

Decomposing Gender Differences in Bankcard Credit Limits

Nathan Blascak

Federal Reserve Bank of Philadelphia
Consumer Finance Institute

Anna Tranfaglia

Federal Reserve Board

WP 23-30

PUBLISHED
December 2023



ISSN: 1962-5361

Disclaimer: This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in these papers are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. Philadelphia Fed working papers are free to download at: <https://philadelphiafed.org/research-and-data/publications/working-papers>.

DOI: <https://doi.org/10.21799/frbp.wp.2023.30>

Decomposing Gender Differences in Bankcard Credit Limits*

Nathan Blascak[†]

Anna Tranfaglia[‡]

December 2023

Abstract

Using linked mortgage application and credit bureau data, we document the existence of unconditional and conditional gender gaps in the distribution of total bankcard limits. We estimate that male borrowers have approximately \$1,300 higher total bankcard limits than female borrowers. This gap is primarily driven by a large gender gap in the right tail of the limit distribution. At the median and in the left tail of the total limit distribution, women have larger limits than men. Results from a Kitagawa-Oaxaca-Blinder decomposition show that 87 percent of the gap is explained by differences in the *effect* of observed characteristics, while 10 percent of the difference is explained by differences in the *levels* of observed characteristics. The gap is persistent across geographies but has varied over time. Overall, these gender gaps are small in economic magnitude and have changed over time favoring women.

Keywords: gender, credit, credit cards, decomposition

JEL Codes: J16, G51 G53

*The authors would like to extend special thanks to Tom Akana for his advice and input throughout this project. This paper has also benefited greatly from comments by Francis Annan, Ken Brevoort, Julia Cheney, Daniel Grodzicki, Robert M. Hunt, Sasha Indarte, Lauren Lambie-Hanson, Jeff Larrimore, Geng Li, Paolina Medina, Slava Mikhed, Sarah Miller, Scott Nelson, Larry Santucci, Keith Wardrip, Douglas Webber, and seminar participants at the Boulder Summer Conference on Consumer Financial Decision Making, the APPAM Fall Research Conference, the SEA Conference, the WEAI Conference, the FRB Philadelphia Workshop on Credit Card Lending and Payments, the Consumer Finance Round Robin, and the Household Finance Brownbag for their helpful feedback. Ian McGroarty and Bryant Wright provided excellent research assistance.

[†]Blascak (Corresponding author): Consumer Finance Institute, Federal Reserve Bank of Philadelphia (email: nathan.blascak@phil.frb.org).

[‡]Tranfaglia: Division of Consumer and Community Affairs, Federal Reserve Board (email: anna.e.tranfaglia@frb.gov).

Disclaimer: This working paper will be published concurrently in the Federal Reserve Bank of Philadelphia Research Working Paper Series and in the Board of Governors of the Federal Reserve System Working Paper Series. This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia, the Federal Reserve Board, or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. No statements here should be treated as legal advice. Philadelphia Fed working papers are free to download at <https://philadelphiafed.org/research-and-data/publications/working-papers>.

1 Introduction

Credit cards are one of the most common debt instruments in the United States. Approximately 76 percent of Americans own at least one credit card (Green and Stavins, 2018), and 44 percent of consumers revolve a credit card debt balance (Bricker et al., 2017). At the end of 2022, U.S. credit card debt totaled over \$986 billion (Federal Reserve Bank of New York, 2022). Given the size of the credit card market and that the existing literature has demonstrated there are gender differences in numerous other financial markets such as housing (Fishbein and Woodall, 2006; Fang and Munneke, 2020; Goldsmith-Pinkham and Shue, 2023), auto loans (Ayres and Siegelman, 1995; Morton, Zettelmeyer, and Silva-Risso, 2003), small business credit (Cavalluzzo and Cavalluzzo, 1998), and investing (Barber and Odean, 2001), we may expect outcomes in this market to differ by gender as well. Motivated by these facts, we examine if there exists a gender gap in bankcard *limits*,¹ a commonly used measure of credit access in the academic literature (e.g., Gross and Souleles (2002), Agarwal et al. (2018), Aydin (2021), and Gross, Notowidigdo, and Wang (2020)).

To assess if there are differences in bankcard limits by gender, we use a large, unique data set of successful mortgage applications (mortgage applications that are subsequently originated) from Home Mortgage Disclosure Act (HMDA) data that are merged with credit bureau data from Equifax. With these data, we observe both credit bureau characteristics, such as credit card limits and accounts, and demographic information that is not available on credit reports.² An important caveat to our analysis is that, in order to isolate the gender of the consumers in our data set, we restrict our analysis to sole mortgage applicants. Despite this limitation, our data set consists of approximately 1 million individuals who originated a mortgage from 2004 to 2014.

Using these merged data, we first document the existence of an unconditional gender difference in total bankcard credit limits, with male borrowers having higher limits than female borrowers on average. This difference persists over time and has increased since the Great Recession. We find that at the start of our study period in January 2006, the average difference between genders for total bankcard credit limit among successful mortgage applicants was approximately \$1,700, or

¹A bankcard is a credit card issued by a bank, bankcard company, national credit card company, or credit union.

²Because the Equal Credit Opportunity Act (ECOA) forbids creditors from discriminating on the basis of race, gender, marital status, national origin, or other demographic information in the underwriting, pricing, and scoring of credit, this long-standing regulation has disincentivized the joint collection of credit attributes and consumer demographics. For an extensive overview of ECOA, see https://files.consumerfinance.gov/f/201306_cfpb_laws-and-regulations_ecoa-combined-june-2013.pdf.

5.7 percent of the male average of \$29,654. This difference declined during the recession, but it subsequently increased to over \$2,300 (7.8 percent of the male average) by the end of 2017. When examining the gender gap at each decile of the bankcard limit distribution across time, we see that the majority of the gap is driven by large gender differences at the 80th and 90th deciles of the bankcard limit distribution. The gender gap at lower deciles favors women, though by a much smaller magnitude.

Using a simple linear regression approach to control for demographic characteristics, geography, income, and credit score,³ we find that the unexplained gender difference in bankcard limits is \$1,312, with male borrowers having higher limits than female borrowers during our period of analysis. This is approximately 4.5 percent of the male sample mean. To put this number in context, this gap is smaller than the 17.5 percent (1.5pp) gender gap in unlevered housing returns estimated by Goldsmith-Pinkham and Shue (2023) but slightly larger than the 2 percent (0.13pp) gap in contract rate estimated in Fang and Munneke (2020). It is important to note that this estimated difference is for an *equilibrium* outcome (total bankcard limit), which is the result of the accumulation of both credit demand and supply decisions over time. Without clear exogenous variation from either the supply or demand side, we are only able to estimate a correlation between gender and total bankcard limits; we are unable to make any causal claims as to why these differences exist.

While we find that the average gender gap is fairly consistent across geographies, we find that the gender gap shrinks when examining individuals living in counties with high unemployment, and the gap closes almost entirely when restricting the analysis to individuals living in counties where a mass layoff has occurred. In these mass layoff counties, we estimate an average marginal effect of -\$278 that is not statistically different from zero. This result suggests that gender differences in employment and income play a role in explaining the gender gap in credit card limits because (1) the decline in the magnitude of the gender gap is consistent with the results in Braxton, Herkenhoff, and Phillips (2023), who estimate that individuals who are laid off in a mass layoff event experience a \$1,000 decline in total bankcard credit limit, and (2) that the majority of mass layoffs result in more men being laid off than women. Since more men than women are laid off, and subsequently have lower incomes, we would expect the gender gap in credit limits to close.

³Our credit score measure is the Equifax Risk Score, which is a proprietary credit score similar to other credit scores used in the industry.

We then attempt to identify what factors explain this gender gap in total bankcard limits by decomposing the gender difference using the standard Kitagawa-Oaxaca-Blinder (KOB) decomposition method. This tool, frequently used in analyzing the gender wage gap, allows us to separately identify how observable and unobservable factors contribute to the male-female difference in total bankcard limits.

Our estimates show that the gender gap in total bankcard limits is primarily driven by the *coefficient* effect (the *effect*⁴ of observable characteristics), which explains approximately 87 percent of the gender gap, while the *endowment* effect (the values of observable characteristics) only explains 10 percent of the difference between genders. This large magnitude for the coefficient effect implies that the effect of unobservable characteristics and the “returns” on observed characteristics (how the effect of observed characteristics affects the gap (not the difference in levels of the observable characteristics)) are *lower* for female borrowers.⁵

The majority of the gender difference is driven by unobservable characteristics consistent with the results from our heterogeneity analyses, which imply that current income and employment (neither of which we observe in our data) play significant roles in explaining the gender gap. Although we are unable to make any causal claims as to why these differences exist, we are able to document that the majority of the gender differences in credit limits cannot be explained by observed differences in demographic, socioeconomic, and credit characteristics alone.

Given that the unconditional gender gap differs across time and across the credit limit distribution, we estimate and decompose the gender gap for each year at each decile of the total bankcard limit distribution from 2006 to 2016. Our results show that the magnitude of the male-female limit difference varies across the distribution of our credit limit variables and across time. At the lower deciles, conditional gender differences are small (\$0 to \$400 up to the 30th decile), with female borrowers having higher limits than male borrowers at some deciles in the later years of the sample. At the higher deciles, gender differences are in favor of male borrowers and larger (\$2,000 and higher at the 80th percentile), with the gap growing over time.

⁴It is standard in the labor economics literature to interpret the coefficient effect as the difference in “returns” to different characteristics since individuals make intentional investments to earn higher wages. In our setting, it is not immediately clear that observed characteristics can be interpreted in this way. Thus, we will use “return” and “effect of” interchangeably through the rest of the paper.

⁵In other words, the large coefficient effect implies that total bankcard limits for women would be *higher* if they had male coefficients.

Our decompositions also reveal that for the upper 70 percent of the credit distribution, women have increasingly *benefited* from having more favorable characteristics over time (the *endowment* effect has become negative). However, the returns to these characteristics are lower, resulting in the credit gap. Additionally, over the same time period, the gap in the returns to those characteristics (the coefficient effect) for the top half of the distribution has increased, which partially explains why the gender gap in limits has grown over time.

Finally, to provide some insights into whether these effects are driven by credit supply or credit demand, we utilize data on direct mail credit card solicitations to examine if there are differences in credit card supply between men and women. Unlike the credit bureau data used in the regression and decomposition analyses, the credit card mail offer data provide us with a more direct measure of credit supply. Using these data, we observe that women receive slightly fewer credit card mail offers than men, that women receive different kinds of offers than men, and that women receive higher *advertised* limits than men.

In the last section of the paper, we discuss three potential mechanisms that would explain the differences we find: (1) differential treatment in the credit market by gender, (2) differences in socioeconomic characteristics, especially at the time of card origination, or (3) differences in preferences for credit and credit cards, or some combination of the three. While we cannot rule out differential treatment in the credit market on the basis of gender, we view this mechanism as unlikely for a number of reasons. In particular, given that the underwriting of credit cards is now highly automated, we may expect there to be limited scope for bias in the assignment of credit limits for bankcards. A combination of differences in (2) and (3) seems most logical, given the established literature on gender differences in both socioeconomic factors, such as the gender pay gap, and in preferences for credit and risk. Our results from both our regression and decomposition analyses are consistent with these mechanisms, and we note that our results also indicate that different combinations of these mechanisms are likely at play when examining the gender differences at different points in the credit limit distribution.

This paper makes contributions to a number of important literatures on gender and economics. First, our research is related to the large body of research that has documented gender differences in a number of different economic settings. Much of the research in this area has focused on outcomes in the labor market, most prominently the gender wage gap (see Blau and Kahn (2017)

for a survey of the literature), and the mortgage market (Fishbein and Woodall, 2006; Cheng, Lin, and Liu, 2011; Goodman, Zhu, and Bai, 2016; Goldsmith-Pinkham and Shue, 2023). Although studies in this literature have also examined gender differences in credit markets with respect to credit card interest rates (Mottola, 2013) and credit scores and credit use (Li, 2018), research on gender disparities in both credit use and credit score has received significantly less attention. We complement this previous research by supplying new empirical evidence of the gender gap in consumer credit markets and documenting how it has changed over time.

Our research also complements the literature on gender differences in financial literacy and confidence. Prior work by Lusardi and Mitchell (2008) and Lusardi, Mitchell, and Curto (2010) document the existence of a pronounced gender gap in financial literacy. Lusardi and Mitchell (2014) summarize the theory and evidence of the economic importance of financial literacy and document the growing body of literature that explains these differences between genders. Fonseca, Mullen, Zamarro, and Zissimopoulos (2012) find that differences in demographic characteristics do not explain the majority of the financial literacy gap between genders. Bucher-Koenen et al. (2017) find a significant financial literacy gap among young women, even though they have greater labor force participation and educational attainment on average than their male counterparts. Survey data have also shown that there might be gender differences with respect to confidence or comfort level when accessing the credit card market. Fernandez and Tranfaglia (2020) show that women are significantly less likely than men to be “very or somewhat” likely to think that a credit card application would be approved if they apply for an additional card, and Canilang et al. (2020) find that 60 percent of men who are college educated and not yet retired are confident managing self-directed retirement accounts compared to 32 percent of women with a bachelor’s degree. While our analyses cannot identify if disparities in financial literacy levels between males and females contribute to differences in bankcard limits, we acknowledge that the connection between these differences warrants further attention.

2 Data

2.1 Data Description

To study the differences in credit card limits between males and females, we use a unique panel data set that combines information from three different data sources. The first data source is a merged data set of anonymized information from Black Knight McDash (McDash) and data from the Home Mortgage Disclosure Act (HMDA) database. The McDash data set contains monthly mortgage servicing information for the largest residential mortgage servicers in the U.S. from 1992 to the present. These data cover approximately two-thirds of the installment-type loans in the residential market, or approximately 151 million loans. The data set contains multiple types of mortgage products and include borrower, property, and loan characteristics.

The HMDA data contain records on mortgage applications, originations, and purchases by depository institutions and certain for-profit, non-depository institutions from 1990 to the present.⁶ Importantly for our paper, this data set contains demographic information on mortgage applicants, including gender, and information on loan characteristics for those loans that are subsequently originated. Certain characteristics of the mortgage application are removed to maintain the anonymity of the applicant. The HMDA data are then matched to the McDash data by the Risk Assessment, Data Analysis, and Research (RADAR) unit of the Federal Reserve Bank of Philadelphia.

This combined data set is then merged with the Equifax Credit Risk Insight Servicing and Black Knight McDash data (CRISM) database using another confidential matching process by the Federal Reserve Bank of Philadelphia. The CRISM data contain anonymized monthly individual-level credit report information for the mortgage borrowers in the McDash data. Since we do not observe account-level information, only information aggregated to the level of the consumer, we have information on the number of accounts and the aggregate account balances for credit cards, auto loans, and student loans. We also observe a consumer's Equifax Risk Score. The merger of the HMDA, McDash, and CRISM (HMC) databases produces a combined data set of over 56 million mortgage loans that were originated from 1992 to the present.

Our primary variable of interest is the total bankcard limit on all bankcard accounts. This variable is the sum of all individual bankcard credit limits that an individual has in a given month.

⁶See Avery, Brevoort, and Canner (2007) for a more detailed discussion of the HMDA data.

Importantly, this variable represents the total accumulation of card limit decisions across the life of active cards and *not* the initial assignment of credit limits when the accounts were opened. From the credit bureau and mortgage servicing data, we utilize variables that contain information on the total number of bankcards that an individual has in a given month, the loan-to-value (LTV) ratio of the individual’s mortgage, the individual’s age, and their credit score. From the HMDA data, we use information on the applicant’s gender, race, ethnicity, state of residence, and income at the time of application.

2.2 Sample Selection

We take a 5 percent random sample of loans in the HMC database that were originated from 1992 to 2014, which produces a data set with approximately 277 million loan-month observations from June 2005 to December 2017. Since we only have six months of data in 2005, we restrict our sample to all full years of data, from January 2006 to December 2017. For each mortgage with a co-applicant, we have CRISM data for each individual applicant of that mortgage. However, the data do not identify which individual is the applicant and which individual is the co-applicant in the HMDA data. Consequently, we drop any co-signed mortgage in the data because, although we know the gender of the applicant and the co-applicant for each mortgage, we cannot accurately assign the applicant/co-applicant gender variables in the HMC sample.⁷ We also drop any individuals missing gender or year of birth. Furthermore, to avoid double counting individuals, we drop observations on mortgages beyond the first mortgage observed in the data for each consumer. The resulting panel data set contains over 55 million observations on approximately 1 million unique individuals.

After implementing this restriction, we create a repeated cross-section data set of the sole mortgage applicants by taking one observation from each individual 24 months after they originate their mortgage in the panel data set. We focus on a period two years after individuals close on their mortgage since it is possible that consumers had adjusted their credit product portfolios in preparation for shopping for a mortgage. The final cross-sectional subsample, which we use for all subsequent regression analyses, consists of 530,122 individuals from 2006 to 2016.

Because of the selected nature of our sample, we briefly discuss how individuals from our sample

⁷While we drop co-signed mortgages, it is important to note that couples can exist in the data in cases where only a single name is recorded on the mortgage. This can happen when the individuals in the household have different credit histories and choose to apply using only the strongest credit history, for example.

differ from individuals in the general U.S. adult populations in Appendix A. Appendix Table A1 includes demographic statistics for the entire U.S. population (not just those with a credit history, a limiting factor for inclusion in the HMC data set). Unsurprisingly, sole mortgage applicants differ from the overall U.S. population in a number of ways: Mortgage applicants have higher incomes and are more likely to be White individuals.

Along with comparing individuals in our sample to the overall U.S. population, we also provide a comparison of our sample to individuals in the credit bureau population in Appendix Table A2. Because sole mortgage applicants are a very specific segment of the entire credit bureau population, it is likely that these individuals are not representative of the credit bureau population. In particular, we expect that the average consumer in our sample would be more creditworthy and would have larger debt balances than an average individual in the overall credit bureau population. We compare summary statistics from our sample and a representative sample of U.S. consumers with a credit report from the FRBNY Consumer Credit Panel/Equifax (CCP) data in Appendix B.⁸ Sole mortgage applicants differ from the overall credit bureau population along a number of dimensions: Sole applicants have higher credit scores, hold more bankcards, and have higher bankcard credit limits.⁹ While the median birth year is the same in both data sets, the standard deviation in the CCP is larger than in the HMC data; this is unsurprising as the total credit bureau population contains younger individuals, who have not yet applied for a mortgage, and older individuals, who have no need for a mortgage or have already paid.

Finally in Appendix C, we compare sole mortgage applicants to dual mortgage applicants (i.e., mortgages with co-applicants) and the overall mortgage application population. As we show in Appendix Table A3, dual mortgage applicants have higher credit scores on average and also have higher total bankcard credit limits.

2.2.1 Selection in Being a Sole Applicant

Even though we focus on sole applicants, we cannot make any assumptions about marital status for the consumers in our sample because we do not observe marital status in our data.¹⁰ Currently,

⁸For a detailed description of the CCP, see Lee and van der Klaauw (2010).

⁹The share of successful mortgage applicants who do have a co-applicant (sole applicants) remained stable throughout our period of study and were approximately 38 percent of our sample overall.

¹⁰While it is very plausible that the majority of dual applicant mortgage holders are not single, the fact that married applicants can apply for mortgages as sole applicants prevents us from determining the marital status of sole applicants.

there is no regulation or law that requires partners to serve as co-applicants, and there are many potential cases where it would be rational for individuals to not include their partner on a mortgage application. For example, if one partner has a significantly lower credit score than another, a household may decide to omit that person from the mortgage application.

If there is selection on who does and does not include a spouse as a co-applicant on a mortgage application by gender, then our analysis may suffer from an endogeneity problem. For example, if married men are more likely to omit their partner on a mortgage application than women, then marital status would explain some of the variation in bankcard limits, resulting in omitted variable bias. In addition, it is likely that the correlation between sole applicant status and marriage is correlated with other variables such as income, in addition to gender.¹¹ As we discuss in the next section, differences in bankcard limits vary along a number of other characteristics besides gender. To understand the magnitude and sign of this bias, we would need a data set with information on marital status, gender, and mortgage application status (sole vs. co-applicant), along with demographic and socioeconomic characteristics. Thus, an important caveat to our analysis is that there is an unknown degree of bias relating to this specific type of selection.

3 Identifying the Gender Credit Gap in Aggregate Data

3.1 Summary Statistics for Sole Mortgage Applicants

To examine the raw differences between males and females in credit card limits, we focus on two different measures of credit limits: the total credit limit on all bankcard accounts and the average credit limit a consumer has on their bankcard accounts. To provide additional context, we produce summary statistics for total bankcard balances and the total number of bankcard accounts, though we do not focus on these variables in the main part of our analysis.¹²

Table 1 provides sample means and quartile values for our cross-section sample of successful sole mortgage applicants and separately female and male applicants. Females hold more bankcards

¹¹For example, it could be the case that male borrowers in high-income households are more likely to omit their partner than low-income households.

¹²It is important to note that one limitation of our data is that we cannot identify different types of credit card that a consumer holds (general purpose, private label, small business, etc.). Therefore, we are unable to examine if females have more of one type of card in particular compared to males, or if the distribution of card types is similar among both genders.

on average than males (3.38 to 3.22). They also have lower total bankcard limits compared to male borrowers (\$28,544 to \$30,079) and lower average bankcard limits (\$9,227 to \$8,359). As also seen in Appendix Table A1, our sample consists mainly of White individuals, who make up 74 percent of our sample, although this percentage is representative of the mortgage-holding population in the U.S.

The results in Table 1 suggest that significant differences exist between male and female borrowers in both credit characteristics and in demographics, though median differences are relatively smaller in terms of economic significance. In absolute terms, mean gender differences tend to be larger than median differences because gender differences tend to be larger in the upper percentiles of their respective distributions. For example, the difference in HMDA incomes¹³ at the 75th percentile for men and women is \$23,000, which is almost double the difference at the median, while gender differences in total bankcard limit at the 75th percentile is \$1,400, which is over three times larger than the median difference. In percentage terms, the gender difference in income at the median is 18.8 percent and increases to 21.7 percent at the 75th percentile, while the difference in bankcard limits is 1.4 percent at the median and 3.4 percent at the 75th percentile.

Looking more closely at the entire income and credit score distributions in Table 1, it is interesting to note that while there are large differences in the income and total bankcard limit distributions by gender, the credit score distribution for each gender is nearly identical. To better understand these relationships among income, credit score, and total bankcard limit for each gender, we create bins for both credit score and income and plot the average total bankcard limit for each bin for each gender. Results for the total bankcard limit and credit score are displayed in Figure 1, while results for total bankcard limit and HMDA income are displayed in Figure 2.

As can be seen in Figure 1, there is a non-linear relationship between credit score and total bankcard limit for individuals in our sample. Total limits are actually higher for individuals with very low credit scores (≤ 500) than for individuals with scores from 500 to 600.¹⁴ However, once scores move out of the deep subprime range, total bankcard limits increase consistently, with average total limits doubling for individuals with scores greater than 800, compared to individuals

¹³These are the applicant incomes reported on the HMDA forms.

¹⁴The presence of individuals with such low credit scores in our sample is due to bankruptcy and/or foreclosure; 75 percent of individuals with credit scores lower than 500 in our sample have ever had either form of financial distress on their credit report.

with scores in the 650-699 range.

We also document that for individuals in our sample, the non-linear relationship between credit score and total bankcard limits differs by gender across different parts of the credit score distribution. For scores less than 600, male borrowers have higher limits than female borrowers in every credit score bin, with a larger separation occurring among consumers with deep subprime scores. In the near-prime to prime score range (600-750), men and women have almost identical total bankcard limits for each credit score category. When scores approach the super prime range (800+), we again observe that men have higher limits relative to women even though they are in the same Risk Score category.

Unlike the relationship between credit scores and total limit, where men had higher total bankcard limits than women across the entire distribution, when looking at the relationship between income and total bankcard limits, women have higher total limits than men in the left tail of the income distribution (Figure 2). We observe that, on average, female borrowers have higher total bank limits than male borrowers for incomes of less than \$100,000; for incomes higher than \$100,000, male borrowers have higher total limits. This result is consistent with two facts: (1) gender differences in income are greater in the right tail of the wage distribution (i.e., “the glass ceiling” effect (Albrecht, Bjorklund, and Vroman, 2003; Fortin, Bell, and Boehm, 2017)) and (2) income is an important factor in determining credit limits at account origination. These relationships among credit scores, income, and total bankcard limits suggest that any analysis examining the gender gap in bankcard limits would be biased if it did not control for these differences.

3.2 Summary Statistics over Time

Along with documenting the existence of differences in bankcard credit limits between male and female sole mortgage applicants using our cross-section sample, we can use the panel nature of the full sample to track these differences over time. Given that our data cover the Great Recession and its recovery period, along with numerous regulatory changes in the form of the Credit Card Accountability Responsibility and Disclosure (CARD) Act and the Dodd-Frank Act, it is possible that these gender differences may have changed over time. To examine if the gender credit gap changed during our study period, we use the full panel version of our data and first plot the monthly means of our bankcard variables for each gender over time in Figure 3. We then plot the time trends

in the gender difference over time by credit score bin to see if the across-time variation in the gender gap varies by credit score in Appendix Figure A2.

We plot three different bankcard variables over time in Figure 3. In panel A of Figure 3, we see that, despite a large reduction in the number of total bankcard accounts for both genders during the Great Recession, female borrowers had more bankcard accounts than male borrowers throughout our period of study, and that this gap has grown over time. There was almost no gap between male and female borrowers at the start of 2006, but this difference increases over our sample period, with females having 0.27 more bankcard accounts than males by the end of 2017.

In panel B of Figure 3, we can see that, at the start of our sample in 2006, the difference in total bankcard limits was \$1,800, or 6.5 percent of the average male total bankcard limit. However, by the end of 2017, this difference had grown to over \$2,300, approximately 7 percent of the male total limit. Panel C shows the gender difference in the average bankcard limit. Throughout the entire sample period, male mortgage holders experienced greater average bankcard credit limits than female mortgage owners.¹⁵ We also note that male borrowers have seen greater growth in average limits than females: Differences in the average bankcard limit were approximately \$400 in 2006 and have grown to \$1,400 by the end of 2017. Overall, Figure 3 illustrates two important stylized facts regarding credit card borrowing in our sample: (1) women hold more bankcard accounts on average than men throughout our sample frame, with the difference widening over time in favor of female borrowers, and (2) men have higher total bankcard credit limits and average bankcard credit limits than women, on average, over this same time period, with the gap only marginally widening in favor of male borrowers.¹⁶

4 Estimating Gender Differences in Bankcard Credit Limits

4.1 Empirical Specification

To better understand the differences between men and women for total bankcard limits, we use our repeated cross-section data set to estimate a series of simple linear regressions of the following form

¹⁵Throughout the analysis, we calculate the average bankcard limit as the consumer's total bankcard limit divided by the number of bankcard accounts.

¹⁶Given that we may also expect to see differences in these gaps by credit score, we also examine how the gap differs across time by credit score category. We discuss the summary statistics for these gaps in Appendix D.

to calculate average differences between genders after controlling for a number of demographic and geographic factors:

$$y_i = \beta_0 + \beta_1 Female_i + \gamma score_i + \theta income_i + \mathbf{\Pi}_i \mathbf{X}_i + \epsilon_i. \quad (1)$$

We define the dummy variable $Female_i$ to be equal to one if individual i 's gender is female. To account for the nonlinear effect of credit score and income on bankcard limits, we follow Han, Keys, and Li (2018) and include the vector $score_i$, which contains dummy variables for 50-point credit score bins, and the vector $income_i$, which contains dummy variables for income bins.¹⁷ \mathbf{X}_i is a vector of control variables that include age fixed effects, race fixed effects, quarter-of-the-year fixed effects, calendar year fixed effects, state fixed effects, state-by-year fixed effects, the loan-to-value (LTV) ratio of the mortgage at origination for individual i , and the number of bankcard accounts.

Although the relationship between the number of bankcards and total bankcard limit is endogenous, the process through which limits change and future limits are assigned is dependent upon the number of cards an individual has. Therefore, when modeling the differences in total bankcard limit between men and women, it is necessary to control for the number of cards an individual has.¹⁸ To properly account for how limits change with the number of bankcards, we instead include the number of cards as a control variable on the right-hand side of Equation (1).¹⁹

As we previously documented in Figures 1 and 2, the non-linear relationship between total bankcard limits, credit score, and income differ by gender. In our preferred specification, we account for these nonlinear differential effects between male and female borrowers by estimating Equation (1) with two sets of gender interaction terms:

$$y_i = \beta_0 + \beta_1 Female_i + \gamma score_i + \theta income_i + \mathbf{\Gamma} Female_i \times score_i + \mathbf{\Theta} Female_i \times income_i + \mathbf{\Pi}_i \mathbf{X}_i + \epsilon_i. \quad (2)$$

¹⁷We use the reported HMDA income variable to measure an individual's income.

¹⁸Another potential way to address this concern would be to use the average bankcard limit (total bankcard limit divided by the total number of bankcards) as our left-hand side variable of interest. This is problematic, however, because the increase in the credit limit on each marginal bankcard is decreasing in the number of cards and using the average limit would not reflect this concave relationship between total limit and total number of bankcards. For example, an individual with four bankcards applying for a fifth card will be offered a different amount of credit than an individual applying for their first card, assuming all other characteristics of the two individuals are equal.

¹⁹We also examine a specification where we include dummy variables for bins of cards instead of a continuous measure of cards. Results are very similar across specifications.

The first set of interaction terms are $Female_i \times income_i$, which capture any non-linear differential effects between males and females in how income affects credit limits. The second set of interaction terms are $Female_i \times score_i$, which capture any nonlinear differential effects between males and female in how credit score affects credit limits. We also include the number of bankcard accounts interacted with the female dummy variable to account for any gender differences in how limits change with the number of cards.

4.2 Results

Results from Equations (1) and (2) for total bankcard limits are presented in Table 2, with the first two columns reporting results for our baseline regression with no interaction effects and the last two columns reporting results from our preferred specification that contains a full set of interaction terms. From our baseline specification, we see that after controlling for race, age, credit score, income, year, and state of residence, female borrowers on average have \$1,272.13 less in total limit relative to male borrowers. However, given that the relationship among income, credit score, and total limit differ by gender, this simple difference from our baseline specification does not fully explain the reasons for the gap.

Results from our preferred specification, reported in the last two columns of Table 2, show that this is indeed the case. When including the full set of interaction terms among number of bankcards, credit score bins, income bins, and the female dummy variable, we can see that the average gender difference in total limit will vary at different income and credit score levels. To properly calculate the gender difference, we first note that the inclusion of the interaction terms means that the coefficient on the female dummy variable of \$4,471.98 is not the average gender difference in total bankcard limits. Second, the coefficient on the number of bankcard accounts is positively correlated with total bankcard limit: One additional bankcard is associated with an increase of \$9,042 in total bankcard limit. However, the coefficient on the interaction term of the female dummy and the number of bankcards indicates that females have \$817 less in total limit than male borrowers per additional bankcard. For the low-income and low-credit score group, which owns four bankcards on average, this implies that our reference group (low credit score, low income, White borrowers, male borrowers) has an average total bankcard limit of \$5,967.18.²⁰ Compared

²⁰Since our estimated constant is -30,203.14 and the coefficient on the total number of bankcards is 9,042.58, this implies that the average limit for the reference group is $-\$30,203.14 + 4 \times \$9,042.58 = \$5,967.18$.

to the reference group, women in the *lowest income and credit score bins* who hold four bankcards have *higher* average total bankcard limits than men by $\$4,471.98 - (4 \times \$817.50) = \$1,202$.²¹

Using the regression coefficients for the interaction between the gender dummy variable and each credit score and income bin variable in Table 2, we can see that the difference reverses in the higher income and credit score categories, with male borrowers having higher limits than female borrowers. For example, the gender difference for women with three bankcards at the 600 to 649 credit score bin and the \$101,00–\$125,000 income bin is \$666.01 (relative to men with the same income and credit score), and it increases to \$1,700.48 when moving up to the highest credit score bin. Similarly, women with three bankcards and a credit score in the 600-649 bin in the highest income bin would see a gender difference of $\$4,471.98 - (3 \times \$817.51) - \$2,979.60 - \$4,173.15 = -\$5,133.3$.

In Figure 4, we plot the average marginal effects for the female dummy variable at each income and credit score level while holding the values of the other covariates at their mean values among those with similar credit and income characteristics. In panel A of Figure 4, we see that female borrowers have a lower total bankcard limit than male borrowers at each credit score bin except for the bottom two categories. This result tells us that if you had two average individuals from our sample, one male and one female, the female borrower would have a higher total limit than a male borrower if they were in the lowest two Risk Score categories, and the female borrower would have a lower total limit than the male borrower in all of the subsequent Risk Score bins. For income, we see in panel B that female borrowers have lower total bankcard limits at every income category. Similar to the credit score marginal effects, these income results tell us that, if you had a male and a female borrower with the same average characteristics, the female borrower would have lower total bankcard limits at every income level. As a summary measure, after controlling for income, credit score, number of bankcard accounts, and demographic and geographic controls, we can calculate an average marginal effect for females of $-\$1,312.21$.

²¹More specifically, the average bankcard limit for females in the lowest income and Risk Score bins with four bankcards is $-\$30,203.14 + (4 \times \$9,042.58) + \$4,471.98 + (4 \times -\$817.50) = \$7,169.16$, which yields $\$7,169.16 - \$5,967.18 = \$1,202$.

4.3 Robustness Checks

As mentioned in Section 3.1, approximately 19 percent of the individuals in our sample have ever had either a bankruptcy or a foreclosure on their credit report. It is plausible that these individuals are sufficiently different than the borrowers who have never had either form of financial distress. To test if our results are sensitive to the inclusion of these individuals in our sample, we first include a dummy variable that is equal to one if an individual ever had a bankruptcy or a foreclosure in our estimating equation. Results are reported in Table 3. The first two columns of Table 3 report results when just including a dummy variable for past financial distress, and the last two columns report results when interacting this financial distress dummy variable with the female dummy variable. The estimated coefficient on the gender dummy is very similar to the results from our preferred specification in Table 2. The coefficient on the financial distress dummy variable indicates that having a prior bankruptcy or foreclosure on a credit report is associated with having \$1,740 less in total bankcard limit.

In the last two columns of Table 3, we interact the financial distress dummy with the female dummy to see if the effects of prior bankruptcy and foreclosure differ by gender. Our results differ slightly from those in the last columns of Table 2, but they are generally in-line with our main result. When we include the interaction between the financial distress dummy variable with the female dummy variable, we see that females are impacted less by the presence of prior financial distress by \$423. Overall, our main results are robust to controlling for prior financial distress in our estimating equation.

In Table 4, we perform a different sensitivity test and use the financial distress dummy variable to restrict our sample to individuals who had no previous financial distress. By restricting our sample in this way, we drop approximately 99,000, about 20 percent, of our observations.²² Compared to our main results, many of our estimated coefficients from using this subsample of individuals are larger in magnitude, especially those on the *Female* \times *Score* interactions. This is unsurprising, given that the nearly 20 percent of our sample that we drop primarily consists of individuals with low credit scores and lower limits. However, when we calculate gender differences at certain values of our covariates, the differences are similar to those we had previously calculated. For example,

²²Note: 57 percent of those individuals dropped are male.

in the previous section, the difference between a man and a woman with three bankcards, in the 600-649 credit score bin, and the highest income bin is \$5,133.3 (in favor of the male borrower); in the non-financial distress sample, the gender gap for the same group is $\$8,179.7 - (3 \times \$914.21) - \$6,376.70 - \$4,049.29 = \$4,988.92$. While our estimates differ when restricting our sample, the implied gender differences that this empirical setup yields are quantitatively similar to our main results.

We also test the robustness of our main results by dropping different sets of fixed effects from our main estimating equation. In the first two columns of Appendix Table A4, we exclude the state-by-year fixed effects, and in the last two columns, we exclude both the year and state-by-year fixed effects. In both cases, our results are very similar to our main results in Table 2. Interestingly, that our results are robust to leaving out the year fixed effects implies that time is not a driving aspect of the average gender difference. Finally, in tests not reported in the tables, we find that our results are robust to the inclusion of gender-specific time trends, gender by year fixed effects, county fixed effects, and county-by-year fixed effects.

4.4 Heterogeneity Analyses

If credit market conditions or preferences for credit vary by geography, it is possible we could see differences in the gender gap in different areas of the U.S. To address this, we split our sample into the four primary census regions and reestimate Equation (2) for each region separately. The average marginal effects are presented in Appendix Figure A3. We find that average marginal effects are generally similar in each region, though the gender gap is smallest in the Midwest and is largest in the Northeast. We also split our sample by the presence of “common law” marital property laws (see Appendix C.1).

To further examine if geographical factors play a role in explaining the gender gap, we divide our sample using county-level variation in labor market conditions. Labor market conditions could play an important role in the bankcard limit gender gap if gender differences in employment lead to differences in income (i.e., ability to pay). To examine if this is the case, we split our sample using two different employment measures. First, we split our sample by the presence of a mass layoff using the mass layoff definition used in Foote et al. (2018). We do this by using Bureau of Labor Statistics (BLS) data on mass layoffs, which was a monthly data set that identified mass

layoff events when more than 50 workers filed an unemployment insurance claim against a single establishment. Unfortunately, the BLS discontinued this data series in 2011, so we are only able to examine gender differences between 2006 and 2011 in these counties. Given that this sample period coincides with the Great Recession and the initial year of the CARD Act, we interpret these results with some caution. We also split our sample by county-level unemployment rate, examining individuals living in counties with very high or very low unemployment rates (counties in either the top or bottom decile of unemployment rate).

Estimates of the average marginal effect of the female dummy variable are presented in panel A of Figure 5, along with the average marginal effect we estimated from our preferred specification in the previous section to serve as a baseline. When we examine individuals who do not live in a mass-layoff county, we can see that our estimate of the average gender difference is similar to our baseline estimate. However, when looking at individuals living in a mass-layoff county, we see that the gender gap nearly disappears: The average marginal effect is approximately -\$280 and is not statistically different than zero. We can also see that the gender gap is smaller in *high* unemployment counties relative to low unemployment counties by about \$400, though both estimates indicate that men have higher limits on average compared to women. This evidence together seems to suggest that labor market conditions explain some of the gender gap in total limits.

Our results showing that the gender gap closes in mass layoff counties are consistent with the results from Braxton, Herkenhoff, and Phillips (2023), who show that individuals who lose their jobs in a mass layoff event experience a decline of approximately \$1,000 in total bankcard credit limit. However, the authors do not look at this result separately by gender. For the gender gap in limits closely following a mass layoff, we would expect that men should be become unemployed at a higher rate than women. To test this, we divide mass layoff counties into two groups: those counties where the percent of men involved in a mass layoff is greater than the percent of women involved and those counties where the percent of women involved in a mass layoff is greater than the percent men involved. We note that 90 percent of mass layoffs resulted in more men being laid off than women. Estimates of the average marginal effects from these two subsamples are presented in panel B of Figure 5. The results show that in the counties where a higher percent of men are involved in mass layoffs, the gender gap shrinks relative to the baseline difference, while in counties

where a higher percent of women are involved in a mass layoff, the gender gap widens. This is the result we should expect if employment status plays a role in the gender gap.

Overall, these results suggest that income and employment are important determinants of the gender gap bankcard limits. While these results are not causal, they do provide additional contextual evidence on the potential drivers of the gender gap.

5 Decomposing the Gender Difference

While the previous results document the gender differences in total bankcard limit and some sources of heterogeneity, we are also interested if observable factors drive this difference. If the gender credit gap is driven completely by differences in observable borrower characteristics, then the gap might be attributable to gender differences in socioeconomic and/or demographic factors. However, if the gap is not due to level differences in borrowers' characteristics, but due to differences in how the characteristics are weighted or due to differences in unobservable factors, this could be suggestive of a number of explanations, including differential treatment in the market. To assess if differences in either observable or unobservable factors explain the gender gap in limits, we employ the three-fold version of the Kitagawa-Oaxaca-Blinder (KOB) decomposition (Kitagawa, 1955; Blinder 1973; Oaxaca 1973).²³ Using this technique, we divide the gender gap into three parts: the part of the gap that can be explained by gender differences in observable characteristics (the endowment effect), the part of the gap that can be attributed to gender differences in the *effect* of the observable characteristics (the coefficient effect), and an interaction term that accounts for the fact that differences in endowments and coefficients can happen simultaneously (the interaction effect).²⁴

5.1 Empirical Decomposition Specification

To analyze and decompose the average difference between genders, we estimate the following reduced form specification for each gender.

$$y_{i\theta} = \beta_{0\theta} + \gamma_{\theta}score_{i\theta} + \delta_{\theta}income_{i\theta} + \mathbf{\Pi}_{\theta}\mathbf{X}_{i\theta} + \epsilon_{i\theta}; \theta \in (A, B), \quad (3)$$

²³See Winsborough and Dickinson (1971), Jones and Kelley (1984), Jann (2008), and Fortin et al. (2011) for details on three-fold variation.

²⁴We summarize the method in Appendix Section E.

where groups A and B represent male and female borrowers, respectively. Similar to Equations (1) and (2) in Section 4, we include the vectors $score_i$ and $income_i$, which contains dummy variables for 50 point credit score bins and \$25,000 income bins, respectively, and the total number of bankcard accounts.²⁵ The control variable vector $\mathbf{X}_{i\theta}$ contains year fixed effects, race fixed effects, month-of-the-year fixed effects, state-fixed-effects, and year-of-birth fixed effects.²⁶ Because we estimate Equation (3) separately for gender, we are unable to include the interaction terms that were previously included in Equation (2). After estimating Equation (3) for each gender, we then obtain the estimated mean decomposition of the difference between genders by taking the difference of the predicted values from each regression in Equation (3) and then algebraically rearranging the difference as expressed in Equation (5) in Appendix E.

There are two caveats to our decomposition analysis. First, because of data limitations, there are a number of unobserved factors that we cannot account for in Equation (3), such as education and marital status, which could also affect the gender gap in credit. However, our results will be consistently estimated if the dependence between the unobservable and observable elements is the same for both males and females (Fortin et al., 2011). Since our sample includes a relatively narrow group of individuals (i.e., successful mortgage borrowers who are sole applicants), it is plausible that this assumption is met. Second, since Equation (3) also suffers from an endogeneity problem, due to omitted variable bias, our decomposition analysis is constructed from correlation measures, not strictly causal ones. Thus, we do not offer a causal interpretation of our results.

5.2 Decomposition Results

The KOB decomposition results from our cross-section sample of sole mortgage applicants for total bankcard limit are displayed in Table 5. We estimate that men’s limits are higher than women’s limits by \$1,499.09, which is 5 percent of the mean male total bankcard limit.²⁷ When decomposing this gap, we find a positive endowment effect of \$162.83, a positive coefficient effect of \$1,312.65, and a positive interaction effect of \$23.61.

²⁵As before, credit score is the Equifax Risk Score.

²⁶We include year-of-birth fixed effects and exclude state-year fixed effects from the main aggregate decomposition for computational efficiency. Our decomposition results are robust to alternate specification where we use age fixed effects instead of year-of-birth fixed effects and including state-by-year fixed effects.

²⁷This estimate differs from our estimates reported in Table 2 because Equation (3) includes different regressors than in Equation (2).

The endowment effect of \$162.83 implies that 10.9 percent of the gender difference in total bankcard limit is due to differences in the values of observed characteristics between male and female borrowers. If female borrowers had the same characteristics as men (i.e., the same values for the covariates), female total bankcard limits would be \$162.83 higher. Results from the detailed decomposition, presented in Appendix Table A6, show that this effect is primarily driven by differences in the level of income between genders, with male borrowers having higher incomes than female borrowers, and differences in the number of bankcards, where women have more bankcards than men.

The estimate of the coefficient effect implies that 87.6 percent of the estimated gap of \$1,499.09 is due to the difference in the values of the coefficients on the observed characteristics. This large, positive coefficient effect indicates that male borrowers receive a larger effect, or return, from the same values of observed characteristics than female borrowers.²⁸ In the KOB framework, this suggests that women’s limits would increase by \$1,312.65 if we applied the men’s coefficients to the women’s characteristics. The detailed decomposition results show that the coefficient effect is driven primarily by differences in returns to (1) the number of bankcards and credit score, with male borrowers having larger effects (higher return) than female borrowers, and (2) income, with women having larger effects. Larger effects for these covariates in the detailed decomposition indicate that these factors play a relatively larger role in explaining the differences in the effects (i.e., returns) between men and women sole applicants.

Overall, these results provide evidence that the coefficient effect is the primary factor in explaining the gender gap in total bankcard limits. Although we cannot draw any causal conclusions regarding the nature of this differential return, as we do not identify if these disparities are derived from demand or supply (nor the relative share between the two), our findings do suggest that the differences in bankcard limits are primarily due to difference in the *returns* to observed characteristics rather than the levels of these characteristics.

²⁸Using the term “return” to explain the coefficient effect is commonplace in labor economics, where workers receive returns on their investments, such a human capital, on their wages.

5.3 Heterogeneity in the Decomposition

While performing this decomposition at the mean of the entire sample provides us with a meaningful analysis of what drives the gender gap for the average male and female sole mortgage applicant borrowers, we may not expect these decomposition results to hold at different points in the credit limit distribution. For example, the summary statistics on both total and average bankcard limits in Table 1 indicate that the gender gap decreases as limits get smaller and grows in magnitude as limits get larger.²⁹ Also, based on the time trends we observe for total bankcard limits in Figure 3, we know that there are differences in the gender gap over time. Given the differences in the gender gap across these dimensions, it is reasonable to assume that performing the decomposition at different points in the distribution at different points in time would yield substantially different results than those reported in Table 5.

To perform these decompositions at other points in the credit limit distribution for each year in our data, we follow Fortin et al. (2011) and use the unconditional quantile regression methodology developed by Firpo, Fortin, and Lemieux (2009) and run a series of regressions of the recentered influence function (RIF) on our observed characteristics for each decile in each year for each gender.³⁰ To decompose the gender differences at different quantiles, we use the same approach as the one outlined in Section 5.1 but now replace the outcome variable with the RIF of the outcome variable and define separate estimating equations for each gender: $RIF(Y_i; Q_\tau)_{tA} = X_A\beta_A + \epsilon_A$ and $RIF(Y_i; Q_\tau)_{tB} = X_B\beta_B + \epsilon_B$. Taking the difference of the fitted values for each gender to create the gender gap $\hat{\Delta}_{\tau,t}$, we can then write an equivalent KOB decomposition for any unconditional quantile τ in time t . To analyze the magnitudes of gender differences across time for different deciles, we present our results in the form of heat maps for more compact presentation.

Panel A of Figure 6 shows the gender differences for our total bankcard limit variable across time for each decile of the limit variable. At the lower end of the total bankcard credit limit distribution, the differences in bankcard credit limit between genders is relatively small, with the gender gap frequently favoring female borrowers with magnitudes varying from \$0 to \$300 up to the 30th percentile. This gender difference in the left tail of the distribution has increased over

²⁹This kind of heterogeneity is also observed in the gender wage gap, where the largest disparity between groups occurs at the right-hand tail of the distribution (Blau and Kahn, 2017).

³⁰We refer readers to the Firpo et al. (2009) and Fortin et al. (2011) studies for the technical details of this methodology.

time, with the gap growing more in favor of female borrowers. For the middle of the distribution, the gender gap at the beginning of our sample in 2006 favors male borrowers and is approximately \$300 to \$900. The gap shrinks starting in 2010, actually favoring female borrowers, with the gap as high as -\$300 at the median of the distribution. The gap favors male borrowers in 2014 and 2015 before becoming negative again in 2016.

However, the differences between male and female borrowers at the higher deciles favor male borrowers and is larger in magnitude: At the 80th percentile, the gender gap ranges from \$700-\$3,300, while at the 90th percentile the difference is over \$3,400 for the majority of years in our sample. Unlike the gender differences in the center and left tail of the distribution, the difference at the 90th percentile has remained relatively constant from 2006 to 2016.

The heat map in panel B of Figure 6 shows the gender gap as a percentage of the male average total bankcard limit. This figure shows the economic significance of the gender gap across deciles and across time. In the middle and right tail of the distribution, the size of the gap is frequently 2 percent to 3 percent of the male total bankcard limit and never exceeds 6 percent, indicating that while the gender gap is statistically significant, its magnitude does not reach the levels of the gender wage gap, which are approximately 18 percent for full-time workers (Hegewisch and Barsi, 2020). We discuss results for the endowment effect and the coefficient effect in Appendix F.

6 Credit Card Supply

Thus far, our analysis has focused on total bankcard credit limits, which is an equilibrium outcome of the amount of credit card debt available to an individual. Because of the limitations we faced in our earlier analyses, it is not possible to identify which portions of the gender differences we estimate for total bankcard limit are due to demand side factors and which are due to supply side factors. To better understand how credit card supply may differ by gender, we follow Firestone (2014) and Han, Keys, and Li (2018) and utilize an anonymized data set of direct mail credit card solicitations from Mintel Comperemedia, Inc. Direct Mail Monitor to test if there are differences in this measure of credit supply. Mintel Comperemedia is a consumer and marketing research firm that collects direct mail credit card offers, among other items, from a set of households every month. Along with information on each piece of mail from the household participants, the data set

includes demographic and other household information from a separate survey and is then merged with credit bureau data from TransUnion. The resulting Mintel Comperemedia, Inc. Direct Mail Monitor and TransUnion LLC Match data set (Mintel/TransUnion data) gives us a snapshot of lenders' offers of credit, as opposed to our main HMC data set, which captures the accumulation of supply and demand decisions of consumer credit over time.

We restrict the Mintel/TransUnion data to mailers that have been sent to consumers who have an active mortgage loan at the time they received the mailer from 2009 to 2017. Note that while we cannot identify the gender of all survey participants, we can identify the gender of individuals in the majority of married households, or approximately 71 percent of the entire data set.³¹ Single heads of households are identified, as are a subset of married households. We provide demographic summary statistics for male and female respondents in the Mintel/TransUnion data in Table 6 and credit card mailer summary statistics in Table 7.

Approximately 46 percent of the individuals in our Mintel/TransUnion data are female and the sample overall heavily is skewed toward White individuals, relative to mortgage holders (see Appendix Table A1). Individuals in our sample are also fairly creditworthy, with an average credit score of almost 800.³² As can be seen in Table 7, female survey respondents received 2.23 mailers on average, compared to 2.52 mailers for male respondents during our sample period. We also observe a difference in advertised credit limits³³ by gender: Male respondents receive slightly lower average advertised limit offers than female respondents (\$5,996.53 compared to \$6,217.43), though this difference is not statistically significant; males and females also have the same median offer of \$1,500. Using the TransUnion credit bureau variables, we examine some recent measures of credit seeking activity: hard inquiries and the success rate³⁴ of the inquiries, prior to the credit card solicitations. While they occur within 6 months of receiving the credit card mailers, they are not a response to receiving solicitations.³⁵ Interestingly, we find that although men receive more credit mailers on average than females in our sample, they have slightly fewer inquiries for credit than females do: 0.495 inquiries for men compared to 0.503 for women. We also observe that females

³¹Specifically, we can accurately identify gender in 142,242 out of 200,253 observations in the full data.

³²Credit score in the Mintel/TransUnion data is the VantageScore 2.0. We use a transformed version of the VantageScore 2.0 so that its range of possible scores is similar to other commercial credit scores.

³³Credit limit is defined as the maximum amount of credit available on the card.

³⁴We define success rate as the number of bankcard trades opened in previous 6 months divided by the number of inquiries in the past 6 months, excluding auto and mortgage inquiries.

³⁵In addition, we are unable to identify if specific inquiries lead to the opening of specific accounts.

have fewer newly opened bankcard accounts within the previous 12 months, compared to males (0.37 for men compared to 0.33 for women).

Similar to total bankcard limit, it is likely that the relationship between the number of credit card offers and credit score is nonlinear. To see how the number of credit card mailers individuals receive varies by credit score, we calculate the average number of mailers received by consumers in each VantageScore 2.0 score bin for both genders. The results are plotted in Figure 7. We can see that men and women with credit scores below 600 receive similar numbers of offers and receive fewer credit card offers than individuals with higher credit scores. Individuals with credit scores greater than 700 receive the most credit offers by mail, with men receiving more mailers than women. Interestingly, individuals in the highest credit score bin receive fewer offers than individuals with credit scores in the 650-799 range.

To better examine these differences between men and women for the number of credit card offers received, we follow Firestone (2014) and estimate a standard Poisson regression model where the dependent variable is the number of mailers received in a month. Our estimating equation takes the following form:

$$Pr(Y_i) = F(\beta_1 Female_i + \mathbf{X}_i \boldsymbol{\Pi}). \quad (4)$$

The vector X_i contains a number of control variables, including race, education, marital status, state of residence, credit score bin, income bin, credit limit utilization, number of inquiries in the past 6 months, a dummy variable equal to one if the individual has ever filed for bankruptcy, and the age of the individual's oldest credit account. Similar to Equation (2), we also include interaction terms between the female dummy and credit score and income bins.

Results from the Poisson regression are reported in Table 8. The coefficient estimate on the *Female* dummy variable shows that female respondents receive fewer offers than male respondents, and the coefficients on interaction terms for *Female* and income categories indicate that female borrowers receive fewer offers in every income category. If we translate these interaction coefficients to incidence rate ratios, females receive offers at a rate of 0.8-0.9 times the rate of male borrowers across each income bin. That women receive fewer offers than male borrowers at higher income levels is consistent with our previous results that show that women had lower total bankcard limits relative to men at the higher income categories. Interestingly, we do not observe statistically

significant differences between men and women in the offers received within different credit score categories, except for the lowest bin, where women receive more offers than men. This would seem to imply that while female borrowers receive fewer offers than male borrowers overall, these differences do not change as credit score changes. Although these results indicate that women receive fewer offers than men, our results are only correlations and they do not suggest what mechanisms may be driving this result.

7 Discussion

The results presented so far in this analysis illustrate that there are gender differences in both total bankcard limits and in the credit supplied to consumers as measured by credit card mail offers. As mentioned previously, the nature of our data and analyses do not allow us to causally identify the specific mechanisms that drive these differences. Despite this limitation, we believe that there are three broad potential explanations for our results and that it is highly likely that some combination of these mechanisms are at play.

The first possible explanation for our results is that banks and lenders are consciously assigning limits differently based on gender. We view this as both an infeasible and highly unlikely reason, given the intense scrutiny given to banks' lending strategies. For example, satisfying ECOA regulations effectively removes gender from directly entering the underwriting decision and banks generally do not collect information on gender and other protected demographic characteristics to avoid even the appearance of treating individuals of specific characteristics differently. We also note that the current system for applying and underwriting bankcards is highly automated, which further reduces lenders' ability to introduce bias into the process. This is in stark contrast with the experience 50 years ago when loan officers would actually know the person they were lending to.

That we find gender differences in total limits even after controlling for a wide variety of factors may indicate instead that automation in the credit card market has its own limitations. Recent research has suggested that use of algorithms, machine learning, and artificial intelligence in credit markets can still lead to biased outcomes (Morse and Pence, 2020; Fuster, Goldsmith-Pinkham, Ramadorai, and Walther, 2021).³⁶ This possibility is further strengthened by the fact that we

³⁶With the increased prevalence of technology and the use of alternate data in lending decisions, it is possible that variables used in lending decisions may serve as a *proxy* for membership in a protected class such as gender. See

do not observe significant differences in the likelihood of delinquency between men and women in Appendix Table A5. We can see that there are no significant differences in the probability of delinquency between male and female borrowers across the income distribution and the credit score distribution. The possibility of the limitations of technology would also be consistent with the broader interdisciplinary research that has found that the use of these kinds of technologies may lead to biased outcomes in a number of other settings.

A second possible explanation is that there are substantial differences in the socioeconomic factors of average male and average female bankcard applicants at the time of origination, and the differences in initial credit limits subsequently do not converge during the account's lifetime. We view this rationale as significantly more probable as women face disparities along a number of dimensions that can potentially affect credit line assignment. For example, it is well documented that on average, women experience worse labor market outcomes than men do, which in turn leads to a gap in wages.³⁷ With regard to wealth, it has been shown that single women own 32 cents on the dollar compared with single men (Baker, Martin-West, and Famakinwa, 2018). Since ability to pay is an important determinant of how much limit a credit card lender will offer, disparities in income will likely impact gender differences in total bankcard limits. Even if men and women receive similar *subsequent* bankcard limit increases over time, this may not be enough to make up for the initial disparity in credit line assignment.

Although we do not have data on income at the time of bankcard origination (which would allow us to more definitively test this mechanism), our regression and decomposition results are consistent with employment and current income likely playing an important role in explaining the gender differences in total bankcard limits. As we show in Figure 5, we observe that gender gap shrinks in geographic areas experiencing economic distress (as measured by employment). For individuals living in counties with high unemployment rates, we estimate the gender gap is -\$750, and for individuals living in counties experiencing mass layoffs (as defined in Foote et al., 2019), the estimate is -\$280 and is not statistically different than zero. That the gender gap shrinks in these counties experiencing distress indicates that employment status, along with income, is an important factor in explaining the gender gap in bankcard limits. It is also likely that employment

Morse and Pence (2020) for an overview of these developments.

³⁷In 2019, full-time, year-round female workers made 82 cents for every dollar men earned (Hegewisch and Barsi, 2020).

and income have important interaction effects, with employment status mediating the effect of income on limits while also affecting limits directly.

Additionally, our detailed decomposition results, shown in Appendix Table A6, show that differences in the *level* of income (the endowment effect) favor male borrowers and play a large role in explaining the part of the gender difference that is due to levels of observed characteristics. Though we document that the endowment effects contribute less to the gender gap than the coefficient effect, these results suggest that the majority of the endowment effect is almost entirely driven by differences in income. For our regression results, we see that even within the same income bins, female borrowers have lower total limits than male borrowers. This may indicate that income information, along with information on employment status and/or occupation and job tenure, may be processed differently by the same algorithm if this information is a proxy for the future ability to pay. This may be the case since there are documented gender differences in occupational choice and career trajectories, and these differences may lead to income information being treated differently by gender.

The third possible explanation is that there is a fundamental difference in credit card seeking, shopping, and management behavior by gender, and this in turn leads to different accumulation patterns of credit card limits over time. This mechanism is also very plausible as prior research has shown that women and men manage housing credit differently (Goodman, Zhu, and Bai, 2016), which may extend to behavior with credit cards. There is also evidence that women have lower levels of financial literacy and are less confident about their math skills, which makes women more likely to engage in more costly credit card behavior (Mottola, 2013). Additionally, there is evidence that shopping behavior can partially explain the different prices individuals pay for credit (Morse and Pence, 2020). While it is technically possible that men and women have *identical* credit card management behaviors, we view this as also unlikely.

Unfortunately, the HMC data do not contain information that would allow us to assess credit card preferences or risk preferences. As mentioned previously, total bankcard limit is a measure of accumulated supply and demand decisions over time, which prevents us from being able to disentangle which portion of total limit is due to risk preferences and which portion is due to other factors. We could use information from the mortgage application data to infer some information regarding risk preferences over loan products, but that would implicitly assume that risk preferences

for housing are a good proxy for risk preferences for credit cards. While we view this as unlikely, given the substantial differences between a mortgage and a credit card, this is a potential area for further study.

8 Conclusion

We present new evidence on the presence of gender differences in bankcard limits in the United States using a unique data set that combines mortgage application information with credit bureau data. After accounting for differences in income, credit score, and other demographic characteristics, we estimate an average marginal effect of -\$1,323 for female borrowers, indicating that female sole mortgage applicants, on average, have lower total bankcard limits relative to male applicants. This gender difference varies along both income and credit score, with the gap increasing in favor of male borrowers at both higher incomes and higher credit scores. Results from a Kitagawa-Oaxaca-Blinder decomposition show that approximately 87 percent of this difference can be explained by differences in the returns to observable characteristics rather than the characteristics themselves. This implies that differences in the *coefficients* of our observed characteristics, not the levels, explain the majority of this gender difference.

Our heterogeneity analyses show that both the gender difference and its decomposition are not consistent over time or across the bankcard limit distribution. Our analyses show that (1) gender difference favors women in the far left tail of the total limit distribution, favors men in the far right tail, and has decreased over time for the bottom 6 deciles of the total bankcard limit distribution, and (2) the factors that drive these gender differences have varied over time for a large portion of the credit limit distribution, with the coefficient effect typically favoring male borrowers. Our analysis of the credit card mail offers indicates that women receive fewer offers than men and receive different kinds of offers, though these differences are not large in economic magnitude. Results from a standard Poisson regression show that this gender gap in offers persists after controlling for a number of demographic and credit variables.

While we are unable to causally identify the mechanism that drives these differences between genders, our results are consistent with the growing literature that show gender disparities exist in a number of other economic settings. Given that income is an important determinant in credit

limits, and that women still face disparities in pay in the U.S., it is likely that socioeconomic characteristics at the time of credit line assignment play an important role in these disparities. To what extent this is the case is an important question to be addressed by future research.

References

- [1] Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. (2015). “Regulating consumer financial products: Evidence from credit cards.” *The Quarterly Journal of Economics*, 130(1), 111–164.
- [2] Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. (2018). “Do banks pass through credit expansions to consumers who want to borrow?” *The Quarterly Journal of Economics*, 133(1), 129–190.
- [3] Albrecht, James, Anders Bjorklund, and Susan Vroman. (2003). “Is there a glass ceiling in Sweden?” *Journal of Labor Economics*, 21(1), 145–177.
- [4] Avery, Robert B., Kenneth P. Brevoort, and Glenn B. Canner. (2007). “Opportunities and issues using HMDA data.” *Journal of Real Estate Research*, 29(4), 351–380.
- [5] Aydin, Deniz. (2021). “Consumption response to credit expansions: Evidence from experimental assignment of 45,307 credit lines.” *American Economic Review*, Forthcoming.
- [6] Ayres, Ian and Peter Siegelman. (1995). “Race and gender discrimination in bargaining for a new car.” *American Economic Review*, 85(3), 304–321.
- [7] Baker, Amy Castro, Stacia Martin-West, and Folasade Famakinwa. (2018). “On shaky ground: Stabilizing the financial security of single women.” Asset Funders Network.
- [8] Barber, Brad M. and Terrance Odean. (2001). “Boys will be boys: Gender, overconfidence, and common stock investment.” *The Quarterly Journal of Economics*, 116(1), 261–292.
- [9] Blau, Francine D. and Lawrence M. Kahn. (2017). “The gender wage gap: Extent, trends, and explanations.” *Journal of Economic Literature*, 55(3), 789–865.
- [10] Blinder, Alan S. (1973). “Wage Discrimination: Reduced Form and Structural Estimates.” *The Journal of Human Resources*, 8(4), 436–455.
- [11] Braxton, J. Carter, Kyle Herkenhoff, and Gordon Phillips. (2023). “Can the unemployed borrowing? Implications for public insurance,” *Journal of Political Economy*. Forthcoming.
- [12] Bricker, Jesse, Lisa J. Dettling, Alice Henriques, Hsu, Joanne W. Hsu, Lindsay Jacobs, Kevin B. Moore, Sarah Pack, John Sabelhaus, Jeffrey Thompson, and Richard A. Windle. (2017). “Changes in U.S. family finances from 2013 to 2016: Evidence from the Survey of Consumer Finances.” *Federal Reserve Bulletin*, vol. 103, no. 3, 1–41.
- [13] Bucher-Koenen, Tabea, Annamaria Lusardi, Rob Alessie, and Maarten Van Rooij. (2017). “How financially literate are women? An overview and new insights.” *Journal of Consumer Affairs*, 51(2), 255–283.
- [14] Canilang, Sara, Cassandra Duchan, Kimberly Kreiss, Jeff Larrimore, Ellen Merry, Erin Troland, and Mike Zabek. (2020). “Report on the economic well-being of U.S. households in 2019, featuring supplemental data from April 2020. Federal Reserve Board of Governors.” <https://www.federalreserve.gov/publications/files/2019-report-economic-well-being-us-households-202005.pdf>.

- [15] Cavalluzzo, Ken S. and Linda C. Cavalluzzo. (1998). “Market structure and discrimination: The case of small businesses.” *Journal of Money, Credit and Banking*, 30(4), 771–792.
- [16] Cheng, Ping, Zhenguo Lin, and Yingchun Liu. (2011). “Do women pay more for mortgages?” *The Journal of Real Estate Finance and Economics*, 43(4), 423–440.
- [17] Fang, Lu and Henry J. Munneke. (2020). “Gender equality in mortgage lending.” *Real Estate Economics*, 48(4), 957–1003. doi:10.1111/1540-6229.12198.
- [18] Federal Reserve Bank of New York. (2022). Quarterly report on household debt and credit: 2022:Q43.
- [19] Fernandez, Jennifer, and Anna Tranfaglia. (2020). “Gender disparities in financial well-being from the Survey of Household Economics and Decisionmaking. Federal Reserve Bank of Philadelphia.” <https://www.philadelphiafed.org/-/media/frbp/assets/community-development/reports/gender-disparities-in-financial-well-being-survey-of-household-economics-and-decisionmaking.pdf>.
- [20] Firestone, Simon. (2014). “Race, ethnicity, and credit card marketing.” *Journal of Money, Credit, and Banking*. 46(6), 1205–1224.
- [21] Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux. (2009). “Unconditional quantile regressions.” *Econometrica*, 77(3), 953–973.
- [22] Fishbein, Allen J., and Patrick Woodall. (2006). “Women are prime targets for subprime lending: Women are diproportionately represented in high-cost mortgage market.” Consumer Federation of America.
- [23] Fonseca, Raquel, Kathleen J. Mullen, Gema Zamarro, and Julie Zissimopoulos. (2012). “What explains the gender gap in financial literacy? The role of household decision-making.” *Journal of Consumer Affairs*, 46(1), 90–106.
- [24] Foote, Andrew, Michel Grosz, and Ann Stevens. (2019). “Locate your nearest exit: Mass layoffs and local labor market response.” *ILR Review*, 72(1), 101–126.
- [25] Fortin, Nicole, Thomas Lemieux, and Sergio Firpo. (2011). “Decomposition methods in economics. In Handbook of labor economics.” (Vol. 4, pp. 1–102). Elsevier.
- [26] Fortin, Nicole, Brian Bell, and Michael Boehm. (2017). “Top income inequality and the gender pay gap: Canada, Sweden, and the United Kingdom.” *Labour Economics*, 47, 107–123.
- [27] Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. (2021). “Predictably unequal? The effects of machine learning on credit markets.” *Journal of Finance*, Forthcoming.
- [28] Goldin, Claudia. (2014). “A grand gender convergence: Its last chapter.” *American Economic Review*, 104(4), 1091–1119.
- [29] Goldsmith-Pinkham, Paul and Kelly Shue. (2023). “The gender gap in housing returns.” *Journal of Finance*, 78(2), 1097–1145.
- [30] Goodman, Laurie, Jun Zhu, and Bing Bai. (2016). “Women are better than men at paying their mortgages.” Urban Institute.

- [31] Green, Claire and Joanna Stavins. (2018). “The 2016 and 2017 surveys of consumer payment choice: Summary results.” (Working Paper 18-3), Federal Reserve Bank of Boston.
- [32] Gross, David B. and Nicholas S. Souleles. (2002). “Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data.” *Quarterly Journal of Economics*, 117(1), 149–185.
- [33] Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. (2020). “The marginal propensity to consume over the business.” *American Economic Journal: Macroeconomics*, 12(2), 351–384.
- [34] Han, Song, Benjamin J. Keys, and Geng Li. (2018). “Unsecured credit supply, credit cycles, and regulation.” *Review of Financial Studies*, 31(3), 1184–1217.
- [35] Hegewisch, Ariane and Zohal Barsi. (2020). “The gender wage gap 2019: Earnings differences by gender, race, and ethnicity.” Fact Sheet, IWPR #C484. Washington, DC: Institute for Women’s Policy Research. <http://iwpr.org/iwpr-issues/employment-and-earnings/the-gender-wage-gap-by-occupation-2019/>.
- [36] Jann, Ben. (2008). “The Blinder-Oaxaca decomposition for linear regression models.” *The Stata Journal*, 8(4), 453–479.
- [37] Jones, F.L. and Jonathan Kelley. (1984). “Decomposing differences between groups: A cautionary note on measuring discrimination.” *Sociological Methods & Research*, 12(3), 323–343.
- [38] Lee, Donghoon and Wilbert van der Klaauw. (2010). “An introduction to the consumer credit panel.” Federal Reserve Bank of New York Staff Report 479.
- [39] Li, Geng. (2018). “Gender-related differences in credit use and credit scores.” (No. 2018-06-22). Board of Governors of the Federal Reserve System.
- [40] Lusardi, Annamaria and Olivia S. Mitchell. (2008). “Planning and financial literacy: How do women fare?” *American Economic Review*, 98(2), 413–17.
- [41] Lusardi, Annamaria, Olivia S. Mitchell, and Vilsa Curto. (2010). “Financial literacy among the youth.” *Journal of Consumer Affairs*, 44(2), 358–380.
- [42] Lusardi, Annamaria and Olivia S. Mitchell. (2014). “The economic importance of financial literacy: Theory and evidence.” *Journal of Economic Literature*, 52(1), 5–44. 10.1257/jel.52.1.5.
- [43] Kitagawa, Evelyn M. (1955). “Components of a difference between two rates.” *Journal of the American Statistical Association*, 50(272), 1168–1194.
- [44] Morse, Adair and Karen Pence. (2020). “Technological innovation and discrimination in household finance.” (NBER Working Paper 26739), National Bureau of Economic Research.
- [45] Morton, Fiona Scott, 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, 65–92.
- [46] Mottola, Gary R. (2013). “In our best interest: Women, financial literacy, and credit card behavior.” *Numeracy*, 6(2), 4.
- [47] Oaxaca, Ronald. (1973). “Male-Female Wage Differentials in Urban Labor Markets.” *International Economic Review*, 14(3), 693–709.

- [48] Ruggles, Steven et al. (2021) IPUMS USA: Version 11.0 [data set]. Minneapolis, MN: IPUMS, 2021. <https://doi.org/10.18128/D010.V11.0>
- [49] Winsborough, H.H. and Peter Dickinson. (1971). “Components of the negro-white income differences.” In *Proceedings of the Social Statistics Section*, 6-8. Washington, DC: American Statistical Association.

Table 1: Summary Statistics for Male and Female Borrowers

	Income (thousands) (\$)	Credit Score	Number of Bankcard Accounts	Total Bankcard Balance (\$)	Average Bankcard Limit (\$)	Total Bankcard Limit (\$)
Total						
Mean	87	723	3.29	8,086	8,854	29,419
Median	63	752	3	3,449	7,650	20,900
25th Percentile	43	680	2	925	4,000	8,800
75th Percentile	96	798	4	9,842	12,100	40,550
Male						
Mean	99	723	3.22	8,340	9,227	30,079
Median	69	752	3	3,491	7,945	21,000
25th Percentile	47	681	2	918	4,100	8,800
75th Percentile	106	798	4	9,973	12,625	41,200
Female						
Mean	72	723	3.38	7,750	8,359	28,544
Median	56	752	3	3,394	7,372	20,700
25th Percentile	39	679	2	933	3,900	8,800
75th Percentile	83	799	4	9,673	11,500	39,800
Percent Female: 42.97%						
Percent White: 73.81%						

Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit Score is the Equifax Risk Score. Income is the HMDA income, reported at the time of mortgage application. Demographic information comes from the HMDA data.

Table 2: OLS Results for Total Bankcard Limit

	Baseline		Preferred	
	Coefficient	(Std. Error)	Coefficient	(Std. Error)
<i>Constant</i>	-27,610.25***	(1,018.88)	-30,203.14***	(1,049.07)
<i>Female</i>	-1,272.14***	(74.40)	4,471.98***	(928.43)
<i>Number of cards</i>	8,684.44***	(136.32)	9,042.58***	(131.56)
<i>Female × Number of cards</i>			-817.51***	(37.78)
Credit Score Category				
<i>Female × I(350 – 399)</i>			-1,205.12	(1,050.35)
<i>Female × I(400 – 449)</i>			-2,719.95***	(907.81)
<i>Female × I(450 – 499)</i>			-2,727.47***	(900.08)
<i>Female × I(500 – 549)</i>			-2,665.04***	(834.10)
<i>Female × I(550 – 599)</i>			-2,993.77***	(923.74)
<i>Female × I(600 – 649)</i>			-2,979.60***	(914.80)
<i>Female × I(650 – 699)</i>			-2,874.40***	(905.34)
<i>Female × I(700 – 749)</i>			-2,934.04***	(914.77)
<i>Female × I(750 – 799)</i>			-3,005.27***	(878.09)
<i>Female × I(800 – 850)</i>			-4,014.08***	(877.88)
Income Categories				
<i>Female × I(21 – 30)</i>			-354.24	(326.32)
<i>Female × I(31 – 40)</i>			-34.29	(285.18)
<i>Female × I(41 – 50)</i>			354.27	(266.26)
<i>Female × I(51 – 75)</i>			404.32	(286.29)
<i>Female × I(76 – 100)</i>			777.59**	(347.95)
<i>Female × I(101 – 125)</i>			294.14	(393.67)
<i>Female × I(126 – 150)</i>			-542.81	(360.66)
<i>Female × I(151 – 200)</i>			-466.08	(449.06)
<i>Female × I(201 – 250)</i>			-839.18	(574.03)
<i>Female × I(> 250)</i>			-4,173.15***	(614.63)
Year FE	Yes		Yes	
State FE	Yes		Yes	
State × Year FE	Yes		Yes	
R^2	0.6299		0.6305	
N	530,122		530,122	

Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit Score is the Equifax Risk Score. Demographic information comes from the HMDA data. Standard errors are clustered at the state level. Reference group consists of White males, with the lowest credit score (280-349), and the lowest income (less than \$20,000). Items omitted from table include uninteracted income and credit score groups. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 3: Robustness Check: Controlling for Prior Financial Distress

	Fin. Dist. Dummy		Fin. Dist. Interaction	
	Point Estimate	(Std. Error)	Point Estimate	(Std. Error)
<i>Constant</i>	-28,972.15***	(1,145.05)	-28,818.27***	(1,126.59)
<i>Female</i>	4,485.47***	(934.87)	4,135.67***	(975.79)
<i>Number of cards</i>	9,036.33***	(132.03)	9,036.22***	(132.00)
<i>Financial distress</i>	-1,740.44***	(124.42)	-1,923.97***	(137.55)
<i>Female × Number of cards</i>	-820.90***	(37.37)	-820.09***	(37.72)
<i>Female × Financial distress</i>			422.94***	(157.04)
Credit Score Category				
<i>Female × I(350 – 399)</i>	-1,243.25	(1,052.37)	-1,224.59	(1,053.30)
<i>Female × I(400 – 449)</i>	-2,748.15***	(913.02)	-2,718.13***	(915.17)
<i>Female × I(450 – 499)</i>	-2,768.96***	(906.78)	-2,708.67***	(912.26)
<i>Female × I(500 – 549)</i>	-2,690.23***	(835.34)	-2,592.26***	(844.46)
<i>Female × I(550 – 599)</i>	-3,012.82***	(930.06)	-2,870.02***	(944.75)
<i>Female × I(600 – 649)</i>	-2,968.25***	(924.26)	-2,774.22***	(939.98)
<i>Female × I(650 – 699)</i>	-2,845.06***	(912.89)	-2,603.14***	(935.33)
<i>Female × I(700 – 749)</i>	-2,912.29***	(922.40)	-2,618.27***	(946.20)
<i>Female × I(750 – 799)</i>	-2,986.25***	(886.18)	-2,650.92***	(916.90)
<i>Female × I(800 – 850)</i>	-3,989.88***	(884.09)	-3,634.20***	(925.27)
Income Categories				
<i>Female × I(21 – 30)</i>	-352.32	(330.63)	-353.87	(330.61)
<i>Female × I(31 – 40)</i>	-44.09	(288.15)	-48.27	(287.13)
<i>Female × I(41 – 50)</i>	335.93	(269.90)	329.25	(268.08)
<i>Female × I(51 – 75)</i>	374.11	(291.24)	362.10	(288.36)
<i>Female × I(76 – 100)</i>	749.62**	(350.52)	730.15**	(345.50)
<i>Female × I(101 – 125)</i>	273.00	(398.58)	245.25	(390.93)
<i>Female × I(126 – 150)</i>	-563.25	(363.53)	-595.23	(357.66)
<i>Female × I(151 – 200)</i>	-462.80	(454.47)	-501.09	(446.73)
<i>Female × I(201 – 250)</i>	-848.90	(571.71)	-888.89	(567.50)
<i>Female × I(> 250)</i>	-4,168.69***	(613.86)	-4,212.5***	(625.01)
Year FE		Yes		Yes
State FE		Yes		Yes
State × Year FE		Yes		Yes
R^2	0.6314		0.6314	
N	530,122		530,122	

Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit Score is the Equifax Risk Score. Demographic information comes from the HMDA data. Specification includes state, year, and state-by-year fixed effects. Standard errors are clustered at the state level. Reference group consists of White males, with the lowest credit score (280-349), and the lowest income (less than \$20,000). *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 4: Robustness Check: Restrict Sample to Individuals Without Prior Financial Distress

	Point Estimate	(Std. Error)
<i>Constant</i>	-36,482.03***	(2,477.83)
<i>Female</i>	8,179.7***	(2,753.69)
<i>Number of cards</i>	9,449.06***	(135.70)
<i>Female × Number of cards</i>	-914.21***	(34.74)
Credit Score Category		
<i>Female × I(350 – 399)</i>	-2,614.25	(2,846.92)
<i>Female × I(400 – 449)</i>	-5,217.79*	(2,780.23)
<i>Female × I(450 – 499)</i>	-6,260.57**	(2,598.56)
<i>Female × I(500 – 549)</i>	-5,956.93**	(2,669.81)
<i>Female × I(550 – 599)</i>	-6,616.15**	(2,906.46)
<i>Female × I(600 – 649)</i>	-6,376.70**	(2,7031.67)
<i>Female × I(650 – 699)</i>	-6,582.52**	(2,804.19)
<i>Female × I(700 – 749)</i>	-6,514.57**	(2,769.85)
<i>Female × I(750 – 799)</i>	-6,565.81**	(2,760.35)
<i>Female × I(800 – 850)</i>	-7,490.78***	(2,723.42)
Income Categories		
<i>Female × I(21 – 30)</i>	-341.78	(329.09)
<i>Female × I(31 – 40)</i>	67.41	(316.68)
<i>Female × I(41 – 50)</i>	525.41	(342.50)
<i>Female × I(51 – 75)</i>	479.21	(337.57)
<i>Female × I(76 – 100)</i>	777.41*	(411.66)
<i>Female × I(101 – 125)</i>	205.61	(427.32)
<i>Female × I(126 – 150)</i>	-499.69	(497.03)
<i>Female × I(151 – 200)</i>	-370.91	(538.33)
<i>Female × I(201 – 250)</i>	-508.17	(758.02)
<i>Female × I(> 250)</i>	-4,049.29***	(666.04)
Year FE		Yes
State FE		Yes
State × Year FE		Yes
R^2	0.6424	
N	431,413	

Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit Score is the Equifax Risk Score. Demographic information comes from the HMDA data. Specification includes state, year, and state-by-year fixed effects. Standard errors are clustered at the state level. Reference group consists of White males, with the lowest credit score (280-349), and the lowest income (less than \$20,000). *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 5: Aggregate Kitagawa-Oaxaca-Blinder Decomposition Results

	Total Bankcard Limit
Mean Male Outcome	\$29,979.91 (55.94)
Mean Female Outcome	\$28,480.82 (57.83)
Mean Gender Differential	\$1,499.09 (80.46)
Endowment Effect	\$162.83 (64.02)
Coefficient Effect	\$1,312.65 (52.04)
Interaction Effect	\$23.61 (25.61)
<i>N</i>	530,125

Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Demographic information comes from the HMDA data. Specification includes year fixed effects, race fixed effects, month-of-the-year fixed effects, state-fixed effects, and age-fixed effects. Standard errors are calculated using the standard Huber/White estimator and reported in parentheses.

Table 6: Demographic and Credit Characteristics of the Mintel/TransUnion Sample

Variable	Mean	Median	Std. Dev
Birth Year	1961	1961	12
Credit Score	767	792	81.4
Number of Credit Cards	8.76	7	6.11
Number of Credit Cards, Balance > \$0	2.27	2	1.94
Total Credit Card Limit	\$49,305	\$34,000	\$55,914

	Overall	Males	Females
Income			
<\$30,000	8.4%	6.5%	10.3%
\$30,000 - \$74,499	36.5%	35.0%	38.2%
\$75,000 - \$149,000	46.9%	49.4%	44.0%
≥\$150,000	8.4%	9.1%	7.5%
Education			
High School or Less	32.7%	33.0%	32.4%
Some College	21.7%	21.6%	21.8%
Bachelor’s Degree or Higher	45.6%	45.4%	45.8%
Race			
White	88.4%	89.4%	88.3%
Black	4.5%	3.9%	5.3%
Asian/Pacific Islander	3.0%	3.1%	3.0%
<i>N</i> = 46,428			

Notes: Authors’ calculations using data from Mintel/TransUnion. Demographic information comes from the Mintel data. Sample includes individuals who had an active mortgage tradeline from 2009 - 2017 and received at least one credit card promotional mailer. We drop any observations where we are unable to determine gender for all cases (e.g., members of a multi-person household where we are unable to determine the gender of the individual who received the mailer). Resulting data set includes both single and married individuals, as well as individuals living by themselves within multi-person households. Credit score is a transformed version of VantageScore 2.0.

Table 7: Card Offer Characteristics of the Mintel/TransUnion Sample

	Males	Females
Average # of Promotional Mailers (per consumer)	2.52	2.23
Breakdown of Card Type		
Affinity	0.039	0.034
Co-Brand	0.205	0.197
Credit	0.691	0.690
Lifestyle	0.065	0.079
Average Credit Limit Advertised (\$)†	5,996.53	6,217.43
Limit by Application Type (\$)		
Affinity	14,883	12,520
Co-Brand	9,343	8,339
Credit	12,118	12,866
Lifestyle	1,818	1,779
Average # of Inquiries	0.495	0.503
Average # of Bankcards Opened, Past 12 Months	0.37	0.33
Success Rate‡	0.246	0.227

Notes: Authors' calculations using data from Mintel/TransUnion. Demographic information comes from the Mintel data. Sample includes individuals who had an active mortgage tradeline from 2009 - 2017 and received at least one credit card promotional mailer. We drop any observations where we are unable to determine gender for all cases (e.g., members of a multi-person household where we are unable to determine the gender of the individual who received the mailer). This translates to dropping 28.95 percent of the full sample. Resulting data set includes both single and married individuals, as well as individuals living by themselves within multi-person households. All calculations made with sample weights provided by the vendor.

† Advertised credit limit is the maximum amount of credit available on the card.


‡ Calculated as the number of bankcard trades opened in the past 6 months divided number of inquiries in the past 6 months, excluding auto and mortgage inquiries.

Table 8: Poisson Results for Number of Credit Card Offers

	Coefficient	(Std. Error)
<i>Female</i>	-0.042	(0.059)
(Household) Income Category		
<i>Female</i> × <i>I</i> (30 – 74)	-0.108***	(0.030)
<i>Female</i> × <i>I</i> (75 – 150)	-0.185***	(0.030)
<i>Female</i> × <i>I</i> (> 150)	-0.186***	(0.038)
Credit Score Category		
<i>Female</i> × <i>I</i> (550 – 599)	0.168**	(0.066)
<i>Female</i> × <i>I</i> (600 – 649)	0.081	(0.061)
<i>Female</i> × <i>I</i> (650 – 699)	0.058	(0.059)
<i>Female</i> × <i>I</i> (700 – 749)	0.011	(0.057)
<i>Female</i> × <i>I</i> (750 – 799)	0.022	(0.056)
<i>Female</i> × <i>I</i> (800 – 850)	0.048	(0.054)
Year FE		Y
State FE		Y
<i>N</i>		46,428

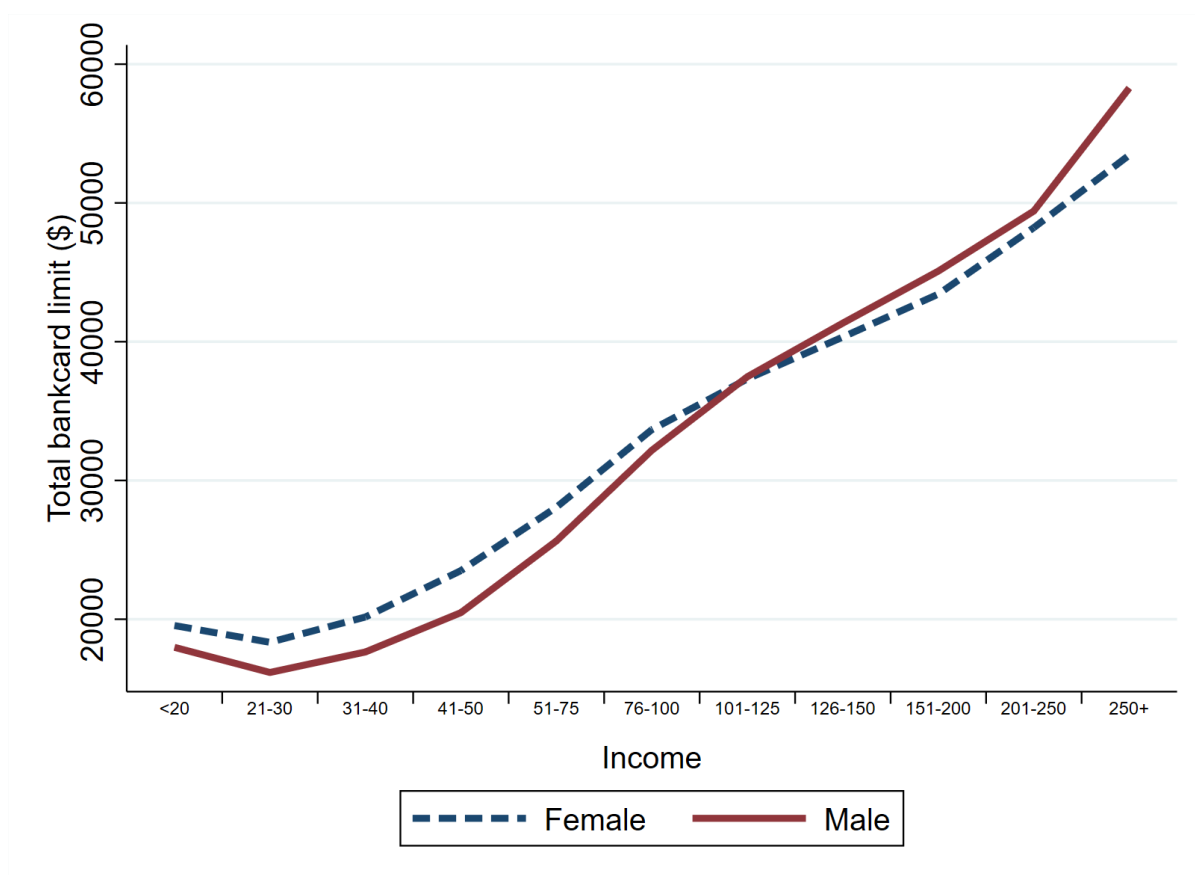
Notes: Authors' calculations using data from Mintel/TransUnion. Demographic information comes from the Mintel data. Sample includes individuals who had an active mortgage tradeline from 2009 - 2017 and received at least one credit card promotional mailer. We drop any observations where we are unable to determine gender for all cases (e.g., members of a multi-person household where we are unable to determine the gender of the individual who received the mailer). This translates to dropping 28.95 percent of the full sample. Resulting data set includes both single and married individuals, as well as individuals living by themselves within multi-person households. All calculations made with sample weights provided by the vendor. Credit score is a transformed version of VantageScore 2.0. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Figure 1: Relationship Between Credit Score and Total Bankcard Limit by Gender



Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit score is the Equifax Risk Score. Demographic information comes from the HMDA data.


Figure 2: Relationship Between Income and Total Bankcard Limit by Gender




Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Income is the HMDA income, reported at the time of mortgage application. Demographic information comes from the HMDA data.

Figure 3: Bankcard Differences over Time by Gender

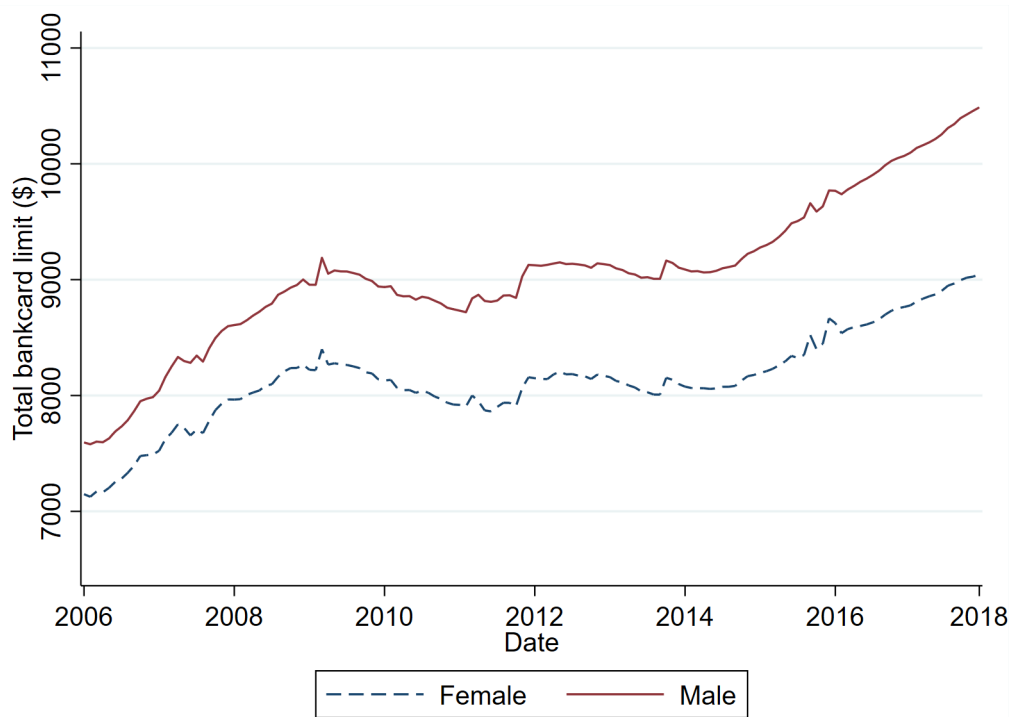
Panel A. Number of Bankcard Accounts



Panel B. Total Bankcard Credit Limit




Panel C. Average Bankcard Credit Limit



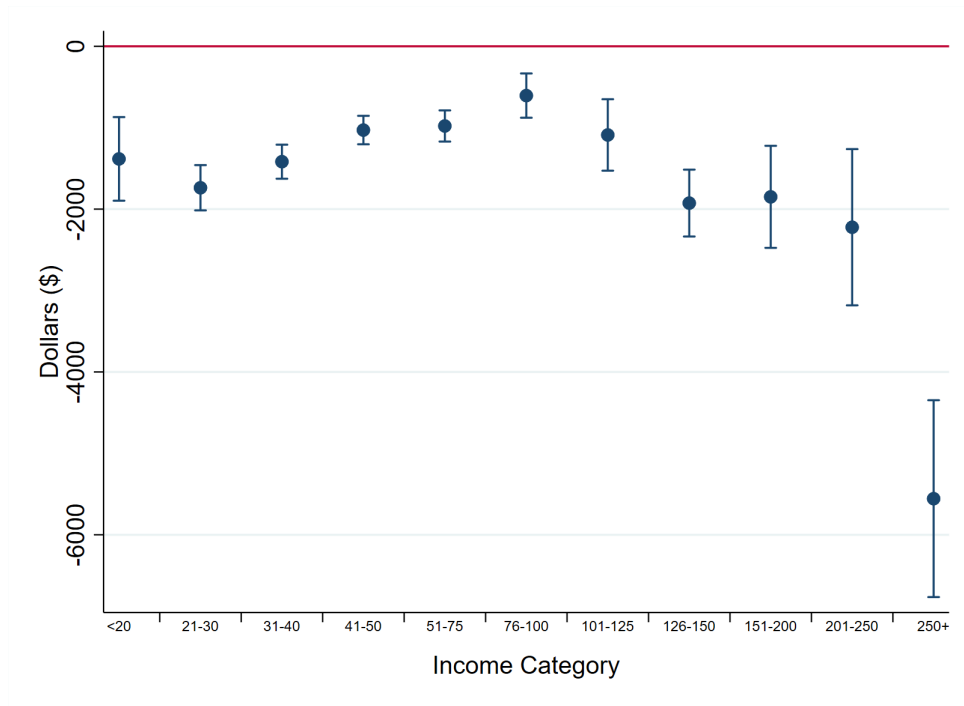
Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Demographic information comes from the HMDA data.

Figure 4: Average Marginal Effects for Credit Score and Income, by Category

Panel A: Credit Score Category



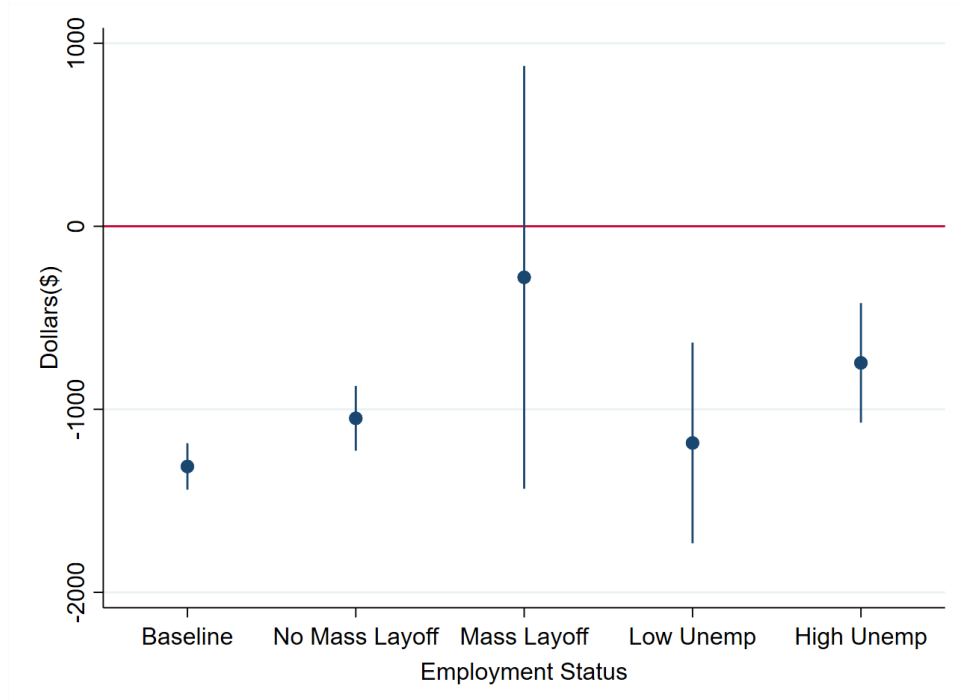
Panel B: Income Category




Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit Score is the Equifax Risk Score. Income is the HMDA income, reported at the time of mortgage application. Demographic information comes from the HMDA data. Bars represent 95% confidence intervals.

Figure 5: Average Marginal Effects by County Employment Status

Panel A: Employment Status




Panel B: Mass Layoff Status




Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Demographic information comes from the HMDA data. Results for mass layoff analyses use data from 2006 to 2011, while results for the baseline and unemployment rate analyses use data from 2006 to 2016.

Figure 6: Heat Maps of Gender Differences (Percent) Across Time by Decile

Panel A. Difference, Measured in Dollars (\$)




Panel B. Difference, as a Percentage of the Male Total Bankcard Limit



Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Demographic information comes from the HMDA data. Reported Z-axis values are the midpoint of each bin.

Figure 7: Relationship Between Credit Score and Bankcard Mail Offers Received by Gender



Notes: Authors' calculations using Mintel/TransUnion data with sample weights provided by the vendor. Demographic information comes from the Mintel data. Includes years 2009-2017. Credit score is transformed version of VantageScore 2.0.

ONLINE APPENDIX

A Selected Features of U.S. Adult Population, 2006-2017

We present summary statistics of the U.S. adult population and adults with a mortgage in Table A1. During our study period, adults age 18 years or older in the U.S. were majority females (51.38 to 48.62) and more than half of U.S. adults were married (55.11). With respect to race and ethnicity, 63 percent of U.S. adults were White, non-Hispanic, 12 percent were Black, non-Hispanic, and 5 percent were Asian; 15 percent were Hispanic. Median income for all adults for this period was \$35,418; median income for female adults was \$17,100 while male median income was \$30,000. The median age among adults was 46, and nearly 40 percent owned their home and had a mortgage (39.4 percent).

However, the race distribution changes when considering the subset of adults with a mortgage. In this group, the vast majority were White, non-Hispanic adults (72.37). Black, non-Hispanic adults and Hispanic adults each represent less than 10 percent of the set of adults with a mortgage (8.5 percent and 9.4 percent, respectively) and during our study time frame, 5 percent of U.S. adults with a mortgage were Asian adults.

B Differences Between Mortgage Borrowers and the Credit Bureau Population (Equifax Data)

In Appendix Table A2, we present summary statistics from our HMC data and from the CCP. We see that individuals with mortgages have greater Risk Scores than the average credit bureau individual (725.86 to 696.89) and that mortgage holders have more bankcard accounts than the average individual (2.77 to 1.87). We also observe that more credit has been extended to mortgage holders than non-mortgage consumers: Mortgage owners have larger bankcard credit limits, on average, than the average consumer (\$31,069 to \$22,404). Interestingly, mortgage owners apply for credit more than the average credit consumer, averaging 2.25 credit inquiries in the past 12 months, compared to the credit bureau average of 1.88. Overall, the summary statistics in Table A2 suggest that mortgage owners have lower credit risk and borrow more than the general population. Broadly speaking, this implies that any results we derive from the HMC sample would have an upward bias relative to the U.S. credit bureau population. While this prevents us from generalizing our result to the entire borrowing population, we argue that the mortgage-holding population is an important subgroup of interest, as approximately 80 million Americans held mortgage debt of \$8.9 trillion by the end of 2017.³⁸

³⁸See Federal Reserve Bank of New York, “Quarterly Report on Household Debt and Credit (Q4:2017).”

C Differences Between Sole Mortgage Applicants and Dual Applicants

As mentioned in the previous subsection, in addition to owning a mortgage, individuals in our sample are sole applicant mortgage owners. Sole applicants make up approximately 39 percent of the set of mortgage holders. While sole mortgage applicants are similar to dual applicant mortgage owners on many dimensions, there are also important differences between the two groups. Appendix Table A3 provides a comparison between these two types of applicants.

Individuals in our sole applicant sample are younger, on average, than dual applicant mortgage holders, as the mean birth year of sole applicants is 1964, as compared to 1960 for dual applicants. In terms of credit market traits, sole applicants have greater credit seeking behavior, as illustrated by a greater volume of inquiries compared to dual applicant mortgage individuals (2.50 to 2.07). Additionally, sole mortgage applicants have lower credit Risk Scores, on average, than dual applicant mortgage holders (707.04 to 738.51). Not surprisingly, dual applicant mortgage holders also have greater credit limits compared to individuals who have mortgages without a co-applicant (\$33,031 to \$28,029, respectively). Overall, we observe differences between sole and dual mortgage applicants, with the summary statistics indicating that our sample of sole applicant mortgage borrowers is less creditworthy and has less total bankcard debt than individual mortgage holders who have a co-applicant.

C.1 U.S. Marital Property Laws

There are two systems of marital property law in the United States: common law and community property law. Property and debts under common law regimes can be owned by one partner or both; ownership of property (and debts) can be separated. Community property law in the U.S. is derived from the Spanish legal system. Therefore, the traditional community property law states are: Arizona, California, Idaho, Louisiana, Nevada, New Mexico, Texas, Washington, and Wisconsin.³⁹ In community property states, spouses are considered to be one economic unit; that is, they have joint ownership over assets and liabilities.

With respect to mortgages in community property law states, if a couple buys a home after they are married, each spouse automatically owns half of it, regardless if there are one or two buyers on the mortgage application. Married home buyers in community property law states will have their spouse's debt added to their debt-to-income ratio calculation, regardless if they are a co-buyer on the mortgage application or not. Likewise, tax liens, judgments, and collections will be taken into account, but generally not the spouse's credit score. To summarize, even though both spouses own the property in community property law states, both individuals do not need to have their names on the mortgage application.

Due to multiple legal systems across our study area, we might expect slightly different rates of sole applicant mortgage takeup by married couples. However, the addition of the state fixed effect alleviates this concern.

³⁹Alaska and Tennessee have optional community property law where individuals can opt in to it.

D Summary Statistics over Time by Risk Score Category

It is possible that the aggregated results in Figure 3 may mask significant heterogeneity in gender differences along a number of dimensions. One of the dimensions where we could expect to see differences is by credit score. Although males and females have similar credit score distributions, as shown in Table 1, it is not clear if this would translate to men and women having similar bankcard limits across time across credit score categories.

To examine if gender differences in bankcard limits differ by credit score, we first calculate each gender's average in each credit score bin in each month of our data, and then take the difference between the two. Appendix Figure A2 tracks these gender differences by credit score bin over time. After splitting our sample by credit score bin, new patterns emerge as consumers at opposite ends of the Risk Score distribution display substantial differences. In panel A of Figure A2, we observe that individuals in the lowest Risk Score bucket have the smallest average difference in the number of bankcard accounts, while individuals in the second-highest Risk Score bucket have the greatest difference. For all Risk Score buckets, the difference grows more negative throughout the study period, indicating that women not only hold more bankcards than men at each Risk Score levels, but that this difference has increased in favor of female borrowers over time.

Panels B and C of Appendix Figure A2 show the gender differences in total and average bankcard credit limits by credit score category. In both of these figures, the largest gender difference between males and females occurs at the highest credit score group (super prime scores of 800 and above), and this difference grows over time. For total bankcard limits, the gender difference among super prime consumers increases from \$2,600 in January 2006 to over \$4,000 by December 2017. For this group, the difference in average bankcard limits almost doubles over the same period. In terms of economic significance, male borrowers went from having 6.7 percent higher limits than females in 2006 to having 10 percent higher limits in 2017.

By the end of the sample, this gender difference for the highest credit score bin is also significantly larger than the next largest difference. In December 2017, the difference between genders for total bankcard limits at the highest credit score category is over two times larger than the next largest difference (\$4,200 to \$1,500); for the average bankcard limit, the difference is approximately 1.5 times as large (\$2,200 to \$1,300). The heterogeneity in trends is evident when looking at gender differences for the lower credit score categories, where differences in total bankcard limits have fallen since the Great Recession and differences in average bankcard limit have increased only modestly.

It is also worth noting that the magnitudes of the gender differences do not monotonically increase with the credit score category. In panel B of Figure A2, the group of individuals with the lowest credit scores exhibits the second-highest gender difference in total bankcard credit limits for a majority of the sample period. For average limits, the second highest gender difference consists of consumers in the second-highest Risk Score bucket. This difference is likely driven by the difference in the average number of bankcard accounts shown in panel A of Figure 3.

E Kitagawa-Oaxaca-Blinder Decomposition Methodology

To assess if observed characteristics are the main driver of the observed disparities in total bankcard limits, we implement a variant of the standard Kitagawa-Oaxaca-Blinder (KOB) decomposition method.⁴⁰ This methodology, introduced by Kitagawa (1955) and popularized in economics by Blinder (1973) and Oaxaca (1973), is based on estimating separate equations for each of the two groups of interest to obtain mean predicted values $E(Y_A)$ and $E(Y_B)$. The gap in the variable of interest, $\hat{\Delta}$, is then calculated as difference in the mean of outcome Y between two groups: $\hat{\Delta} = E(Y_A) - E(Y_B)$. This difference is then algebraically decomposed into “explained” and “unexplained” components, which then can be used to demonstrate how much of a difference can or cannot be explained by observable characteristics. Our analysis uses a variation of the method, first introduced in the sociology literature, that decomposes the group difference into three components, commonly referred to as a “threefold decomposition” (Winsborough and Dickinson, 1971; Jones and Kelley, 1984; Jann, 2008; Fortin et al., 2011).

In the threefold decomposition case, if we assume that estimating equations for each group are $Y_A = X_A\beta_A + \epsilon_A$ and $Y_B = X_B\beta_B + \epsilon_B$, the outcome gap $\hat{\Delta}$ can be represented as:

$$\hat{\Delta} = [E(X_A) - E(X_B)]'\beta_B + E(X_B)'(\beta_A - \beta_B) + [E(X_A) - E(X_B)]'(\beta_A - \beta_B). \quad (5)$$

Following Jann (2008), we take each term from Equation (3) and specify the following equation:

$$\hat{\Delta} = E + C + I. \quad (6)$$

The first component of Equation (5), E , corresponds to

$$E = [E(X_A) - E(X_B)]'\beta_B.$$

This is often called the “endowment effect,” and it represents the portion of the gap that is generated by group (gender) differences in the explanatory variables. E measures the expected change in group B’s mean outcome if group B had group A’s levels of explanatory variables.

The next segment of Equation (5) is

$$C = E(X_B)'(\beta_A - \beta_B),$$

which corresponds to the “coefficient effect.” This part of the gap is generated by differences in the two sets of coefficients for both the explanatory variables and the intercepts. C measures how much group B’s mean outcome would change if group B experienced the returns to endowments (coefficients) that group A did. By allowing group B to take on group A’s endowments and coefficients, this part of the decomposition effectively creates a counterfactual to compare with the original two groups.

Finally, the last component,

$$I = [E(X_A) - E(X_B)]'(\beta_A - \beta_B),$$

represents the “interaction effect.” This effect exists because a portion of the differences in explanatory variables and beta coefficients occurs simultaneously. In other words, the interaction effect measures how much the mean outcome for group B would change if group B had the same

⁴⁰Frequently used in labor applications, the Kitagawa-Oaxaca-Blinder decomposition is a multivariate regression analysis technique that estimates counterfactual outcomes of variables of interest.

endowments as group A and if the additional amount of the endowments for group B had the same coefficients as group A.⁴¹

In our setting, female and male mortgage holders make up the two mutually exclusive groups, A and B, our estimating equations $Y_A = X_A\beta_A + \epsilon_A$ and $Y_B = X_B\beta_B + \epsilon_B$ are credit limit equations for each gender, and $\hat{\Delta}$ is the gender credit gap for total bankcard limits.

⁴¹For example, in the gender wage gap literature, the interaction effect would be interpreted as the additional wage that women would earn if they worked the same hours as men *and* if those additional hours were paid at the male wage (Jones and Kelley, 1984).

F Heterogeneity in the Decomposition

Appendix Figure A4 displays the endowment effects of our decompositions, which identifies the part of the gender gap due to differences in observable characteristics. In panel A, which shows the magnitude of the endowment effect, we observe significant heterogeneity both across deciles and across time. We first note that for the first five years of our sample, the endowment effect is positive across the entirety of the credit limit distribution, implying that portions of the gender gap were driven by male borrowers having an advantage in the values of their observed characteristics. For smaller limits (at or below the 20th percentile, or approximately \$5,500 to \$8,000), the endowment effect was between \$0 and \$400, while at higher limits, it ranged from \$1,000 to over \$2,000.

Starting in 2010, we see that the endowment effect decreases and by 2011 becomes *negative* for all deciles. A negative endowment effect implies that women have an advantage in the levels of their endowments relative to men.⁴² In other words, we would expect the gender gap to be smaller (e.g., less in favor of men or more in favor of women) if female borrowers had the same observed characteristics as male borrowers. This effect is largest in the right tail of the distribution, where the endowment effect ranges from -\$1,300 to -\$1,700 starting at the 70th percentile. Although the negative endowment effects shrink in magnitude in 2014, they remain negative for most of the distribution. The exception is for the top two deciles, where the endowment effects have favored male borrowers in the final years of our sample.

Panel B of Figure A4 shows the endowment effect as a percentage of the average male total bankcard limit. We can see that from 2006 to 2010, the magnitude of the endowment effect in panel A ranges from 0 percent to 6 percent of the average male total bankcard limit, with the endowment effect being relatively larger in economic terms in the left tail of the distribution. When the endowment effect becomes negative in 2011, the size of the endowment effect ranges from -1.5 percent to -4 percent of the male average total limit and is consistently in that range until the end of the sample in 2016.

Appendix Figure A5 documents how the coefficient effect, which explains the part of the gender gap that is due to differences in the returns on observable characteristics of borrowers, differs across deciles and over time. In panel A of Figure A5, we can see that for almost all years in our sample, the coefficient effect is negative for the deciles below the median of the total limit distribution. This implies that female borrowers would have lower bankcard limits than male borrowers if they had male coefficients. In other words, in this part of the distribution, female borrowers receive higher returns on their observed characteristics than male borrowers. For the deciles above the median of the limit distribution, we can see that the coefficient effect is growing positive over time, going from negative values from the 50th to 70th percentiles in 2006, to positive values by 2009. For the 70th to 90th deciles, the coefficient effect is consistently positive for almost all years in our sample, with large values at the highest decile. This indicates that total bankcard limits for women would be higher if they had male coefficients.

We also see that the coefficient effect for deciles below the median have been relatively stable over time. However, the effect has become less negative starting in 2011. For the deciles above the median, the coefficient effect has become more positive over time; at the 70th percentile of the total bankcard limit distribution, the coefficient effect was between -\$200 and \$0 in 2006 and increases to \$600-\$800 by 2016. At the highest bankcard limits in the 90th decile, the coefficient effect has consistently been over \$3,000 since 2010, indicating that male borrowers receive very high returns on their observed characteristics relative to female borrowers. In panel B of Figure A5, we present the coefficient effect as a percentage of the male total bankcard limit. When compared to

⁴²While uncommon, negative values for the parts of the KOB decomposition are possible. We provide a stylized example of a situation graphically in Appendix Figure A6.

the endowment effect, the size of the coefficient effect is larger in comparison. For example, in the latter years of our sample, the coefficient effect at the 80th percentile ranges from \$1,800 to \$3,200, which translates to 3 percent to 6 percent of the male total bankcard limit.

Overall, we observe that gender differences in total bankcard limit vary across the bankcard limit distribution and across time. In the left tail of the total limit distribution, the gender gap favors female sole applicants, where differences range from \$0 to \$300. In the right tail of the total limit distribution, the gender gap favors male sole applicants, with differences over \$3,400 at the 90th decile, which represent up to 7.5 percent of the male total bankcard limit at each decile and year. The decomposition analyses reveal that the factors driving the gender gap in total bankcard limits vary over time and by decile of the limit variable. In particular, we note that the large coefficient effect we documented in Table 5 is driven completely by the upper tail of the credit limit distribution. Our results also show that the relatively large changes in both the gender gap and in the endowment effect occur in 2010 and 2011, which was the period immediately after the implementation of the CARD Act, which was a comprehensive reform of the credit card industry.⁴³

⁴³For a more detailed discussion of the CARD Act, see Agarwal, Chomsisengphet, Mahoney, and Stroebel (2015).

Table A1: Census Summary Statistics: Overall U.S. Adult Population

Characteristic	Percent
Sex	
Male	48.62
Female	51.38
Race	
White, non-Hispanic	63.22
Black, non-Hispanic	12.23
Asian, non-Hispanic	5.22
Hispanic or Latino	14.54
Marital Status	
Married	51.11
Education	
High school or less	38.89
Some college/associate degree	28.67
Bachelor's degree	18.25
Graduate or professional degree	10.95
Owns Home and Has Mortgage	39.36
Race and Ethnicity Among Adults 18+ with a Mortgage	
Race/Ethnicity	Percent
White, non-Hispanic	72.37
Black, non-Hispanic	8.52
Asian, non-Hispanic	4.34
Hispanic or Latino	9.41
Other	5.36
Median Age (in years)	46
Income, Adults 18+	
Mean	\$36,418
Median	\$23,000
25th Percentile	\$8,500
75th Percentile	\$46,100

Notes: Authors' calculations using IPUMS USA data (University of Minnesota) from 2006 to 2017. Statistics are for adults 18 years and older, except education estimates (25 years and older).

Table A2: Summary Statistics: HMC Sample vs. CCP Sample

Variable	Mean	Median	Std. Dev	Obs.
Birth year: HMC	1962	1963	13	168,885,478
Birth year: CCP	1961	1963	19	2,434,769
Credit score: HMC	725.86	754	94.90	167,279,935
Credit score: CCP	696.89	718	106.21	2,176,960
Number of credit cards: HMC	2.77	2	2.31	167,549,933
Number of credit cards: CCP	1.87	1	2.11	2,194,070
Total credit card limit: HMC	31,069.39	22,547	30,537.59	147,082,985
Total credit card limit: CCP	22,404.07	13,700	30,822.11	1,509,530
Number of inquiries, past 12 months: HMC	2.25	2	2.45	138,061,352
Number of inquiries, past 12 months: CCP	1.884	1	2.291	1,492,699

Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, Equifax Credit Risk Insight Servicing data, and FRBNY Consumer Credit Panel/Equifax data. Credit Score is the Equifax Risk Score. Demographic information comes from the HMDA data.

Table A3: HMC Summary Statistics: Sole Applicants vs. Dual Applicants

Variable	Mean	Median	Std. Dev	Obs.
Birth year: total	1962	1963	13	168,885,478
Birth year: sole applicants	1964	1965	13	65,915,540
Birth year: co-applicants	1960	1961	13	99,382,295
Credit score: total	725.86	754	94.90	167,279,935
Credit score: sole applicants	707.04	731	100.34	65,301,769
Credit score: co-applicants	738.51	767	88.86	98,421,911
Number of credit cards: total	2.77	2	2.31	167,549,933
Number of credit cards: sole applicants	2.80	2	2.35	65,462,821
Number of credit cards: co-applicants	2.74	2	2.28	98,523,985
Total card limit: total	31,069	22,547	30,537	147,082,985
Total card limit: sole applicants	28,029	19,500	29,298	57,184,245
Total card limit: co-applicants	33,021	24,800	31,054	86,704,287
Inquiries, past 12 months: total	2.25	2	2.45	138,061,352
Inquiries, past 12 months: sole applicants	2.50	2	2.70	56,146,290
Inquiries, past 12 months: co-applicants	2.07	1	2.23	78,872,446

Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit Score is the Equifax Risk Score. Demographic information comes from the HMDA data.

Table A4: Robustness Checks: Exclude Fixed Effects

	(No State \times Year FEs)		(No Time FEs)	
	Coefficient	(Std. Error)	Coefficient	(Std. Error)
<i>Constant</i>	-30,256.78***	(1,028.25)	-29,489.44***	(1,149.50)
<i>Female</i>	4,518.63***	(938.31)	4547.90***	(935.80)
<i>Number of cards</i>	9,044.73***	(132.05)	9,072.55***	(129.18)
<i>Female \times Number of cards</i>	-818.46***	(38.02)	-830.93***	(39.25)
Credit Score Category				
<i>Female \times I(350 – 399)</i>	-1,229.04	(1,048.59)	-1,227.04	(1,055.14)
<i>Female \times I(400 – 449)</i>	-2,744.54***	(909.99)	-2,730.23***	(906.50)
<i>Female \times I(450 – 499)</i>	-2,766.75***	(905.67)	-2,756.94***	(900.90)
<i>Female \times I(500 – 549)</i>	-2,701.67***	(844.80)	-2,692.56***	(842.07)
<i>Female \times I(550 – 599)</i>	-3,022.59***	(929.40)	-2,986.67***	(924.17)
<i>Female \times I(600 – 649)</i>	-3,026.46***	(921.31)	-3,009.95***	(921.87)
<i>Female \times I(650 – 699)</i>	-2,919.57***	(910.77)	-2,887.81***	(905.86)
<i>Female \times I(700 – 749)</i>	-2,974.21***	(918.41)	-2,932.04***	(917.83)
<i>Female \times I(750 – 799)</i>	-3,056.18***	(882.38)	-3,006.41***	(875.25)
<i>Female \times I(800 – 850)</i>	-4,058.17***	(882.42)	-4,030.36***	(876.31)
Income Categories				
<i>Female \times I(25 – 49)</i>	-370.37	(323.30)	-383.85	(323.39)
<i>Female \times I(50 – 74)</i>	-40.37	(282.62)	-63.58	(285.21)
<i>Female \times I(75 – 99)</i>	358.19	(263.77)	326.57	(262.31)
<i>Female \times I(100 – 124)</i>	398.45	(284.69)	353.6	(287.76)
<i>Female \times I(125 – 149)</i>	783.25**	(344.75)	754.79**	(349.11)
<i>Female \times I(150 – 174)</i>	290.44	(391.44)	289.27	(398.24)
<i>Female \times I(175 – 199)</i>	-567.98	(353.49)	-560.71	(361.66)
<i>Female \times I(200 – 224)</i>	-506.73	(436.58)	-481.37	(437.48)
<i>Female \times I(225 – 249)</i>	-853.98	(567.62)	-806.80	(565.28)
<i>Female \times I(\geq 250)</i>	-4,216.94***	(619.46)	-4,162.39***	(627.21)
Year FE	Yes		No	
State FE	Yes		Yes	
State \times Year FE	No		No	
R^2	0.6305		0.6298	
N	530,122		530,122	

Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit Score is the Equifax Risk Score. Demographic information comes from the HMDA data. Standard errors are clustered at the state level. Reference group consists of White males, with the lowest credit score (280-349), and the lowest income (less than \$25,000). *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A5: Probability of Having a Delinquent Bankcard Account

	Point Estimate	(Std. Error)
<i>Constant</i>	-0.946***	(0.007)
<i>Female</i>	-0.009	(0.01)
<i>Number of cards</i>	0.006	(0.0002)
<i>Female × Number of cards</i>	-0.0004*	(0.0002)
Credit Score Category		
<i>Female × I(350 – 399)</i>	-0.0011	(0.0107)
<i>Female × I(400 – 449)</i>	-0.0115	(0.0127)
<i>Female × I(450 – 499)</i>	-0.008	(0.0136)
<i>Female × I(500 – 549)</i>	-0.0121	(0.0093)
<i>Female × I(550 – 599)</i>	-0.0054	(0.0105)
<i>Female × I(600 – 649)</i>	-0.0164*	(0.0094)
<i>Female × I(650 – 699)</i>	-0.0152	(0.0096)
<i>Female × I(700 – 749)</i>	-0.0111	(0.0091)
<i>Female × I(750 – 799)</i>	-0.01	(0.0092)
<i>Female × I(800 – 850)</i>	-0.01	(0.0092)
Income Categories		
<i>Female × I(21 – 30)</i>	-0.0011	(0.0041)
<i>Female × I(31 – 40)</i>	-0.0014	(0.0038)
<i>Female × I(41 – 50)</i>	-0.0029	(0.0034)
<i>Female × I(51 – 75)</i>	-0.0034	(0.0036)
<i>Female × I(76 – 100)</i>	-0.0039	(0.0033)
<i>Female × I(101 – 125)</i>	-0.0015	(0.0039)
<i>Female × I(126 – 150)</i>	0.0004	(0.0041)
<i>Female × I(151 – 200)</i>	0.0002	(0.004)
<i>Female × I(201 – 250)</i>	0.0005	(0.0051)
<i>Female × I(> 250)</i>	0.0015	(0.0052)
Year FE		Yes
State FE		Yes
State × Year FE		Yes
R^2	0.4251	
N	582,023	


Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit Score is the Equifax Risk Score. Demographic information comes from the HMDA data. Specification includes state, year, and state-by-year fixed effects. Standard errors are clustered at the state level. Reference group consists of White males, with the lowest credit score (280-349), and the lowest income (less than \$20,000). *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A6: Detailed Aggregate KOB Decomposition

Detailed Decomposition			
	Endowment Effects	Coefficient Effects	Interaction Effects
Number of cards	-1,334.30*** (51.55)	2,778.83*** (75.75)	-133.66*** (6.31)
Income	1,789.19*** (20.47)	-596.15*** (54.65)	169.16*** (20.23)
Risk Score	26.26 (21.90)	425.65*** (107.11)	-6.840*** (1.86)
Race	13.46** (6.56)	-60.44 (185.25)	14.86 (9.06)
Age	-266.09*** (8.82)	-1,182.79 (718.86)	-33.38*** (11.31)
Calendar year	5.144*** (3.08)	-39.87*** (11.09)	-0.504 (2.02)
State	-7.946** (3.61)	-105.02 (73.17)	10.13** (4.45)
<i>N</i>	530,125		

Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Risk Score is the Equifax Risk Score. Demographic information comes from the HMDA data. Results for the detailed decomposition results for categorical variables are based on "normalized" effects so our results do not rely upon our choice for the omitted category. For more information, see Jann (2008). We report the sum of coefficients for each variable category for tractability. For example, *Income* is the sum of all coefficients for each income bin in the regression Specification includes year fixed effects, race fixed effects, month-of-the-year fixed effects, state-fixed effects, and age. Standard errors are calculated using the standard Huber/White estimator and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.


Figure A1: Unconditional Gender Differences in Total Bankcard Limit by Year and Limit Decile



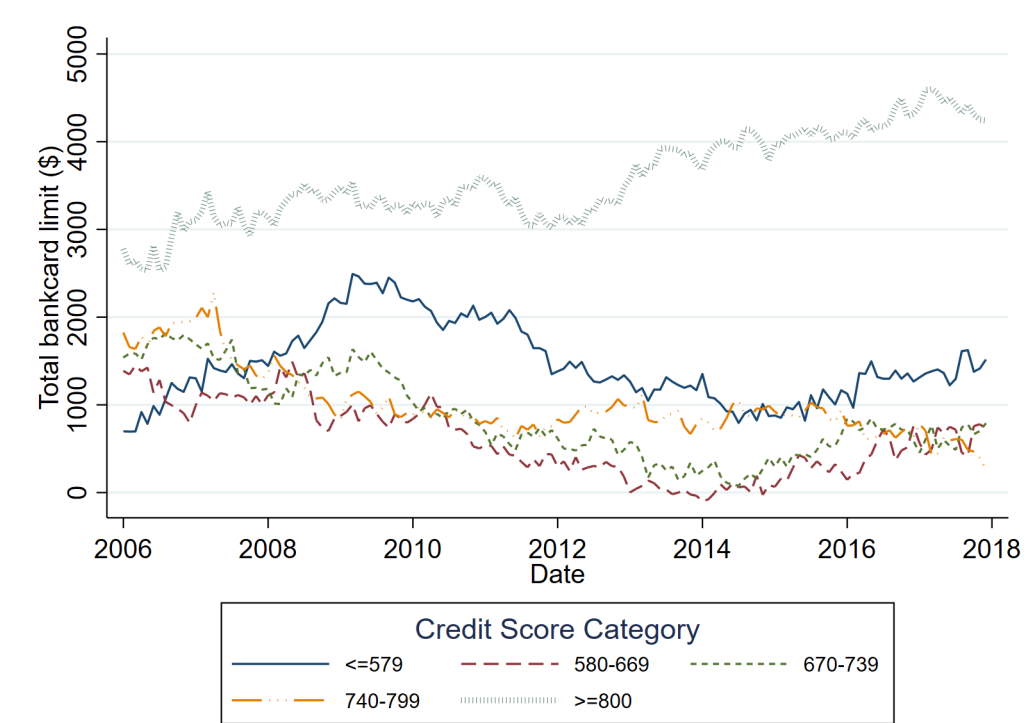
Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit score is the Equifax Risk Score. Demographic information comes from the HMDA data.

Figure A2: Bankcard Differences by Credit Score Category


Panel A. (Male-Female) Difference in the Number of Bankcard Accounts



Panel B. (Male-Female) Difference in Total Bankcard Credit Limit

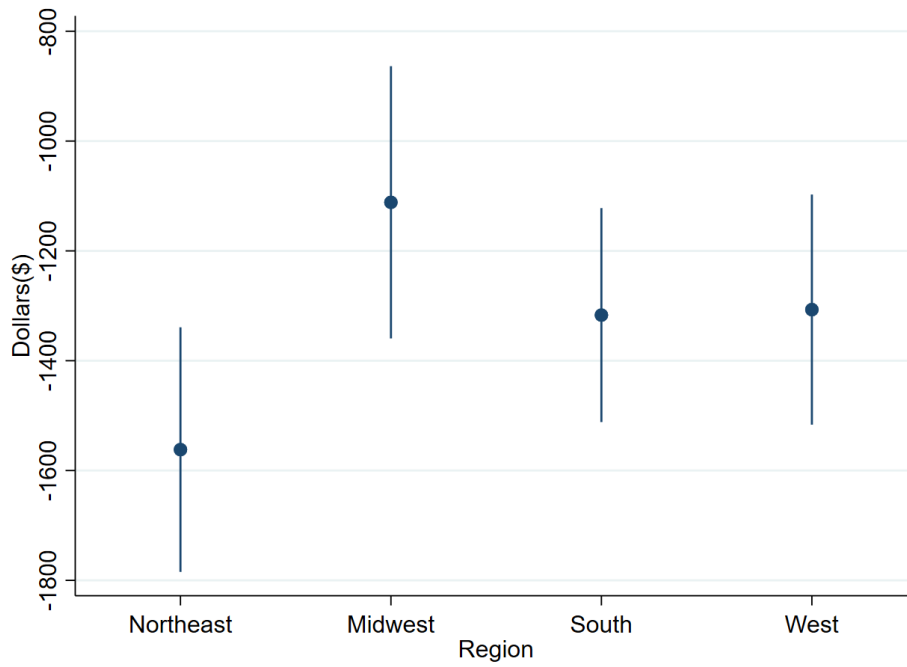


Panel C. (Male-Female) Difference in Average Bankcard Credit Limit



Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit Score is the Equifax Risk Score. Demographic information comes from the HMDA data.

Figure A3: Average Marginal Effects by Census Region





Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Credit score is the Equifax Risk Score. Demographic information comes from the HMDA data.

Figure A4: Heat Maps of the Endowment Effect Across Time by Decile

Panel A. Difference Due to the Endowment Effect, Measured in Dollars (\$)

Panel B. Difference, as a Percentage of the Male Total Bankcard Limit





Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Demographic information comes from the HMDA data. Reported Z-axis values are the midpoint of each bin.

Figure A5: Heat Maps of the Coefficient Effect Across Time by Decile

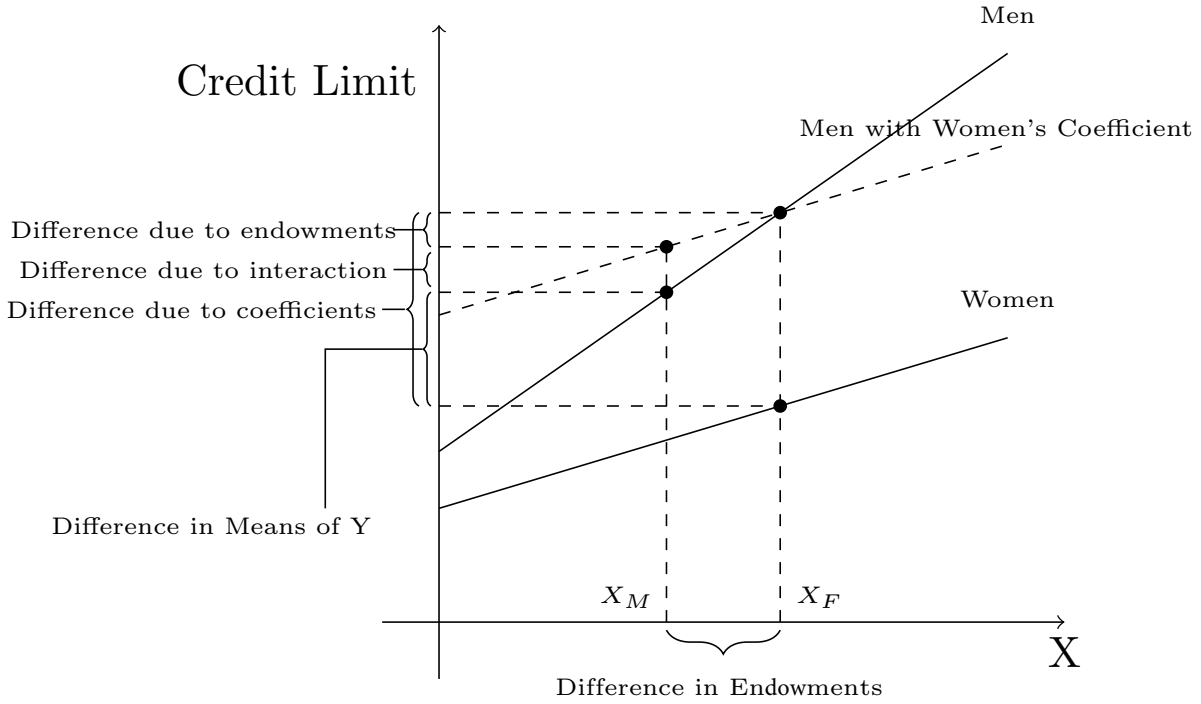
Panel A. Difference Due to the Coefficient Effect, Measured in Dollars (\$)

Panel B. Difference, as a Percentage of the Male Total Bankcard Limit



Notes: Authors' calculations using Home Mortgage Disclosure Act (HMDA) data, Black Knight McDash loan servicing data, and Equifax Credit Risk Insight Servicing data. Demographic information comes from the HMDA data. Reported Z-axis values are the midpoint of each bin.

Figure A6: Example KOB Decomposition



Notes: This example produces a KOB decomposition where the interaction and endowment effects are negative, but the coefficient effect is positive. To generate this result, female endowments are greater than male endowments, while men have greater average values for the Y variable than women. If the relationship for men is given by $Y_M = a_M + \beta_M X_M$ and the relationship for women is given by $Y_W = a_W + \beta_W X_W$, then the coefficient effect is $(\beta_M - \beta_W)X_M$, the endowment effect is given by $\beta_W(X_M - X_W)$, and the interaction effect is given by $(\beta_M - \beta_W)(X_M - X_W)$. For a canonical example in this format, see Jones and Kelley (1984).