

Credit Failures and Entrepreneurial Risk Aversion*

William Fuchs[†] Ganesh Iyer[‡] Przemyslaw Jeziorski[§]

March 7, 2024

Abstract

This paper examines how adverse experiences of credit use arising post the adoption of a new credit technology impacts entrepreneurial risk aversion. In a randomized controlled trial which deployed a new credit technology to small retail entrepreneurs in Kenya we show that the experience of a failed use of a new credit line significantly increases their risk aversion. By separating out the causal effect of credit from selection effects we show that accounting for selection, especially selection into credit, matters – those who adopt credit have substantially lower ex-ante risk aversion. Thus, their post adoption treatment effect which increases risk aversion is even more pronounced. Consequently, the more risk loving entrepreneurs in the population adopt the credit, but they also end up becoming substantially more risk averse upon experiencing a credit failure. A heterogeneous treatment effects analysis identifies important demographic moderators and shows that male entrepreneurs who are younger and who run smaller businesses are more likely adopt the new credit product, but they are also particularly susceptible to credit failure experiences and the resulting increase in their risk aversion. We also show that the increased risk aversion leads to a counterfactual credit adoption of only 19% in comparison to the actual adoption rate of 26%.

*We would like to thank 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

Preferences are typically considered an immutable foundation for economic decision making. It is standard in economic models to assume that agents are endowed with stable risk preferences that are unchanged by past outcomes and experiences. Yet the recent evidence from neuroscience shows that the manner in which we respond to a set of risky choices is mediated by how our brains process the outcomes of these choices (see Christopoulos, Tobler, Bossaerts, Dolan, and Schultz, 2009). Thus the wiring in our brains influences these choices.¹ Importantly, as documented in the literature on brain plasticity, the wiring of our brains is not fixed and is affected by past experiences. This provides the rationale for possible “experience effects” through which our past experiences can affect our current choices.

This paper asks a basic yet surprisingly unexplored question in the area of entrepreneurship and technology adoption: Do past experiences matter and how consequential are the failures from past business decisions for risk taking by entrepreneurs? In the context of a randomized controlled trial which deployed a new credit technology to small retail entrepreneurs we show that this is indeed the case. A failed use of a new credit line made the subjects significantly more risk averse. The magnitude being even larger than the ex-ante differences in risk aversion across gender in the control group.

The credit product, Jaza Duka², was introduced in Kenya in 2019 as a collaboration between Mastercard and Unilever, as first of a kind modern and massively accessible retail credit line. The goal of the credit product was to alleviate financial constraints of small shop owners by providing working capital to grow their sales. The study was conducted in the Malindi region on a sample of 998 retailers, a majority of whom (approximately 80%) did not have any experience with formal credit in the past year. A random sample of 50% of these retailers received an offer of credit in September 2019, the treatment group. The remaining retailers never received an offer, the control group. Of the retailers in the treatment group, approximately 40% accepted the credit line and used it at least once. Unfortunately, 69% of

¹The wiring may depend on age and gender and both 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“Jaza Duka” means “fill up your store” in Swahili.

those that used the credit line, in particular those that used it more intensively and bought new products with it, ended up having difficulty selling those products and failed to repay their loans on time.

Between November and December 2019, we measured risk preferences by asking our subjects to choose between a certain payoff of 100 Shillings or a series of fictitious 50/50 lotteries that payed 0 when they lose or π when they win. Thus, the higher the π needed to accept the lottery the more risk averse are the subjects. The design enables us to measure the causal impact of the introduction of the Jaza Duka credit on the attitudes towards risk. Compared to the control group, we observe that those that received the offer, have a significantly lower probability of being willing to accept the risky hypothetical gamble. We find that 7.8% more subjects reject a fair gamble in the treatment group as compared to the control and this difference is significant at the 1% level. This difference is consequential. It is even greater in magnitude to the differences across gender (5.2%) and the difference between older and younger subjects (5.3% based on the median split in age), in the control group.

Accepting and using the credit line is an endogenous choice. Given the natural assumption that the act of receiving the offer has, in and of itself, no effect over preferences, we can then see that those that rejected the offer are actually more risk averse than the control group. This makes sense since more risk averse people are less likely to try a new financial product. Thus, those that actively participated in the program were initially less risk averse ex-ante than those in the control group. This implies that simply comparing agents that accepted the credit offer to those that rejected it, absent the control group, would underestimate the effect of credit on risk preferences. Indeed in our context, comparing these two populations would produce no effect, despite the fact that unbiased estimates of causal effects are large.

The role of “experience effects” in the economics literature was first proposed by Malmendier and Nagel (2011). Our paper makes three distinct contributions relative to this literature. First, we study experience effects in the context of a randomized controlled trial which allows us to clearly establish causality. Unlike the existing literature that relies on panel data we can 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 approximately of the same magnitude (and opposite 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, provides new evidence about how past failures might stunt efficient entrepreneurial risk taking and innovation in the future.

Lastly, most of the previous work has focused mainly 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 might simply be available but which the individual in question did not directly experience or suffer from. This is well documented in several papers looking at the formation of expectations on future inflation (or choices such as fixed versus variable rate mortgages that would depend on individual agent's beliefs about future inflation). Instead, since we are presenting our retailers with a series of objective gambles, we argue that the effect we document speaks to actual 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 as compared to the control. Related to the literature, we also find that the effects are stronger for the young compared to the old. This could be because with fewer experiences the brain reacts more to a new stimulus. This is consistent with the neuroscience notion that the brain plasticity is decreasing with age (see Burke and Barnes, 2006).

In an important paper Guiso, Sapienza and Zingales (2018) describe a related analysis based on interesting 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, i.e. before and after the financial crises.³ They show that this measure of

³See 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 to identify if the changes are due to beliefs about future prospects from preferences.

risk aversion increases significantly between both waves. They argue that changes in wealth or future expectations are not consistent with the data and thus changes must be attributed to a change in preferences or to the salience of negative outcomes. They too cannot rule out that for some exogenous reason there was change in risk-aversion which in turn helped cause the crisis.⁴ Indeed, the empirical asset pricing literature basically assumes that the causality goes the other way. They argue “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.” Cochrane (2011). As 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.

The analysis uncovers several results which are important for the understanding of entrepreneurial risk taking. The estimates reveal both causal and selection effects which have important implications. 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; i.e., compared to the control group, 23.5% fewer adopters were risk-averse prior to obtaining credit, conversely, among non-adopters, there was a 8.5% higher incidence of risk aversion. The impact of the treatment effect of credit adoption which increases risk aversion is even more pronounced and it overwhelms the effect of selection into credit. In particular, 29.6% of adopters switch from risk-loving to risk-averse as a result of adopting credit. Thus the more risk loving entrepreneurs in the population are the ones that adopt the credit line and are exposed to potential failures. But post the experience of a failure it is precisely these entrepreneurs that end up becoming overly risk averse.

The analysis also reveals interesting demographic effects. Consistent with the previous literature we find that females and older subjects are more risk averse. Males are more likely to adopt credit and, upon adoption, modify their preferences more compared to females. The analysis of heterogeneous treatment effects reveals several demographic moderators of the treatment effect of credit. Female entrepreneurs show a treatment effect on risk aversion

⁴See Asriyan, Fuchs and Green (2019) for a rational model that could provide such type of rationalization based on beliefs rather than preferences.

that is about half the size observed for males. Male retailers with smaller stores experience a relatively larger treatment effect when compared to those with larger stores. This effect is even more pronounced for younger retailers. Thus, the effects of adverse experiences on risk aversion are especially salient for young male entrepreneurs who run smaller businesses. We also isolate the impact on risk aversion after credit adoption from potential wealth effects and show that the influence of wealth changes is negligible.

We identify an interesting channel of new SKU adoption that can influence the impact of negative experience due to credit adoption. 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 the new SKU increases risk aversion above and beyond adopting credit. In particular, adopting credit without the adoption of a new SKU increases the risk premium by 10 Shillings, while adopting both credit and the new SKU increases the risk premium by 19 Shillings. This difference can be explained by the fact that new SKU adoption is correlated with a larger default rate.

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 can 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 26%, underscoring the idea that experiences of past failures have the potential significantly stunt future entrepreneurial risk taking. In addition, we show that the 8% less entrepreneurs (3% less females and 10% less males) adopt new SKUs under the ex-post, more risk-averse preferences.

2 Institutional Details and Market Setting

The study involves manipulating the availability of Jaza Duka, a credit program that is a collaborative effort between Mastercard, Unilever, and the Kenya Central Bank. 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 instead of smaller sachets, 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.

Jaza Duka had a significant impact because the Kenyan retail sector constitutes a large portion of the nation's economy. The retail sector amounts to \$28 billion or over a third of the GDP, and the wholesale and retail trade sectors combined account for about 29% of urban jobs in Kenya. Women play a prominent role, comprising around 60% of those engaged in wholesale and retail trade. Micro-retailers, who are the subject of our study, constitute a substantial portion of the retail sector in Kenya, representing about 70% of the total, and they are a dominant form of small entrepreneurship in the country. These micro-retailers are crucial in providing essential goods and services to local communities, particularly in rural and underserved areas. There are approximately 250,000 micro-retailers, supporting around 1.5 million individuals through employment and related support services.⁵

The deployment of Jaza Duka was an unprecedented effort due to its sheer scale, reaching over 40,000 retailers throughout Kenya. Of these, approximately 13,000 retailers successfully opened a credit line through the program. A significant aspect of Jaza Duka was that the credit qualification process for retailers was solely based on their purchasing history with Unilever. 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.

One notable advantage of Jaza Duka, compared to other existing financing options like group loans, is that no co-signing or group formation was required. This feature aimed to promote financial inclusion by enabling individuals from ethnic minorities and women to access individual loans in cases where forming a lending group may be impractical or challenging. 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

⁵Source AT Kearney 2017 Global Retail Development Index; Kenya Vision 2030

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.

Jaza Duka functions akin to a modern credit card. 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 avoid their credit line 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. Given that Unilever constitutes a significant portion (15-20%) of the revenue for merchants, the threat of being cut off can pose a non-trivial business cost.

All credit operations are conducted through mobile money, specifically the widely used platform called M-Pesa. Retailers are responsible for keeping track of their credit balance and ensuring they make sufficient payments to avoid restrictions. Due to possible income shocks and the perceived complexity of the credit product for the retailers (many of whom were first-time users of formal credit), many retailers in our sample, specifically more than 60%, experienced restrictions on their credit line at some point.

To understand the financial sophistication in the market, we conducted measurements. 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 complexity of the Jaza Duka program's rules, leads us to categorize the adoption of this credit product as a risky business strategy.

Moreover, among less sophisticated and inexperienced entrepreneurs with limited exposure to formal credit technologies, risk preferences are still being formed. This situation provides us with a unique opportunity to study the initial formation of risk preferences in response to the adoption and experience with a novel financial instrument. This advantage of studying the early development of risk preferences in response to a new financial instrument is particularly pronounced in developing markets like Kenya. In more developed markets

where the use of modern financial instruments is already more established and ingrained, individuals may have already formed their risk preferences either based on prior experience or from information from peers or the news media. However, in a market like Kenya, with a low penetration of business credit and limited familiarity with financial instruments, studying the formation of risk preferences based on initial experiences is valuable. It allows us to capture the dynamics and influences that shape risk preferences in a relatively uncharted financial landscape, providing insights that may be less readily available in more developed markets. Furthermore, by randomizing the adoption of this financial instrument, we can establish causal relationships and draw robust conclusions about its impact on risk preferences.

3 Experimental Design

To investigate the impact of adopting a risky business strategy on subsequent risk aversion, we conducted a randomized experiment involving access to the Jaza Duka retail credit line. This experiment’s subject pool consisted of retailers eligible for the Jaza Duka program and affiliated with a single distributor called Banjara. This distributor operates in Kenya’s coastal region, specifically in Malindi, which was the market chosen for the experiment. Banjara was chosen as the distributor following discussions with Unilever and Mastercard, with an emphasis on ensuring equitable access to credit. We were particularly cautious not to exclude any retailers from access to credit that they might have otherwise secured without our experimental intervention.

Mastercard implemented the Jaza Duka program in a phased manner because of the technological and logistical costs of introducing the product in a new market. Jaza Duka was initially launched in major markets like Nairobi and Mombasa, and the expansion to Malindi was scheduled for the latter half of 2020. For our experimental design, Mastercard agreed to hasten the introduction of Jaza Duka in Malindi, advancing its start to the latter half of 2019.

Our study population consists of all Banjara retailers 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.

The timing of data collection for this study, completed before the unforeseen COVID-19 pandemic, circumvented the confounding factors introduced by the crisis. After the study concluded, the pandemic led to an unforeseen extension of the control period for the group originally slated to receive credit later at the originally planned roll-out date in 2020. Mastercard, responded to the crisis by placing a hold on the expansion of the Jaza Duka program. Thus, in the absence of our experiment, none of the retailers in Malindi would have received Jaza Duka until several years later. But, due to our experiment the retailers in the treatment group were able to get the credit offer before the onset of the crisis.

Before proceeding with the data analysis, it is helpful to frame our experimental design using the potential outcomes framework as outlined by Angrist, Imbens, and Rubin (1996). For each subject i , the dependent variables, which in our case indicate risk aversion measures, are labeled as Y_i . Whether or not the subject is offered credit is denoted by Z_i , and we use the dummy variable D_i to reflect whether individual i has adopted credit before the risk preferences were measured.⁶

To further specify the framework, define the function $Y_i(\mathbf{Z}, \mathbf{D})$, which determines the outcome for an individual who is in treatment arm Z_i and who makes a credit adoption

⁶Given that our sample includes panel data on credit adoption, we are able to identify individuals who adopted credit after their risk aversion levels were assessed. Obviously, such a decision could not impact the measured risk preferences; thus, our analysis of causal impact of credit adoption on preferences does not use adoption after the risk preferences were measured. Nonetheless, we use overall adoption rates when relevant, for example, when assessing the overall default rates, or when describing adoption rates across segments. Additionally, although this particular segment of the population might offer intriguing perspectives, it is too limited in size to conduct a statistical analysis with sufficient power.

decision D_i . Additionally, we have the function $D_i(\mathbf{Z})$, which represents the credit adoption decision of i based on the treatment arm they are assigned to. Following the convention, at this stage, we allow dependence of both functions on the entire vectors, \mathbf{Z} and \mathbf{D} .

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 that the credit would have enabled. The salespeople were instructed not to aggressively promote credit sales as it could divert attention from selling actual products. As a result, the adoption rate of credit, defined as buying at least one product on credit, was relatively modest, amounting to 26% before risk aversion was measured, and eventually peaking at 38%. 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.

Under the assumptions of SUTVA and Unconfoundedness, we can denote potential risk attitudes as $Y_i(D_i)$ depending on the credit adoption decision. Specifically, we refer to $Y_i(0)$ as the ex-ante or baseline risk aversion and $Y_i(1)$ as the ex-post risk aversion. The ex-ante risk aversion represents the risk aversion that would prevail in our experimental population without introducing the credit product. Notably, the risk aversion of the control group can be seen as equivalent to ex-ante risk aversion due to the strict control features of our experimental design.

In the next section, we discuss the descriptive statistics of the data, including the variation

in ex-ante risk aversion among the participants.

4 Data

The research uses data collected as part of a large RCT program in Malindi that evaluated Jaza Duka credit product and accompanying business training offered by Mastercard. The entire 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 COVID and was the same as baseline with added questions about risk aversion, time value of money, and psycho-metrics. The follow-up end-line telephone survey was conducted after COVID and was focused on the responses to the pandemic. We also obtained individual level whole-sale ledgers from Unilever at the transaction level that contains aggregate pre-study data and product-level 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 whole-sale ledgers), patterns of leveraged purchases (wholesale ledgers).

As previously noted, 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 the Online Appendix. 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. The Appendix Table 9 presents an attrition analysis comparing the 150 stores to the remaining 780 establishments. Initially, we scrutinize the 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 midline 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. The Appendix Table 10 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. However, it is worth noting that owners tended to be slightly older and predominantly male. 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 will revisit our analysis concerning the co-owner and employee sub-populations in order to empirically test the aforementioned hypothesis.

Table 1 contains summary statistics of our sample, entrepreneurs that open, credit eligible, completed midline survey and reported as sole owners, across both control and treatment arms. The sample is 65% males, as measured in a post-experiment survey. Age was measured in a follow up telephone survey and averaged to 39 years old. 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 instance, Statista⁷ reported that only 3.5% of Kenyan residents have college degrees.

The third panel of Table 1 contains distribution of store size according to the volume of 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 categorized as small, while 12% are large. We also measured self-reported pre-treatment revenue and find that an average store generates 11,665 Kenyan Shillings per day which amounts to circa

⁷<https://www.statista.com/statistics/1237796/distribution-of-population-in-kenya-by-highest-level-of-education-completed/>

\$100. The median store revenue is approximately 7,000 Shillings.

The Table confirms that 50% of subjects were offered credit. Of those, 38% signed up and successfully made at least one purchase on credit. Males had a 41% credit adoption rate, while the female adoption rate was 34%. 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 26%, 29% and 20% for the entire population, males and females, respectively.

The online Appendix A contains randomization checks using observable characteristics. We find that out of 15 coefficients, only age is significant at 10% level. However, the economic significance of this difference is minimal. In particular, the treatment group has an average age of 28.9 years, compared to 27.2 years in the control group.⁸ Additionally, an F-test, which considers all the variables in the regression together, delivers a p-value of 0.65, which confirms that the two groups are well-balanced.

4.1 Risk-Taking and Default Behavior

In this section we provide descriptive and model-free evidence of risk taking and credit default behavior. This would set 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 our focal market. 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 are unaware of the interest rate they need 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.”

Apart from default rates, we have a measure of risky behavior facilitated by credit. Interviews with Banjara, the local Unilever distributor serving all the stores in our sample,

⁸The difference does not survive a t-test (p-value of 0.26). The initial finding of age being significant at the 10% level in a joint regression is likely attributed to the issue of multiple comparisons.

revealed that an important intended function of credit was to allow the stores to try products they otherwise would not have been able to try due to liquidity constraints. Consequently, substantial sales force effort was dedicated to promoting new products to stores that adopted the credit card. Notably, the salesperson knew the remaining credit card balance of a retailer while deciding whether to push new products.

An initial observation from the data reveals that, for those who adopt credit, a typical credit purchase aligns closely in size with a non-credit purchase. Nonetheless, 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 tend to 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 Stock Keeping Units (SKUs) with their first swipe—items that were previously not part of their inventory. To put this into context, in 2019, a mere 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 instance, 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 6.2.

	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	11		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 1: Descriptive statistics

4.2 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 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.⁹ We used simple 50/50 gambles instead of a collection of gambles with varying probabilities, because we suspected that the concept of probability may be difficult to understand.

The first gamble in the sequence sets π to 100 Shillings. A rational decision maker should always reject this gamble which is Pareto dominated. Therefore this choice serves as detection of subject comprehension or possible mis-coding by an enumerator. We find that 7.41% of subjects fail the comprehension test. Subjects that pass and fail a comprehension test are not statistically different regarding credit adoption rates, SKU adoption, or default. In the subsequent analysis, we work with the full sample, if possible. We drop subjects that fail the comprehension test if the analysis is infeasible without assuming some degree of rationality. In other words, some of our analysis that assumes rationality cannot be redone incorporating irrational subjects who fail the comprehension test. The analysis that can be redone, such as regressions using percentage of risk-averse subjects as dependent variable, was successfully replicated when including irrational subjects. Note that some of the descriptive statistics numbers in the subsequent sections may differ slightly from the numbers reported in Table 1, if the analysis was done using only the rational sample.¹⁰

Another important aspect of assessing decision-making coherence is the consistency of

⁹The exact wording was: “which one of the following two options will you choose to receive: 100 vs π - 50% of the time.” The enumerators were instructed to answer any questions about the choice options.

¹⁰As previously noted, 139 establishments did not participate in the baseline survey but did respond to the post-treatment survey. This attrition is inconsequential for certain demographics, such as gender, as we utilize measurements from the post-treatment survey. Concerning education, we define the dummy variables as “responded to baseline survey and reported given education level.” As for age and revenue, we initially gathered 439 and 500 data points in the baseline survey, respectively, and we present the averages derived from the collected data. The attrition does not affect most of our conclusions since these variables are utilized only sparingly. For the examination of the determinants of ex-ante risk aversion outlined in Table 2, where the baseline survey data is employed for the control group, we eliminate the absent observations. This culling procedure yields 277 observations, while the total count of the control group stands at 303. In our analysis of heterogeneous treatment effects, we substitute missing age data with the mean age value.

preferences across different gambles. If an individual's preferences are stable and consistent, they should consistently accept all gambles that are Pareto superior to a specific gamble they have accepted, and similarly, they should consistently reject all gambles that are Pareto inferior to a gamble they have rejected. Our findings reveal that a mere 5% of participants exhibit inconsistency in their choices. This proportion is balanced across both the treatment and control groups, as well as between those who adopted credit and those who did not (with p-values of 0.99 and 0.30, respectively). This distribution of inconsistencies hints at a "trembling-hand" explanation, which attributes the occasional inconsistent choices to small, random errors rather than systematic errors in judgment or understanding. Despite the existence of these inconsistencies, we decided to retain these individuals in our sample. To account for their inconsistencies in a systematic manner, we will employ the "envelope approach," which we will elaborate on in the following section. Taking into account the possibility of coding errors, as well as the hypothetical nature of the gambles, we regard the small proportion of inconsistencies as an indicator that participants were indeed paying attention and engaging with the questionnaire.

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 will introduce other more complicated measures that aggregate information from many gambles at the individual level; however, in this section we will 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 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 these 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 can help to alleviate this concern. Since we can attribute the change to risk attitudes rather than beliefs.

As mentioned earlier, the gambles are hypothetical and low-stakes. This design prompts a valid question regarding whether the recovered risk preferences can predict actual risky business decisions, such as adopting loans or new products. The 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) 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 further been emphasized 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.

Furthermore, we proceed to correlate our measure of risk aversion with the pertinent responses obtained from the survey. Given our hypothesis that risk aversion measures in the treatment group might shift, a similar transformation might also affect the survey responses, leading to simultaneity. To address this issue, we confine our focus to the control group, which corresponds to *ex-ante* risk aversion, as defined in the previous section.

	(1)	(2)
	Risk aversion	Risk aversion
Have you taken a loan in the last 12 months	-0.101* (0.0574)	-0.106* (0.0604)
Do you have cash savings?		0.123* (0.0685)
Save at the bank		-0.111** (0.0502)
Do you keep financial records?		0.00171 (0.0487)
Would you consider mobile money for loan payments		-0.0563 (0.0468)
Customer credit		-0.0430 (0.0463)
Daily revenue		0.00000187 (0.00000152)
Credit arm	277	277

Standard errors in parentheses

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

Table 2: Correlation between business practices from pre-treatment survey and rejection of an actuarially fair gamble (risk aversion).

To explore the drivers of our ex-ante risk aversion measure, we utilize a pre-treatment survey that was conducted over eight months before we elicited risk preferences. This significant time gap allows us to assess the validity and time consistency of our measure of risk aversion.

Initially, we reproduce the finding that risk-averse entrepreneurs tend to have a lower demand for credit, focusing this time on loans other than Jaza Duka. In column (1) of Table 2, the data indicates that subjects with higher risk aversion reported fewer loan uptakes in the past 12 months compared to their risk-loving counterparts. We then broaden our analysis to include various measures of business and financial practices. We observe that risk-averse entrepreneurs tend to save more money, however, they tend to save more in cash rather than in bank savings. This pattern supports the idea that storing cash at home may

be considered as a more preferred option and a safer practice as compared to saving in a bank and this may speak to a general lack of exposure to the formal financial sector in our sample.

Next, we explore the correlation of our measure of risk aversion and proxies of wealth. We find no correlation of risk aversion with daily revenue. We also tried daily profits and found similar results. We also explore demographic variables, such as store size, gender, religion, and education.¹¹ We find some evidence that owners of larger stores, older owners, and Muslims, are more risk-averse.¹²

5 Results

Our analysis is conducted 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, business size, and big5 personality metrics. Additionally, we provide both model-free and structural estimates of the ITT effects.

Moving to the second step, we aim to estimate the effect of 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 credit line was not utilized. Additionally, we propose a mechanism that influences risk aversion by analyzing credit default data and SKU-level purchase data. By utilizing reduced form and structural partial compliance frameworks, we demonstrate that the ATT effect is more substantial for individuals who purchased new SKUs during their initial credit purchase and for those who eventually defaulted on the credit loan.

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

¹¹The results are contained in the Online Appendix A.

¹²We also collected psychometric measures of big5 personality traits as in Rammstedt and John (2007). We find that scores related to Agreeableness (reverse of “tends to find fault with others”), Conscientiousness (reverse of “tends to be lazy”), Neuroticism (reverse of “is relaxed and handles stress well”) are correlated with risk aversion (see Nigel Nicholson and Willman, 2005).

5.1 Causal Effect of the Credit Offer

Since the credit offer Z_i 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). As mentioned earlier, 92.6% of participants reject the Pareto dominated gamble, which we interpret as rationality (or coding error) test. Notably, the pass rate for this test does not exhibit significant difference between the treatment and control groups. Moving to the second gamble, it is also rejected a

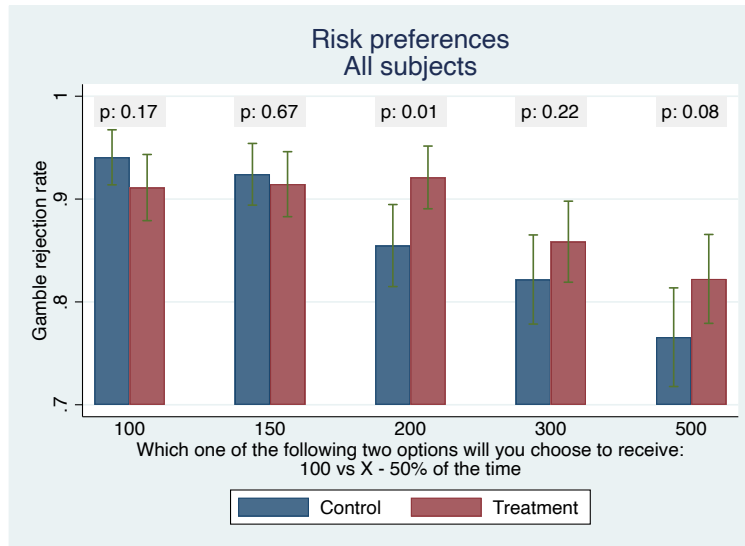


Figure 1: Gamble rejection rate. Average of the raw data of responding to gamble questions. Treatment is the offer of credit. Brackets are 95% CIs for each bar.

large, 91.9%, likelihood in both arms. Acceptance of this gamble would imply a considerable level of risk-loving behavior. However, since very few individuals accept this gamble, it is subject to ceiling effects. In literal terms, the treatment does not impact the percentage of highly risk-loving individuals.

The third 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 85.5% rejection rate, whereas the treatment group showed a 92.1% rate, resulting in a statistically significant 6.6% increase among those offered credit (p-value of 0.01). This shift suggests that access to credit alters risk preferences, converting 6.6% of individuals from risk-loving or -neutral to risk-averse preferences. The remaining two gambles in Figure 1 do not reveal statistically significant differences in rejection rate at the 5% significance level in the overall population, possibly due to ceiling effects. However, the last gamble demonstrates significant difference at the 10% level, with a p-value of 8.4%.

Moving to the regression analysis, we will estimate models using data on the individual gambles as well as data aggregated across all gambles at the individual level. As explained later, aggregation of choices requires some degree of rationality; thus, we have excluded 7.4% of participants, who accept the Pareto dominated gamble. For consistency we estimate all the regressions with this smaller sample. We will also re-estimate all the models that do not rely on rationality using the entire sample, and report any relevant differences.

Column (1) of Table 1 presents the regression outcomes where 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 – i.e., 6.6% treatment effect in the entire sample. When utilizing the rational sample, the effect is slightly larger. However, it does not significantly diverge in qualitative terms from the effect observed in the entire sample. We extended our analysis to encompass the entire sample, incorporating sole owners, co-owners, and employees. As predicted, we observed no discernible effect of credit on either co-owners or employees. This may partially be due to a smaller sample size; however, when pooling the populations of owners and co-owners, we detected a statistically significant albeit smaller effect of 4.7% (significant at the 5% level), which suggests smaller effect or absent effect for co-owners. Detailed regression analyses supporting these findings are provided in the Appendix Table 11.¹³ As mentioned earlier, approximately 5% of total participants, and less than 3% of participants in the rational sample, exhibit inconsistency in their choices, either accepting dominated gambles or rejecting dominant gamble, as compared to their previous choices. Such behavior also

¹³The regressions are performed on the sample of 755 establishments that furnished the answer to the ownership question.

	(1)	(2)	(3)	(4)	(5)	(6)
	Risk aversion	Risk premium LB	Risk aversion	Risk aversion	Risk premium LB	Risk premium LB
Credit	0.0785*** (0.0215)	11.51** (4.799)				
Not adopted credit			0.0848*** (0.0234)		14.53*** (5.210)	
Adopted credit			0.0611* (0.0332)	0.296*** (0.0865)	3.157 (7.408)	43.35** (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 3: 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.

deviates from the rationality assumptions. We considered various approaches to address this empirical challenge, all of them involving carefully curating the data. The most drastic approach involves excluding users who exhibit irrational behavior, while a more nuanced alternative entails removing specific gambles that elicit inconsistent responses. Below we describe our preferred way to prune the data using induction from above.

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 agent, define rejection threshold, $\bar{\pi}$, so that they reject the gamble if and only if $\pi_t \leq \bar{\pi}$. For simplicity, we assume that users 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, $\text{RPAT} = 0.5\bar{\pi} - 100$.

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. Risk premium for the fair gamble is formally

defined and will be used in the structural analysis in the next section. The advantage of RPAT is that it is partially observable in the data, as we will later see, through the obtained bounds on the threshold. Consequently, it serves as a useful monetary and model-agnostic measure of risk preferences. In contrast, the risk premium for the fair gamble is more challenging to obtain directly and it requires a model for determination. A disadvantage of RPAT is that it may reach infinity, particularly if the agent rejects any gamble where they have a 50% chance of obtaining π and a 50% chance of zero.

Denote each gamble by t . We solicit responses to 6 gambles with different values of π_t . Thus, our data partially identifies the threshold $\bar{\pi}$. 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 $\bar{\pi} \in [300, 500)$. By induction, we can partially identify the CDF of $\bar{\pi}$.

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 agents 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 depicts the empirical CDF of the threshold in the control and treatment group obtained using the empirical distribution of $\bar{\pi}$. Each point on the X-axis represents possible rejection thresholds, and corresponding risk premiums, 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 0.021, which indicates that the offer of credit affects the CDF of rejection thresholds; thus, altering the risk preferences.

Column (2) in Table 3 contains regressions of the lower bound of risk premium on the




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

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.

5.2 Selection into and Causal Effect of Credit

In this section, we estimate causal impact of adopting credit for adopters on their risk aversion. In order to satisfy the exclusion restriction for instrumental variables, we make the assumption that simply receiving the credit offer does not impact risk aversion for non-adopters. Formally, $Y_i(1,0) - Y_i(0,0) = 0$. Following the convention, we define $Y_i(D_i) = Y_i(Z_i, D_i)$.

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.,

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

We refer to the measured risk aversion of the individual before they take the credit offer $Y_i(0)$ as *ex-ante* risk aversion.

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

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 control 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 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 3 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 8.5% and 6.1% 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. The validity of this estimate hinges on the assumption that the mere offer of credit does not influence risk aversion – the exclusion restriction. 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 about -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 the 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 [3](#).

We conducted the same analysis using a pooled sample and using the risk premium as the dependent variable. The results are presented in Columns (5) and (6) of Table [3](#). Our findings reveal that credit adopters have *ex-post* risk premium that is 11.37 Shilling lower than that of the non-adopters. The IV estimate suggests 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 selection out of credit equal to 14.53 Shillings and selection into credit of -40.19 Shillings. Overall, the findings of the risk premium analysis are similar to the one with the risk aversion measure. The selection into credit is substantial but the treatment effect is even more pronounced.

Comparing findings from the fair gamble acceptance rates and those derived from risk premiums suggests that the latter 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.

5.3 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, you’ll find a regression of CATE on age, which indicates a negative and statistically significant relationship with age, confirming our hypothesis.

To further rank the demographic moderators in terms of their influence on the treatment

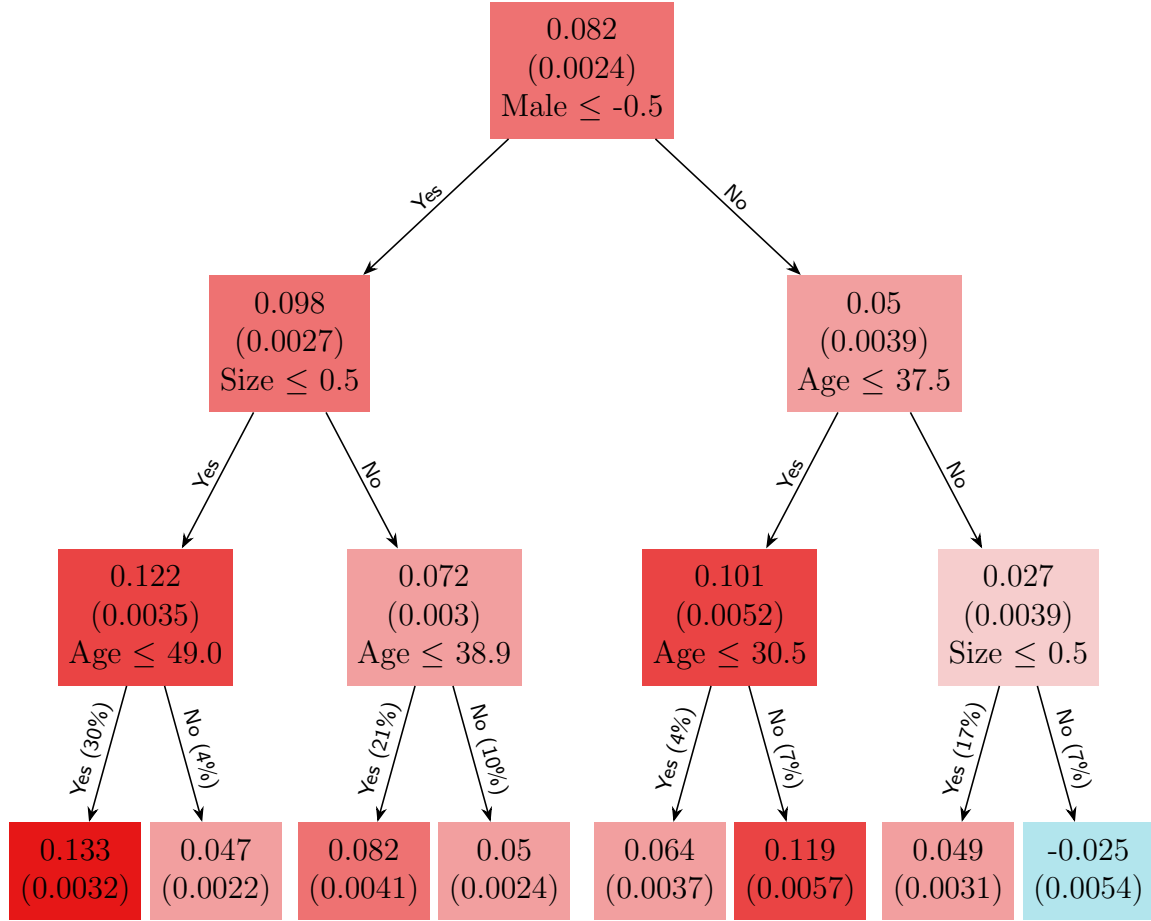


Figure 4: Decision tree depicting decomposition of Conditional Average Treatment Effects.

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, as will be demonstrated 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 74% compared to a 54% 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 50% of women (with a p-value over 0.1%).

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 very similar again across both genders but the relevant cutoff age is higher for males than females. The younger group shows a larger effect. There seems to be a reversal for very young women but the sample size here is very small. This is consistent with the evidence in neuroscience that as we age, a new experience has a progressively smaller effect on the wiring of our brains. It is also worth pointing out that if we split the sample on those 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, most of the results align closely with those using 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. In other words, greater introspection may make individuals more open to changing their behavior in the future.

The individuals that "find fault in others" are classified by big-5 scale as less "agreeable." According to Enrick and Langford (2000) based on this scale managers tend to be more agreeable than entrepreneurs. Thus, individuals exhibiting entrepreneurial personality traits are less likely to internalize failures and are less affected by the negative experience than individuals with lesser entrepreneurial personality [CITE THIS PAPER YOU MENTIONED].

	(1) Risk aversion	(2) Risk aversion	(3) Risk aversion	(4) Risk aversion	(5) Risk aversion	(6) Risk aversion	(7) Risk aversion
Credit	0.0292 (0.0365)	0.105*** (0.0271)	0.0292 (0.0348)				
Male	-0.0517 (0.0318)						
Male/Credit	0.0754* (0.0451)						
Not adopted credit				0.101*** (0.0300)		0.0542 (0.0366)	
Adopted credit				0.113*** (0.0400)	0.349*** (0.0976)	-0.0708 (0.0599)	0.146 (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 4: 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 4, 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, however, the degree to which their risk preferences are altered post-adoption varies.

6 Utility Model

This section evaluates the determined effects of credit adoption through various structural models that assume utility-maximizing behavior. This will allow 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 6.1, we analyze a model that endogenizes the gamble 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 CARA and CRRA utility models and show that accounting for wealth effects does not change our main findings. In Section 6.2, 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.

6.1 Credit Adoption Model

In this subsection, we develop and estimate two structural models of credit adoption and the subsequent risk behavior associated with credit usage based on expected utility maxi-

mization. The first, more straightforward model, relates the change in preferences to credit adoption and gauges both the ex-ante and ex-post distributions of utility parameters. The strength of this model lies in its close connection to our descriptive analysis, and it leans solely 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 simply a result of wealth effects without underlying changes in the utility function. 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 informed by post-experience preferences.

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

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

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

In this model, the risk aversion 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.¹⁴

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

¹⁴We 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).

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).

It is helpful to link the structural parameters of the model to the treatment effects identified in the previous section. The simplest mapping is for ATT, which is given by Δ . Similarly, the ITT is given by the difference in the distribution of γ depending on the experimental arm, represented by $F(\gamma|Z_i = 0)$ and $F(\gamma|Z_i = 1)$. Concretely, ITT can be expressed as in the previous section

$$ITT = E[\gamma_i|Z_i = 1] - E[\gamma_i|Z_i = 0] = \Delta E[D_i|Z_i = 1].$$

In an analogous fashion average selection (SEL) is expressed as

$$SEL = E[\gamma_i|D_i = 1, Z_i = 1] - E[\gamma_i|D_i = 0, Z_i = 1].$$

To ease the interpretation, we report ITT, ATT and SEL as an implied difference in the average risk premium of the fair gamble. The risk premium is defined as

$$u^{-1} \left(\frac{1}{2}u(200; \gamma_i, w_i) + \frac{1}{2}u(0; \gamma_i, w_i) - u(100; \gamma_i, w_i); \gamma_i, w_i \right).$$

In addition to capturing the average ITT, the model also accounts for variance shifts, which are determined by the correlation between ϵ_i and ν_i . For instance, we demonstrated

that risk-loving individuals adopt credit more frequently, which implies a negative correlation between ϵ_i and ν_i . In such case, the variance of γ in the treatment group may be dampened. This occurs because the positive treatment is more commonly applied to individuals in the lower tail of the γ distribution. While these effects can be identified without relying on a specific model in a non-parametric manner (such as through quantile regressions), it may be necessary in practice to randomize credit on a very large scale to ensure sufficient statistical power. Given our sample size we follow a parametric approach.

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

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

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 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.,

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

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 when writing the choice likelihood. Unfortunately, we do not have wealth information. But we do have information on daily profits which is likely the most relevant driver of differences in wealth induced by credit. Thus, this measure should capture the relevant variation in wealth-driven differences in risk aversion between treatment and control.

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). The unit of observation is a single entrepreneur and the data is given by the observed lower and upper bounds of the implied acceptance thresholds, denoted by $\bar{\pi}_i^L$, and $\bar{\pi}_i^U$, respectively. As a result, standard errors are clustered at the unit of randomization.

To compute the likelihood $L_i = \Pr(\bar{\pi}_i^L < \bar{\pi}_i \leq \bar{\pi}_i^R | X_i; \theta)$ we simulate $R = 10^7$ draws from the implied distribution $F(\gamma_i; \theta)$ and count gamble choices that are consistent with the threshold being in the observed interval. We subsequently find $\hat{\theta}$ that maximizes the population log-likelihood. We obtain standard errors by using a non-parametric bootstrap which samples agents with replacement from the empirical distribution.

6.1.1 CARA Utility Model

Table 5 presents the results of the estimation. Column (1) displays the outcomes of the CARA model that does not allow for wealth effects (and hence, does not use data on wealth). This approach closely mirrors the model-free analysis, attributing all experimental variations in gamble acceptance rates to inferred differences in risk aversion.

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 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.

The second section of the table outlines causal effect and selection estimates for the male population. These effects are quantified by examining the average differences in risk premiums, which are derived from the participants' demonstrated preferences. The risk premium calculation is based on making subjects indifferent between engaging in an actuarially fair

		(1) CARA No wealth effect	(2) CRRA Constant wealth	(3) CRRA Linear in daily profits	(4) CRRA Linear in daily averaged profits	(5) CRRA Quadratic in daily averaged profits
RA intercept, $\bar{\gamma}_{\text{FEMALE}}$		1.76*** (0.17)	3.17*** (0.25)	2.93*** (0.46)	2.88*** (0.44)	2.07*** (0.22)
RA intercept, $\bar{\gamma}_{\text{MALE}}$		1.73*** (0.16)	3.07*** (0.28)	2.87*** (0.45)	2.83*** (0.43)	2.04*** (0.20)
RA Dispersion, σ_{γ}		1.21*** (0.13)	2.13*** (0.16)	2.03*** (0.29)	1.99*** (0.28)	1.44*** (0.13)
RA-adoption correlation, ρ		-0.39*** (0.13)	-0.45*** (0.12)	-0.52*** (0.13)	-0.50*** (0.13)	-0.52*** (0.12)
Impact of adoption, Δ_{FEMALE}		0.16* (0.09)	0.10*** (0.02)	0.31 (0.36)	0.32 (0.42)	0.45 (0.39)
Impact of adoption, Δ_{MALE}		0.67** (0.29)	1.36*** (0.40)	1.35*** (0.39)	1.26*** (0.39)	0.96*** (0.30)
Wealth intercept, \bar{w}		-	59.31*** (10.89)	51.59*** (16.94)	46.06*** (15.81)	14.82*** (4.76)
Wealth slope, ϕ		-	-	0.00*** (0.00)	0.00* (0.00)	0.00 (0.00)
Wealth quadratic		-	-	-	-	0.00 (13.11)
Risk premium for the fair gamble	ITT	2.99* (1.73)	4.13*** (1.52)	5.24** (2.28)	5.00** (2.23)	5.90*** (2.15)
	SEL out of	5.12** (2.24)	6.51*** (2.17)	7.81*** (2.68)	7.50*** (2.61)	8.04*** (2.58)
	SEL into	-14.00** (5.81)	-17.90*** (5.36)	-21.58*** (7.57)	-20.70*** (7.21)	-22.46*** (7.16)
	ATT	9.26*** (2.17)	13.35*** (4.83)	17.97** (8.98)	17.21** (8.69)	21.15** (8.55)
Credit adoption (males)	Baseline	27%	27%	27%	27%	27%
	Counterfactual	20%	19%	16%	17%	17%
Risk premium for the fair gamble (males)	ITT	5.07** (2.29)	6.52*** (2.34)	7.30*** (2.78)	6.91** (2.74)	7.57*** (2.67)
	SEL out	5.69** (2.53)	7.24*** (2.50)	8.67*** (3.05)	8.32*** (2.98)	8.93*** (2.97)
	SEL in	-12.95** (5.50)	-16.55*** (5.14)	-19.74*** (6.94)	-18.94*** (6.60)	-20.39*** (6.54)
	ATT	16.60** (7.33)	21.44*** (7.43)	23.92*** (9.22)	22.64** (8.99)	24.85*** (8.77)
Credit adoption (males)	Baseline	26%	26%	26%	26%	26%
	Counterfactual	19%	18%	16%	16%	16%
N		587	587	587	587	587

Standard errors in parentheses

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

Table 5: 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.

gamble — winning 200 Shillings with a 50% chance, or receiving a certain 100 Shillings. By expressing these effects in Shillings, the study offers a tangible measure of the economic impact of the structural parameters.

Males in the treatment group are inclined to pay an additional 5.07 Shillings to bypass the actuarially fair gamble. Regarding selection, entrepreneurs who opted out of credit have a 5.69 Shillings higher risk premium than the general population before the offer of credit. On the other hand, those selecting into credit exhibited ex-ante 12.95 Shillings lower risk premium. Furthermore, post-credit adoption, entrepreneurs exhibit a significant shift in preferences, with an increased willingness to pay an extra 16.6 Shillings to avoid taking the gamble, than they would have without adopting credit.

These findings corroborate the conclusions drawn from the reduced form analysis, particularly the observation that selection into credit (i.e., initially lower risk aversion among those who choose to take credit) is more pronounced than the selection out of credit. Moreover, the post-adoption preference change towards greater risk aversion is notably larger than selection into credit. It is the more risk loving entrepreneurs in the population who drive the adoption of new innovations (in our case the new credit technology). But the failure experience affects the risk preferences of precisely these entrepreneurs more and may make them unwilling to adopt future credit offers that would be valuable.

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 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 26%, whereas the counterfactual adoption rate decreases to 19%.

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.

6.1.2 CRRA Utility Model

Columns (2) to (5) of Table 5 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.¹⁵ 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 5 additional percentage points to 17%.

¹⁵For 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).

The next step is to enhance the specification to consider the potential influence of credit on wealth. We make use of our data on daily profits (solicited in the same survey as the risk aversion measure) and consider the following equation for wealth:

$$w_i = \bar{w} + \phi \times (\text{DAILY PROFITS})_i.$$

This approach assumes that an entrepreneur’s daily profits can serve as a direct representation of their wealth because for many entrepreneurs the store is the only source of income. In particular, according to our survey only 18% of entrepreneurs report having other income besides the store. Column (3) offers results from this model. All the values are very close to those in Column (2) indicating that after partialling out wealth there remains a large unexplained variation in risk taking across the three relevant groups.

Note that the treatment effect metrics in the table’s concluding three rows also accommodate wealth disparities, as they derive from the updated structural estimates. Moreover, to pinpoint the influence of intrinsic risk preferences, all risk premiums are calculated utilizing the mean daily profit level in the control group. Little difference between treatment effect 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 the method of using individual-level daily profits provides a granular view of financial standings, it may also capture inherent variability in individual daily profit reports from the survey. This may lead to attenuation of the estimate of ϕ and underestimation of wealth effects. To address this, instead of using individual response we use average reported daily profits computed separately for three distinct groups: the control group, the treatment group without credit, and the treatment group with credit. In essence, such a specification should adjust for average cross-sectional wealth differences between credit adopters, non-adopters, and control group. Column (4) shows that averaging daily profits has no impact on the results, which suggests that survey noise in individual daily profits is also not consequential. Finally, for robustness, in (5) we report our findings when w_i is allowed to be a quadratic function of daily profits. Larger treatment effect values were derived, implying that a more complex functional form does not account for the observed result.

In the bottom panel of Columns (3)-(5), we reassess the credit adoption rates fixing the daily profits to the average in the control group. Our aim here is to isolate the impact of changes in risk aversion resulting from credit adoption from any potential effects of credit on wealth. We observe no significant differences between Columns (2)-(5). This lack of variation again suggests that the influence of wealth in driving lower adoption rates is negligible. In other words, a population of “already treated” entrepreneurs would adopt less credit because they have adjusted their utility function, not because their wealth changed.

6.1.3 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 Online Appendix B.

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 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 online Appendix C. 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).

6.2 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, which are probably contributing to the shift in preferences. To gain a better comprehension of this impact, it is beneficial to zoom into the wholesale ledgers at the level of SKU transactions. As mentioned in Section 4.1, 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 merely adopting credit and buying familiar SKUs. To study this, we extend our model to endogenize SKU adoption on the first credit transaction in addition to modeling 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 Heckman and Robb Jr, 1985), beyond the decision to adopt credit. Thus, our analytical structure has to 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 the credit and SKU adoption on credit, modifying a model to include the endogenization of SKU adoption.

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. For example,

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. Another observation is that 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 car ID and Unilever market shares within the sample of credit adopters. Therefore, we elected to formulate an interaction instrument. We define a “low type” as a store that is served by a less successful salesperson *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 car IDs. However, other combinations result in an imbalanced distribution of types and a reduction in statistical power. We proceed with this definition, while being aware that we are testing the limits of our data.

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:

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

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

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

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 6 displays the results of our estimation. The average values, dispersion, and correlations are nearly identical to those observed in the baseline CARA (Constant Absolute Risk Aversion) 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 on credit. It is worth mentioning that the correlation between risk aversion and SKU adoption

		CARA
Risk premium for the fair gamble	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, ρ_1	-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)
	ITT	4.80*** (1.47)
	SEL out of credit	7.12*** (2.02)
	SEL into credit, out of new SKU	-15.87*** (5.19)
	SEL into credit and new SKU	-21.44*** (4.99)
	ATT adoption and no new SKU	10.69*** (4.02)
	ATT adoption and new SKU	19.44*** (5.50)
N		587

Standard errors in parentheses

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

Table 6: Structural Estimates with Credit and SKU Adoption

is smaller than the correlation with credit adoption. This observation may initially appear counterintuitive; however, it can be explained 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 by computing the relevant selection measures later on.

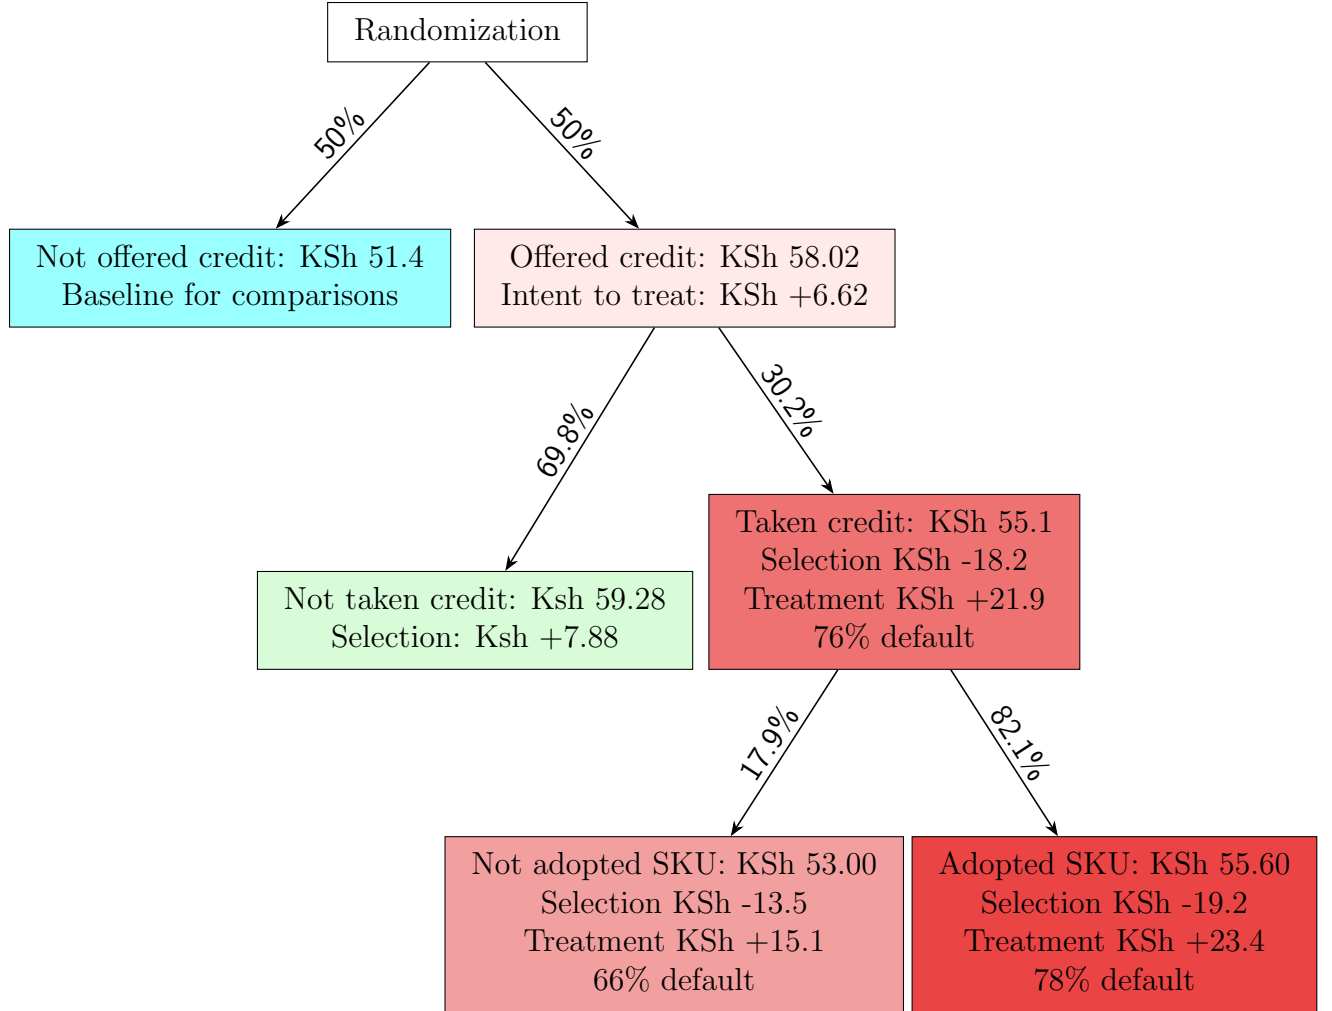


Figure 5: Decision tree depicting the average risk premium for each group and selection (relative to the control) and treatment effects for the high-type males.

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 itself modifies risk preferences, albeit to a lesser extent than using credit for riskier endeavors. This additional modification is captured by the parameter Δ^2 , which is positive.

We detect treatment effects for women that are statistically significant at the 1% level. Yet, the effects are small in magnitude; thus, we concentrate the remaining analysis on the effects for males contending that impact of credit on female preferences is small. The implied treatment effects for males are presented in the second section of the table. The intent-to-treat value is equal to KSh 6.40, slightly larger than in the baseline model, although the

difference is economically small. This variation likely arises from the different functional form of the model and the additional instruments used in this analysis.

To provide a clearer understanding, let’s visualize the decision process and treatment effects for a particular segment – high-type males, depicted as a tree in Figure 5. In the initial section of the tree, we illustrate 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. Additionally, the treatment group shows an average intent-to-treat value of 6.62 Shillings, suggesting a roughly 20% increase in the risk premium in the treatment group.

Moving forward, the treatment group is divided into credit adopters and non-adopters. Among high-type males, 30.2% adopted credit, and 69.8% did not adopt credit before our survey. Credit non-adopters exhibited, ex-ante, a 7.88 Shillings larger risk premium than the general population. Conversely, credit adopters had 18.2 Shillings smaller risk premiums. These selection effects replicate our earlier results. In addition, the model produces selection estimates when adopting SKUs on credit. For instance, new SKU non-adopters are 13.5 Shillings, while SKU adopters are 19.2 Shillings less ex-ante risk averse than the average. These, numbers constitute 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.

Furthermore, we demonstrate that the impact of credit on preferences depends on the nature of its usage. Prudent credit usage generates a shift towards risk aversion, but the shift is not as pronounced as it is for those users who expand their offerings. Part of this effect can be explained by the fact that SKU adopters initially exhibit lower levels of risk aversion, allowing for more significant changes in their preferences. However, because SKU adoption is correlated with defaults, one might expect it to generate a more negative experience and result in increased risk aversion as a consequence. The corresponding set of numbers for all males, both high and low types, are contained in Table 6.

To gauge the impact of these effects on business decisions, we computed counterfactual scenarios, analogous to credit adoption counterfactuals in the previous sections. This time, we considered three sets of preferences: the baseline with ex-ante risk aversion, preferences

modified by credit experience without SKU adoption, and preferences modified by credit with SKU adoption. The results are presented in Table 7.

		Preferences (female)		
		Baseline	Credit without new SKU	Credit with new SKU
High type	Credit adoption	20%	19%	15%
	New SKU if adopted credit	82%	82%	82%
	New SKU	16%	15%	13%
Low type	Credit adoption	20%	19%	16%
	New SKU if adopted credit	50%	50%	49%
	New SKU	10%	9%	8%

		Preferences (male)		
		Baseline	Credit without new SKU	Credit with new SKU
High type	Credit adoption	30%	22%	19%
	New SKU if adopted credit	82%	81%	80%
	New SKU	25%	18%	15%
Low type	Credit adoption	30%	22%	19%
	New SKU if adopted credit	50%	49%	48%
	New SKU	15%	11%	9%

Table 7: Counterfactual business decisions for the male population.

In our study, we separately evaluated the effects for males and females and for High and Low types. For female High types, the baseline rate of credit adoption is 20%, which decreases to 19% for more prudent credit usage scenarios (i.e., no new SKUs adopted) and further to 15% for risky usage scenarios (i.e., new SKUs adopted). For males, we observe a stronger effect by which the credit usage drops from 30% to 22% and 19%, respectively. The rates of SKU adoption also undergo changes, albeit they are very small for both males and females. Initially, over 82% of high types adopting credit also opted for new SKUs. This rate slightly decreases to approximately 81% when considering preferences altered by experience, indicating that high types exhibit a low elasticity in their SKU adoption already adopted credit. This pattern is influenced by the fact that credit adoption already selects individuals according to certain risk preferences (i.e., with lower risk aversion), thus, the populations of credit adopters would exhibit more homogeneous risk preferences, albeit the effect of experience is significantly smaller. This pattern may also imply that the product adoption decisions are predominantly influenced by liquidity constraints. When analyzing the overall rate of SKU adoption, we observe a reduction in SKU adoption from 16% to 13%

for females, and from 25% to 15% for males. For Low types, the observed effect is comparable in magnitude, yet their conditional adoption of new SKUs shows a more significant decline than that of the high types.

The large incremental impact of SKU adoption, over and above adopting credit, suggests that a less-than-ideal experience with adopting new SKUs may backfire and lead to friction when adopting SKUs in the future. Importantly, this effect isolates the impact of changing risk preferences resulting from adopting a risky activity post credit as distinct from other channels that do not rely on preference changes, such as learning about demand. This demonstrates that considering changing risk preferences is an important channel when examining the dynamics of retail entrepreneurial decisions.

7 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, experienced failures arising from a business decision to adopt a new credit technology, can increase risk aversion and influence future risk taking. In the context of a randomized controlled trial which 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. 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 identify some crucial demographic moderators of the selection and treatment effects which have decision making and policy relevance. Males show significantly less ex-ante risk aversion than females, but they also have a treatment effect of credit adoption on risk aversion which is almost twice as that of females. Similarly, younger entrepreneurs and those with smaller businesses are relatively more risk loving prone to adopting credit, but post adoption the effect of a failure also transforms their preferences and makes them substantially more risk averse. Thus from the policy perspective it is younger male entrepreneurs who run smaller businesses who are initially more likely adopt the new credit product, but these are precisely the entrepreneurs who may be particularly susceptible to credit failures and its effects on risk aversion. In so far as these segments of the population are salient in many economies, our findings may help in the design of policies to optimally manage entrepreneurial risk taking and innovation.

References

- 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.
- 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.
- BURKE, S. N., AND C. A. BARNES (2006): “Neural plasticity in the ageing brain,” *Nature reviews neuroscience*, 7(1), 30–40.
- 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.
- 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.
- ELMINEJAD, A., T. HAVRANEK, AND Z. IRSOVA (2022): “Relative risk aversion: a meta-analysis,” .
- ENVICK, B. R., AND M. LANGFORD (2000): “The five-factor model of personality: Assessing entrepreneurs and managers,” *Academy of Entrepreneurship Journal*, 6(1), 6–17.
- 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.
- 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.

- HECKMAN, J. J., AND R. ROBB JR (1985): “Alternative methods for evaluating the impact of interventions: An overview,” *Journal of econometrics*, 30(1-2), 239–267.
- HOLT, C. A., AND S. K. LAURY (2002): “Risk aversion and incentive effects,” *American economic review*, 92(5), 1644–1655.
- 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.
- MALMENDIER, U., AND S. NAGEL (2011): “Depression babies: Do macroeconomic experiences affect risk taking?,” *The quarterly journal of economics*, 126(1), 373–416.
- NIGEL NICHOLSON, EMMA SOANE, M. F., AND P. WILLMAN (2005): “Personality and domain-specific risk taking,” *Journal of Risk Research*, 8(2), 157–176.
- PALACIOS-HUERTA, I., AND R. SERRANO (2006): “Rejecting small gambles under expected utility,” *Economics Letters*, 91(2), 250–259.
- 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.
- RAMMSTEDT, B., AND O. P. JOHN (2007): “Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German,” *Journal of research in Personality*, 41(1), 203–212.
- 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.

A Attrition analysis

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

Standard errors in parentheses.

Table 8: Characteristics of stores that closed or became credit ineligible during the study.

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

Standard errors in parentheses.

Table 9: Characteristics of stores that did not complete the survey.

B Store ownership

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

Standard errors in parentheses.

Table 10: Characteristics of owners vs. non-owners.

	(1)	(2)	(3)	(4)	(5)
	Risk aversion	Risk aversion	Risk aversion	Risk aversion	Risk aversion
Credit Arm	0.0663*** (0.0256)	0.0663*** (0.0255)	-0.0735 (0.0629)	-0.0429 (0.0987)	0.0465** (0.0237)
Credit Arm \times Co-owner	-0.140** (0.0682)				
Credit Arm \times 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 11: Effect of credit offer on risk aversion for co-owners and employees.