381	201	201

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 14: Analysis of the difference in treatment effects of male vs female.

To zoom further into the gender differences, we compare selection and ATT estimates across males and females. The analysis utilizes the rational sample, aligning it with previous regressions. In Column (1) of Table 14, we show ITT estimates that include the interaction between a treatment dummy and gender. Notably, we find that the difference between the ITT for males and females is statistically significant at the 10% level (the lower statistical significance is likely due to noise in the female sample). Columns (2) and (3) present the ITT estimated separately on the male and female samples. The point estimate indicates that approximately 10% of the male population switches from risk-loving to risk-averse. This is in contrast to the 7.85% shift observed within the combined-gender sample. For females, the point estimate is near zero, suggesting negligible change in their risk preference in response to the treatment. Albeit, that estimate is considerably noisy, likely due to a smaller sample size and lower credit uptake.

When estimating the selection and average treatment effects on the treated (ATT) for male and female sub-samples separately, the data reveals distinct patterns. For males, we obtain a treatment effect of 34.9%, selection out of credit of 10.1% and selection into credit of -23.6%. In the case of females, the point estimate for the causal effect is 14.6%, with a selection out of credit of 5.4% and selection into credit of 21.7%. Although the estimates for females are statistically less precise, it appears that the inclination to adopt credit is comparable between genders, while the magnitude of the treatment effects is markedly less for females. This differential suggests that while the decision-making process for adopting credit may be similar across gender, the degree to which their risk preferences are altered post-adoption varies.