

NBER WORKING PAPER SERIES

THE AGE OF REASON:
FINANCIAL DECISIONS OVER THE LIFECYCLE

Sumit Agarwal
John C. Driscoll
Xavier Gabaix
David Laibson

Working Paper 13191
<http://www.nber.org/papers/w13191>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2007

Gabaix and Laibson acknowledge support from the National Science Foundation (Human and Social Dynamics program). Laibson acknowledges financial support from the National Institute on Aging (R01-AG-1665). The views expressed in this paper are those of the authors and do not represent the policies or positions of the Board of Governors of the Federal Reserve System or the Federal Reserve Bank of Chicago. For their helpful comments we thank Marco Bassetto, Stephane Bonhomme, David Cutler, Ray Fair, Luigi Guiso, Gur Huberman, Erik Hurst, Ulrike Malmendier, Mitch Petersen, Rich Rosen, Timothy Salthouse, Fiona Scott Morton, Jesse Shapiro, William Sharpe, Paolo Sodini, Nick Souleles, Jon Zinman, and participants at the AEA meetings, the Bank of Spain, the Chicago Fed, the Federal Reserve Board, Princeton, the Institute for Fiscal Studies, the NBER (Aging and Behavioral Economics groups), the Bank of the Netherlands, the Society of Actuaries, and the Yale Behavioral Economics conference. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Sumit Agarwal, John C. Driscoll, Xavier Gabaix, and David Laibson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Age of Reason: Financial Decisions Over the Lifecycle
Sumit Agarwal, John C. Driscoll, Xavier Gabaix, and David Laibson
NBER Working Paper No. 13191
June 2007, Revised February 2008
JEL No. D1,D8,G2,J14

ABSTRACT

In cross-sectional data sets from ten credit markets, we find that middle-aged adults borrow at lower interest rates and pay fewer fees relative to younger and older adults. Fee and interest payments are minimized around age 53. The measured effects are not explained by observed risk characteristics. We discuss several leading factors that may contribute to these effects, including age-related changes in experience and cognitive function, selection effects, and cohort effects.

Sumit Agarwal
Financial Economist
Federal Reserve Bank of Chicago
230 South LaSalle Street
Chicago, IL 60604-1413
sumit.agarwal@chi.frb.org

John C. Driscoll
Mail Stop 85
Federal Reserve Board
20th and Constitution Ave., NW
Washington, DC 20551
John.C.Driscoll@frb.gov

Xavier Gabaix
New York University
Finance Department
Stern School of Business
44 West 4th Street, 9th floor
New York, NY 10012
and NBER
xgabaix@stern.nyu.edu

David Laibson
Department of Economics
Littauer M-14
Harvard University
Cambridge, MA 02138
and NBER
dlaibson@harvard.edu

1 Introduction

Performance tends to rise and then fall with age. Baseball players peak in their late 20s (Fair 2005, James 2003). Mathematicians, theoretical physicists, and lyric poets make their most important contributions around age 30 (Simonton 1988). Chess players achieve their highest ranking in their mid-30s (Charness and Bosnian 1990). Autocratic rulers are most effective in their early 40s (Simonton 1988). Authors write their most influential novels around age 50 (Simonton 1988).¹

The present paper studies an activity that is less august, though it is directly relevant to more people: financial decision making. We find that financial choices also have a hump-shaped performance pattern. We document cross-sectional variation in the prices that people pay for financial services. We study several proprietary datasets from multiple financial institutions, containing detailed information on the choices and characteristics of consumers. We present evidence on all the data that is available to us, including (a) interest rates in six different collateralized and non-collateralized consumer credit markets, (b) consumers' ability to optimally exploit balance transfers, and (c) three kinds of credit card fees.

We find that younger adults and older adults borrow at higher interest rates and pay more fees than middle-aged adults controlling for all observable characteristics, including measures of risk. Fee and interest payments are at their minimums for borrowers in their early 50s.

Factors idiosyncratic to each market likely contribute to the observed patterns that we observe. The recurrence of this pattern in many different markets suggests, however, that there may be common underlying explanations. We argue that the most likely explanations are age-related cognitive effects, selection effects, and cohort effects. We also discuss several other theories that are unlikely to play an important role in explaining the observed patterns in our data, including age-varying risk effects, age-varying opportunity costs of time, age discrimination and other supply factors.

The paper has the following organization. Section 2 describes the basic structure of the empirical analysis. The next two sections present results for interest rates on six different financial products, three different kinds of credit card fee payments, and on the use of balance transfer credit card offers. Section 5 uses all ten sets of results to estimate the age of peak sophistication. Section 6 discusses various common explanations for the pattern. Section 7 discusses market equilibrium, the economic magnitude of the effects, related work, and opportunities for future research. Section 8 concludes.

2 Overview

In the body of the paper, we document a U-shaped age-related curve in the prices people pay for financial services. We study ten separate contexts: home equity loans and lines of credit; auto loans; credit card interest rates; mortgages; small business credit cards; credit card late payment fees; credit

¹What about economists? Oster and Hamermesh (1998) find that economists' output in top publications declines sharply with age. This may simply reflect lower motivation with age. More optimistic data are reported in Weinberg and Galenson (2005)'s study of Nobel (Memorial) Prize winners. They find that "conceptual" laureates peak at age 43, and "experimental" ones at age 61.

card over limit fees; credit card cash advance fees; and use of credit card balance transfer offers. We discuss three forms of prices paid: higher APRs (Annual Percentage Rates, i.e., interest rates); higher fee payments; and suboptimal use of balance transfer offers.

For each application, we conduct a regression analysis that identifies age effects and controls for observable factors that might explain patterns of fee payments or APRs by age. Thus, unless otherwise noted, in each context we estimate a regression of the type:

$$(1) \quad F = \alpha + \beta \times \text{Spline}(\text{Age}) + \gamma \times \text{Controls} + \epsilon.$$

Here F is the level of the APR paid by the borrower (or the frequency of fee payment), Controls is a vector of control variables intended to capture alternative explanations in each context (for example, measures of credit risk), and $\text{Spline}(\text{Age})$ is a piecewise linear function that takes consumer age as its argument (with knot points at ages 30, 40, 50, 60 and 70).² We then plot the fitted values for the spline on age. Regressions are either pooled panel or cross-sectional, depending on the context.

Below, we discuss the nature of the products and the prices paid, briefly document the datasets used, and present the regression results and graphs by age. We provide summary statistics for the data sets in the Appendix.

3 Three Financial Choices

In this section, we present evidence for a U-shaped pattern in prices paid for three financial choices: (1) borrowing through home equity loans, (2) borrowing through home equity lines of credit and (3) the use of credit card balance transfer offers. For the first two choices, we are able to tease out the mechanism leading to higher interest payments (namely, mistakes made in estimating house values). The U-shaped pattern by age for the third choice is a relatively clean example and thereby merits greater emphasis.

3.1 Home Equity Loans

3.1.1 Data Summary

We use a proprietary panel dataset constructed with records from a national financial institution that has issued home equity loans and home equity lines of credit. The lender has not specialized in subprime loans or other market segments. Between March and December 2002, the lender offered a menu of standardized contracts for home equity credits.³ Consumers chose (a) either a loan or a credit line; (b) either a first or second lien; and (c) an incremental loan amount and an estimate of her property value, corresponding to a loan-to-value (LTV) ratio of less than 80 percent, between 80 and 90 percent, or

²For instance, in Table 1, the “Age 30-40” spline is: $\max(30, \min(40, \text{Age}))$, the “Age < 30” spline is $\min(30, \text{Age})$, and the “Age > 70” spline is $\max(70, \text{Age})$.

³Other interest rates in the economy varied considerably during this time period. One might therefore ask whether the age results we report below are an artifact of borrowers of different ages happening to disproportionately borrow earlier or later in the sample. We observe no pattern in the distribution by month of borrowing by age over the sample.

between 90 and 100 percent. In effect, the lender offered twelve different contract choices.⁴ For 75,000 such contracts, we observe the contract terms, borrower demographic information (age, years at current job, home tenure), financial information (income and debt-to-income ratio), and risk characteristics (credit (FICO) score, and LTV)⁵ We also observe borrower estimates of their property values and the loan amount requested.

3.1.2 Results

Table 1 reports the results of estimating regressions of APRs (interest rates) on home equity loans on a spline for age and control variables. As controls, we use all borrower-related variables observed by the financial institution that might affect loan pricing, including credit risk measures, house and loan characteristics, and borrower financial and demographic characteristics. The control variables all have the expected sign, and most are statistically significant, although some of them lack economic significance.⁶

The measure of credit risk, the log of the FICO score (lagged three months because it is only updated quarterly), is statistically significant but with a negligible magnitude. Discussions with people who work in the industry reveal that financial institutions generally use the FICO score to determine whether a loan offer is made, but conditional on the offer being made, do not use the score to do risk-based pricing. The results here, and for the other consumer credit products discussed below, are consistent with this hypothesis.

Loan APRs do depend strongly on the absence of a first mortgage (reducing the APR) and whether the property is a second home or a condominium. The absence of a first mortgage reduces the probability of default and raises the amount that might be recovered conditional on a default. Second homes and condominiums are perceived as riskier properties. Log income and log years on the job also have large and negative effects on APRs, as expected, since they indicate more resources available to pay off the loan and perhaps less risk in the latter case. The largest effects on APRs come from dummy variables for LTV ratios between 80 and 90 percent and for ratios greater than 90 percent. This is consistent with different LTV ratios corresponding to different contract choices.⁷

Even after controlling for these variables, we find that the age splines have statistically and economically

⁴We focus on APR payment across contracts for four reasons. First, contracts do not differ in points charged or in other charges to the borrower. Second, we verify that, even conditioning on contract choice, some borrowers pay higher APRs than others. Third, we control for borrower risk characteristics. Fourth, in section 3.3, we show that the residual variation in APRs is explained by the propensity to make an identifiable mistake in the loan acquisition process.

⁵We do not have internal behavior scores (a supplementary credit risk score) for these borrowers. Such scores are performance-based, and are thus not available at loan origination.

⁶Note that although we include all observed variables on the borrower, R-squareds are not 100 percent. In part, this reflects the fact that bank loan pricing models also depend on other variables external to the borrower, such as the cost of funds. Banks may also reassess their lending standards, depending on macroeconomic or other factors. As long as such factors are not correlated with consumer age, the regression coefficients on age will correctly report the impact of age on APR.

⁷We estimate three variants as a specification check. First, we allow the FICO scores, income, and LTV ratios to have quadratic and cubic terms. This allows us to make sure that the nonlinear effects with age that we see are not a consequence of omission of potential nonlinear effects of other control variables. Second and third, we allow the splines to have knot points at every five years, and have a dummy for each age, to ensure that the smoothing caused by the use of ten-year splines does not artificially create a U-shape. In all three cases, our results are not qualitatively or quantitatively changed.

Home Equity Loan APR		
	Coefficient	Std. Error
Intercept	8.1736	0.1069
Log(FICO Score)	-0.0021	0.0001
Loan Purpose-Home Improvement	0.0164	0.0138
Loan Purpose-Rate Refinance	-0.0081	0.0113
No First Mortgage	-0.1916	0.0097
Log(Months at Address)	0.0021	0.0039
Second Home	0.3880	0.0259
Condominium	0.4181	0.0165
Log(Income)	-0.0651	0.0077
Debt/Income	0.0034	0.0002
Log(Years on the Job)	-0.0246	0.0039
Self Employed	0.0106	0.0161
Home Maker	-0.0333	0.0421
Retired	0.0355	0.0225
Age < 30	-0.0551	0.0083
Age 30-40	-0.0336	0.0043
Age 40-50	-0.0127	0.0048
Age 50-60	0.0102	0.0039
Age 60-70	0.0174	0.0076
Age > 70	0.0239	0.0103
LTV 80-90	0.7693	0.0099
LTV 90+	1.7357	0.0111
State Dummies	YES	
Number of Observations	16,683	
Adjusted R-squared	0.7373	

Table 1: The first column gives coefficient estimates for a regression of the APR of a home equity loan on a spline with age as its argument, financial control variables (Log(FICO) credit risk score, income, and the debt-to-income-ratio), and other controls (state dummies, a dummy for loans made for home improvements, a dummy for loans made for refinancing, a dummy for no first mortgage on the property, months at the address, years worked on the job, dummies for self-employed, retiree, or homemaker status, and a dummy if the property is a condominium).




Figure 1: Home equity loan APR by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.

significant effects. Figure 1 plots the fitted values on the spline for age for home equity loans. The line has a pronounced U-shape.⁸ For this and the nine other studies, we present in section 5.2 a formal hypothesis test for the U-shape. To anticipate those results, we reject the null hypothesis of a flat age-based pattern in 9 out of 10 cases.

3.2 Home Equity Lines of Credit

3.2.1 Data Summary

The dataset described in the previous section is also used here.

3.2.2 Results

Table 2 reports a regression of the APRs from home equity lines on a spline for age and the same control variables used for the home equity loans regression. The control variables have similar effects on home equity line APRs as they did on home equity loan APRs.

Fitted values on the age splines, plotted in Figure 2, continue to reveal a pronounced U-shape.

⁸Mortgage and other long-term interest rates were generally falling during this period. Thus, another possible explanation for the observed pattern is that younger and older adults disproportionately borrowed at the beginning of the sample period. However, we found no time-variation in the age distribution of borrowers over the sample period.

Home Equity Line APR		
	Coefficient	Std. Error
Intercept	7.9287	0.0570
Log(FICO Score)	-0.0011	0.0000
Loan Purpose-Home Improvement	0.0551	0.0051
Loan Purpose-Rate Refinance	-0.0386	0.0047
No First Mortgage	-0.1512	0.0054
Log(Months at Address)	-0.0160	0.0019
Second Home	0.3336	0.0132
Condominium	0.4025	0.0079
Log(Income)	-0.1474	0.0037
Debt/Income	0.0044	0.0001
Log(Years on the Job)	-0.0164	0.0020
Self Employed	0.0135	0.0073
Home Maker	-0.0818	0.0215
Retired	0.0139	0.0109
Age < 30	-0.0529	0.0050
Age 30-40	-0.0248	0.0023
Age 40-50	-0.0175	0.0022
Age 50-60	0.0152	0.0035
Age 60-70	0.0214	0.0064
Age > 70	0.0290	0.0154
LTV 80-90	0.6071	0.0050
LTV 90+	1.8722	0.0079
State Dummies	YES	
Number of Observations	66,278	
Adjusted R-squared	0.5890	

Table 2: The first column gives coefficient estimates for a regression of the APR of a home equity lines of credit on a spline with age as its argument, financial control variables (Log(FICO) credit risk score, income, and the debt-to-income-ratio), and other controls (state dummies, a dummy for loans made for home improvements, a dummy for loans made for refinancing, a dummy for no first mortgage on the property, months at the address, years worked on the job, dummies for self-employed, retiree, or homemaker status, and a dummy if the property is a condominium).




Figure 2: Home equity credit line APR by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.

3.3 One Mechanism: Borrower Misestimation of Home Values

The amount of collateral offered by the borrower, as measured by the loan-to-value (LTV) ratio, is an important determinant of loan APRs. Higher LTVs imply higher APRs, since the fraction of collateral is lower. At the financial institution that provided our data, borrowers first estimate their home values, and ask for a credit loan or credit line falling into one of three categories depending on the implied borrower-generated LTV estimate. The categories correspond to LTVs of 80 percent or less; LTVs of between 80 and 90 percent; and LTVs of 90 percent or greater. The financial institution then independently verifies the house value using an industry-standard methodology. The bank then constructs a bank-generated LTV based on the bank’s independent verification process. The bank-LTV can therefore differ from the borrower-LTV.⁹

Loan pricing depends on the LTV category that the borrower falls into and not on the specific LTV value within that category; for example, a loan with an LTV of 60 has the same interest rate as a loan with an LTV of 70, holding borrower characteristics fixed, since the LTVs of both loans are less than 80.¹⁰ If the borrower has overestimated the value of the house, so that the bank-LTV is higher than borrower-LTV, the financial institution will direct the buyer to a different loan with a higher interest rate corresponding to the higher bank-LTV. In such circumstances, the loan officer is also given some discretion to depart from

⁹Bucks and Pence (2006) present evidence that borrowers do not generally have accurate estimates of their house values.

¹⁰We have verified this practice in our dataset by regressing the APR on both the level of the bank-LTV and dummy variables for whether the bank-LTV falls into one of the three categories. Only the coefficients on the dummy variables were statistically and economically significant. Ben-David (2007) also shows that there are discrete jumps in lending rates at LTV cutoff points.




Figure 3: Home equity loan APRs for borrowers who do not make a rate-changing mistake. The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.

the financial institution’s normal pricing schedule to offer a higher interest rate than the officer would have offered to a borrower who had correctly estimated her LTV. If the borrower has underestimated the value of the house, however, the financial institution need not direct the buyer to a loan with a lower interest rate corresponding to the bank-LTV (which is lower in this case than the borrower-LTV); the loan officer may simply choose to offer the higher interest rate associated with the borrower-LTV, instead of lowering the rate to reflect the lower bank-LTV.¹¹

Since the APR paid depends on the LTV category and not the LTV itself, home value misestimation leads to higher interest rate payments if the category of the bank-LTV differs from the category of the borrower-LTV. If, in contrast, the borrower’s estimated LTV was 60, but the true LTV was 70, the borrower would still qualify for the highest quality loan category ($\text{LTV} < 80$) and would not suffer an effective interest rate penalty. We define a *Rate-Changing Mistake (RCM)* to have occurred when the borrower-LTV category differs from the bank-LTV category – for instance, when the borrower estimates an LTV of 85 but the bank calculates an LTV of 75 (or vice versa).¹² We find that, on average, making a RCM increases the APR by 125 basis points for loans and 150 basis points for lines (controlling for other variables, but not age).

To highlight the importance of RCMs, we first study the APR for consumers who do not make a Rate-Changing Mistake. Figures 3 and 4 plot the fitted values from re-estimating the regressions in Table 1 and 2, but now conditioning on borrowers who do not make a RCM. The plots show only slight differences in

¹¹ Even if the financial institution’s estimate of the true house value is inaccurate, that misestimation will not matter for the borrower as long as other institutions use the same methodology.

¹² Recall that the categories are less than 80, 80 to 90, and greater than 90.

**Home Equity Credit Line APRs for Borrowers Who Do Not
Make a Rate-Changing Mistake**




Figure 4: Home equity credit line APRs for borrowers who do not make a rate-changing mistake. The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.

APR paid by age. The APR difference for a home equity loan for a borrower at age 70 over a borrower at age 50 has shrunk from 36 basis points to 8 basis points; for a home equity line of credit, it has shrunk from 28 basis points to 4 basis points. For a borrower at age 20, the APR difference over a borrower at age 50 has shrunk to 3 basis points for home equity loans and 3 basis points for home equity lines of credit. We conclude that, conditional on not making a RCM, the APR is essentially flat with age. So the U-shape of the APR is primarily driven by the Rate-Changing Mistakes.

We next study who makes a RCM. Figures 5 and 6 plot the probability of making a rate-changing mistake by age for home equity loans and home equity lines, respectively. The figures show U-shapes for both. Borrowers at age 70 have a 16 (19) percentage point greater chance of making a mistake than borrowers at age 50 for home equity loans (lines); borrowers at age 20 have a 35 (41) percentage point greater chance of making a mistake than borrowers at age 50. The unconditional average probability of making a rate-changing mistake is 24 percent for loans and 18 percent for lines.

This age effect is consistent with the cost of a RCM calculated above and the additional probability of making a RCM by age. For example, a 70-year old has a 16 and 19 percent additional chance of making a RCM for loans and lines, respectively. Multiplying this by the average APR cost of a RCM for home equity lines and loans of about 150 and 125 basis points, respectively, gives an expected incremental APR paid of about 26 and 23 basis points. These differences are very close to the estimated differences of about 23 basis points for loans (reported in Figure 1) and of about 28 basis points for lines (reported in Figure 2).

We conclude that in the example of home equity lines and loans, we have identified the channel for the U-shape of the APR as a function of age (as always, controlling for other characteristics). Younger and




Figure 5: Propensity of making a Rate Changing Mistake on home equity loans by borrower age. We define a Rate Changing Mistake to have occurred when a borrower’s misestimation of house value causes a change in LTV category and potentially a change in interest rate paid (see the text for a full definition). The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.




Figure 6: Propensity of making a Rate Changing Mistake on home equity credit lines by borrower age. We define a Rate Changing Mistake to have occurred when a borrower’s misestimation of house value causes a change in LTV category and potentially a change in interest rate paid (see the text for a full definition). The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.

older consumers have a greater tendency to misestimate the value of their house, which leads to a Rate-Changing Mistake, which leads them to borrow at an increased APR. On the other hand, for consumers who do not make a Rate-Changing Mistake, the APR is essentially independent of age. Hence, this channel explains quantitatively the higher APR paid by younger and older adults.

Given the large costs associated with a Rate-Changing Mistake, one might ask why borrowers do not make greater effort to more accurately estimate their house values. One possibility is that potential borrowers may not be aware that credit terms will differ by LTV category; or, even if they are aware of this fact, they may not know how much the terms differ by category. This particular aspect of loan pricing may thus be a shrouded attribute, in the sense of Gabaix and Laibson (2006).

3.4 “Eureka” Moments: Balance Transfer Credit Card Usage

3.4.1 Overview

Credit card holders frequently receive offers to transfer account balances on their current cards to a new card. Borrowers pay substantially lower APRs on the balances transferred to the new card for a six-to-nine-month period (a ‘teaser’ rate). However, *new* purchases on the new card have high APRs. The catch is that *payments* on the new card *first* pay down the (low interest) transferred balances, and only subsequently pay down the (high interest) debt accumulated from new purchases.

The optimal strategy during the teaser-rate period, is for the borrower to make all new purchases on her *old* credit card and to make all payments to her old card. The optimal strategy implies that the borrower should make no new purchases with the new card to which balances have been transferred (unless she has already repaid her transferred balances on that card).

Some borrowers will identify this optimal strategy immediately – before making any purchases with the new card. Some borrowers will never identify the optimal strategy. Some borrowers may not initially identify the optimal strategy, but will discover it after one or more pay cycles as they observe their (surprisingly) high interest charges. Those borrowers will make purchases for one or more months, then have a “eureka” moment, after which they will implement the optimal strategy.¹³ We categorize account holders by the speed with which they converge on the optimal strategy (and stop using the “balance transfer” card for new purchases).

3.4.2 Data Summary

We use a proprietary panel data set from several large financial institutions, later acquired by a single financial institution, that made balance transfer offers nationally. The data set contains 14,798 individuals who accepted such balance transfer offers over the period January 2000 through December 2002. The bulk of the data consists of the main billing information listed on each account’s monthly statement, including total payment, spending, credit limit, balance, debt, purchases, cash advance annual percentage

¹³We thank Robert Barro for drawing our attention to this type of potentially tricky financial product.

rates (APRs), and fees paid. We also observe the amount of the balance transfer, the start date of the balance transfer teaser rate offer, the initial teaser APR on the balance transfer, and the end date of the balance transfer APR offer. At a quarterly frequency, we observe each customer’s credit bureau rating (FICO) and a proprietary (internal) credit ‘behavior’ score. We have credit bureau data about the number of other credit cards held by the account holder, total credit card balances, and mortgage balances. We have data on the age, gender, and income of the account holder, collected at the time of account opening. In this sample, borrowers did not pay fees for the balance transfer. Further details on the data, including summary statistics and variable definitions, are available in the Appendix.

3.4.3 Results

About one third of all customers who make a balance transfer do no spending on the new card, thus implementing the optimal strategy immediately. Slightly more than one third of customers who make a balance transfer spend on the new card every month during the promotional period, thus never experiencing a “Eureka” moment. The remaining third of customers experience “Eureka” moments between the first and sixth months.

Figure 7 plots the frequency of Eureka moments for each age group. The plot of those who never experience a “Eureka” moment – that is, who never implement the optimal strategy – is a pronounced U-shape by age. The plot of those who implement the optimal strategy immediately (the “Month One” line) is a pronounced inverted U-shape by age. Plots for Eureka moments in the interior of the time space (that is Eureka moments that occur strictly after Month One) are flat.¹⁴ The No Eureka line implies that the groups with the greatest frequency of maximal confusion are younger adults and older adults. The group with the greatest frequency of optimality is middle-aged adults.

Table 3 reports the results of a regression of a dummy variable for ever having a Eureka moment on a spline for age and controls for credit risk ($\log(\text{FICO})$), education, gender, and $\log(\text{income})$.¹⁵ Credit risk is included because higher scores may be associated with greater financial sophistication. Similarly, we would expect borrowers with higher levels of education to be more likely to experience Eureka moments. The coefficients on the age spline imply that young adults and older adults are less likely to experience Eureka moments.

Figure 8 plots the fitted values of the age splines for the propensity of ever experiencing a “Eureka” moment. Note that, unlike the other figures, higher values indicate a smaller propensity to make mistakes. Consistent with the evidence so far, we observe a performance peak in middle age.

¹⁴Although the average percent of borrowers for each of the intermediate categories is small—on the order of five percent—summing over all the months yields a fraction of borrowers equal to the one-third of total borrowers.

¹⁵Although we report an OLS regression for ease in interpreting the coefficients, we have also run the regression as a logit and found similar results.




Figure 7: Fraction of borrowers in each age group experiencing specific delays. For example, the dashed line plots the fraction of borrowers experiencing no delay to a Eureka moment. These sophisticated borrowers represent a large fraction of middle-aged households and a much smaller fraction of younger and older households.




Figure 8: Propensity of ever experiencing a “Eureka” moment by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$, education, and credit-worthiness.

Propensity of ever experiencing a “Eureka” Moment		
	Coefficient	Std. Error
Intercept	0.2587	0.0809
Age < 30	0.0134	0.0026
Age 30-40	0.0019	0.0005
Age 40-50	-0.0001	0.0000
Age 50-60	-0.0029	0.0009
Age 60-70	-0.0035	0.0008
Age > 70	-0.0083	0.0072
Some High School	-1.6428	0.9570
High School Graduate	-0.6896	0.8528
Some College	-0.4341	0.8944
Associate’s Degree	-0.2439	0.4537
Bachelor’s Degree	0.3280	0.5585
Graduate Degree	0.6574	0.3541
Log(FICO)	0.0102	0.0019
Log(Limit)	0.0120	0.0022
Log(Income)	-0.0044	0.0067
Number of Observations	3,622	
Adjusted R-squared	0.1429	

Table 3: This table reports estimated coefficients from a panel regression of the month in which the borrower did no more spending on the balance transfer card (the “Eureka” moment) on a spline with age as its argument and other control variables.

4 Seven Other Choices

In this section, we present results on all seven other financial choices we studied.

4.1 Credit Cards

4.1.1 Data Summary

We use a proprietary panel dataset from several large financial institutions that offered credit cards nationally, later acquired by a larger financial institution. The dataset contains a representative random sample of about 128,000 credit card accounts followed monthly over a 36 month period (from January 2002 through December 2004). The bulk of the data consists of the main billing information listed on each account’s monthly statement, including total payment, spending, credit limit, balance, debt, purchases and cash advance annual percent rates (APRs), and fees paid. At a quarterly frequency, we observe each customer’s credit bureau rating (FICO) and a proprietary (internal) credit ‘behavior’ score. We have credit bureau data about the number of other credit cards held by the account holder, total credit card balances, and mortgage balances. We have data on the age, gender and income of the account holder, collected at the time of account opening. Further details on the data, including summary statistics and variable definitions, are available in the data Appendix.

Credit Card APR		
	Coefficient	Std. Error
Intercept	14.2743	3.0335
Age < 30	-0.0127	0.0065
Age 30-40	-0.0075	0.0045
Age 40-50	-0.0041	0.0045
Age 50-60	0.0023	0.0060
Age 60-70	0.0016	0.0184
Age > 70	0.0016	0.0364
Log(Income)	-0.0558	0.0803
Log(FICO)	-0.0183	0.0015
Home Equity Balance	0.0003	0.0022
Mortgage Balance	-0.0000	0.0000
Number of Observations	92,278	
Adjusted R-squared	0.0826	

Table 4: This table gives coefficient estimates for a regression of the APR of a credit card on a spline with age as its argument, financial control variables (Log(FICO) credit risk score, income, total number of cards, total card balance, home equity debt balance and mortgage balance).

4.1.2 Results

Table 4 reports the results of regressing credit card APRs on a spline with age as its argument and other control variables. As controls, we again use information observed by the financial institution that may influence pricing. As before, we find that credit scores have little impact on credit card APRs. APRs rise with the total number of cards, though the effect is not statistically significant. Other controls, including the total card balance, log income, and balances on other debt, do not have economically or statistically significant effects on credit card APRs.

Figure 9 plots the fitted values on the spline for age. A U-shape is present, though it is much weaker than the age-based patterns that we document for other financial products.

4.2 Auto Loans

4.2.1 Data Summary

We use a proprietary data set of auto loans originated at several large financial institutions that were later acquired by another institution. The data set comprises observations on 6,996 loans originated for the purchase of new and used automobiles. We observe loan characteristics including the automobile value and age, the loan amount and LTV, the monthly payment, the contract rate, and the time of origination. We also observe borrower characteristics including credit score, monthly disposable income, and borrower age.




Figure 9: Credit card APR by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.

4.2.2 Results

Table 5 reports the results of regressing the APR paid for auto loans on an age-based spline and control variables. FICO credit risk scores again have little effect on the loan terms. Higher incomes lower APRs and higher debt-to-income ratios raise them, though the magnitudes of the effects are small. We also include car characteristics, such as type and age, as one of us has found those variables to matter for APRs in other work (Agarwal, Ambrose, and Chomsisengphet, forthcoming)—though we note that the financial institutions do not directly condition their loans on such variables. We also include loan age and state dummies.

Figure 10 plots the fitted values on the spline for age. The graph shows a pronounced U-shape.

4.3 Mortgages

4.3.1 Data Summary

We use a proprietary data set from a large financial institution that originates first mortgages in Argentina. Using data from one other country provides suggestive evidence about the international applicability of our findings. The data set covers 4,867 owner-occupied, fixed-rate, first mortgage loans originated between June 1998 and March 2000 and observed through March 2004. We observe the original loan amount, the LTV and appraised house value at origination, and the APR. We also observe borrower financial characteristics (including income; second income; years on the job; wealth measures, such as second house ownership, car ownership, and car value), borrower risk characteristics (Veraz score—a credit score similar to the U.S. FICO score—and mortgage payments as a percentage of after-tax income), and

Auto Loan APR		
	Coefficient	Std. Error
Intercept	11.4979	1.3184
Age < 30	-0.0231	0.0045
Age 30-40	-0.0036	0.0005
Age 40-50	-0.0054	0.0005
Age 50-60	0.0046	0.0007
Age 60-70	0.0031	0.0017
Age > 70	0.0091	0.0042
Log(Income)	-0.3486	0.0176
Log(FICO)	-0.0952	0.0059
Debt/Income	0.0207	0.0020
Japanese Car	-0.0615	0.0270
European Car	-0.0127	0.0038
Loan Age	0.0105	0.0005
Car Age	0.1234	0.0031
State Dummies	YES	
Quarter Dummies	YES	
Number of Observations	6,996	
Adjusted R-squared	0.0928	

Table 5: This table gives coefficient estimates from a regression of the APR of an auto loan on a spline with age as its argument, financial control variables (Log(FICO) credit risk score, income, and the debt-to-income ratio), and other controls (state dummies, dummies for whether the car is Japanese or European, loan age, and car age).




Figure 10: Auto loan APR by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as log(income) and credit-worthiness.




Figure 11: APR for Argentine mortgages by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.

borrower demographic characteristics (age, gender, and marital status).

4.3.2 Results

Table 6 reports results of regressing the mortgage APR on an age-based spline and control variables. As controls, we again use variables observed by the financial institution that may affect loan pricing, including risk measures (credit score, income, mortgage payment as a fraction of income, and LTV), and various demographic and financial indicators (gender, marital status, a dummy variable for car ownership, and several others – these coefficients are not reported to save space). The coefficients on the controls are again of the expected sign and generally statistically significant, though of small magnitude.

The coefficients on the age spline are positive below age 30, then negative through age 60 and positive thereafter. Figure 11 plots the fitted values on the spline for age. The figure provides only partial support for the U-shape hypothesis.

4.4 Small Business Credit Cards

4.4.1 Data Summary

We use a proprietary data set of small business credit card accounts originated at several large institutions that issued such cards nationally. The institutions were later acquired by a single institution. The panel data set covers 11,254 accounts originated between May 2000 and May 2002. Most of the business are very small, owned by a single family, and have no formal financial records. The dataset has all information collected at the time of account origination, including the business owner’s self-reported personal

Mortgage APR		
	Coefficient	Std. Error
Intercept	12.4366	4.9231
Age < 30	0.0027	0.0046
Age 30-40	-0.0023	0.0047
Age 40-50	-0.0057	0.0045
Age 50-60	0.0127	0.0093
Age 60-70	0.0155	0.0434
Age > 70	0.0234	0.0881
Log(Income)	-0.2843	0.1303
Log(Credit Score)	-0.1240	0.0217
Debt/Income	0.0859	0.2869
Loan Term	-0.0114	0.0037
Loan Term Squared	-0.0000	0.0000
Loan Amount	-0.0000	0.0000
Loan to Value	0.1845	0.0187
Years on the Job	-0.0108	0.0046
Second Home	0.1002	0.1014
Auto	0.1174	0.0807
Auto Value	0.0000	0.0000
Gender (1=Female)	0.0213	0.0706
Married	-0.0585	0.0831
Two Incomes	-0.1351	0.1799
Married with Two Incomes	-0.0116	0.1957
Employment: Professional	-0.0438	0.1174
Employment:Non-Professional	0.0853	0.1041
Merchant	-0.1709	0.1124
Bank Relationship	-0.2184	0.1041
Number of Observations	4,867	
Adjusted R-squared	0.1004	

Table 6: This table reports the estimated coefficients from a regression of mortgage APR on a spline with age as its argument and financial and demographic control variables.

Small Business Credit Card APR		
	Coefficient	Std. Error
Intercept	16.0601	0.6075
Age < 30	-0.0295	0.0081
Age 30-40	-0.0068	0.0040
Age 40-50	-0.0047	0.0038
Age 50-60	-0.0017	0.0055
Age 60-70	0.0060	0.0209
Age > 70	0.0193	0.0330
Years in Business 1-2	-0.5620	0.1885
Years in Business 2-3	-0.7463	0.1937
Years in Business 3-4	-0.2158	0.1031
Years in Business 4-5	-0.5100	0.0937
Years in Business 5-6	-0.4983	0.0931
Log(FICO)	-0.0151	0.0008
Number of Cards	0.1379	0.0153
Log(Total Card Balance)	<0.0001	<0.0001
Log(Total Card Limit)	<0.0001	<0.0001
Number of Observations	11,254	
Adjusted R-squared	0.0933	

Table 7: This table reports the estimated coefficients from a regression of the APR for small business credit cards on a spline with the business owner’s age as its argument and other control variables (dummies for years in business, log(FICO) credit risk score, number of cards, total card balance, and total card limit).

income, the number of years the business has been in operation, and the age of the business owner. We observe the quarterly credit bureau score of the business owner.

4.4.2 Results

Table 7 reports the results of regressing the APR for small business credit cards on an age-based spline and control variables. As with individual credit card accounts, we control for the FICO score of the business owner, the total number of cards, card balance, and card limit. We also include dummy variables for the number of years the small business has been operating – we expect APRs to fall for businesses with longer operating histories. All control variables are statistically significant and have the expected sign, though only the dummies for years in business have substantial magnitudes.

The APRs are decreasing in the age of the borrower through age 60 and increasing thereafter. Figure 12 plots the fitted values on the spline for age. The graph shows a pronounced U-shape.




Figure 12: Small business credit card APR by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.

4.5 Credit Card Fee Payments: Late Fees

4.5.1 Overview

Certain credit card uses involve the payment of a fee. Some kinds of fees are assessed when terms of the credit card agreement are violated. Other fees are assessed for use of services.

In the next three sections, we focus on three important types of fees: late fees, over limit fees, and cash advance fees.¹⁶ We describe the fee structure for our dataset below.

1. **Late Fee:** A late fee of between \$30 and \$35 is assessed if the borrower makes a payment beyond the due date on the credit card statement. If the borrower is late by more than 60 days once or by more than 30 days twice within a year, the bank may also impose ‘penalty pricing’ by raising the APR to over 24 percent. The bank may also choose to report late payments to credit bureaus, adversely affecting consumers’ FICO scores. If the borrower does not make a late payment during the six months after the last late payment, the APR will revert to its normal (though not promotional) level.
2. **Over Limit Fee:** An over limit fee – also between \$30 and \$35 – is assessed the first time the borrower exceeds his or her credit limit. Over limit violations generate penalty pricing that is analogous to the penalty pricing that is imposed as a result of late fees.
3. **Cash Advance Fee:** A cash advance fee – which is the greater of 3 percent of the amount advanced,

¹⁶Other types of fees include annual, balance transfer, foreign transactions, and pay by phone. All of these fees are relatively less important to both the bank and the borrower. Few issuers (the most notable exception being American Express) continue to charge annual fees, largely as a result of increased competition for new borrowers (Agarwal et al., 2005). The cards in our data do not have annual fees. We study balance transfer behavior using a separate data set below. The foreign transaction fees and pay by phone fees together comprise less than three percent of the total fees collected by banks.

or \$5 – is levied for each cash advance on the credit card. Unlike the first two fees, this fee can be assessed many times per month. It does not cause the imposition of penalty pricing. However, the APR on cash advances is typically greater than the APR on purchases, and is usually 16 percent or more.

Payment of these fees is not generally a mistake. For example, if a card holder is vacationing in Tibet, it may not be optimal to arrange a credit card payment for that month. However, payments of fees are sometimes mistakes, since the fee payment can often be avoided by small and relatively costless changes in behavior. For instance, late fees are sometimes due to memory lapses that could be avoided by putting a reminder in one’s calendar.

We use the same data set as that used for the credit card APR case study discussed above.

4.5.2 Results

Table 8 presents panel regressions for each type of fee. In each of the three regressions, we regress a dummy variable equal to one if a fee is paid that month on an age-based spline and control variables. Hence, the coefficients give the conditional effects of the independent variables on the propensity to pay fees.

The control variables differ from those of the preceding six examples. Now we control for factors that might affect the propensity to pay a fee, which are not necessarily the same as factors that might lead borrowers to default or otherwise affect their borrowing terms. “Bill Existence” is a dummy variable equal to one if a bill was issued last month; borrowers will only be eligible to pay a late fee if a bill was issued. “Bill Activity” is a dummy variable equal to one if purchases or payments were made on the card; borrowers will only be eligible to pay over limit or cash advance fees if the card was used. “Log(Purchases)” is the log of the amount purchased on the card, in dollars; we would expect that the propensity to pay over limit and cash advance fees would be increasing with the amount of purchases. “Log(FICO)” is the credit risk score, and “Log(Behavior)” is an internal risk score created by the bank to predict late and delinquent payment beyond that predicted by the FICO score. Higher scores mean less risky behavior. The scores are lagged three months because they are only updated quarterly. We would expect the underlying behavior leading to lower credit risk scores would lead to higher fee payment. “Debt/Limit” is the ratio of the balance of credit card debt to the credit limit; we would expect that having less available credit would raise the propensity to pay over limit fees, and possibly other fees.

For late fee payments – column one of the table – all control variables have the expected signs and are statistically significant, though they are also small in magnitude. Note that some control variables may partly capture the effects of age-related cognitive decline on fees. For example, if increasing age makes borrowers more likely to forget to pay fees on time, that would both increase the propensity to pay late fees and decrease credit and behavior scores. Hence, the estimated coefficients on the age splines may understate some age-related effects.

Coefficients on the age splines are uniformly negative for splines through age 50; negative or weakly

	Late Fee		Over Limit Fee		Cash Adv. Fee	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Intercept	0.2964	0.0446	0.1870	0.0802	0.3431	0.0631
Age < 30	-0.0021	0.0004	-0.0013	0.0006	-0.0026	0.0011
Age 30-40	-0.0061	0.0003	-0.0003	0.0001	-0.0004	0.0002
Age 40-50	-0.0001	0.0000	-0.0002	0.0000	-0.0002	0.0000
Age 50-60	-0.0002	0.0000	-0.0002	0.0000	-0.0003	0.0000
Age 60-70	0.0004	0.0002	0.0003	0.0001	0.0004	0.0000
Age > 70	0.0025	0.0013	0.0003	0.0001	0.0004	0.0000
Bill Existence	0.0153	0.0076	0.0104	0.0031	0.0055	0.0021
Bill Activity	0.0073	0.0034	0.0088	0.0030	0.0055	0.0021
Log(Purchases)	0.0181	0.0056	0.0113	0.0023	0.0179	0.0079
Log(Behavior)	-0.0017	0.0000	-0.0031	0.0012	-0.0075	0.0036
Log(FICO)	-0.0016	0.0007	-0.0012	0.0003	-0.0015	0.0005
Debt/Limit	-0.0066	0.0033	0.0035	0.0013	0.0038	0.0012
Acct. Fixed Eff.	YES		YES		YES	
Time Fixed Eff.	YES		YES		YES	
Number of Obs.	3.9 Mill.		3.9 Mill.		3.9 Mill.	
Adj. R-squared	0.0378		0.0409		0.0388	

Table 8: This table reports coefficients from a regression of dummy variables for credit card fee payments on a spline for age, financial control variables (log(FICO) credit risk score, internal bank behavior risk score, debt over limit) and other control variables (dummies for whether a bill existed last month, for whether the card was used last month, the dollar amount of purchases, and account- and time- fixed effects).

positive for the spline between age 50 and 60; and positive with increasing slope for splines above age 50.

The top line in Figure 13 plots fitted values for the age splines for the late fee payment regression.¹⁷

4.6 Credit Card Fee Payments: Over Limit Fees

The second column of Table 8 presents regression results for the over limit fee, on the same controls and age splines that were used for the late fee. Results are similar to those generated in analysis of the late fee.

The bottom line in Figure 13 plots fitted values of the age splines for the over limit fee payment regression.

4.7 Credit Card Fee Payments: Cash Advance Fees

The second column of Table 8 presents regression results for the cash advance fee, on the same controls and age splines that were used for the late fee. Results are similar to those generated in analysis of the late fee and the over limit fee.

¹⁷In Agarwal, Driscoll, Gabaix, and Laibson (2008), we study this propensity of paying fees as the interaction of learning from the payment of past fees and forgetting.




Figure 13: Frequency of fee payment by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.

The middle line in Figure 13 plots fitted values of the age splines for the cash advance fee payment regression.

5 The Peak of Performance

5.1 Locating the Peak of Performance

Visual inspection of the age splines for the ten case studies suggests that fees and interest rates paid are minimized in the late 40s or early 50s. To estimate the minimum more precisely, we re-estimate each model, replacing the splines between 40 and 50 and 50 and 60 with a single spline running from 40 to 60, and the square of that spline. This enables us to more precisely estimate the local properties of the performance curve.

In other words, we run the following regression, where F is the outcome associated respectively with each of the 10 studies:

$$(2) \quad F = \alpha + \beta \times \text{Spline}(\text{Age})_{\text{Age} \notin [40,60]} + \gamma \times \text{Controls} + \epsilon \\ + a \times \text{Spline}(\text{Age})_{\text{Age} \in [40,60]} + b \cdot \text{Spline}(\text{Age})_{\text{Age} \in [40,60]}^2.$$

Here $\text{Spline}(\text{Age})$ is a piecewise linear function that takes consumer age as its argument (with knot points at ages 30, 40, 60 and 70). $\text{Spline}(\text{Age})_{\text{Age} \notin [40,60]}$ represents the splines outside of the $[40, 60]$ age range, while $\text{Spline}(\text{Age})_{\text{Age} \in [40,60]}$ is the linear spline with knot points at 40 and 60. Hence, for age between 40

	Age of Peak Performance	Standard Error
Home Equity Loans-APR	55.9	4.2
Home Equity Lines-APR	53.3	5.2
Eureka Moment	45.8	7.9
Credit Card-APR	50.3	6.0
Auto Loans-APR	49.6	5.0
Mortgage-APR	56.0	8.0
Small Business Credit Card-APR	61.8	7.9
Credit Card Late Fee	51.9	4.9
Credit Card Over Limit Fee	54.0	5.0
Credit Card Cash Advance Fee	54.8	4.9
Average over the 10 Studies	53.3	

Table 9: Age at which financial mistakes are minimized, for each case study

and 60, the above formulation is implicitly quadratic in age:

$$F = Controls + a \times Age + b \times Age^2.$$

The peak of performance is defined as the value that minimizes the above function:

$$(3) \quad Peak = -a / (2b).$$

We calculate the asymptotic standard errors on $Peak$ using the delta method, so that the standard error of $Peak$ is the standard error associated with the linear combination: $-1/(2b) \cdot (\text{Coefficient on age}) + a/(2b^2) \cdot (\text{Coefficient on } age^2)$.

In Table 9, we report the location of the “age of reason”: the point at which financial mistakes are minimized. The mean age of reason appears to be at 53.3 years. The standard deviation calculated by treating each study as a single data point is 4.3 years.

Formal hypothesis testing ($H_0: a + 2b \times 53 = 0$) shows that only the location of the Eureka moment is statistically different from 53 years. Interestingly, the Eureka task is arguable the most most dependent on analytic capacity and least dependent on experience (since the kinds of balance transfer offers that we study were new financial products when our data was collected). It is not surprising that the peak age for succeeding at that task would be earlier than the peak age for the other tasks. However, since we do not have a rigorous measure of the “difficulty” of a task, the interpretation of the Eureka case remains speculative.

5.2 Formal Test of a Peak of Performance Effect

Table 9 allows us do a formal test for a peak effect. In regression (2), the null hypothesis of a peak effect is: (i) $b > 0$, and (ii) $Peak = -a / (2b) \in [40, 60]$. Together these conditions imply that mistakes

follow a U-shape, with a peak that is between 40 and 60 years of age.

For criterion (i), we note that the b coefficients are positive for all 10 studies. For 9 of the 10 studies, b is significantly different from zero (the credit card APR study is the exception).¹⁸ For criterion (ii), Table 9 shows that a peak in the 40-60 age range can not be rejected for all ten studies.

6 Possible Explanations

Each credit market has idiosyncratic factors that contribute to the hump-shaped age patterns that we have measured. The recurrence of that hump-shaped pattern across all ten outcomes suggests that the regularities may also have some common underlying explanations. In this section, we discuss such common explanations. We argue that the leading candidates are cognitive age effects, selection effects, and cohort effects. However, our data does not enable us to identify the respective contributions of these factors.

6.1 Cognitive Age Effects

One possible explanation for the U-shaped pattern of mistakes is a combination of diminishing returns to learning and age-based declines in analytic function. Relatively young borrowers have low levels of experience and a high degree of analytic function, while older borrowers have high levels of experience but relatively lower levels of analytic function. We discuss these mechanisms below and explain how these offsetting lifecycle trends produce a hump-shaped pattern in financial sophistication.

Analytic cognitive function can be measured in many different ways, including tasks that evaluate working memory, reasoning, spatial visualization, and cognitive processing speed (see Figure 14). Analytic function shows a robust age pattern in cross-sectional datasets. Analytic function is strongly negatively associated with age in adult populations (Salthouse, 2005; and Salthouse, forthcoming). On average, analytic function falls by 2 to 3 percent of one standard deviation¹⁹ with every incremental year of age after age 20. This decline is steady from age 20 to age 90 (see Figure 15).

The measured age-related decline in analytic performance results from both age effects and cohort effects, but the available panel data implies that the decline is primarily driven by age effects (Salthouse, Schroeder, and Ferrer, 2004).²⁰ Medical pathologies represent one important pathway for age effects. For instance, dementia is primarily attributable to Alzheimer’s Disease (60%) and vascular disease (25%). The prevalence of dementia doubles with every five additional years of lifecycle age (Fratiglioni, De Ronchi, and Agüero-Torres, 1999). There is a growing literature that identifies age-related changes in cognition (see Park and Schwarz, 1999; and Denburg, Tranel, and Bechara 2005).²¹

¹⁸To save space, we only report the t -statistics associated with the b coefficients. Following the order of Table 9, they are: 2.20, 4.55, 7.80, 8.77, 17.05, 1.61, 4.57, 2.91, 3.08, 2.67.

¹⁹This is a standard deviation calculated from the entire population of individuals.

²⁰See Flynn (1984) for a discussion of cohort effects.

²¹Mather and Carstensen (2005) and Carstensen (2006) identify a different type of age-variation in cognitive preferences. Subjects with short time horizons or older ages attend to negative information relatively less than subjects with long time horizons or younger ages.




Figure 14: Four tasks used to measure cognitive function. Source: Salthouse (forthcoming).




Figure 15: Age-normed results from four different cognitive tests. The Z-score represents the age-contingent mean, measured in units of standard deviation relative to the population mean. More precisely, the Z-score is (age-contingent mean minus population mean) / (population standard deviation). Source: Salthouse (forthcoming).




Figure 16: Hypothesized relation between general task performance and age. Analytic capital declines with age and experiential capital increase with age. This generates the hypothesis that general task performance (which uses both analytic and experiential capital) first rises and then declines with age.

Age-driven declines in *analytic* function are partially offset by age-driven increases in *experience*.²² Most day-to-day tasks rely on both analytic and experiential human capital – e.g., buying the right amount of milk at the grocery store. For most tasks, we hypothesize that performance is hump-shaped with respect to age. Formally, this would result from the following conditions (1) general task performance is determined by the sum of analytic capital and experiential capital, (2) experiential capital is accumulated with strictly diminishing returns over the lifecycle, and (3) analytic capital falls linearly (or concavely) over the lifecycle (see Figure 15). Then general task performance will be hump-shaped with respect to age. Figure 16 illustrates this case.

A Possible Interpretation of the Location of the Performance Peak This hypothesis above also provides us with a possible explanation of the location of the peak of performance. We hypothesize that peak performance reflects a trade-off between experience (that is accumulated with diminishing returns) and analytic function (that declines roughly linearly after age 20). If so, the sooner people start experimenting with the product, the earlier the peak of performance should be. For instance, suppose Analytic Capital declines linearly with age, so that Analytic Capital = $\alpha - age/\beta$. Suppose that Experiential Capital is accumulated with diminishing returns – for instance, Experiential Capital = $\ln(age - \gamma age_0)$, where age_0 is the actual age at which people start using the product, and $\gamma age_0 < age_0$ is the effective age at which people start using the product (so $\gamma < 1$). The effective age is less than the actual age, since consumers get indirect experience (observation and advice) as a result of their interactions with slightly older individuals who use the product. The model implies that peak performance occurs at $Peak = \beta + \gamma age_0$. Hence,

²²Experience may either be directly acquired or it may be indirectly acquired from peers. As social networks are built up over the lifecycle, external sources of experience become better and better developed. However, such social networks tend to fray as individuals retire and leave well-developed work-based social networks. Peer mortality also contributes to a late-life decline of social networks. All of these channels suggest that experiential knowledge embodied in social networks follows a hump-shaped lifecycle pattern.

peak performance is later when people start using the product later in life.

To evaluate this hypothesis for each financial product, we first construct the distribution of the ages of the users of this product in our data set and calculate the age at the 10th percentile of the distribution, which we call “age_{10%}”. It is a crude proxy for the age at which people start using the product. We then regress the location of the peak of performance on age_{10%}. We find: $Peak = 33 + 0.71 \times age_{10\%}$, ($R^2 = 0.62, n = 10$; the s.e. on the coefficients are respectively 5.7 and 0.19).²³ We reject the null hypothesis of no relationship between *Peak* and *age*_{10%}. Products that are used later in life tend to have a later performance peak.

This analysis only provides suggestive evidence. It is important to measure this correlation and to test the hypothesized mechanism with other data sets.

6.2 Selection effects

The cross-sectional age effects that we measure are probably also partially attributable to differences in the pool of borrowers by age group: a selection effect. For example, in the *total* population of US households, retirees borrow less than other adults (matching a prediction of the lifecycle consumption model). Older adults who are using home equity loans and lines of credit may therefore be unrepresentative of the population of *all* older adults. Likewise, older adults who are using home equity loans and lines of credit might be less financially savvy than 20- to 60-year-old borrowers (since borrowing is not a “bad signal” at these lower ages).²⁴

In this subsection, we look at measurable financial characteristics by age in our sample. We first ask whether the older adults in our sample have comparable socio-economic characteristics to the other adults in our sample. Figure 17 shows that credit-worthiness (FICO) scores on home equity loans and lines show a U-shape by age distribution. In other words, older and younger borrowers in our sample, are less risky than middle-aged borrowers in our sample. Figure 18 shows that LTV ratios decline substantially with age, indicating that older borrowers in our sample are devoting a relatively smaller fraction of their assets to servicing home equity loans and lines. Likewise, Appendix Table A8 shows that debt levels rise from age 25 to 50 and then decline to age 75.²⁵ Finally, in section 6.4 (below), we report that default rates are *lower* for older borrowers in our sample relative to younger borrowers. These analyses suggest that the older borrowers in our sample compare favorably (in terms of risk characteristics) to the younger borrowers in our sample.

We also find that average income for home equity loan borrowers rises from \$76,000 for those aged less than 30, and about the same to those between 30 and 40 to a peak of \$88,500 for those between 40 and 50, and then declines to about \$69,000 for those between 60 and 70 and \$62,000 for those over 70. These

²³The effect is robust to the choice of the 10th percentile. For instance, the correlation between Peak age and Median age (of users for the product, in our data set) is 0.83.

²⁴They could also be riskier, a hypothesis that we consider (and reject) in section 6.4.

²⁵In comparing the debt levels with those from survey data, one should bear in mind that these data, from the lender, may be higher than those reported by individuals. Gross and Souleles (2002a, 2002b) document the under-reporting of credit card debt by individuals.




Figure 17: This figure plots the FICO (credit-worthiness) scores of home equity loan and line of credit borrowers by age. A high FICO score means a high credit-worthiness.

relative income levels are consistent with the pattern of earnings measured in studies of representative populations of US households (e.g. Gourinchas and Parker, 2001).

All in all, the characteristics that we observe do not point to strong negative selection effects for the older adults in our sample. This is probably due to the fact that our sample only includes *prime* borrowers, and is thus truncated in a way that reduces some selection effects that would otherwise arise. In our sample average FICO score levels are much higher than those for the population as a whole; default rates are lower; and income levels are higher.

Figure 19 shows the results of re-estimating the regressions for home equity loans and lines of credit, now dropping data on all borrowers over the age of 60. There is less reason to believe that the pool of borrowers below 60 are subject to the sample selection issues discussed above. The results still show a U-shape, albeit a somewhat less pronounced one.²⁶

These analyses do not lend support to selection effects, but the analyses also do not rule selection effects out. For example, it is possible that only unsophisticated older adults borrow; and such unsophisticated older adults might still have good credit scores and low default rates. Hence, we believe that selection effects probably do contribute to our results, but it is hard to quantify that impact. However, we do not believe that selection effects explain the entire U-shaped pattern, particularly the left-hand arm of the U. Indeed, the kinds of selection effects that we have discussed above would predict a monotonic decline in financial performance over the entire lifecycle.

²⁶This graph also reinforces the arguments above that potential higher riskiness of borrowers above age 60 is likely not responsible for the results.




Figure 18: This figure plots the loan-to-value (LTV) ratio of home equity loan and line of credit borrowers by borrower age.




Figure 19: This figure plots the residual effect of age on home equity loan and line APRs, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness. Observations on borrowers over age 60 have been dropped.

6.3 Cohort Effects

Older borrowers in our cross-section are likely to make relatively less sophisticated financial choices because they belong to cohorts that have less human capital than younger cohorts (e.g. Flynn, 1984). For example, older cohorts may be less mathematically literate than younger cohorts. In addition, older cohorts may use less sophisticated search technologies – for instance, older cohorts may be less inclined to use the internet to compare financial products. Finally, older cohorts may have grown up with different financial products than the products that are now available from financial intermediaries.

In the absence of a true panel dataset with information of twenty years or more, we cannot measure the role of cohort effects to explain the U-shape relative to other explanations. However, several facts make us think that cohort effects do not provide a *complete* explanation for the U-shaped patterns in our data.

First, education-based cohort effects do not explain the pattern of *declining* mistakes that we observe over the *first* half of the adult lifecycle. Second, we observe the U-shaped pattern over a broad range of products; while some of these products, such as mortgages, have seen substantial changes in their institutional characteristics over time, others, such as auto loans, have not. Third, if cohort effects were dominant, we might expect to see differences in APRs between male and female borrowers on the grounds that the current cohort of older female borrowers has tended to be less involved in financial decision-making than their male contemporaries. Figures 20 and 21 plot the residual effects of age on home equity line and loan APR for female and male borrowers, respectively. Both show a U-shaped pattern by age, with no substantive difference between the two groups.

Finally, for two products—auto loans and credit cards—we have data from 1992, ten years earlier than the data used for our other studies. Figures 22 and 23 replicate the plots of the fitted values of the effects of age on APR for this earlier dataset. Both plots show the same U-shape, with the minimum in the early 50s (like our results using *later* cross-sections). If our findings were driven by cohort effects, the U-shape should not reproduce itself in cross-sections from different years.

In summary, cohort effects are probably present in our data, though we doubt that they explain all of the U-shaped pattern. Cohort effects are most likely to make some contribution to the decline in performance that we measure after middle age. The improvement in performance up to middle age is harder to explain with cohort stories, though some preference-based cohort story might be generating this pattern.²⁷

6.4 Risk effects

Some of our results could be driven by unobserved variation in default risk that is not reflected in the risk measures – like FICO – that we use as control variables. For instance, the U-shape of APRs could be due to a U-shape of default by age. We test this alternative hypothesis by analyzing default rates of credit cards, auto loans, and home equity loans and credit lines. Specifically, we estimate a linear regression in which the default rate is modeled as a weighted sum of an age spline, log income, and all of the standard risk measures that are in our data.

²⁷See Malmendier and Nagel (2007) for examples of preference based cohort effects.




Figure 20: This figure plots the residual effect of age on home equity loan and line APRs for women, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.




Figure 21: This figure plots the residual effect of age on home equity loan and line APRs for men, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.




Figure 22: Auto loan APR by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as log(income) and credit-worthiness. Data is from 1992.




Figure 23: Credit card APR by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as log(income) and credit-worthiness. Data is from 1992.




Figure 24: Default frequency by borrower age. The figure plots the residual effect of age, after controlling for other observable characteristics, such as $\log(\text{income})$ and credit-worthiness.

We plot fitted values in Figure 24. None of the graphs is U-shaped. On the contrary, home equity loans and lines show a pronounced *inverted* U-shape, implying that the young and old have *lower* default rates. Credit cards and auto loans also show a slight inverted U-shape. Hence, Figure 24 contradicts the hypothesis that our APR results are driven by an unmeasured default risk. Finally, note that age-dependent default risks would not explain the observed patterns in credit card fee payments or suboptimal use of balance transfers.

6.5 Opportunity Cost of Time

Some age effects could be generated by age-variation in the opportunity cost of time (Aguilar and Hurst, forthcoming). However, standard opportunity-cost effects would predict that retirees pay lower prices, which is not what we observe in our data. Nevertheless, our findings and those of the Aguilar and Hurst article are not contradictory. Shopping for a familiar commodity – for instance, a gallon of milk – is much less analytically demanding than shopping for a complicated and somewhat unfamiliar product that can differ across many dimensions – for instance, a mortgage. Hence, we are not surprised to see older adults shop more effectively for food at the same time that they lose ground in relatively complex domains – like shopping for mortgage. In addition, shopping at stores and supermarkets may be a more pleasant activity than shopping at banks and other lenders, leading consumers to do more intensive shopping for food than for loans.

6.6 Discrimination and Other Supply Factors

The presence of age effects might also be interpreted as evidence for some kind of age discrimination. Banks may explicitly choose to charge older and younger borrowers higher APRs, or may simply market products that happen to have higher APRs or fees more aggressively to the young or old. We believe these explanations to be unlikely for two reasons. First, the U-shaped pattern shows up in contexts such as fee payments and failures to optimally use balance transfer offers in which discrimination is not relevant (since the products are the same and all card holders face the same rules). Second, firms avoid age discrimination for legal reasons. Penalties for age discrimination from the Fair Lending Act are substantial (like the resulting negative publicity).²⁸ We discuss the issue of supply factors further in the section on market equilibrium below.

7 Discussion

7.1 Market Equilibrium

The markets we describe may seem paradoxical. First, the markets appear to be competitive, since many firms compete to sell commodity credit products. However, consumers with ostensibly identical risk characteristics, fare differently, implying that the goods being sold are somehow de-commodified.

Markets like this have been described in the industrial organization literature. A first generation of models (e.g., Salop and Stiglitz 1977, Ellison 2005, and the citations therein) emphasizes heterogeneous search costs, which are costs of discovering the products of different firms. A second generation (sometimes under the name of “behavioral industrial organization,” e.g. Ellison 2005 and Gabaix and Laibson, 2006) emphasizes heterogeneous levels of consumer rationality. For instance, a balance transfer offer provides a rent that only some consumers are smart enough to exploit. Some consumers unravel the shrouded attribute – the “catch” that they should transfer balances to the card but make no purchases with it – and some consumers never get it. In the market equilibrium (with competition and free entry), the naive consumers end up paying above marginal cost, subsidizing the sophisticated consumers, who pay below marginal cost. From an ex ante point of view, the market is fully competitive, since expected firm profits are zero.²⁹

²⁸Charles, Hurst and Stephens (2006) show that racial differences in lending rates exist at auto finance companies, but not at banks.

²⁹One may ask how such a potentially inefficient equilibrium can persist in a competitive environment. An answer is proposed in Gabaix and Laibson (2006): the cross-subsidy from naives to sophisticates makes the market more “sticky.” The sophisticates may not have an incentive to switch from the firms with shrouded attributes (at which they are getting cross-subsidies). Such stickiness explains why these equilibria are robust even when the equilibria are inefficient. For example, in the home equity loan and line of credit markets, shrouding may take the following form. The degree to which the pricing depends on the loan-to-value ratio may well not be apparent to the consumer; the bank will generally not give them a full schedule of APRs by LTV bucket. Thus, consumers may be unaware of the cost of a rate-changing mistake.

7.2 On the Economic Magnitude of the Effects

The effects we find have a wide range of dollar magnitudes. Appendix Table A7 provides estimates of the magnitudes for borrowers aged 75 and 25 relative to borrowers aged 50.³⁰ For instance, for home-equity lines of credit, 75-year-olds pay about \$265 more each year than 50-year-olds, and 25-year-olds pay about \$295 more. For other quantities, say, credit card fees, the implied age differentials are small – roughly \$10-\$20 per year for each kind of fee.³¹

None of the economic decisions that we study is of significant economic relevance on its *own*, but rather that there is a U-shaped pattern of payments that may merit economists’ attention because it points to a phenomenon that might apply to many decision domains. An important question is whether the U-shape of mistakes translates into other decision domains, including savings choices, asset allocation choices, and healthcare choices.

7.3 Related Work

Other authors have studied the effects of aging on the use of financial instruments. Korniotis and Kumar (2007) examine the performance of investors from a major U.S. discount brokerage house. They use census data to impute education levels and data from the Survey of Health, Aging and Retirement in Europe to estimate a model of cognitive abilities. They find that investors with cognitive declines earn annual returns between 3-5 percentage points lower on a risk adjusted basis.

In their work on financial literacy, Lusardi and Mitchell find evidence consistent with an inverse-U shape of financial proficiency. Lusardi and Mitchell (2006) find a decline in financial knowledge after age 50. Lusardi and Mitchell (2007) also find an inverse U-shape in the mastery of basic financial concepts, such as the ability to calculate percentages or simple divisions.

After some of our presentations, other researchers have offered to look for age patterns of financial mistakes in their own data sets. Lucia Dunn has reported to us that the Ohio State Survey on credit cards shows a U-shaped pattern of credit card APR terms by age (Dunn, personal communication). Fiona Scott Morton has reported that in her data set of indirect auto loans (made by banks and finance companies using the dealer as an intermediary; see Scott Morton et al., 2003), loan markups show a U-shaped pattern (Scott Morton, personal communication). Luigi Guiso finds that, when picking stocks, consumers achieve their best Sharpe ratios at about age 43, and this effect appears to be entirely driven by the participation margin (Guiso, personal communication). Ernesto Villanueva finds that mortgage APRs in Spanish survey data (comparable to the U.S. Survey of Consumer Finances) are U-shaped by age (Villanueva, personal communication).

³⁰We use the average levels of debt at ages 25 and 75, which differ from those at age 50. We provide those figures in Appendix Table A8.

³¹As shown in the columns in Table A7, a difference in late fee probability of 2% per month, and a fee amount \$35, leads to a total extra yearly expense of \$8.40. Note, however, that some of these fees, if paid too often, can trigger “penalty pricing,” in which interest rates ten percentage points or higher are levied on card balances, thus greatly increasing the cost of fee payment. See Agarwal, Driscoll, Gabaix and Laibson (2006) for further discussion.

A relationship between earning and performance has been noted in many nonfinancial contexts. Survey data suggests that labor earnings peak around age 50 (Gourinchas and Parker, 2002) or after about 30 years of experience (Murphy and Welch, 1990). Shue and Luttmer (2006) find that older and younger voters disproportionately make more errors in voting.

Aguiar and Hurst (2007, forthcoming) demonstrate that older adults find lower prices for everyday items by spending more time shopping around. In contrast, we find that older adults seem to make more mistakes in personal financial decision-making. We reconcile these findings by noting that financial products require more analytic ability than everyday items (like food or clothing). Moreover, shopping for financial products may be less pleasurable.

Turning to purely noneconomic domains, there is a literature on estimating performance peaks in professional athletics and other competitive areas. Fair (1994, 2005, 2007) estimates the effects of age declines in baseball and chess, among other sports. Simonton (1988) provides a survey.

A new literature in psychology and economics reports systematic differences in “rationality” between groups of people. Benjamin, Brown, and Shapiro (2006) find that subjects with higher test scores, or less cognitive load, display fewer behavioral biases. Frederick (2005) identifies a measure of “analytical IQ”: people with higher scores on cognitive ability tasks tend to exhibit fewer/weaker psychological biases. While this literature is motivated by experimental data (where it is easier to control for unobservables), we rely on field data in our paper. Similarly, Massoud, Saunders, and Schnolnick (2006) find that more educated people make fewer mistakes on their credit cards, and Stango and Zinman (2007) find evidence that more naive consumers make mistakes across a range of financial decisions.

Several researchers have looked at the response of consumers to low, introductory credit card rates (‘teaser’ rates) and at the persistence of otherwise high interest rates. Shui and Ausubel (2004) show that consumers prefer credit card contracts with low initial rates for a short period of time to ones with somewhat higher rates for a longer period of time, even when the latter is *ex post* more beneficial. Consumers also appear ‘reluctant’ to switch contracts. DellaVigna and Malmendier (2004) theorize that financial institutions set the terms of credit card contracts to reflect consumers’ poor forecasting ability over their future consumption.

Many of those effects are discussed in “behavioral industrial organization,” a literature that documents and studies markets with behavioral consumers and rational firms: examples from this literature include DellaVigna and Malmendier (2004), Gabaix and Laibson (2006), Heidhues and Koszegi (2006), Malmendier and Devin Shanthikumar (2005), Mullainathan and Shleifer (2005), Oster and Scott Morton (2005), Spiegler (2006). In some of those papers, it is important to have both naive and sophisticated consumers (Campbell 2006). Our paper suggests that those naive consumers will disproportionately be younger and older adults.

Bertrand et al. (2006) find that randomized changes in the “psychological features” of consumer credit offers affect adoption rates as much as variation in the interest rate terms. Ausubel (1991) hypothesizes that consumers may be over-optimistic, repeatedly underestimating the probability that they will borrow, thus possibly explaining the stickiness of credit card interest rates. Callem and Mester (1995) use the 1989 Survey of Consumer Finances (SCF) to argue that information barriers create high switching costs

for high-balance credit card customers, leading to persistence of credit card interest rates, and Calem, Gordy, and Mester (2005) use the 1998 and 2001 SCFs to argue that such costs continue to be important. Kerr and Dunn (2002) use data from the 1998 SCF to argue that having large credit card balances raises consumers' propensity to search for lower credit card interest rates. Kerr, Cosslett, and Dunn (2004) use SCF data to argue that banks offer better lending terms to consumers who are also bank depositors (and about whom the bank would thus have more information).

A literature analyzes heuristics and biases in financial decision-making. For instance, Benartzi and Thaler (2002) show that investors prefer the portfolios chosen by other people rather than the ones chosen by themselves, a pattern that suggests that task difficulty prevents people from reaching an optimal decision. Benartzi and Thaler (forthcoming) also document the use of a number of sometimes inappropriate heuristics. Our findings imply that the U-shaped pattern of financial mistakes should also be found in the examples that Bernatzi and Thaler document.

A number of researchers have written about consumer credit card use. Our work most closely overlaps with that of Agarwal et al. (2005), who use another large random sample of credit card accounts to show that, on average, borrowers choose credit card contracts that minimize their total interest costs net of fees paid. About 40 percent of borrowers initially choose suboptimal contracts. While some borrowers incur hundreds of dollars of such costs, most borrowers subsequently switch to cost-minimizing contracts. The results of our paper complement those of Agarwal et al. (2005), since we find evidence of learning to avoid fees and interest costs given a particular card contract. Other authors have used credit card data to evaluate more general hypotheses about consumption. Agarwal, Liu, and Souleles (forthcoming) use credit card data to examine the response of consumers to the 2001 tax rebates. Gross and Souleles (2002a) use credit card data to argue that default rates rose in the mid-1990s because of declining default costs, rather than a deterioration in the credit-worthiness of borrowers. Gross and Souleles (2002b) find that increases in credit limits and declines in interest rates lead to large increases in consumer debt. Ravina (2005) estimates consumption Euler equations for credit card holders and finds evidence for habit persistence.

Finally, from a methodological perspective our work is related to recent research that studies age variation along other dimensions. For example, Blanchflower and Oswald (2007) report that well-being is U-shaped over the lifecycle controlling for observable demographic characteristics. The trough occurs in the 40s.

7.4 Some Open Questions for Future Research

Our findings suggest several directions for future research.

First, it would be useful to measure such effects in other decision domains. We have described a simple procedure for this: (1) identify the general shape of age effects, as in equation (1), using controls and age splines; (2) estimate a linear-quadratic form to localize the peak of performance, as in equations (2) and (3).

Second, it may be possible to develop models that predict the location of peak performance. There is a growing consensus that analytically intensive problems – like mathematics – are associated with younger

peak ages (see Simonton 1988, Galenson 2005, and Weinberg and Galenson 2005). Analogously, problems that require more experiential training have older peak ages. For instance, Jones (2006) finds that the peak age for natural scientists has drifted higher over the twentieth century. Relative to 100 years ago, more experience now needs to be accumulated to reach the cutting edge of scientific fields.

In our last case study, we found that what is arguably the most analytically demanding task – deducing the best way to exploit “interest-free” balance transfers – is associated with the youngest age of peak performance. It would be useful to assess the generality of this association between analytically demanding problems and young peak ages.

Advice markets may solve many of the potential problems identified in this paper. On the other hand, advice markets may not function efficiently because of information asymmetries between the recipients and the providers of advice (Dulleck and Kerschbamer, 2006). It may be particularly interesting to study the advice market for retirees who control historically unprecedented amounts of financial wealth.

8 Conclusion

We find that middle-aged adults borrow at lower interest rates and pay lower fees in ten financial markets. Our analysis suggests that this fact is not explained by age-dependent risk factors. For example, FICO scores and default rates would predict the opposite pattern of age variation.

We believe that our findings are driven by three complementary factors: age-related cognitive effects, selection effects, and cohort effects. We are unable to disentangle the contributions of each of these factors. We speculate that, to the extent they are present, cohort and selection effects have larger effects on the reported outcomes for older borrowers, since the natural cohort- and selection-stories for younger borrowers imply that those groups should on average be paying lower prices (controlling for observable risk characteristics).

Whatever the mechanisms, there is a robust relationship between age and financial sophistication in cross-sectional data, after controlling for available consumer characteristics. Future research should untangle the different forces that give rise to these effects. Measuring age-related cognitive effects will be particularly important because cognitive decline over the lifecycle has significant implications for the financial decisions that older adults are now routinely expected to make.

References

- Agarwal, Sumit, Brent Ambrose, and Souphala Chomsisengphet (forthcoming) “Information Asymmetry and the Automobile Loan Market.”
- Agarwal, Sumit, Souphala Chomsisengphet, Chunlin Liu and Nicholas S. Souleles (2005). “Do Consumers Choose the Right Credit Contracts?” Mimeo, The Wharton School, University of Pennsylvania.

- Agarwal, Sumit, John C. Driscoll, Xavier Gabaix and David I. Laibson (2008). "Learning in the Credit Card Market." Mimeo, Harvard University.
- Agarwal, Sumit, Chunlin Liu and Nicholas S. Souleles (forthcoming). "The Reaction of Consumer Spending and Debt to Tax Rebates— Evidence from Consumer Credit Data." *Journal of Political Economy*.
- Aguiar, Mark and Eric Hurst (2005). "Consumption versus Expenditure," *Journal of Political Economy*, 113(5), pp. 919-948.
- Aguiar, Mark and Eric Hurst (forthcoming). "Lifecycle Prices and Production," *American Economic Review*, forthcoming December 2007.
- Ausubel, Lawrence M. (1991). "The Failure of Competition in the Credit Card Market." *American Economic Review*, 81(1), pp. 50-81.
- Benartzi, Shlomo, and Richard Thaler (2002), "How Much Is Investor Autonomy Worth?" *Journal of Finance*, , 57(4), pp. 1593-1616.
- Benartzi, Shlomo, and Richard Thaler (forthcoming), "Heuristics and Biases in Retirement Savings Behavior," *Journal of Economic Perspectives*.
- Ben-David, Itzhak (2007). "Manipulation of Collateral Values by Borrowers and Intermediaries." Working Paper, University of Chicago Graduate School of Business.
- Benjamin, Daniel, Sebastian A. Brown and Jesse Shapiro (2006) "Who is Behavioral? Cognitive Ability and Anomalous Preferences." Working Paper, University of Chicago.
- Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman (2006). "What's Psychology Worth? A Field Experiment in the Consumer Credit Market." Mimeo, University of Chicago Graduate School of Business.
- Blanchflower, David G. and Andrew Oswald (2007), "Is Well-being U-Shaped over the Life Cycle?," NBER Working Paper No. 12935
- Bucks, Brian and Karen Pence (2006), "Do Homeowners Know their House Values and Mortgage Terms?" Finance and Economics Discussion Paper Series, Federal Reserve Board.
- Calem, Paul S. and Loretta J. Mester (1995). "Consumer Behavior and the Stickiness of Credit-Card Interest Rates." *American Economic Review*, 85(6) , pp. 1327-1336.
- Calem, Paul S., Michael B. Gordy and Loretta J. Mester. (2005) "Switching Costs and Adverse Selection in the Market for Credit Cards." Mimeo, Federal Reserve Bank of Philadelphia.
- Campbell, John Y. (2006). "Household Finance," *Journal of Finance*, 61 (4), pp. 1553–1604.
- Carstensen, Laura L. (2006) "The Influence of a Sense of Time on Human Development," *Science*, 312, pp. 1913-15.

- Charles, Kerwin Kofi, Erik Hurst and Melvin Stephens, Jr. (2006). “Explaining Racial Differences in Vehicle Loan Rates.” Working Paper, University of Chicago Graduate School of Business.
- Charness, N., and Bosnian, E. A. (1990). “Expertise and aging: Life in the lab.” In T. H. Hess (Ed.), *Aging and cognition: Knowledge organization and utilization* (pp. 343-385). Amsterdam: Elsevier.
- DellaVigna, Stefano and Ulrike Malmendier (2004). “Contract Theory and Self-Control: Theory and Evidence.” *Quarterly Journal of Economics*, 119(2), pp. 353-402.
- Denburg, Nathalie L., Daniel Tranel, Antoine Bechara, “The ability to decide advantageously declines prematurely in some normal older persons,” *Neuropsychologia* 43 (2005) pp. 1099–1106
- Dulleck, Uwe and Rudolf Kerschbamer (2006) “On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods,” *Journal of Economic Literature* 44(1), pp. 5-42.
- Ellison, Glenn, “A Model of Add-on Pricing,” *Quarterly Journal of Economics*, 120 (2005), 585-638.
- Fair, Ray C. (1994). “How Fast Do Old Men Slow Down?” *Review of Economics and Statistics*, 76(1), pp. 103-118.
- (2007). “Estimated Age Effects in Athletic Events and Chess.” *Experimental Aging Research*, 37-57.
- (2008). “Estimated Age Effects in Baseball.” forth., BE Press.
- Flynn, J. R. (1984). “The Mean IQ of Americans: Massive Gains 1932 to 1978”. *Psychological Bulletin*, 95, pp. 29-51.
- Fratiglioni Laura, Diana De Ronchi and Hella Agüero Torres (1999) , “Worldwide Prevalence and Incidence of Dementia,” *Drugs & Aging*, 15(5), pp. 365-375
- Frederick, Shane (2005). “Cognitive Reflection and Decision Making”, *Journal of Economic Perspectives*, 19(4), pp. 25–42.
- Gabaix, Xavier and David I. Laibson (2006). “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets.” *Quarterly Journal of Economics*, 121(2), pp. 505-540.
- Gabaix, Xavier, David I. Laibson, Guillermo Moloche and Stephen Weinberg (2006), “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model,” *American Economic Review*, 96 (4), pp. 1043-1068.
- Galenson, David (2005), *Old Masters and Young Geniuses: Two Life Cycles of Artistic Activity*. Princeton, NJ: Princeton University Press.
- Gourinchas, Pierre-Olivier, and Jonathan Parker (2002). “Consumption Over the Life Cycle.” *Econometrica*, 70(1), pp. 47-89.
- Gross, David B. and Nicholas S. Souleles (2002a). “Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data.” *Quarterly Journal of Economics*, 117(1), pp. 149-185.

- (2002b). “An Empirical Analysis of Personal Bankruptcy and Delinquency.” *Review of Financial Studies*, 15(1), pp. 319-347.
- Heidhues, Paul, and Botond Koszegi (2006), “Loss Aversion, Price Stability, and Sales,” Working Paper, University of California, Berkeley.
- James, Bill (2003). *The New Bill James Historical Baseball Abstract*. New York, NY: Free Press.
- Jones, Benjamin F.(2006), “Age and Great Invention” Mimeo, Northwestern University.
- Kamenica, Emir (2006) “Contextual Inference in Markets: On the Informational Content of Product Lines”, Mimeo, University of Chicago.
- Kerr, Sougata and Lucia Dunn (2002). “Consumer Search Behavior in the Changing Credit Card Market.” Mimeo, Ohio State University.
- Kerr, Sougata, Stephen Cosslett, and Lucia Dunn (2004). “Do Banks Use Private Information from Consumer Accounts? Evidence of Relationship Lending in Credit Card Interest Rate Heterogeneity.” Mimeo, Ohio State University.
- Korniotis, George and Alok Kumar (2007) “Does Investment Skill Decline due to Cognitive Aging or Improve with Experience?” Mimeo, Federal Reserve Board.
- Lusardi, Annamaria and Olivia S. Mitchell (2006), “Financial Literacy and Planning: Implications for Retirement Wellbeing,” Mimeo, Dartmouth.
- Lusardi, Annamaria and Olivia S. Mitchell (2007), “Financial Literacy and Retirement Preparedness: Evidence and Implications for Financial Education Programs,” Mimeo, Dartmouth
- Malmendier, Ulrike and Stephan Nagel (2007) “Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?” mimeo.
- Malmendier, Ulrike, and Devin Shanthikumar (2005), “Are Small Investors Naive about Incentives?” Mimeo, Harvard University and Stanford University.
- Massoud, Nadia, Anthony Saunders and Barry Scholnick (2006). “Who Makes Credit Card Mistakes?,” Mimeo, University of Alberta.
- Mather, Mara, and Laura Carstensen (2005). “Aging and motivated cognition: The positivity effect in attention and memory,” *Trends in Cognitive Science*, 9(10), pp. 496-502.
- Mullainathan, Sendhil, and Andrei Shleifer (2005), “The Market for News,” *American Economic Review*, 95(6), pp. 1031-1053.
- Murphy, Kevin M. and Finis Welch (1990). “Empirical Age-Earnings Profiles.” *Journal of Labor Economics*, 8(2), pp. 202-229.

- Oster, Sharon M. and Daniel S. Hamermesh, "Aging and Productivity Among Economists," *Review of Economics and Statistics*, 1998, 80(1), pp. 154-156.
- Oster, Sharon M., and Fiona M. Scott Morton (2005), "Behavioral Biases Meet the Market: The Case of Magazine Subscription Prices," *Advances in Economic Analysis and Policy*, 5.
- Park, Denise, and Norbert Schwarz eds (1999), *Cognitive Aging: A Primer*. Psychology Press.
- Ravina, Enrichetta (2005). "Habit Persistence and Keeping Up with the Joneses: Evidence from Micro Data." Mimeo, New York University.
- Salop, Steven, and Joseph Stiglitz (1977), "Bargains and Ripoffs: A Model of Monopolistically Competitive Price," *Review of Economic Studies*, 44, pp. 493-510.
- Salthouse, Timothy (2005), "Effects of aging on reasoning", *Cambridge Handbook of Thinking and Reasoning*. Cambridge, England: Cambridge University Press.
- Salthouse, Timothy (forthcoming), "Executive Functioning," D.C. Park & N. Schwarz (Eds.), *Cognitive Aging: A Primer, 2nd edition*, Psychology Press.
- Salthouse, Timothy, David H. Schroeder and Emilio Ferrer (2004). "Estimating retest effects in longitudinal assessments of cognitive functioning in adults between 18 and 60 years of age," *Developmental psychology*, 40(5), pp. 813-822.
- Scott Morton, Fiona, Florian Zettelmeyer and Jorge Silva-Risso (2003). "Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?" *Quantitative Marketing and Economics* 1(1), pp. 65-92.
- Shue, Kelly and Erzo F. P. Luttmer (2006). "Who Misvotes? The Effect of Differential Cognition Costs on Election Outcomes." NBER Working Paper No. 12709.
- Shui, Haiyan and Lawrence M. Ausubel (2004) "Time Inconsistency in the Credit Card Market." Mimeo, University of Maryland.
- Simonton, D. K. (1988). "Age and outstanding achievement: What do we know after a century of research?" *Psychological Bulletin*, 104, pp. 251-267.
- Spiegler, Rani (2006), "The Market for Quacks," *Review of Economic Studies*, 73(4), pp. 1113-1131.
- Stango, Victor and Jonathan Zinman (2007) "Fuzzy Math and Red Ink: When the Opportunity Cost of Consumption is Not What it Seems," Mimeo, Dartmouth.
- Weinberg, Bruce A. and David A. Galenson (2005), "Creative Careers: The Life Cycle of Nobel Laureates in Economics" NBER Working Paper No. 11799.

Appendix: Data Summary Statistics

Table A1: Home Equity Loans and Credit Lines				
	Loans		Credit Lines	
Description (Units)	Mean	Std. Dev.	Mean	Std. Dev.
APR(%)	7.96	1.16	4.60	0.88
Borrower Age (Years)	43	14	46	12
Income (\$, Annual)	78,791	99,761	90,293	215,057
Debt/Income (%)	40	18	41	19
FICO (Credit Bureau Risk) Score	713	55	733	49
Customer LTV (%)	66	26	62	24
Appraisal LTV (%)	69	29	64	23
Borrower Home Value Estimate (\$)	196,467	144,085	346,065	250,355
Bank Home Value Estimate (\$)	186,509	123,031	335,797	214,766
Loan Requested by Borrower (\$)	43,981	35,161	61,347	50,025
Loan Approved by Bank (\$)	42,871	33,188	60,725	51,230
First Mortgage Balance (\$)	79,496	83,560	154,444	112,991
Months at Address	92	122	99	129
No First Mortgage (%)	29	45	15	42
Second Home (%)	3	14	3	12
Condo (%)	8	18	6	17
Refinancing (%)	66	47	39	49
Home Improvement (%)	18	39	25	44
Consumption (%)	16	39	35	35
Self Employed (%)	7.9	27	7.8	27
Retired (%)	9.5	29	7.7	27
Homemaker (%)	1.4	12	1.3	11
Years on the Last Job	6.3	8.1	7.6	9.1

Table A2: Credit Cards			
Account Characteristics	Frequency	Mean	Std. Dev.
Purchase APR	Monthly	14.40	2.44
Interest Rate on Cash Advances (%)	Monthly	16.16	2.22
Credit Limit (\$)	Monthly	8,205	3,385
Current Cash Advance (\$)	Monthly	148	648
Payment (\$)	Monthly	317	952
New Purchases (\$)	Monthly	303	531
Debt on Last Statement (\$)	Monthly	1,735	1,978
Minimum Payment Due (\$)	Monthly	35	52
Debt/Limit (%)	Monthly	29	36
Fee Payment			
Total Fees (\$)	Monthly	10.10	14.82
Cash Advance Fee (\$)	Monthly	5.09	11.29
Late Payment Fee (\$)	Monthly	4.07	3.22
Over Limit Fee (\$)	Monthly	1.23	1.57
Extra Interest Due to Over Limit or Late Fee (\$)	Monthly	15.58	23.66
Extra Interest Due to Cash Advances (\$)	Monthly	3.25	3.92
Cash Advance Fee Payments/Month	Monthly	0.38	0.28
Late Fee Payments/Month	Monthly	0.14	0.21
Over Limit Fee Payments/Month	Monthly	0.08	0.10
Borrower Characteristics			
FICO (Credit Bureau Risk) Score	Quarterly	731	76
Behavior Score	Quarterly	727	81
Number of Credit Cards	At Origination	4.84	3.56
Number of Active Cards	At Origination	2.69	2.34
Total Credit Card Balance (\$)	At Origination	15,110	13,043
Mortgage Balance (\$)	At Origination	47,968	84,617
Age (Years)	At Origination	42.40	15.04
Income (\$)	At Origination	57,121	114,375

Notes: The “Credit Bureau Risk Score” is provided by Fair, Isaac, and Company (FICO). The greater the score, the less risky the consumer is. The “Behavior Score” is a proprietary score based on the consumer’s past payment history and debt burden, among other variables, created by the bank to capture consumer payment behavior not accounted for by the FICO score.

Table A3: Auto Loan APRs		
Description (Units)	Mean	Std. Dev.
APR(%)	8.99	0.90
Borrower Age (Years)	40	21
Income (\$, Monthly)	3416	772
LTV(%)	44	10
FICO (Credit Bureau Risk) Score	723	64
Monthly Loan Payment (\$)	229	95
Blue Book Car Value (\$)	11,875	4,625
Loan Amount (\$)	4172	1427
Car Age (Years)	2	1
Loan Age (Months)	12	8

Table A4: Mortgage Loans		
	Loans	
Description (Units)	Mean	Std. Dev.
APR(%)	12.64	2.17
Borrower Age (Years)	40.54	9.98
Income (\$)	2,624	2,102
Monthly Mortgage Payment/Income (%)	22.84	12.12
Veraz (Credit Bureau Risk) Score	686	253
LTV (%)	61	17
Loan Amount (\$)	44,711	27,048
Years at Current Job	9.43	8.01
Second House (%)	15.54	5.18
Car Ownership (%)	73.56	44.11
Car Value (\$)	5,664	13,959
Gender (Female=1)	30.96	46.24
Second Income (%)	20.44	40.33
Married (%)	71.32	45.23
Married with Two Incomes (%)	16.75	37.34
Self Employed (%)	13.87	34.57
Professional Employment (%)	15.78	36.46
Nonprofessional Employment (%)	52.78	49.93
Relationship with Bank (%)	10.40	30.52

Table A5: Small Business Credit Cards APRs		
Description (Units)	Mean	Std. Dev.
APR(%)	13.03	5.36
Borrower Age (Years)	47.24	13.35
Line Amount (\$)	9,623.95	6,057.66
Total Unsecured Debt	12,627.45	17,760.24
FICO (Credit Bureau Risk) Score	715.86	55.03
Mortgage Debt (\$)	102,684.70	160,799.57

Table A6: Age Distribution by Product					
Product	Age Percentile				
	10%	25%	50%	75%	90%
Home Equity Loans	34	40	48	59	71
Home Equity Lines	32	40	47	58	70
“Eureka”	24	34	44	53	63
Credit Card	25	34	44	57	68
Auto Loans	27	35	45	57	67
Mortgage	34	42	49	60	69
Small Business Credit Card	37	43	53	62	72
Credit Card Late Fee	25	35	45	58	67
Credit Card Over Limit Fee	26	34	43	56	65
Credit Card Cash Advance Fee	25	36	46	58	68

Table A7: Cost by Product				
Product	APR (bp) or Probability (%) Difference		Annual Cost Difference	
	Age 25-Age 50	Age 75-Age 50	Age 25-Age 50	Age 75-Age 50
Home Equity Loans	73	40	\$284	\$146
Home Equity Lines	68	51	\$296	\$265
“Eureka”	8	11	\$37	\$13
Credit Card APR	17	5	\$2	\$1
Auto Loans	20	12	\$8	\$4
Mortgage	6	15	\$25	\$62
Small Business Credit Card	26	14	\$3	\$2
Credit Card Late Fee	2	2	\$8	\$8
Credit Card Over Limit Fee	1	1	\$4	\$4
Credit Card Cash Advance Fee	2	1	\$8	\$4

Notes: This table computes the difference in annual costs of each product borne by 25-year-old and 75-year-old borrowers relative to 50-year-old borrowers. For home equity loans and lines of credit, credit cards, auto loans, mortgages, and small business credit cards, the annual cost difference is the product of the APR difference and the average debt levels by age given in Table A8. For “Eureka,” it is the probability difference of experiencing a “eureka” moment multiplied by the APR difference multiplied by the amount of the balance transferred. For the three types of credit card fees, it is the probability difference of paying the fee multiplied by a fee amount of \$35. For these last three, the calculations likely understate the amount of the fee, since they do not incorporate interest rate changes that may be triggered by multiple fee payments or the interest paid on cash advance balances.

Table A8: Average Debt Levels by Age			
Product	Age		
	25	50	75
Home Equity Loans	\$38,879	\$46,057	\$36,601
Home Equity Lines	\$43,477	\$56,891	\$52,031
Balance Transferred	\$2,723	\$3,123	\$2,422
Credit Card	\$1,426	\$1,778	\$1,203
Auto Loans	\$3,782	\$4,031	\$3,554
Mortgage	\$40,645	\$47,337	\$41,403
Small Business Credit Card	\$1,321	\$1,479	\$1,275