The Human Costs of Business Cycles*

Francesco D'Acunto[†] Georgetown University

Pulak Ghosh[‡] IIM Bangalore

Bin Li[§] Wuhan University

Alberto G. Rossi[¶] Georgetown University

This version: August 2025

Preliminary and Incomplete: Comments Welcome

Abstract

Consumers are willing to pay more to avoid business cycle fluctuations than the resulting loss in consumption would suggest. Are agents systematically incorrect in estimating these costs, or do they face costs that standard economic models overlook? By matching over one million photographs with consumers' financial accounts across local business cycles from 2016 to 2023, we find that consumers facing higher local unemployment rates age faster relative to their actual age and to others. This "excess aging" also occurs within the same consumers observed across different cycles. Moreover, we find little scope for negative excess aging ("rejuvenation") during good times. Excess aging in bad times is not explained by health shocks or by financial and demographic characteristics that capture the direct consumption effects of business cycles, such as income levels or employment status. It is more pronounced during volatile cycles and among middle-aged consumers, women, and those with high-responsibility jobs and higher incomes. Access to finance mitigates excess aging in bad times, revealing an overlooked benefit of financial inclusion. We interpret excess aging as a proxy for the human costs of business cycles that standard economic data fail to capture. Alternative data can facilitate a holistic understanding of the effects of economic shocks on consumers and inform the design and assessment of stabilization policies.

Keywords: Behavioral Macroeconomics, Economic Fluctuations, AI, Alternative Data, Access to Finance, Micro-to-Macro.

*For very helpful comments, we thank XXX. All errors are our own.

[†]McDonough School of Business, Georgetown University. e-Mail: francesco.dacunto@georgetown.edu

[‡]Indian Institute of Management, Bannerghatta Road- 560076, Bangalore, India. e-Mail: pulak.ghosh@iimb.ac.in

[§]Wuhan University, Economics and Management School, Wuhan, PRC. e-Mail: binli.whu@whu.edu.cn

[¶]McDonough School of Business, Georgetown University, Washington, DC, USA. Fellow of the Luohan Academy. e-Mail: agr60@georgetown.edu.

"Although I do not suppose that either of us knows anything really beautiful and good, I am better off than he is, for he knows nothing, and thinks that he knows. I neither know nor think that I know."

Plato, *The Apology of Socrates*

1 Introduction

Across the world, consumers express a strong desire to avoid business cycle downturns—so strong that their willingness to pay for stability far exceeds the consumption losses predicted by standard models (Deaton, 2011; Georgarakos et al., 2025; Pfajfar and Winkler, 2025; Stantcheva, 2024; Wolfers, 2003). This empirical fact stands in sharp contrast to the assumptions and conclusions of conventional economic theory, which finds only modest welfare losses from cyclical fluctuations (Lucas Jr, 2003). The discrepancy implies one of two possibilities: either most agents systematically misjudge the consumption costs of business cycle fluctuations and fail to learn from experience over their lifetime, or they face additional costs that standard economic models do not capture. In short, either the consumers are wrong—or the economists are.

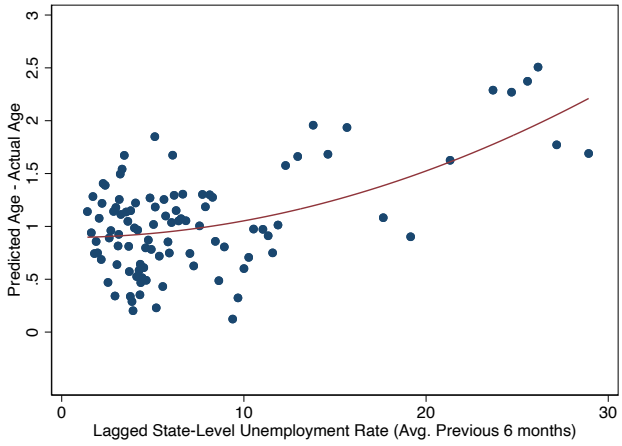
We define the potentially overlooked and elusive burdens that consumers bear—and economists might miss—as the *human costs* of business cycles. Identifying and measuring these costs in real-world settings is challenging because, by definition, they are not reflected in the standard observables that economic models link to consumption. Moreover, self-reported subjective estimates conflate actual objective costs with agents' errors and biases in assessing standard consumption-related costs.These challenges make it essential to evaluate the existence of human costs using externally measurable, objective outcomes.

In this paper, we use alternative data to propose a field proxy for the human costs of business cycles, test for their existence, and examine their properties. Specifically, we match image data for more than one million users of a large Indian bank from 2016 to 2023—a period that spans disparate local business cycles. For a subset of users, we observe multiple photographs taken at different points in time, allowing for a within-individual analysis of how the same person responds to varying macroeconomic conditions. Using state-of-the-art machine learning techniques trained on large-scale image data (Guenzel et al., 2025), we predict each individual's appearance age. We then combine these predictions with bank account information and demographic characteristics, including

users' dates of birth. Subtracting a user's calendar age at the time the photo was taken from their predicted appearance age yields a measure of excess age (Borgschulte et al., ming). Excess age is positive when a person appears older than their actual age, and negative when they appear younger. Across cohorts and years, the mode and average of predicted appearance ages closely track actual age, providing reassurance about the accuracy and reliability of our measurement approach.

We propose using excess age—after controlling for users' economic and demographic characteristics—as a proxy for human costs. In our sample, excess age is, on average, positive and increases more sharply during periods of adverse local economic conditions, even after controlling for users' income, employment status, medical debt, and other proxies of consumption ability and health conditions we observe directly in our data. This pattern supports the interpretation of excess age as a measure of the persistent, non-monetary costs individuals accumulate beyond standard consumption losses.

Figure 1: Excess Age and Local Unemployment Rate



The y-axis plots the difference between bank users' predicted age based on their pictures and their actual calendar age at the time the picture was taken. The x-axis sorts bank users based on the average unemployment rate in the Indian state where they reside over the 6 months preceding the picture.

Figure 1, based on our raw data, motivates the rest of the analysis. It presents a binscatter plot of bank users' excess age (y-axis) against the lagged average unemployment rate in the Indian state where each user resided during the six months prior to when their photo was taken. The figure reveals three baseline facts. First, excess age increases with the local unemployment rate at the time and place the photo was taken. In other words, individuals exposed to worse local business cycles appear older relative to their actual age than those facing more favorable conditions. Second, the relationship is highly nonlinear and exhibits increasing marginal effects: the more severe the downturn, the greater the excess age gap compared to individuals facing less adverse conditions. Third, the relationship flattens as the unemployment rate approaches low levels, suggesting limited scope for "rejuvenation"

when the economy improves. These patterns remain similar when computing average unemployment rates over different horizons. However, the longer the horizon, and thus the more persistent the local unemployment rate, the stronger the observed marginal effects.

The relationship between excess age and lagged local unemployment rates is a robust feature of the data: it persists even after accounting for common shocks that affect all users over time—such as COVID-19 and other time-varying, aggregate fluctuations—as well as time-invariant, location-specific factors, including local labor regulations and economic institutions. Importantly, this relationship holds not only in the cross-section but also within individuals, by exploiting variation among those with multiple images taken at different points in time.

Among users for whom we observe detailed demographic characteristics, the relationship between excess age and local unemployment remains both economically and statistically significant. Importantly, we observe direct proxies for users' consumption ability at the time their pictures were taken—such as income, employment status, and whether they incurred medical debt due to major health expenses. The fact that the association persists even after controlling for these variables reinforces our interpretation: excess age captures objectively measurable human costs of business cycles that go beyond the consumption effects typically captured by standard economic variables in administrative data. These standard metrics reflect economists' and statisticians' attempts to map the concept of consumption onto observable transactions, but they may miss the broader, non-monetary burdens that individuals experience during downturns.

Having established the robustness of the relationship between unemployment levels and excess age, we next assess the combined effect of the level and volatility of local business cycles. To capture business cycle volatility, we compute the standard deviation of monthly unemployment rates in the locations where users reside over the six months preceding the date their picture was taken. In our setting, this standard deviation varies substantially—ranging from about 0.6 percentage points for observations in the bottom third of the distribution to nearly 5 percentage points for those in the top third. We find that both the average level and the volatility of recent unemployment rates are significant predictors of excess age. Moreover, the two dimensions interact: the relationship between average unemployment and excess age is stronger when volatility is high than when the same average rate is more stable. This interplay between business cycle level and volatility holds both in the cross-section and when relying solely on within-individual variation.

We then assess potential heterogeneities in the relationship between local business cycles and excess age across users' demographic characteristics. We find that the relationship is more pronounced for women, who also display, on average, unconditionally higher excess age in our sample compared to men. In the Indian context, women are more economically vulnerable—facing substantial gaps in wages, wealth, and self-employment relative to men (Jose and Sivaraman, 2023; Naaraayanan, 2019)—and often bear greater responsibility for childcare and eldercare (Albanesi and Kim, 2021; Brody, 2003). These factors may expose them to elevated non-consumption burdens—such as emotional strain—during economic downturns.

We also uncover a hump-shaped pattern across the working life cycle: the relationship is strongest for users in mid-career and is neither economically nor statistically significant for younger and older individuals. These groups may be less exposed to the non-consumption effects of business cycles—for example, because they are students or retirees, or, in the case of older users, because they may have accumulated sufficient savings to buffer against adverse macroeconomic shocks.

Interestingly, we find that the relationship between business cycles and human costs is more pronounced among individuals in higher-ranked occupations, such as top leadership and senior management roles, compared to lower-ranked employees and manual workers. This finding aligns with the notion that high-level executives may face substantial non-monetary costs associated with job-related shocks, as documented by Borgschulte et al. (ming), whereas mid- and low-ranked employees experience weaker effects. Sorting users by income—which is highly correlated with occupational rank—yields a similar pattern of heterogeneity.

In the final part of the paper, we leverage our rich bank-account-level administrative data to ask a question of both research and policy relevance: Does access to finance mitigate or exacerbate the human costs of business cycles? Ex ante, the answer is ambiguous. On the one hand, consumers who accumulate debt during periods of negative or volatile business cycles may face additional burdens, such as anxiety, psychological stress, and worry, knowing they will have to manage interest payments and meet repayment deadlines in the future. These burdens may contribute to the human costs that excess age captures. On the other hand, access to finance allows consumers to smooth the adverse effects of business cycle fluctuations, potentially offering peace of mind and thereby reducing the human costs of downturns above and beyond its effect on limiting consumption costs.[1]

[1]Recall that in our setting, we observe users' income and occupational status, allowing us to control for the direct effects

We find that access to finance reduces the accumulation of excess age as local business cycles worsen at both the extensive margin (i.e., whether users have access to credit lines from the bank) and the intensive margin (i.e., the amount of debt users take on as a share of their income). Importantly, this attenuation of excess age during downturns does not occur in the case of medical debt, which suggests that excess age is not simply capturing untreated medical conditions during bad times.

These findings reveal an overlooked benefit of access to finance: it reduces not only the negative impact of economic fluctuations on consumption, but also the non-consumption-related human costs of business cycles, as captured by our measure of excess age. In turn, this result implies that denying credit access based on factors unrelated to actual repayment risk may have an additional discriminatory effect by unnecessarily exacerbating the human costs of business cycles for financially rationed consumers.

Our paper contributes to at least four areas of research in economics. First, we contribute to the research in empirical macroeconomics and applied microeconomics on assessing and estimating the costs consumers face from business cycle fluctuations. In a long-lived tradition, macroeconomists have tackled this question using standard models of business cycles (Lucas Jr, 1987), as well as incorporating refinements in otherwise standard representative-agent and rational-expectations models, which have consistently delivered negligible estimates of the consumption costs of business cycle fluctuations (Atkeson and Phelan, 1994; Barlevy, 2003, 2004; Imrohoroğlu, 1989; Imrohoroglu, 2008; Ramey and Ramey, 1991; Tallarini Jr, 2000). These results have led macroeconomists to argue that the welfare consequences of stabilization policies are minimal (Lucas Jr, 2003).

At the same time, these results and conclusions have long been at odds with survey-based evidence showing that consumers tend to perceive non-negligible declines in their well-being when exposed to business-cycle fluctuations (Davis and vonWachter, 2011; Deaton, 2011; Guerreiro et al., 2024; Wolfers, 2003). And, indeed, when eliciting consumers' willingness to pay to avoid business cycle fluctuations using surveys across many countries, Georgarakos et al. (2025) estimate it to be substantially higher than what would be implied by standard models. Moreover, Pfajfar and Winkler (2025) find that households' preferences for monetary policy tradeoffs do not align with what would be implied by them facing minimal costs of business cycles.

This evidence suggests that either consumers' preferences and beliefs depart from those assumed

of business cycle fluctuations on users' ability to consume.

by standard models, or consumers face costs that are not captured by standard economic variables, or both. We contribute to this debate by proposing an empirical setting in which to test whether consumers face any objectively measurable, non-standard costs when facing business cycle fluctuations—costs that elude standard micro- or macro-level variables used in economics. We do this by combining data on a large sample of consumers' actual and appearance ages with detailed administrative information on their economic characteristics from bank accounts. Our evidence suggests that consumers face overlooked costs of business cycles—costs that merit consideration not only in measurement and simulation but also in the foundations of economic theory—alongside the already ongoing introduction of non-standard preferences and beliefs in macroeconomics (Akerlof, 2002; Bordalo et al., 2020; De Grauwe, 2012; Gabaix, 2020; Hommes, 2021; Pfäuti and Seyrich, 2022).

Second, from a methodological perspective, our work contributes to the increasingly popular literature that adopts a micro-to-macro approach, whereby researchers use detailed administrative micro data at the consumer level to document the drivers of decision-making and belief formation, and to test directly for alternative microfoundations of representative- and heterogeneous-agent macroeconomic models (Kaplan and Violante, 2018). In particular, we add to contributions in this area that study households' consumption, investment, and liability management decisions, including, among others, Aguiar et al. (2013); Bornstein and Indarte (2023); D'Acunto et al. (2024); Fagereng et al. (2020); Fuster et al. (2021); Guiso et al. (1996); Indarte et al. (2025); Kaplan et al. (2014)

Third, we contribute to a growing literature in economics and finance that leverages alternative data to measure outcomes typically unobserved in administrative sources and to test economic mechanisms more directly. This research harnesses nontraditional sources—such as satellite imagery, web-scraped prices, social networks, and mobile metadata—to capture latent dimensions of economic activity. For instance, satellite night light intensity has been used to proxy for subnational growth in regions lacking reliable GDP data (Henderson et al., 2012; Michalopoulos and Papaioannou, 2014), while satellite-based deforestation measures have helped identify the political economy of environmental degradation (Burgess et al., 2012). Web-scraped price data have refined inflation measurement (Cavallo, 2017), and search volume indices and social media content have been shown to predict investor behavior and macroeconomic sentiment (Baker et al., 2016; Da et al., 2011). Some studies combine such alternative measures with quasi-experimental designs or instrumental variables—for example, using Facebook friendship networks to estimate peer effects in home-buying (Bailey et al.,

2018). A common thread in this literature is the ability of novel data to uncover heterogeneity, dynamics, and subjective burdens—such as attention, sentiment, or social influence—that traditional economic metrics cannot measure.

Our paper extends this tradition by introducing an image-based proxy—excess age—for the non-monetary costs of macroeconomic fluctuations. In doing so, we build on recent work using facial appearance as a field outcome (Borgschulte et al., ming; Guenzel et al., 2025), and expand the scope of alternative data in economics. Just as night lights reveal economic growth, and search data capture attention, facial aging offers a scalable, objectively measurable indicator of the human costs of business cycles—broadening the economist's toolkit to include dimensions of well-being long considered unobservable.

Finally, our results on the mitigating role of access to finance in reducing the human costs of business cycles speak to the extensive literature on the causes and consequences of access to finance and credit rationing in consumer lending markets. Theoretically, collateral and liquidity constraints can amplify the effects of macroeconomic shocks by restricting households' ability to smooth consumption (Kiyotaki and Moore, 1997). Empirically, easing access to credit—especially in developing and financially underserved contexts—has been shown to improve consumption smoothing and support microenterprise formation (Banerjee and Duflo, 2010; Townsend, 1994). Importantly, however, access to finance is not universally granted and is often shaped by cultural biases, structural exclusion, and discrimination. A large body of research documents that even after controlling for observable credit risk, minority and marginalized borrowers are more likely to be denied credit or offered worse terms. For example, Bayer et al. (2018) find persistent racial disparities in mortgage pricing, while Bartlett et al. (2022) use algorithmic models to document large gaps in mortgage approval and pricing by race, even within fintech lenders. Fuster et al. (2022) further demonstrate that algorithmic credit scoring can entrench racial biases, as digital lenders exhibit discriminatory behavior despite claims of neutrality.

Our paper contributes to this literature by introducing a novel, objectively measurable proxy—excess age—for the non-consumption-related costs of business cycles. We show that access to formal financial products substantially mitigates the rise in excess age during downturns, holding constant income, employment, and health shocks. This suggests that finance helps not only by smoothing financial constraints, but by alleviating psychological and physiological burdens unobserved in stan-

dard metrics. By highlighting this channel, we expand the welfare implications of financial access and highlight the potentially compounding costs of exclusion and discrimination in lending markets.

2 Institutional Setting and Data

We collaborate with a leading Indian digital lending platform that provides short-term, unsecured personal loans through a mobile app. Founded in 2016, the platform operates as a non-bank financial company and primarily serves urban, salaried millennials and Gen Z professionals—many of whom lack formal credit histories. To underwrite these borrowers, it employs a proprietary credit scoring algorithm that uses hundreds of covariates, including smartphone metadata, education, income, employment history, and social media activity.

Loan applications are initiated and completed through the platform's mobile app, which supports a fully digital onboarding process. As part of the required Know Your Customer (KYC) protocol, users must upload government-issued identification documents and take a live selfie using their smartphone. These self-portraits—typically one per user—are timestamped and stored in the platform's system. A minority of users (approximately 10%) later update their profile photo, resulting in two images on record. Because these photos are generally captured during the loan application process and are time-stamped, we can align them precisely with each user's financial status and the prevailing macroeconomic conditions at that point in time.

Our empirical analysis draws on three distinct samples from the company's user base between 2016 and 2023. The first, which we refer to as the *Full Sample*, comprises approximately 1 million users who uploaded at least one photo alongside loan application data. Each observation includes a timestamped image, geolocation, and linked administrative account information. By design, all users in this sample are smartphone owners with digital literacy and an active demand for credit.

The second sample, the *Two-Image Sample*, includes approximately 93,000 users for whom the company recorded two distinct photos taken at different times. This feature enables us to track within-individual variation in appearance across different phases of the business cycle.

Finally, for a smaller subset of approximately 67,000 users, we also observe self-reported demographic characteristics, including gender, income, education, and occupation. While these variables are not required to use the platform, they enrich our ability to assess heterogeneity in outcomes across

observable characteristics.

Although the sample is not representative of the broader Indian population, it captures a key segment of digitally connected, credit-seeking, working-age consumers concentrated in urban and semi-urban areas. The company's paperless, mobile-first infrastructure ensures that all user images are collected under reasonably uniform conditions—typically using smartphone cameras during loan onboarding—though image quality may still vary.

Each user photo is linked to that individual's date of birth, financial data (such as income, loan approvals, and debt levels), and local economic conditions proxied by state-level unemployment rates in the months leading up to the image timestamp. This rich, high-frequency dataset allows us to examine whether individuals appear to age more—or less—than expected in response to local macroeconomic fluctuations, using predicted appearance age as a proxy for the human costs of business cycles.

2.1. Measuring Excess Age

To quantify the visual impact of macroeconomic stress, we construct a measure of excess age—the difference between a user's predicted appearance age (based on facial features in a selfie) and their calendar age (based on date of birth). This measure builds on recent advances in computer vision and leverages the DeepFace library in Python, a widely used framework for facial attribute analysis.[2]

DeepFace is an open-source facial recognition and analysis library that integrates all components of the facial analysis pipeline—face detection, alignment, normalization, feature extraction, and classification—into a streamlined interface. Age prediction, in particular, is performed using deep convolutional neural networks trained via transfer learning atop the VGG-Face architecture. In benchmark evaluations, the age estimation module achieves a mean absolute error (MAE) of approximately ± 4.6 years,[3] making it suitable for detecting systematic deviations in appearance across economic environments.

Each user selfie is analyzed through a standardized facial analysis pipeline to estimate appearance-based age. The pipeline proceeds in two stages. In the primary pass, the image is scanned for detectable faces using a selected backend (e.g., RetinaFace or MTCNN). If a face is found, DeepFace

[2]We use Deepface version 0.0.81 and modify it to allow for age predictions at the two-digit decimal level. Please refer to the following link for additional details: `https://github.com/serengil/deepface`.

[3]See DeepFace documentation: `https://github.com/serengil/deepface`. The MAE reflects the average deviation between predicted and actual ages on benchmark datasets.

predicts the age using enforced detection. If no face is detected—often due to poor image quality, lighting, or pose—the image is rotated up to four times, and DeepFace attempts detection without enforcement, allowing estimation under looser constraints. Each image is tagged with metadata about detection quality, estimated appearance age, and whether fallback procedures were triggered. In cases where multiple faces are detected, we retain only the most prominent one (typically the largest and most centered).

To mitigate heterogeneity in selfie quality—such as lighting, angle, and resolution—DeepFace applies a set of automatic preprocessing steps prior to age estimation.[4] Alignment significantly improves accuracy by correcting for facial geometry (e.g., tilt, asymmetry), while normalization reduces visual noise from lighting and contrast variation.[5]

To illustrate the concept of excess age, we present two example images below. Both images show men who were 41 years old at the time the photos were taken. As noted above, the original images have been cropped to the same size for consistency. Our procedure predicts different appearance ages: 44 years for the subject on the left and 38 years for the subject on the right. Despite having the same calendar age, this results in excess age values of $+3$ and -3 years, respectively.



Calendar Age: 41
Predicted Age: 44
Excess Age: +3



Calendar Age: 41
Predicted Age: 38
Excess Age: -3

This example underscores the distinction between calendar age and appearance-based age. While such stark contrasts are not typical in our sample, they highlight how facial features—such as groom-

[4]The DeepFace pipeline automatically performs (i) face alignment using facial landmarks (e.g., eyes, nose, mouth) to correct for tilt and angle; (ii) normalization by resizing images to a standard 224×224 pixel input and adjusting color and brightness to stabilize pixel distributions across inputs. See: `https://github.com/serengil/deepface`.

[5]By default, DeepFace avoids geometric distortion by padding with black borders to preserve facial aspect ratio. It also rescales pixel intensities to match the expectations of the VGG-based CNN.

ing, skin tone, facial tension, and symmetry—influence the model's predictions. The individual on the right, who appears younger, has fuller hair, a well-kept beard, even skin tone, and a relaxed expression. In contrast, the individual on the left presents with visible under-eye shadows, a strained expression, and graying hair, all of which likely contribute to a higher predicted age.

At the same time, it is important to recognize that some variation in predicted appearance age reflects not intrinsic facial characteristics, but technical features of the selfie itself. These images are self-portraits taken during the digital loan application process, under non-standardized conditions that introduce heterogeneity in lighting, pose, and angle. For instance, the subject on the left is photographed under cooler lighting and from a slightly lower angle, exaggerating certain facial features; the subject on the right is photographed frontally, with warmer lighting that softens contrast and improves symmetry. These differences can subtly affect the model's predictions—even when both users look directly into the camera with similarly open, centered eyes.

Although DeepFace's preprocessing is explicitly designed to reduce such noise—and generally improves robustness—residual sensitivity to image quality remains. According to the DeepFace documentation, poor lighting, unusual angles, or visual artifacts may still influence the model's output when facial features are not cleanly resolved. Nevertheless, the combination of rigorous preprocessing and two-stage detection ensures a high degree of consistency and reliability in estimating appearance-based age across large samples.

As a further validation of the procedure, we construct "pseudo-individuals" by pooling observations across users who share the same cohort, gender, education, and occupation. This design averages out the idiosyncratic variation that individual users may exhibit in appearance age due to transient factors such as health, stress, grooming, or image quality, and instead isolates systematic patterns that should characterize narrowly defined demographic groups. Tracking these pseudo-individuals across years allows us to test whether predicted appearance ages evolve consistently with calendar age.

Figure 2 illustrates this validation exercise for one pseudo-individual: males born between 1990 and 1992, with a high school degree, employed as manual workers. Three features emerge. First, in both years considered, 2019 and 2022, the modal predicted age closely aligns with the actual calendar age of the cohort, confirming that the procedure captures demographic regularities—in this case, calendar age. Second, the distribution shifts rightward between 2019 and 2022 in line with the natural aging of the cohort, while the overall shape remains stable. Third, the dispersion of predicted ages is nearly

identical across years, indicating that the model yields consistent estimates for the same demographic profile over time.

Together, these findings provide reassurance that our measure of excess age is not driven by noise or technical artifacts, but rather reflects persistent demographic patterns at the group level. This validation supports its use as a reliable proxy for studying the human costs of business cycles.

2.2. Samples and Summary Statistics

Our empirical analysis is based on three samples. First is the *Full Sample*, which includes all users for whom the bank has one or two pictures on record, accompanied by geolocation and timestamp information indicating when each picture was taken. Recall that these pictures are self-portraits that users are required to take and upload to their accounts as part of the loan application process. This method of sourcing images inevitably introduces selection in ways that may cause the sample to differ from the broader Indian population. First, because the company operates entirely through digital channels, all included users must either own a personal device capable of taking photographs and connecting to the internet, or have access to one through acquaintances. Given the rapid and widespread adoption of smartphones and digital financial services in India following the demonetization shock of November 2016 (Crouzet et al., 2023)—which aligns with the beginning of our sample period—we do not believe this technological filter produces substantial divergence from the average Indian urban population. Second, all users are active loan applicants at the time their picture is taken, which ensures in-sample homogeneity in credit demand but may also introduce selection on unobservable characteristics—such as financial stress, credit constraints, or risk preferences—that distinguish individuals in our sample from the broader population.

Our second sample, the *Two-Image Sample*, is a subset of the "Full Sample" comprising users for whom the company recorded two pictures rather than just one, enabling us to conduct within-individual analyses by comparing the excess age of the same person across different business cycles. The potential sources of selection in this sample are the same as for the "Full Sample," with the additional caveat that these users chose to update their profile picture on the platform at some point, while others did not. Although we do not have any compelling hypotheses about how this feature might introduce selection related to the outcomes and choices we study in this paper, this additional caveat warrants acknowledgment.

Third, we consider the *Full Sample with Demographics*, which is the subsample of the *Full Sample* for which we observe the full set of demographics the company collects in their user account profile. In addition to the selection criteria to enter the *Full Sample*, consumers in this third sample decided to provide their full set of demographics on the account profile. It is important to note that, although encouraged, providing such information is not required to open an account. While application timing and user location are recorded via geolocation services, demographic attributes—such as gender, occupational status, and job rank—are entirely self-reported. About one in ten users include this information on their profiles. Once again, although the precise ways in which this additional selection may correlate with key outcomes are unclear, the restriction should be kept in mind when interpreting results.

Table 1 reports summary statistics for the main variables that enter our analysis for each of the three samples. Starting with the *Full Sample*, we note that, although largely overlapping, the distribution of predicted appearance ages is slightly shifted to the right relative to the distribution of actual calendar ages, with a sample mean of 30 years versus 29 years. Calendar age also displays greater dispersion, with a standard deviation of 7.02 years versus 4.52 years for predicted age. The most pronounced difference between the two distributions is the heavier lower tail of calendar age, whereas their 75th percentiles are nearly identical.

Turning to local economic conditions, we find that lagged unemployment rates—measured over one-, three-, and six-month horizons—display substantial variation across users. Regardless of the horizon, the standard deviations are nearly as large as the means, and the interquartile ranges span values consistent with both near-full employment and pronounced economic downturns. This high degree of dispersion in local business cycle exposure—across both geography and time—is a distinguishing feature of the Indian setting and is central to our empirical strategy.

Finally, we consider two outcomes related to access to finance: the extensive margin—whether a user's loan application is approved—and the intensive margin—the debt-to-income ratio for those who obtain credit. About 16% of users are approved for a loan, while the average debt-to-income ratio among borrowers is 0.96, indicating that, on average, users borrow nearly their full monthly income. The interquartile range of this ratio exceeds 80 percentage points, and a substantial share of users carry debt loads well above their monthly income, reflecting significant heterogeneity in borrowing

intensity.[6]

Moving on to the *Two-Image Sample*, Table 1 shows that we observe two images for approximately 93K users. The predicted appearance age in this group averages 29.7 years, closely aligned with their calendar age of 29.2 years. The distributional properties remain similar to the *Full Sample*, with modest shifts that are economically insignificant. Similarly, while the average lagged unemployment rates are slightly lower across all horizons (e.g., 6.89% over the prior month versus 7.25% in the *Full Sample*), the standard deviations and interquartile ranges remain sizable—ranging from near full employment to rates exceeding 8%—preserving the substantial variation crucial to our empirical strategy.

In terms of financial access, this subsample exhibits a notably higher loan approval rate, with 24% of users obtaining credit compared to 16% in the *Full Sample*. For those approved, average debt-to-income ratios are slightly lower (0.88 versus 0.96), with an interquartile range of 0.68, suggesting substantial heterogeneity in borrowing intensity but slightly more conservative leverage among this group. This pattern is consistent with a more creditworthy or financially stable profile among users who engage with the platform repeatedly.

Finally, we provide statistics for the *Full Sample with Demographics*. First and foremost, we note that the subsample of users who fill the demographic information fully in their profile is, unfortunately, a small minority of the *Full Sample*—it represents about 7.4% of the full set of users whose portraits we observe. This limited size constrains its use in our core empirical analyses, which are instead based on the broader *Full Sample* and the larger *Two-Image Sample*. The latter, combined with the fact that several demographic attributes are time-invariant, allows us to control for individual fixed effects in panel regressions and absorb unobserved heterogeneity more effectively.

In terms of age and local unemployment rates, the statistics for this subsample are broadly similar to those already discussed: average predicted appearance age is 30.2 years, while calendar age is 30.5 years; unemployment rates average around 6.5% to 6.9% depending on the horizon, with substantial variation across individuals.

This sample does, however, allow us to examine the socioeconomic composition of the user base. First, 88% of users identify as male, consistent with broader evidence that digital finance platforms tend to over-represent men, particularly in developing country contexts. The average log monthly in-

[6]Debt-to-income is computed only for users whose loan application is approved.

come is 11.0, with a standard deviation of 0.97—less than 9% of the mean—indicating limited income dispersion within this group. Roughly 18% of users report a college degree, and 62% have completed high school, implying that approximately 80% have attained at least secondary education. This is a much higher educational attainment rate than in the general Indian population: in 2022, approximately 13% of Indian adults held a college degree and about 51% had completed upper secondary school (OECD, 2022).

Finally, our data allow observing users' occupation rank, which is self reported and selected by users from a four-category multiple choice list: (i) Employee, which includes workers that have applied jobs and work under somebody else's direction without any managerial responsibilities and with nobody reporting to them; Middle Management—we label users in this group "Managerial Role" users; Senior Management and Leadership, which we pool into a single category labeled "Leadership Role." Overall, we find that a bit less than half of the sample belongs to the lowest category, whereas 25% self-categorize themselves as middle management and 30% as having leadership roles. This occupational distribution, consistent with patterns in education and income, suggests that our sample is positively selected toward more educated and professionally senior individuals relative to the broader Indian population.

3 Excess Age and Business Cycles

We begin by examining the univariate relationship between users' excess age and the average local unemployment rate in the location where each user resides, measured over different windows. Figure 3 illustrates this relationship. Panel A plots excess age against the unemployment rate in the month immediately preceding the date the picture was taken. Panel B uses the three-month average of local unemployment prior to the image, while Panel C replicates the six-month horizon already shown in the Introduction, included again here to facilitate comparison across timeframes. Presenting all three panels side by side allows us to assess the stability of the relationship between macroeconomic stress and facial aging across different business cycle horizons.

As discussed in Figure 1 in the Introduction, the univariate analyses across different horizons reveal three remarkably consistent baseline patterns. First, excess age increases with the local unemployment rate at the time and place each user's image is recorded. This finding suggests that, in cross-section,

users exposed to worse local business cycle conditions tend to appear older than their actual age—relative to peers who experienced more favorable economic environments.

Second, the relationship between excess age and local unemployment is nonlinear and exhibits increasing marginal effects: the worse the business cycle, the larger the gap between predicted appearance age and actual age, relative to users facing more favorable economic conditions. This nonlinearity becomes especially pronounced as the persistence of local unemployment increases. While the relationship is approximately linear when using the previous month's unemployment rate, it becomes markedly nonlinear when considering the six-month average. In other words, sustained exposure to elevated unemployment—rather than short-term fluctuations—is strongly associated with greater signs of accelerated aging among users.

Third, excess age remains positive on average even during periods of strong local economic conditions—when unemployment rates are low and the economy is near full employment. Moreover, the relationship between business cycles and excess age tends to flatten in good times, indicating diminishing returns to improved economic conditions. This pattern suggests that while adverse business cycles are associated with accelerated aging in appearance, the potential for "rejuvenation" during favorable periods is limited. That is, moving leftward along the unemployment distribution—toward better economic conditions—does not yield proportionally lower excess age.

Because Figure 3 plots average values within groups of users, it cannot provide insights into the distributional properties of the relationship between local business cycles and excess age. A key open question is whether this relationship is driven primarily by a small subset of users experiencing pronounced aging during economic downturns, or whether it reflects broader shifts across the entire distribution. To investigate this, Figure 4 overlays the full density distributions of excess age—computed as the residuals from a regression of appearance age on actual calendar age—for two subsamples: users in the bottom quartile of the local unemployment rate distribution (solid lines) and users in the top quartile (dashed lines). Across all panels and time horizons, we observe a clear rightward shift in the distribution for users exposed to higher unemployment, indicating that the relationship is not concentrated in the tail, but instead holds across the full distribution.

This result reveals that the baseline association is not driven by outliers or the tails of the excess age distribution. Instead, facing worse business cycles is associated with higher excess age for the majority of users, irrespective of the initial level of their excess age.

The widespread nature of this relationship has two important implications. First, it suggests that the human costs of business cycles affect a large part of the population—not just a small subset of especially vulnerable or psychologically sensitive individuals. This implies that these costs should be modeled as a general feature of consumer behavior, rather than as rare or idiosyncratic exceptions, and thus deserve inclusion in standard macroeconomic frameworks, including representative-agent models.

Second, the fact that the effect is pervasive rather than concentrated implies that policy responses to economic downturns—such as forward guidance or optimism-inducing communication—should be directed at the population as a whole, not only at traditionally vulnerable groups like the unemployed or financially fragile. This departs from the typical rationale behind transfer-based fiscal policies and underscores the potential value of broader, population-wide interventions aimed at alleviating the widespread psychological and physiological toll of recessions.

3.1. Multivariate Analysis

After documenting the univariate relationship between local unemployment rates and consumers' excess age, we now leverage the rich administrative data from users' bank accounts to test the robustness of this relationship and explore its underlying mechanisms in a multivariate setting.

Our baseline multivariate analysis is based on variations of the following most parsimonious linear specification:

$$\begin{aligned} \textit{Predicted Age}_{i,t,s} = \alpha + \beta \textit{Calendar Age}_{i,t,s} + \gamma \textit{Avg. Unemployment Rate}_{t-k \to t-1,s} \\ + \eta_t + \eta_s + \eta_i + \epsilon_{i,t,s}, \end{aligned} \tag{1}$$

where $\textit{Predicted Age}_{i,s,t}$ is the predicted appearance age of consumer i whose image is taken in state s and time t; $\textit{Calendar Age}_{i,s,t}$ is the actual age of consumer i in state s and time t based on their date of birth; $\textit{Avg. Unemployment Rate}_{t-k \to t-1,s}$ is the average lagged unemployment rate in state s, where k=2 captures the lagged unemployment rate one month before i's image is observed (Panel A), k=4 captures the 3-month average lagged unemployment rate before i's image is observed (Panel B), and k=7 captures the 6-month average lagged unemployment rate before i's image is observed (Panel C); η_s, η_i, and η_t capture full set of Indian state, individual, and year fixed effects. The coefficient γ captures the marginal effect of local macroeconomic conditions on predicted age. A positive and

significant γ indicates that individuals exposed to higher unemployment in their region tend to look older than their actual age, even after controlling for observable characteristics and fixed effects. In other words, γ quantifies the average appearance-based aging penalty associated with economic downturns.

In our baseline specification, we use predicted appearance age as the dependent variable while controlling directly for individuals' actual calendar age, rather than using excess age (the difference between the two) as the outcome. This approach allows us to estimate the correlation between appearance and actual age explicitly, and to isolate how local macroeconomic conditions correlate with predicted aging after accounting for chronological age. For statistical inference, we cluster standard errors at both the individual and year-month levels. This double-clustering strategy accounts for potential autocorrelation in repeated observations of the same individuals over time, as well as for common shocks affecting all users in a given period—such as the COVID-19 pandemic or other time-varying national events.

Table 2 report the results for estimating several versions of equation (*3.1.*). Before discussing our main coefficients of interest, note that the estimate of the conditional correlation between calendar age and predicted appearance age is positive but substantially lower than 1, ranging between 22% and 30% across specifications and confirming the fact that predicted ages are not noise but at the same time substantial variation exists in that dimension relative to individuals' calendar age, which provides scope for our analysis.

Our coefficient of interest is γ, which captures the association between lagged local unemployment rates in the states in which individuals reside and their predicted appearance age, after partialling out calendar age and other characteristics across specifications. Like in the univariate analysis, we consider three horizons for the computation of average local unemployment rates—1 month (Panel A), 3 months (Panel B), and 6 months (Panel C). To allow comparison across panels, we standardize the unemployment rates computed at different horizons to have a mean of zero and a standard deviation of 1.

Column (1) shows that the baseline association between lagged average unemployment rates and predicted age—controlling for actual age—is consistently between 0.07 and 0.08. In other words, individuals exposed to a one-standard-deviation higher local unemployment rate appear 0.07–0.08 years older than others of the same calendar age who experienced more favorable labor market conditions

over the same period. This effect is economically meaningful, especially considering that the relationship appears markedly flatter during periods of low unemployment—suggesting that the perceived aging effect is more pronounced in downturns than in recoveries, as we further illustrate in Figure 3. Moreover, consistent with the univariate patterns, the relationship may be nonlinear across the unemployment distribution, with disproportionately larger effects for individuals experiencing deeper local recessions. We investigate these nonlinearities below.

Columns (2)–(4) assess the robustness of our baseline estimates by progressively tightening the sources of identifying variation. Column (2) includes state fixed effects to absorb time-invariant characteristics shared by all consumers within a state—such as labor market institutions, cultural norms, or social networks defined along racial, ethnic, or religious lines. Column (3) introduces year fixed effects to account for nationwide time-varying shocks, including the COVID-19 pandemic and broader macroeconomic cycles. Column (4) includes both sets of fixed effects simultaneously. Across all specifications, the estimated coefficients on local unemployment remain economically meaningful and statistically significant, underscoring the robustness of the relationship.

3.2. Within-Individual Analysis and Demographics

The evidence reported so far relies on cross-sectional variation in predicted appearance age across individuals who share the same calendar age but were exposed to different local business cycles. However, our sample also includes a subset of users for whom we observe two images taken at different points in time. This allows us to conduct a within-individual analysis, comparing how predicted appearance age varies for the same person under different macroeconomic conditions, while controlling for changes in calendar age.

This analysis is important because it allows us to test the robustness of our baseline results to unobserved, time-invariant consumer characteristics—such as cognitive ability, stable time and risk preferences, human capital, and inherited wealth—that are either unavailable for the full sample or entirely unmeasured.

Columns (5)-(6) of Table 2 report the results for this within-individual analysis. Before commenting on the results, note that the subsample of users for whom the bank observes two pictures taken at different points in time includes 92,629 consumers ($92{,}629 \times 2 = 185{,}258$), which is one order of magnitude smaller than our baseline sample. This limitation is the primary reason we do not rely

solely on the within-individual design, which necessarily focuses on a specific subset of users who may differ systematically from the broader population of account holders in unobservable ways.

The within-individual estimates closely mirror the results obtained from the full cross-sectional sample, both in magnitude and direction. In fact, the estimated association between lagged local unemployment rates and predicted appearance age is slightly larger in this subsample, ranging from 0.07 to 0.13 years. We can always reject the null that the estimate equals zero at any plausible level of significance.

The within-individual analysis alleviates concerns that observed or unobserved time-invariant individual characteristics—such as preferences, cognitive ability, or socioeconomic background—may be correlated with local business cycles and spuriously drive the relationship between unemployment and predicted appearance age. However, this approach does not account for time-varying individual-level factors, which could still act as confounders in our baseline analysis.

To address this concern, we turn to a third subsample of the data for which we observe a set of individual-level demographic characteristics. This subsample is significantly smaller than the main sample ($N = 67,682$) and potentially selected, which is why we do not rely on it for our primary analysis. However, it allows us to control for both time-invariant and time-varying characteristics across individuals. For completeness, we report estimates from specifications that include all available covariates—even though, as noted above, the within-individual analysis already accounts for time-invariant heterogeneity.

In Table A.1 of the Online Appendix, we report results from estimating versions of equation (*3.1.*) that include the full set of available individual-level observables. Importantly, this includes time-varying variables such as consumers' income—arguably the most direct proxy for shocks to consumption capacity during business cycles—and employment status, which captures the user's exposure to labor market fluctuations. We also observe time-invariant characteristics such as gender and education level.

Across the board, the estimates obtained from this subsample—controlling directly for individual-level characteristics—are very similar, both qualitatively and quantitatively, to the cross-sectional results reported in columns (1)–(4) of Table 2. Notably, our baseline estimates remain virtually unchanged even after controlling for the logarithm of income and employment status. We interpret this as evidence that the variation in predicted appearance age beyond calendar age is unlikely to reflect

the direct consumption effects of business cycles, and instead may capture non-monetary consequences such as stress, uncertainty, or psychological burden.

3.3. Stable vs. Volatile Business Cycles

Thus far, our analysis has focused on documenting the relationship between the level of local business cycles—captured by lagged local unemployment rates—and consumers' excess age. However, economic theory suggests that volatility in business cycle conditions is an equally important driver of the welfare costs associated with macroeconomic fluctuations. In this section, we extend the analysis by comparing the roles of both the average level and the volatility of local unemployment in shaping predicted appearance age.

To capture volatility, we focus on lagged 6-month local business cycles and compute the standard deviation of monthly unemployment rates over this window, alongside the average unemployment rate used in the previous analysis. To facilitate interpretation of both level and interaction coefficients—and to explore potential nonlinearities—we construct three indicator variables corresponding to the bottom, middle, and top terciles of the volatility distribution. The variation across these groups is substantial: the average standard deviation of unemployment ranges from 0.6 percentage points in the bottom tercile to 4.85 percentage points in the top tercile.

Table 3 reports results from estimating a modified version of equation (*3.1.*) that includes indicator variables for being in the middle or top tercile of the volatility distribution, as well as their interactions with the average six-month lagged unemployment rate. Each column replicates the specifications from Table 2, with these additional terms included. In this setup, the coefficient on the average unemployment rate reflects the estimated association for consumers exposed to the least volatile business cycles—those in the bottom tercile, which serves as the omitted reference group.

The results on business cycle volatility yield two key insights. First, when examining the coefficients on the indicators for medium and high volatility, we find that consumers exposed to more volatile local business cycles have significantly higher predicted appearance age relative to those in low-volatility environments—even after controlling for the average level of unemployment. As shown in Table 3, this pattern is consistent and robust across all specifications, including both cross-sectional and within-individual analyses. These findings suggest that volatility itself—independent of average economic conditions—contributes to excess aging, even when comparing the same individuals at different points

in time.

The second key insight from Table 3 concerns the interaction terms between the average six-month lagged unemployment rate and the indicators for business cycle volatility. Focusing first on the cross-sectional results in columns (1)–(4), we find that volatility plays a critical role in amplifying the relationship between unemployment and predicted appearance age. Among consumers exposed to low-volatility business cycles, higher average unemployment rates are not associated with increased appearance-based aging. In contrast, the positive baseline relationship documented in Table 2 is entirely driven by consumers facing moderate to high volatility. That is, unemployment leads to a higher predicted age only when accompanied by unstable labor market conditions.

The importance of business cycle volatility in amplifying the effects of unemployment on excess aging is further confirmed by the within-individual results in columns (5)–(6) of Table 3, both qualitatively and quantitatively. These specifications also show that exposure to volatile business cycles is positively associated with predicted appearance age, regardless of the average unemployment rate faced over the preceding six months.

3.4. Heterogeneity: Life Cycle, Gender, Occupations, and Income

Our baseline multivariate results may mask important sources of heterogeneity in the sensitivity of excess aging to business cycles. One key dimension is consumers' age: individuals at different stages of their working life may experience the human costs of macroeconomic fluctuations differently. Younger consumers, for example, may be more adaptable in responding to downturns, face fewer financial or caregiving responsibilities, and in many cases may still be studying—making them less directly affected by adverse labor market conditions.

Based on these considerations, Table 4 examines whether the relationship between adverse business cycles and excess aging is attenuated among younger consumers. We construct an indicator variable equal to one for individuals aged 25 or younger and interact it with the average lagged local unemployment rate. While our results are robust to alternative age thresholds, we use 25 as a natural cutoff to distinguish college-aged individuals and young workers from the broader sample. Columns (1)–(3) of Table 4 show that younger consumers not only have lower predicted appearance age conditional on calendar age, but are also less sensitive to increases in local unemployment. Specifically, the estimated association between average unemployment and predicted age is between 13% and 29% smaller for

consumers under age 25, depending on the specification.

The binary age split in Table 4 may mask important nonlinearities in how the human costs of business cycles vary across the life cycle. In particular, older consumers may also be less sensitive to local economic conditions—though for reasons distinct from those affecting younger individuals. Many older consumers have accumulated stable sources of income or wealth, or may be retired and thus receive predictable pension payments that are insulated from labor market fluctuations. If anything, recessions may benefit retirees by reducing prices for goods and services, thereby increasing their real purchasing power.

To assess potential nonlinearities in the relationship between age and the human costs of business cycles, we estimate a version of equation (*3.1.*) that includes interactions between consumers' average local lagged unemployment rate and six age group indicators: under 20, 20–24, 25–29, 30–34, 35–39, and 40 or older. Results are robust to alternative age thresholds. To facilitate interpretation, we plot the estimated interaction coefficients in Panel A of Figure 5, using the specification that includes the full sets of state and year fixed effects. The results reveal a clear hump-shaped relationship: the association between unemployment and predicted appearance age is statistically insignificant for the youngest and oldest consumers, but strong and positive for individuals in the middle of the age distribution—particularly those in the early to mid-stages of their working lives.

The second demographic dimension we examine for heterogeneous effects is gender. This focus is motivated by two key considerations. First, in the Indian context, women are generally more economically vulnerable than men due to persistent income and wealth gaps (Blau and Kahn, 2017), weaker bargaining power in the labor market (Card et al., 2016), limited access to business ownership and independent asset accumulation (Naaraayanan, 2019), and disproportionate responsibility for managing household finances and caregiving for children and the elderly (Kleven et al., 2019). These factors can increase the psychological toll women face during economic downturns, when such roles and resources are more likely to be strained. Second, a large body of research in psychology and neuroscience finds that women are, on average, more sensitive than men to negative shocks in terms of mental health outcomes Kessler et al. (1993). Taken together, these considerations lead us to test whether women experience greater excess aging than men following exposure to adverse business cycles, holding other factors constant.

Columns (4)–(6) of Table 4 present results from extending equation (*3.1.*) to include an interaction

between users' gender and the average lagged local unemployment rate. As noted earlier, demographic information is only available for a subset of bank users, which reduces the sample size for this analysis by roughly an order of magnitude relative to the main specification.

First, we find that—unconditionally—men exhibit substantially lower excess age than women. The estimated intercept for male users at a hypothetical 0% local unemployment rate is approximately -0.7 years. This baseline gender gap is sizable: for comparison, it is equivalent to the difference in predicted appearance age between same-gender consumers exposed to local unemployment rates that differ by six standard deviations. Assuming no systematic gender bias in the model predicting appearance age—a view supported by the literature (**?**)—this result is consistent with the interpretation that women face higher average human costs in daily life than men, even in the absence of adverse macroeconomic conditions.

Second, turning to conditional effects, we find that the sensitivity of predicted appearance age to local business cycle conditions is significantly higher among women than among men. Across specifications, the estimated effect of unemployment on predicted age is between 22% and 60% lower for men than for women, as captured by the coefficient on the interaction between gender and the average lagged local unemployment rate. This pattern holds even when comparing men and women in the same geographic locations and time periods, including during large aggregate shocks such as the COVID-19 pandemic. While the statistical significance of this result varies across specifications—likely due to limited power from the smaller subsample with gender data—the point estimates consistently support the view that women bear greater human costs in response to negative business cycles.

We next examine whether the baseline association between business cycles and excess aging varies across occupational ranks. The bank provides an internal classification of users into four categories: Executive (40% of the sample), referring to lower-ranked employees without managerial responsibilities; Middle Management (26%); Senior Management (25%); and Leadership (9%). Ex ante, the direction of heterogeneity across these groups is unclear. On the one hand, lower-ranked workers may be more vulnerable to the economic and psychological consequences of downturns due to greater income instability, limited financial buffers, and increased exposure to job market competition. On the other hand, prior work by Borgschulte et al. (ming) documents significant appearance-based aging among CEOs—one of the highest-ranked occupational groups—around major corporate events such as mergers and acquisitions. While that study does not examine lower-ranked workers, it highlights

that individuals in high-responsibility roles may experience acute stress and visible aging as a result of professional pressure and accountability. This motivates our empirical exploration across the occupational hierarchy.

Columns (7)–(9) of Table 4 report results from interacting average lagged local unemployment rates with an indicator variable for users in the top two occupational categories—Senior Management and Leadership. Across specifications, we find a consistent pattern: higher-ranked employees are substantially more sensitive to local unemployment fluctuations in terms of excess aging. The additional effect for this group is as large as the baseline estimate for the full sample—approximately 0.06 to 0.07 years. In contrast, the effect for lower-ranked workers is smaller in magnitude and statistically indistinguishable from zero across all specifications. Notably, in stable periods (captured by the intercept for the high-rank dummy at a hypothetical zero unemployment rate), excess aging does not differ across occupational groups. This reinforces the interpretation that higher-ranked individuals bear greater human costs during adverse business cycles. These findings underscore the relevance of occupational status in shaping vulnerability to macroeconomic conditions, with potential implications for targeted workplace and health interventions.

Finally, we examine heterogeneity based on users' income. We define a dummy variable, *High Income*, equal to 1 for individuals in the top income quintile and 0 otherwise. While the threshold is arbitrary, our results are robust to alternative definitions. Overall, we find that high-income users exhibit substantially greater sensitivity to local business cycles: the estimated effect is between two and seven times larger than that for the rest of the sample, depending on the specification. At the same time, in stable periods (i.e., low unemployment), high-income individuals tend to exhibit lower excess age relative to others. This suggests that, despite their greater sensitivity to macroeconomic shocks, the absolute level of human costs faced by high- and lower-income users might be close to each other.

The estimated coefficients for high-income users are consistently larger than those for the rest of the sample, though we also find economically and statistically significant effects for other income groups. To explore this heterogeneity further, Panel B of Figure 5 plots the interaction coefficients between average local unemployment and five income quintiles, using the specification with full state and year fixed effects. The results confirm a pronounced effect for the top quintile. Among the remaining four quintiles, the estimated coefficients are relatively similar in magnitude and significance, with slightly

larger values for those in the lowest income group—though these differences are not statistically significant.

Taken together, these findings suggest that the association between excess age and business cycles is broadly similar across most of the income distribution, but rises sharply among top earners—highlighting that exposure to economic stress may not decline monotonically with income.

4 Access to Finance: Does It Moderate the Human Costs of Business Cycles?

In the final part of the paper, we investigate whether consumers' access to finance moderates or amplifies the relationship between local business cycles and excess aging.

Tackling this question helps to understand the mechanisms behind the relationship we document and carries important policy implications. The net effect of access to finance is ex-ante ambiguous. On the one hand, credit access can help consumers smooth consumption during downturns, alleviating financial stress and shielding well-being from the adverse effects of negative shocks (Banerjee and Duflo, 2010; Haushofer and Fehr, 2014; Karlan and Zinman, 2010)—beyond what can be explained by observed changes in income or employment status. On the other hand, taking on new debt or accumulating liabilities relative to income may itself introduce stress, elevate anxiety about repayment obligations, and raise the perceived risk of future financial distress. In this sense, access to finance may buffer or exacerbate the human costs of economic fluctuations, depending on how individuals use credit and how they experience its psychological consequences (Deckard et al., 2022).

4.1. Raw-Data Evidence

As with the baseline relationship, we begin by assessing the role of access to finance using raw data. We focus on the extensive margin—whether consumers' loan applications were approved or rejected—which lends itself to the intuitive graphical comparison in Figure 6. This figure replicates the baseline analysis in Panel C of Figure 3, but plots the relationship between excess age (y-axis) and the average lagged local unemployment rate over the previous six months separately for two groups: users whose loan applications were rejected (green line) and those for whom at least one loan application was

accepted (orange line). For visual clarity, we omit the underlying binscatter points.

Figure 6 reveals two raw-data facts that deserve further empirical investigation. First, there is a substantial level difference in excess age between users based on credit access, irrespective of local economic conditions. We emphasize this by plotting horizontal lines at the average excess age for each group in the lowest unemployment bin (gray dashed lines). Users whose loan applications were rejected exhibit excess ages more than one year higher than those with at least one accepted application. This gap suggests that limited access to finance may elevate the human costs consumers face even in the absence of adverse local business cycles.

Second, the slope of the relationship between excess age and lagged unemployment differs markedly by credit access. For users whose loan applications were rejected, excess age rises steadily as unemployment increases. In contrast, for users who gained credit access, the relationship remains virtually flat across most of the unemployment distribution—up to 16% average unemployment, a value well above the 75^{th} percentile (8.76%) in our sample. This suggests that, under most macroeconomic conditions, access to finance buffers consumers against the aging effects of local downturns. Finally, during extreme downturns—sustained periods of very high unemployment—excess aging accelerates for both groups, but the gap between them widens further. Specifically, the difference in excess age between users with and without credit access is 1.1 years in the lowest-unemployment bin and grows to 1.7 years in the highest. This pattern reinforces the notion that access to finance not only dampens the human costs of typical economic fluctuations but also provides increasing protection as local labor markets deteriorate more severely.

4.2. Multivariate Analysis

The graphical raw-data evidence in Figure 6 motivates us to further examine whether access to finance alleviates the human costs of business cycle fluctuations. We therefore move to a multivariate analysis, estimating specifications inspired by equation (*3.1.*) that compare users with and without access to credit. For completeness, we consider both the extensive margin—whether a loan application was approved—and the intensive margin—the amount of debt obtained relative to income for approved borrowers.

We present the results in Table 5, covering three forecasting horizons: one-month-ahead (columns 1–4), average three-months-ahead (columns 5–8), and average six-months-ahead unemployment rates

(columns 9–12). For each horizon, the first two columns estimate the baseline specification in equation (*3.1.*) separately for consumers whose loan applications were rejected (columns 1, 5, and 9) and those whose applications were accepted (columns 2, 6, and 10). Across all horizons, the estimated relation between local business cycles and excess aging is consistently larger and statistically significant for rejected applicants. By contrast, the estimates for accepted applicants are smaller and, with the exception of the three-month horizon, not statistically distinguishable from zero.

To formally test whether the relationship between local business cycles and excess aging differs by credit access, we turn to columns (3), (7), and (11) of Table 5, where we pool all loan applicants and augment the baseline specification with an interaction between the local unemployment rate and a dummy equal to one if the consumer's loan application was accepted. Two key patterns emerge. First, the difference in coefficients between accepted and rejected applicants is both economically meaningful—ranging from 33% to 50% of the baseline effect for rejected applicants—and statistically significant at standard levels, except at the three-month horizon. For this intermediate case, the estimated difference is of similar magnitude but not statistically different from zero. Second, we find that accepted applicants appear significantly younger—conditional on calendar age—even in the absence of local labor market stress (i.e., when unemployment is zero). To the extent that this association captures cases in which consumers apply for credit because they face individual shocks and needs unrelated to the local business cycle, this result suggests that having a loan application accepted rather than rejected—conditional on applying—does lead consumers to appear younger relative to their calendar age and others.

We then move on to the intensive margin of access to finance: within the sample of approved loan applicants, we assess whether being able to raise a higher amount of debt as a share of income moderates or mitigates the relationship between local business cycles and excess age. Columns (4), (8), and (12) report the estimates for each horizon we consider. Consistent with the extensive margin, we find that raising more debt to income among those who raise any debt reduces the positive association between local business cycles and excess age. Economically, the baseline effect is reduced between 38% and 75% if we move from an approved loan applicant who obtains a negligible share of their income as credit to one who obtains 100% of their income as credit. All the interactions are statistically different from zero at the 5% or 10% levels of significance.

Despite detecting a moderating role of access to finance at both margins for the baseline association

between local business cycles and excess aging, interestingly, we find divergent directions on how accessing unconditionally more debt intensively relates to excess aging relative to the effect of obtaining a loan application acceptance. Indeed, across all horizons, we find that obtaining an additional amount of debt equal to one's income, irrespective of the local business cycle, is associated with 0.07 additional predicted appearance years. This result is consistent with the accumulation of large amounts of debt relative to one's income, leading consumers to face higher human costs.

Overall, our results consistently document that access to finance, at both the extensive and intensive margin, reduces the positive association between facing negative local business cycles and excess age. This result reveals an overlooked benefit of access to finance, which appears to not only loosen consumers' financial and liquidity constraints and hence allow consumption smoothing in times of negative economic shocks, but also to reduce the human costs of business cycles. Symmetrically, these results also reveal an additional cost of rationing consumers from lending markets, which adds especially to the concerning cases in which access to finance is rationed to vulnerable consumers and/or due to explicit or implicit discriminatory practices (Bartlett et al., 2022).

5 Conclusions

This paper provides an empirical framework to assess why consumers are systematically willing to pay more to avoid business cycle fluctuations than what standard measures of consumption losses would suggest. Specifically, our setting aims to test for the existence of potentially overlooked costs that consumers may face during downturns, which are not captured by standard economic variables used as proxies for consumption. Evidence of these costs would suggest that consumers' "excessive" willingness to pay cannot be fully explained by their inability to accurately estimate the consumption costs of business cycles.

To address this question, we match image data from the pictures of approximately one million bank customers with their bank account and demographic information, as well as external data on the local business cycles consumers faced when their photographs were taken. For each image, we use state-of-the-art machine learning techniques to predict appearance age and compare it to individuals' actual calendar age, thereby generating a measure of "excess age" at the individual-time level. We interpret this excess age, controlling for direct proxies of individuals' consumption ability, as a proxy

for the bundle of non-consumption-related costs consumers may face, or *human costs*.

We find that consumers exposed to more adverse local business cycles also exhibit higher human costs than others. This relationship holds both across and within individuals and is highly nonlinear, with marginal effects increasing as local unemployment rates rise and persist. The association is particularly pronounced during periods of heightened macroeconomic volatility, measured by the standard deviation of monthly local unemployment rates. Importantly, the baseline relationship remains robust after controlling for a range of time-invariant and time-varying characteristics—including income and employment status, which proxy for individuals' consumption capacity. We also document heterogeneity in the strength of this relationship across demographic groups: it is notably stronger among middle-aged individuals, women, and those in higher-responsibility occupations and income brackets. Finally, our setting allows us to assess whether access to finance moderates or mediates these effects. We find that both at the extensive margin (i.e., loan approval) and the intensive margin (i.e., the share of income borrowed), access to finance mitigates the link between local unemployment and excess aging.

Our results open several avenues for future research. While our proxy enables the detection and field measurement of human costs, identifying their underlying drivers—whether psychological, biological, or social—and quantifying their relative importance remain open challenges. Future work could also explore the long-term consequences of elevated human costs, such as reduced labor market participation, delayed investment in human capital, or increased economic inertia. A deeper understanding of these mechanisms would offer a more comprehensive view of the burdens consumers face during business cycle downturns.

On the theoretical side, standard models tend to overlook the existence and relevance of human costs by focusing primarily on the consumption consequences of business cycles. Advancing theory to incorporate a more holistic view of the costs consumers face—and to allow for comparisons across different types of costs—remains an important direction for future research.

In terms of policy implications, our results suggest that policymakers may underestimate the costs consumers face during negative and volatile business cycles. In particular, relying solely on consumption-based measures may yield a misleading view of the cross-sectional distribution of these costs, given the systematic heterogeneity we document across demographic groups. A more comprehensive assessment—one that accounts for both consumption-based and human costs—would offer a

fuller understanding of which populations are most vulnerable and help inform the design of more targeted policy interventions. Moreover, our findings point to an overlooked benefit of access to finance: beyond easing liquidity constraints and enabling consumption smoothing during downturns, credit access also mitigates the non-consumption-related human costs of economic stress. These results imply that rationing access to debt markets may impose greater welfare losses than previously recognized, underscoring the need for improved measurement and deeper analysis of these impacts in future research.

References

Aguiar, M., E. Hurst, and L. Karabarbounis (2013). Time use during the great recession. *American Economic Review 103*(5), 1664–1696.

Akerlof, G. A. (2002). Behavioral macroeconomics and macroeconomic behavior. *American Economic Review 92*(3), 411–433.

Albanesi, S. and J. Kim (2021). The gendered impact of the covid-19 recession on the us labor market. Technical report, National Bureau of Economic Research.

Atkeson, A. and C. Phelan (1994). Reconsidering the costs of business cycles with incomplete markets. *NBER macroeconomics annual 9*, 187–207.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2018). The economic effects of social networks: Evidence from the housing market. *Journal of Political Economy 126*(6), 2224–2276.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The quarterly journal of economics 131*(4), 1593–1636.

Banerjee, A. V. and E. Duflo (2010). Giving credit where it is due. *Journal of Economic Perspectives 24*(3), 61–80.

Barlevy, G. (2003). Credit market frictions and the allocation of resources over the business cycle. *Journal of monetary Economics 50*(8), 1795–1818.

Barlevy, G. (2004). The cost of business cycles under endogenous growth. *American Economic Review 94*(4), 964–990.

Bartlett, R., A. Morse, R. Stanton, and N. Wallace (2022). Consumer-lending discrimination in the fintech era. *Journal of Financial Economics 143*(1), 30–56.

Bayer, P., F. Ferreira, and S. L. Ross (2018). What drives racial and ethnic differences in high-cost mortgages? the role of high-risk lenders. *The Review of Financial Studies 31*(1), 175–205.

Blau, F. D. and L. M. Kahn (2017). The gender wage gap: Extent, trends, and explanations. *Journal of economic literature 55*(3), 789–865.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). Overreaction in macroeconomic expectations. *American Economic Review 110*(9), 2748–2782.

Borgschulte, M., M. Guenzel, C. Liu, and U. Malmendier (Forthcoming). Ceo stress, aging, and death. *Journal of Finance*.

Bornstein, G. and S. Indarte (2023). The impact of social insurance on household debt. *Available at SSRN 4205719*.

Brody, E. M. (2003). *Women in the middle: Their parent-care years*. Springer Publishing Company.

Burgess, R., M. Hansen, B. A. Olken, P. Potapov, and S. Sieber (2012). The political economy of deforestation in the tropics. *The Quarterly journal of economics 127*(4), 1707–1754.

Card, D., A. R. Cardoso, and P. Kline (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly journal of economics 131*(2), 633–686.

Cavallo, A. (2017). Are online and offline prices similar? evidence from large multi-channel retailers. *American Economic Review 107*(1), 283–303.

Crouzet, N., A. Gupta, and F. Mezzanotti (2023). Shocks and technology adoption: Evidence from electronic payment systems. *Journal of Political Economy 131*(11), 3003–3065.

Da, Z., J. Engelberg, and P. Gao (2011). In search of attention. *The journal of finance 66*(5), 1461–1499.

D'Acunto, F., M. Weber, and X. Yin (2024). Subjective income expectations and household debt cycles. Technical report, National Bureau of Economic Research.

Davis, S. and T. vonWachter (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity*.

De Grauwe, P. (2012). *Lectures on behavioral macroeconomics*. Princeton University Press.

Deaton, A. (2011). The financial crisis and the well-being of america. In *Investigations in the Economics of Aging*, pp. 343–368. University of Chicago Press.

Deckard, F. M., B. J. Goosby, and J. E. Cheadle (2022). Debt stress, college stress: Implications for black and latinx students' mental health. *Race and Social Problems 14*(3), 238–253.

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2020). Heterogeneity and persistence in returns to wealth. *Econometrica 88*(1), 115–170.

Fuster, A., P. Goldsmith-Pinkham, T. Ramadorai, and A. Walther (2022). Predictably unequal? the effects of machine learning on credit markets. *The Journal of Finance 77*(1), 5–47.

Fuster, A., G. Kaplan, and B. Zafar (2021). What would you do with $500? spending responses to gains, losses, news, and loans. *The Review of Economic Studies 88*(4), 1760–1795.

Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review 110*(8), 2271–2327.

Georgarakos, D., K. H. Kim, O. Coibion, M. Shim, Y. Gorodnichenko, G. Kenny, S. Han, and M. Weber (2025). How costly are business cycle volatility and inflation? a vox populi approach. *NBER Working Paper Series No. 33476*.

Guenzel, M., S. Kogan, M. Niessner, and K. Shue (2025). Ai personality extraction from faces: Labor market implications. *Available at SSRN*.

Guerreiro, J., J. Hazell, C. Lian, and C. Patterson (2024). Why do workers dislike inflation?: Wage erosion and conflict costs.

Guiso, L., T. Jappelli, and D. Terlizzese (1996). Income risk, borrowing constraints, and portfolio choice. *The American economic review*, 158–172.

Haushofer, J. and E. Fehr (2014). On the psychology of poverty. *science 344*(6186), 862–867.

Henderson, J. V., A. Storeygard, and D. N. Weil (2012). Measuring economic growth from outer space. *American economic review 102*(2), 994–1028.

Hommes, C. (2021). Behavioral and experimental macroeconomics and policy analysis: A complex systems approach. *Journal of Economic Literature 59*(1), 149–219.

Imrohoroğlu, A. (1989). Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political economy 97*(6), 1364–1383.

Imrohoroglu, A. (2008). Welfare costs of business cycles. In *The new Palgrave dictionary of economics*, pp. 1–5. Springer.

Indarte, S., R. Kluender, U. Malmendier, and M. Stepner (2025). Consumption wedges: Measuring and diagnosing distortions.

Jose, I. and S. Sivaraman (2023). Gender inequality and gender gap: An overview of the indian scenario. *Gender Issues 40*(2), 232–254.

Kaplan, G. and G. L. Violante (2018). Microeconomic heterogeneity and macroeconomic shocks. *Journal of Economic Perspectives 32*(3), 167–194.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity 2014*(1), 77–138.

Karlan, D. and J. Zinman (2010). Expanding credit access: Using randomized supply decisions to estimate the impacts. *The Review of Financial Studies 23*(1), 433–464.

Kessler, R. C., K. A. McGonagle, M. Swartz, D. G. Blazer, and C. B. Nelson (1993). Sex and depression in the national comorbidity survey i: Lifetime prevalence, chronicity and recurrence. *Journal of affective disorders 29*(2-3), 85–96.

Kiyotaki, N. and J. Moore (1997). Credit cycles. *Journal of political economy 105*(2), 211–248.

Kleven, H., C. Landais, and J. E. Søgaard (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics 11*(4), 181–209.

Lucas Jr, R. E. (1987). Models of business cycles. *New York: Wiley*.

Lucas Jr, R. E. (2003). Macroeconomic priorities. *American economic review 93*(1), 1–14.

Michalopoulos, S. and E. Papaioannou (2014). National institutions and subnational development in africa. *The Quarterly journal of economics 129*(1), 151–213.

Naaraayanan, S. L. (2019). Women's inheritance rights and entrepreneurship gender gap. Technical report.

OECD (2022). Education at a glance 2022.

Pfajfar, D. and F. Winkler (2025). Households' preferences over inflation and monetary policy trade-offs.

Pfäuti, O. and F. Seyrich (2022). A behavioral heterogeneous agent new keynesian model.

Ramey, G. and V. A. Ramey (1991). Technology commitment and the cost of economic fluctuations.

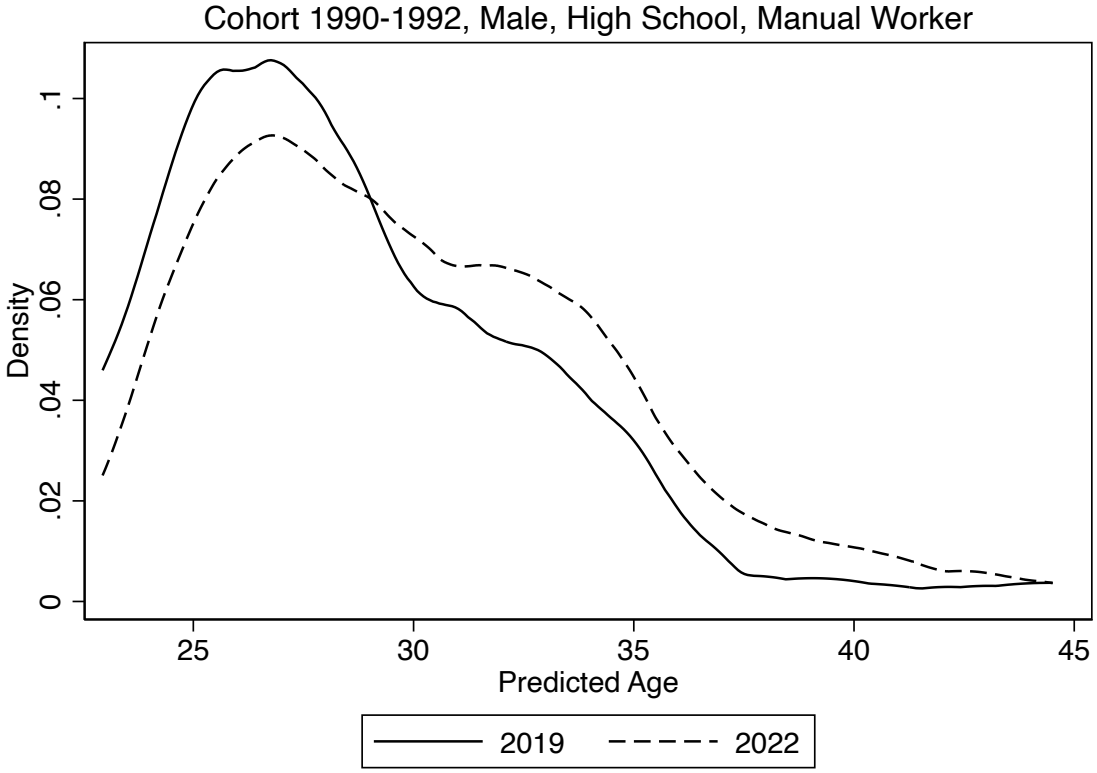
Stantcheva, S. (2024). Why do we dislike inflation? *Brookings Papers on Economic Activity 2004*(1), 1–46.

Tallarini Jr, T. D. (2000). Risk-sensitive real business cycles. *Journal of monetary Economics 45*(3), 507–532.

Townsend, R. M. (1994). Risk and insurance in village india. *Econometrica*, 539–591.

Wolfers, J. (2003). Is business cycle volatility costly? evidence from surveys of subjective well-being. *International finance 6*(1), 1–26.

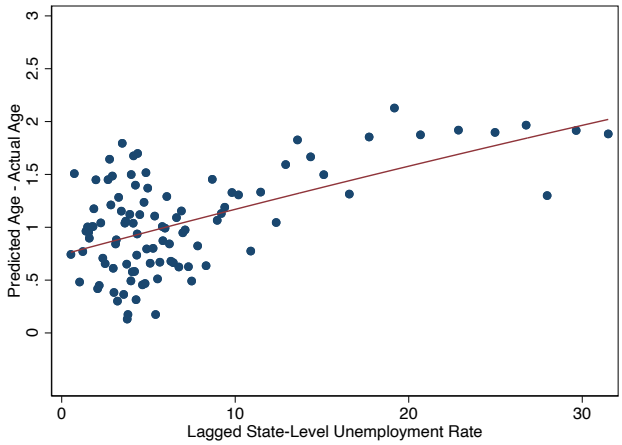
Figure 2: Is Predicted Age Noise? Evolution within Sample Cohorts



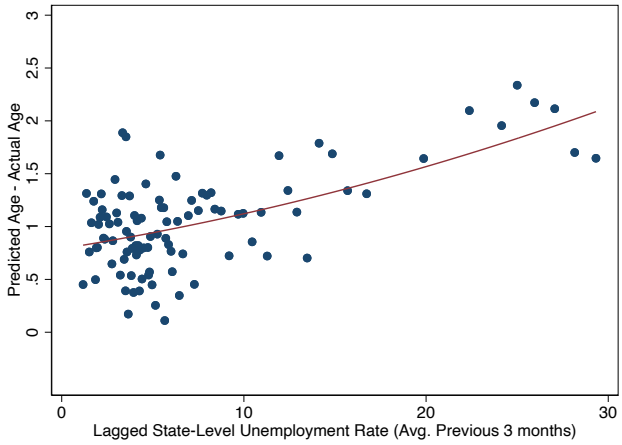
This figure plots kernel density distributions of predicted appearance age for consumers in the 1990–1992 cohort who are male, hold a high school degree, and work in manual occupations. The distributions are shown separately for 2019 (solid line) and 2022 (dashed line).

Figure 3: Excess Age and Local Unemployment Rate by Horizons

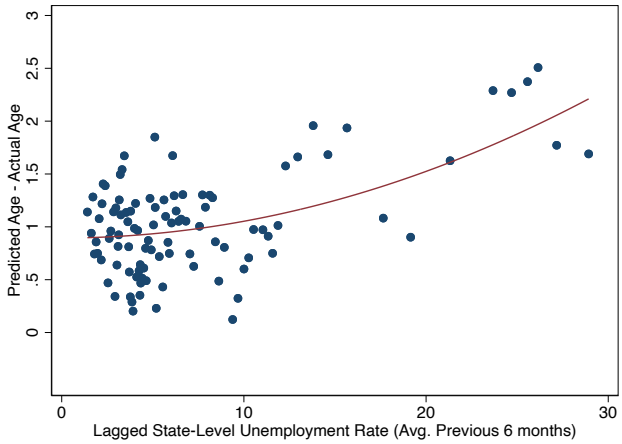
Panel A. Previous-Month Unemployment Rate



Panel B. Previous 3-Month Unemployment Rate



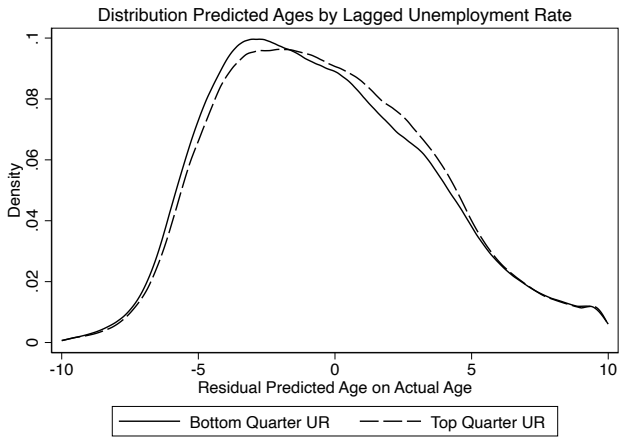
Panel C. Previous 6-Month Unemployment Rate



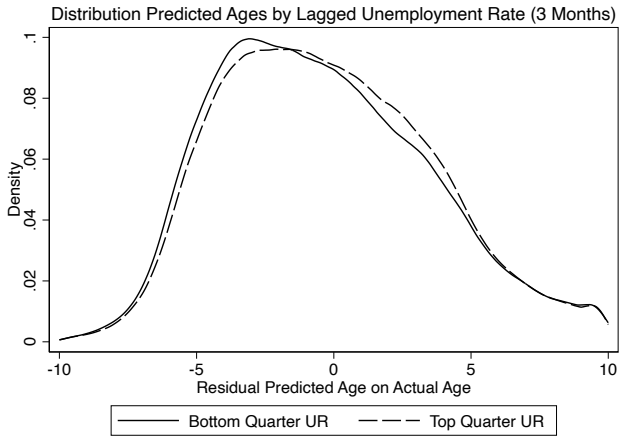
Panels A–C show the relationship between excess age and local unemployment rates measured over three different horizons: (A) 1-month, (B) 3-month average, and (C) 6-month average. The y-axis plots the difference between bank users' predicted age based on their pictures and their actual calendar age at the time the picture was taken. The x-axis sorts bank users based on the average unemployment rate in the Indian state where they reside over the 6 months preceding the picture.

Figure 4: Distribution of Excess Age by Top vs. Bottom Local Unemployment Rates

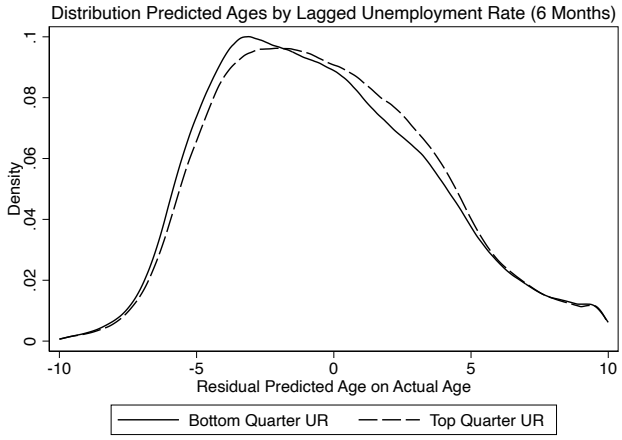
Panel A. Previous-Month Unemployment Rate



Panel B. Previous 3-Month Unemployment Rate

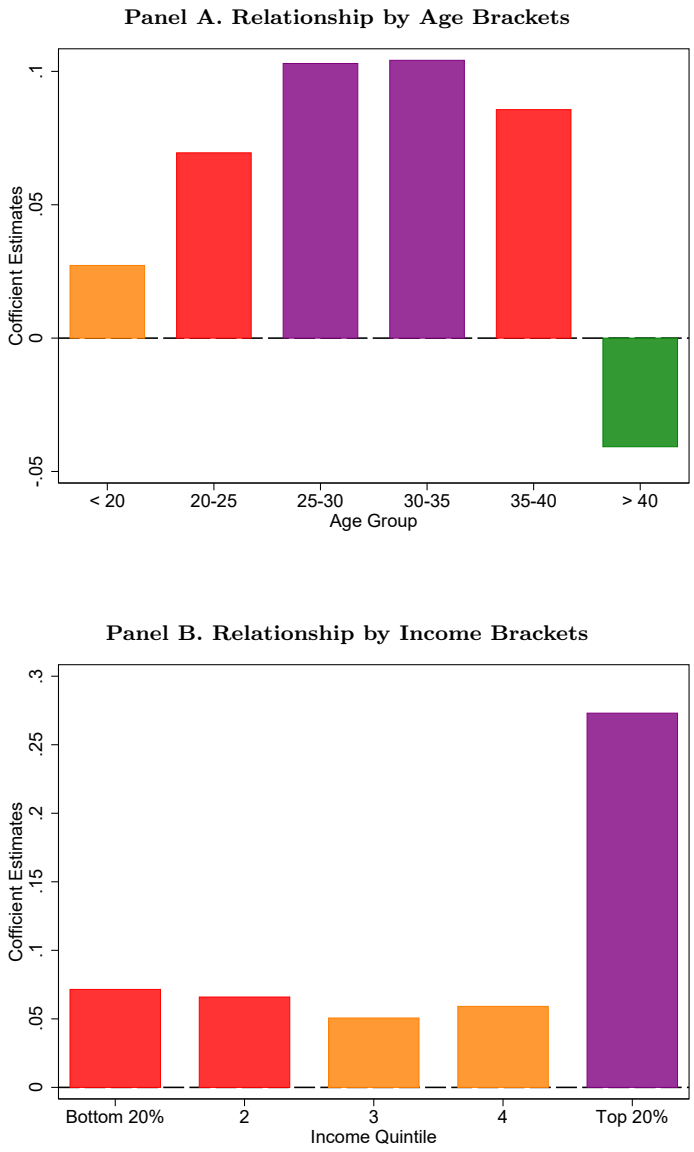


Panel C. Previous 6-Month Unemployment Rate



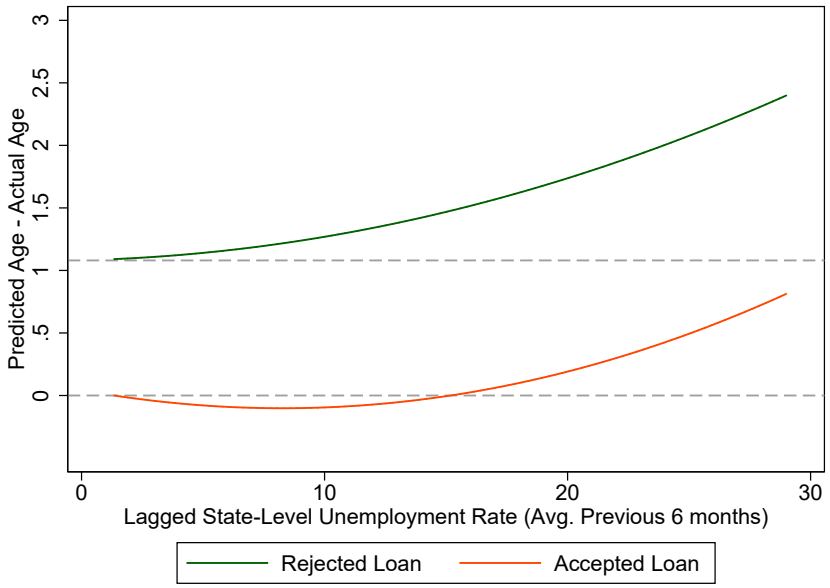
Panels A–C compare the distribution of excess age across different levels of local unemployment exposure. The y-axis plots the estimated density of residuals from a regression of predicted age on actual age. The solid lines represent users residing in Indian states with unemployment rates in the bottom quartile, while the dashed lines represent users in the top quartile of the unemployment distribution. Panel A uses 1-month lagged unemployment, Panel B uses a 3-month average, and Panel C uses a 6-month average.

Figure 5: Excess Age and Local Unemployment Rate: Heterogeneity



Panels A–B show the estimated interaction coefficients between excess age and average local unemployment rates across different demographic groups. In both panels, the y-axis plots coefficient estimates obtained from regressions of predicted appearance age on calendar age, average local unemployment, demographic group indicators, and their interactions. Panel A displays interaction coefficients by age group: under 20, 20–24, 25–29, 30–34, 35–39, and 40 or older. The x-axis corresponds to the six age bins. Panel B displays interaction coefficients by income quintile. The x-axis corresponds to increasing income levels.

Figure 6: Excess Age and Local Unemployment Rate by Access to Finance



The figure shows the relation between excess age and average local unemployment rates over the previous six months, separately for users whose loan applications were accepted (orange line) and rejected (green line). The y-axis plots the difference between users' predicted appearance age based on their photographs and their actual calendar age at the time the picture was taken. The x-axis sorts users based on the six-month average unemployment rate in the Indian state where they reside. Each line is a smoothed binscatter fit.

Table 1. Summary Statistics

	N. obs.	Mean	St. dev.	25*th* p.	Median	75*th* p.
Full Sample						
Predicted Appearance Age	915,696	30.09	4.52	26.58	29.45	32.90
Calendar Age	915,696	29.04	7.02	23.83	27.63	32.78
Unemployment Rate (Previous Month)	914,465	7.25	6.95	3.17	4.76	8.11
Avg. Unemployment Rate (3 Months)	914,572	7.38	6.69	3.48	4.84	8.33
Avg. Unemployment Rate (6 Months)	915,696	7.51	6.47	3.50	5.12	8.76
Approved Loan	915,696	0.16	0.37	0	0	0
Debt-to-Income	149,847	0.96	0.65	0.45	0.83	1.26
Two-Image Sample						
Predicted Appearance Age	185,286	29.73	4.53	26.21	28.98	32.53
Calendar Age	185,286	29.18	6.38	24.49	28.10	32.70
Unemployment Rate (Previous Month)	185,143	6.89	6.45	3.12	4.75	7.70
Avg. Unemployment Rate (3 Months)	185,132	6.96	6.20	3.35	4.74	7.98
Avg. Unemployment Rate (6 Months)	185,286	7.03	6.01	3.42	4.89	8.24
Approved Loan	185,286	0.24	0.43	0	0	0
Debt-to-Income	44,886	0.88	0.56	0.46	0.77	1.14
Full Sample with Demographics						
Predicted Appearance Age	67,682	30.21	4.58	26.64	29.52	33.03
Calendar Age	67,682	30.53	6.18	25.96	29.63	34.09
Unemployment Rate (Previous Month)	67,664	6.53	6.39	2.94	4.37	7.00
Avg. Unemployment Rate (3 Months)	67,669	6.74	6.14	3.19	4.57	7.23
Avg. Unemployment Rate (6 Months)	67,682	6.94	5.93	3.17	4.89	8.18
Log(Income)	67,682	11.00	0.97	10.28	10.86	11.56
Male	67,682	0.88	0.32	1	1	1
High School Degree	67,682	0.62	0.49	0	1	1
College Degree	67,682	0.18	0.38	0	0	0
Employee	67,682	0.45	0.50	0	0	1
Managerial Role	67,682	0.25	0.44	0	0	1
Leadership Role	67,682	0.30	0.46	0	0	1

Table 1 presents summary statistics for the main variables used in the analysis, reported separately for the three samples we examine. The *Full Sample* comprises all users for whom the platform has recorded one or two photographs. The *Two-Image Sample* is a subset of the Full Sample and includes users with two distinct pictures taken at different points in time. The *Full Sample with Demographics* is another subset of the Full Sample, consisting of users who voluntarily provided a complete set of demographic information in their profile.

Table 2. Excess Age and Local Unemployment Rate: Across and Within Individuals

Dependent variable: Predicted Age	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A. Previous Month					
	Across Individuals				Within Individuals	
Calendar Age	0.22***	0.22***	0.22***	0.22***	0.30***	0.30***
	(121.4)	(122.0)	(124.4)	(124.6)	(26.87)	(5.47)
Unemployment Rate	0.07***	0.07***	0.06***	0.03**	0.07***	0.09***
	(13.89)	(5.00)	(12.78)	(2.32)	(2.65)	(3.61)
Observations	914,691	914,691	914,691	914,691	185,258	185,258
	Panel B. Previous 3 Months					
	Across Individuals				Within Individuals	
Calendar Age	0.22***	0.22***	0.22***	0.22***	0.30***	0.29***
	(121.7)	(122.44)	(124.8)	(124.95)	(25.94)	(5.32)
Average Unemployment Rate	0.08***	0.13***	0.07***	0.08***	0.10***	0.12***
	(15.04)	(8.25)	(15.19)	(4.58)	(3.33)	(4.25)
Observations	914,781	914,781	914,781	914,781	185,232	185,232
	Panel C. Previous 6 Months					
	Across Individuals				Within Individuals	
Calendar Age	0.22***	0.22***	0.22***	0.22***	0.30***	0.27***
	(121.5)	(122.2)	(124.6)	(124.9)	(25.71)	(5.07)
Average Unemployment Rate	0.08***	0.14***	0.06***	0.08***	0.11***	0.13***
	(12.83)	(6.95)	(13.96)	(3.70)	(3.33)	(4.16)
Observations	915,696	915,696	915,696	915,696	185,286	185,286
State Fixed Effects		✓		✓		
Year Fixed Effects			✓	✓		✓
Individual Fixed Effects					✓	✓

Table 2 reports the results of estimating the following specification by ordinary least squares:

$$Predicted\ Age_{i,t,s} = \alpha + \beta Calendar\ Age_{i,t,s} + \gamma Avg.\ Unemployment\ Rate_{t-k\to t-1,s} + \mathbf{X}'\delta + \eta_t + \eta_s + \eta_i + \epsilon_{i,t,s}$$

where $Predicted\ Age_{i,s,t}$ is the appearance age of consumer i in state s and time t predicted from i's image(s) at time t; $Calendar\ Age_{i,s,t}$ is the actual age of consumer i in state s and time t based on their date of birth; $Avg.\ Unemployment\ Rate_{t-k\to t-1,s}$ is the average lagged unemployment rate in state s, where k=2 captures the lagged unemployment rate one month before i's image is observed (Panel A), k=4 captures the 3-month average lagged unemployment rate before i's image is observed (Panel B), and k=7 captures the 6-month average lagged unemployment rate before i's image is observed (Panel C); η_t, η_s, and η_i capture full set of year, state, and individual fixed effects. Standard errors are double-clustered at the individual and month levels.

Table 3. Excess Age and Local Unemployment Rate: Stable vs. Volatile Business Cycles

Dependent variable: Predicted Age	(1)	(2)	(3)	(4)	(5)	(6)
	Across Individuals				Within Individuals	
Calendar Age	0.22***	0.22***	0.22***	0.22***	0.30***	0.26***
	(122.1)	(122.7)	(124.6)	(125.0)	(25.81)	(5.06)
Average Unemployment Rate	-0.06	-0.12	-0.06	-0.14**	-0.49***	-0.31**
	(-0.73)	(-1.25)	(-0.83)	(-2.31)	(-3.49)	(-2.38)
Average Unemployment Rate × Top Volatility	0.14	0.27***	0.12	0.21***	0.60***	0.43***
	(1.57)	(2.92)	(1.57)	(3.48)	(4.43)	(3.45)
Average Unemployment Rate × Med. Volatility	0.21**	0.22***	0.24***	0.23***	0.42***	0.38***
	(2.44)	(3.06)	(2.87)	(4.09)	(3.27)	(3.12)
Top Volatility	0.07	0.12*	0.09	0.15***	0.39***	0.28***
	(1.17)	(1.94)	(1.54)	(2.87)	(4.69)	(3.63)
Medium Volatility	0.10*	0.11**	0.11**	0.11***	0.28***	0.19***
	(1.68)	(2.21)	(2.08)	(3.11)	(3.69)	(2.75)
Observations	915,681	915,681	915,681	915,681	185,284	185,284
State Fixed Effects		✓		✓		
Year Fixed Effects			✓	✓		✓
Individual Fixed Effects					✓	✓
Mean St. Dev. UR Top Volatility	4.85 pp					
Mean St. Dev. UR Medium Volatility	1.39 pp					
Mean St. Dev. UR Bottom Volatility	0.60 pp					

Table 3 reports the results of estimating the following specification by ordinary least squares:

$$\begin{aligned} \textit{Predicted Age}_{i,s,t} = &\alpha + \beta \textit{Calendar Age}_{i,s,t} + \zeta \textit{Avg. Unemployment Rate}_{s,t-6} \\ &+ \gamma_T \textit{Avg. Unemployment Rate}_{s,t-6} \times \textit{Top Volatility}_{s,t-6} \\ &+ \gamma_M \textit{Avg. Unemployment Rate}_{s,t-6} \times \textit{Medium Volatility}_{s,t-6} \\ &+ \delta_T \textit{Top Volatility}_{s,t-6} + \delta_M \textit{Medium Volatility}_{s,t-6} \\ &+ \eta_s + \eta_i + \eta_t + \epsilon_{i,s,t}, \end{aligned}$$

where $\textit{Predicted Age}_{i,s,t}$ is the appearance age of consumer i in state s and time t predicted from i's image(s) at time t; $\textit{Calendar Age}_{i,s,t}$ is the actual age of consumer i in state s and time t based on their date of birth; $\textit{Avg. Unemployment Rate}_{s,t-6}$ is the 6-month average lagged unemployment rate before i's image is observed; $\textit{Top Volatility}_{s,t-6}$ and $\textit{Medium Volatility}_{s,t-6}$ are dummy variables that equal 1 if the standard deviation of the unemployment rates in state s over the 6 months before the image was observed are in the top third or middle third of the distribution, and zero otherwise; η_s, η_i, and η_t capture full set of state, individual, and year fixed effects. Standard errors are double-clustered at the individual and month levels.

Table 4. Excess Age and Local Unemployment Rate: Heterogeneity by Demographics

Dependent variable: Predicted Age	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Age			Gender			Occupation Rank			Income		
Calendar Age	0.21***	0.21***	0.21***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***
	(87.18)	(88.37)	(87.70)	(32.80)	(32.28)	(32.76)	(19.65)	(19.48)	(19.69)	(32.16)	(31.63)	(32.24)
Average Unemployment Rate	0.08***	0.15***	0.07***	0.11***	0.18***	0.10***	0.02	0.08	0.01	0.05***	0.12***	0.03***
	(11.84)	(7.01)	(13.46)	(4.15)	(4.41)	(3.62)	(1.52)	(1.49)	(0.45)	(3.89)	(3.86)	(2.71)
Average Unemployment Rate × Young	-0.02***	-0.02***	-0.02***									
	(-3.07)	(-3.09)	(-2.70)									
Young	-0.06**	-0.06***	-0.09***									
	(-2.61)	(-2.81)	(-4.09)									
Average Unemployment Rate × Male				-0.05	-0.04	-0.06**						
				(-1.63)	(-1.43)	(-1.99)						
Male				-0.70***	-0.72***	-0.66***						
				(-19.07)	(-20.18)	(-20.15)						
Average Unemployment Rate × High Rank							0.07**	0.07**	0.06**			
							(2.55)	(2.47)	(2.04)			
High Rank							0.01	0.02	0.08			
							(0.24)	(0.26)	(1.50)			
Average Unemployment Rate × High Income										0.27**	0.26***	0.20**
										(3.52)	(3.37)	(2.52)
High Income										-0.28***	-0.26***	0.20
										(-3.52)	(-3.23)	(2.52)
Observations	915,696	915,696	915,696	149,847	149,847	149,847	67,687	67,687	67,687	149,835	149,835	149,835
State Fixed Effects		✓			✓			✓			✓	
Year Fixed Effects			✓			✓			✓			✓

Table 4 reports the results of estimating the following specification by ordinary least squares:

$$Predicted\ Age_{i,s,t} = \alpha + \beta Calendar\ Age_{i,s,t} + \zeta Avg.\ Unemployment\ Rate_{s,t-6} + \eta_s + \eta_i + \eta_t + \epsilon_{i,s,t},$$

where $Predicted\ Age_{i,s,t}$ is the appearance age of consumer i in state s and time t predicted from i's image(s) at time t; $Calendar\ Age_{i,s,t}$ is the actual age of consumer i in state s and time t based on their date of birth; $Avg.\ Unemployment\ Rate_{s,t-6}$ is the 6-month average lagged unemployment rate in state s before i's image is observed; $Young$ is a dummy variable that equals 1 for users who are 25 years old or younger, and zero otherwise; $Male$ is a dummy variable that equal 1 for male users, and zero otherwise; $High\ Rank$ is a dummy variable that equals 1 for users who are in Senior Management or Leadership occupation ranks, and zero for those in Middle Management or No Responsibility Employee occupation ranks; $High\ Income$ is a dummy variable that equals 1 for users whose income belongs to the top quintile of the distribution in our sample, and zero otherwise; η_i and η_t capture full set of individual and year fixed effects. Standard errors are double-clustered at the individual and month levels.

Table 5. Excess Age and Local Unemployment Rate: The Role of Access to Finance

Dependent variable: Predicted Age	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	1 Month				3 Months				6 Months			
	Rejected	*Accepted*	*All*	*Intensive*	*Rejected*	*Accepted*	*All*	*Intensive*	*Rejected*	*Accepted*	*All*	*Intensive*
Calendar Age	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***
	(126.28)	(32.23)	(124.78)	(31.28)	(126.74)	(32.19)	(125.18)	(31.23)	(126.33)	(32.19)	(125.10)	(31.25)
(Average) Unemployment Rate	0.04***	0.02	0.04**	0.04*	0.08***	0.06**	0.06***	0.08***	0.10***	0.03	0.09***	0.06
	(2.43)	(0.82)	(2.54)	(1.73)	(4.75)	(2.08)	(4.72)	(2.66)	(4.59)	(0.72)	(3.89)	(1.43)
(Average) Unemployment Rate × Accepted			-0.02*				-0.02				-0.03**	
			(3.82)				(-1.51)				(-2.08)	
Accepted			-0.17***				-0.17***				-0.17***	
			(-5.38)				(-5.37)				(-5.43)	
(Average) Unemployment Rate × Mean DTI				-0.03*				-0.03**				-0.03*
				(-1.91)				(1.96)				(-1.79)
Mean DTI				0.07***				0.07***				0.07***
				(3.07)				(3.05)				(3.08)
Observations	764,865	149,826	914,691	149,826	764,952	149,829	914,781	149,829	765,849	149,847	915,696	149,847
State Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 5 reports the results of estimating the following specification by ordinary least squares:

$$Predicted\ Age_{i,s,t} = \alpha + \beta Calendar\ Age_{i,s,t} + \zeta Avg.\ Unemployment\ Rate_{s,t-x} + \eta_s + \eta_i + \eta_t + \epsilon_{i,s,t},$$

where $Predicted\ Age_{i,s,t}$ is the appearance age of consumer i in state s and time t predicted from i's image(s) at time t; $Calendar\ Age_{i,s,t}$ is the actual age of consumer i in state s and time t based on their date of birth; (*Avg.*) $Unemployment\ Rate_{s,t-x}$ is the one-month lagged unemployment rate in state s in columns (1)-(4), and the average lagged unemployment rate in state s otherwise, where x=3 captures the 3-month average lagged unemployment rate before i's image is observed (columns (5)-(8)), and x=6 captures the 6-month average lagged unemployment rate before i's image is observed (columns (9)-(12)); *Accepted* is a dummy variable that equals 1 for users for whom at least one loan application was accepted, and zero otherwise (note that users might file more than one loan application at each point in time); $Mean\ DTI$ is the within-individual average debt over income for the amount of debt raised and income at the time of the loan application for the loan applications we observe, and is only defined for users for whom at least one loan application was accepted; η_i and η_t capture full set of individual and year fixed effects. Standard errors are double-clustered at the individual and month levels.

Online Appendix:
The Human Costs of Business Cycles

Francesco D'Acunto, Pulak Ghosh, Bin Li, Alberto G. Rossi

Not for Publication

Table A.1. Excess Age and Local Unemployment Rate: Subsample with Demographic Characteristics

Dependent variable: Predicted Age	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	1 Month			3 Months			6 Months		
Calendar Age	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***	0.22***
	(16.97)	(16.83)	(17.75)	(16.96)	(16.82)	(17.74)	(16.88)	(16.77)	(17.67)
(Average) Unemployment Rate	0.06***	0.07**	0.04**	0.07***	0.14***	0.05***	0.06***	0.09*	0.03**
	(3.36)	(1.95)	(2.25)	(3.82)	(2.79)	(2.69)	(3.44)	(1.79)	(2.24)
Log(Income)	-0.03	-0.02	0.16***	-0.03	-0.02	0.16***	-0.03	-0.02	0.16***
	(-0.51)	(-0.36)	(2.89)	(-0.50)	(-0.34)	(2.89)	(-0.51)	(-0.35)	(2.89)
Male	-0.68***	-0.70***	-0.65***	-0.68***	-0.70***	-0.65***	-0.67***	-0.70***	-0.65***
	(-16.28)	(-17.16)	(-15.63)	(-16.25)	(-17.12)	(-15.60)	(-16.37)	(-17.14)	(-15.72)
High-School Degree	0.07	0.06	0.07	0.07	0.06	0.07	0.07	0.06	0.07
	(1.38)	(1.28)	(1.39)	(1.41)	(1.29)	(1.42)	(1.43)	(1.32)	(1.45)
College Degree	0.13**	0.11*	0.12**	0.13**	0.11*	0.12**	0.13**	0.11*	0.12**
	(2.12)	(1.83)	(2.05)	(2.14)	(1.85)	(2.07)	(2.14)	(1.84)	(2.08)
Employee	0.04	0.03	0.04	0.04	0.03	0.04	0.04	0.03	0.04
	(0.70)	(0.61)	(0.85)	(0.70)	(0.62)	(0.85)	(0.70)	(0.61)	(0.84)
Managerial Role	-0.03	-0.04	0.09	-0.03	-0.03	0.09	-0.03	-0.04	0.09
	(-0.36)	(-0.51)	(1.26)	(-0.35)	(-0.45)	(1.26)	(-0.34)	(-0.47)	(1.27)
Leadership Role	0.12	0.12*	0.05	0.117	0.121*	0.06	0.118	0.12*	0.06
	(1.62)	(1.68)	(0.78)	(1.63)	(1.69)	(0.79)	(1.63)	(1.69)	(0.79)
Observations	67,670	67,670	67,670	67,673	67,673	67,673	67,682	67,682	67,682
State Fixed Effects		✓			✓			✓	
Year Fixed Effects			✓			✓			✓

Table A.1 reports the results of estimating the following specification by ordinary least squares:

$$Predicted\ Age_{i,s,t} = \alpha + \beta Calendar\ Age_{i,s,t} + \zeta Avg.\ Unemployment\ Rate_{s,t-x} + \eta_s + \eta_i + \eta_t + \epsilon_{i,s,t},$$

where $Predicted\ Age_{i,s,t}$ is the appearance age of consumer i in state s and time t predicted from i's image(s) at time t; $Calendar\ Age_{i,s,t}$ is the actual age of consumer i in state s and time t based on their date of birth; $Avg.\ Unemployment\ Rate_{s,t-x}$ is the average lagged unemployment rate in state s, where x=3 captures the 3-month average lagged unemployment rate before i's image is observed (columns (1)-(3)), and x=6 captures the 6-month average lagged unemployment rate before i's image is observed (columns (4)-(6)); $St.dev.\ Unemployment\ Rate_{s,t-x}$ is the standard deviation of the monthly unemployment rates in state s computed over the same horizons as above; η_i and η_t capture full set of individual and year fixed effects. Standard errors are double-clustered at the individual and month levels.