Learning in the Credit Card Market

Sumit Agarwal, John C. Driscoll, Xavier Gabaix, and David Laibson*

April 23, 2011

Abstract

Agents with more experience make better choices. We measure learning dynamics using a panel with four million monthly credit card statements. We study add-on fees, specifically cash advance, late payment, and overlimit fees. New credit card accounts generate fee payments of $15 per month. Through negative feedback – i.e. paying a fee – consumers learn to avoid triggering future fees. Paying a fee last month reduces the likelihood of paying a fee in the current month by about 40%. Controlling for account fixed effects, monthly fee payments fall by 75% during the first three years of account life. We find that learning is not monotonic. Knowledge effectively depreciates about 10% per *month*, implying that learning displays a strong recency effect. The speed of net learning is about twice as great for higher-income borrowers than it is for lower-income borrowers; the rate of knowledge depreciation, or forgetting, is about half as fast for high- relative to low-income borrowers. Middle-aged borrowers have the same advantageous learning dynamics relative to older borrowers. (JEL: D1, D4, D8, G2)

*Agarwal: Federal Reserve Bank of Chicago, sagarwal@frbchi.org. Driscoll: Federal Reserve Board, john.c.driscoll@frb.gov. Gabaix: New York University and NBER, xgabaix@stern.nyu.edu. Laibson: Harvard University and NBER, dlaibson@harvard.edu. Gabaix and Laibson acknowledge support from the National Science Foundation (DMS-0938185). Laibson acknowledges financial support from the National Institute on Aging (R01-AG-1665). The views expressed in this paper are those of the authors and do not represent the policies or positions of the Board of Governors of the Federal Reserve System or the Federal Reserve Bank of Chicago. Ian Dew-Becker, Keith Ericson, Mike Levere, Tom Mason and Thomas Spiller provided outstanding research assistance. The authors are grateful to the editor and the referees, and to Murray Carbonneau, Stefano Dellavigna, Joanne Maselli, Nicola Persico, Devin Pope, Matthew Rabin, and seminar participants at the AEA, Berkeley, EUI and the NBER (Law and Economics) for their suggestions. This paper previously circulated under the title "Stimulus and Response: The Path from Naiveté to Sophistication in the Credit Card Market."

1 Introduction

Economists often motivate optimization and equilibrium as the outcome of learning. Learning is a key mechanism that underpins economic theories of rational behavior. Accordingly, many economic studies have analyzed learning in the lab,[1] and in the field.[2]

Because of data limitations, only a few papers measure learning with household-level panel data. Such household studies, usually find that households learn to optimize over time. For example, Miravete (2003) and Agarwal, Chomsisengphet, Liu and Souleles (2006) respectively show that consumers switch telephone calling plans and credit card contracts to minimize monthly bill payments. A few papers are able to identify the specific information flows that elicit learning. For instance, Fishman and Pope (2006) study video stores, and find that renters are more likely to return their videos on time if they have recently been fined for returning them late. Ho and Chong (2003) use grocery store scanner data to estimate a model in which consumers learn about product attributes. Their learning model has greater predictive power, with fewer parameters, than forecasting models used by retailers.[3]

In the current paper, we study individual households that learn to avoid add-on fees in the credit card market.[4] We analyze a panel dataset that contains three years of credit card statements, representing 120,000 consumers and 4,000,000 credit card statements. We focus our analysis on credit card fees — late payment, over limit, and cash advance fees. Some observers argue that account holders do not optimally minimize such fees.[5] We want to know whether credit card holders change the way they use their credit cards – e.g., paying

[1]For example, Van Huyck, Cook and Battalio (1994), Crawford (1995), Roth and Erev (1995), Camerer (2003), and Wixted (2004).

[2]For example, see Bahk and Gort (1993), Marimon and Sunder (1994), Thornton and Thompson (2001).

[3]Lemieux and MacLeod (2000) study the effect of an increase in unemployment benefits in Canada. They find that the propensity to collect unemployment benefits increases as a consequence of a previous unemployment spell. Odean, Strahlevitz and Barber (2010) find evidence that individual investors tend to repurchase stocks that they previously sold for a gain. Dellavigna (2009) surveys the field evidence on behavioral phenomena.

[4]Ausubel (1991, 1999), Shui and Ausubel (2004) and Kerr and Dunn (2008) analyze the magnitude of interest payments and fees in the credit card market.

[5]For example, Frontline reports that "The new billions in revenue reflect an age-old habit of human behavior: Most people never anticipate they will pay late, so they do not shop around for better late fees." (http://www.pbs.org/wgbh/pages/frontline/shows/credit/more/rise.html) There is also a nascent academic literature that studies how perfectly rational firms interact in equilibrium with imperfectly rational consumers. See Shui and Ausubel (2004), DellaVigna and Malmendier (2004), Mullainathan and Shleifer (2005), Oster and Morton (2005), Gabaix and Laibson (2006), Jin and Leslie (2003), Koszegi and Rabin (2006), Malmendier and Shanthikumar (2007), Grubb (2009), and Bertrand et al. (2010). See Spiegler (2011) for an overview.

fewer fees – as they gain experience.

We find that fee payments are very large in the first few months after the opening of an account. We find that new accounts generate direct monthly fee payments (*not* including interest payments) that average \$15 *per month*.[6] However, these payments fall by 75 percent during the first four years of account life.

These learning effects may be driven by many different channels. Consumers learn more about the existence and magnitude of fees when they knowingly or accidentally trigger them. Painful fee payments may also train account holders to be more vigilant in their card usage. As a result of many different learning pathways, card holders sharply cut their fee payments over time.

We find that the learning dynamics are not monotonic. Card holders act as if their knowledge depreciates – i.e., learning patterns exhibit a recency effect.[7] A late payment charge from the *previous* month engenders vigilant fee avoidance this month, and this response is much stronger than the vigilance engendered by a late payment charge that was paid further back in time. We estimate that the learning effect of a fee payment effectively depreciates at a rate of between 10 and 20 percent per month. At first glance, such depreciation may seem counter-intuitive. However, if attention is a scarce resource, attention may wander as the salience of a past fee payment fades. After making any significant mistake (e.g., getting a speeding ticket), people are likely to pay attention and avoid the mistake; however as the key event recedes into history, vigilance fades.

There are several papers that have also documented forgetting effects, though the settings of these papers are quite different from our credit card application. For instance, Benkard (2000) finds evidence for both learning and forgetting – that is, depreciation of productivity over time – in the manufacturing of aircraft, as do Argote, Beckman and Epple (1990), in shipbuilding.

We analyze the mechanisms that may explain the fee dynamics that we measure. We

[6]Moreover, this understates the impact of fees, since some behavior — e.g. a pair of late payments — not only triggers direct fees but also triggers an interest rate increase, which is *not* captured in our \$15 calculation. Suppose that a consumer is carrying \$2,000 of debt. Changing the consumer's interest rate from 10% to 20% is equivalent to charging the consumer an extra \$200. Late payments also may prompt a report to the credit bureau, adversely affecting the card holder's credit accessability and creditworthness. The average consumer has 4.8 cards and 2.7 actively used cards.

[7]See Lehrer (1988), Aumann, Hart and Perry (1997) and Besanko et al. (2010) for some theoretical models of forgetfulness.

first explore several explanations that are not consistent with our preferred explanation of learning/forgetting–for example, that card usage might be negatively autocorrelated–and find that these explanations are not consistent with the data. On the other hand, we find support for mechanisms that support the learning/forgetting interpretation. Notably, we find that in the month after paying a late fee, account holders are especially likely to make their next payment more than two weeks before the due date. This suggests that a late payment fee acts as a wake up call that induces earlier fee payment. We also find that the speed of (i) net learning, (ii) the magnitude of the recency effect, and (iii) the speed of forgetting all differ across borrower characteristics. Higher-income borrowers learn more than twice as fast, have a recency effect double the size, and forget about three-times as slowly as lower-income borrowers. Likewise, middle-aged borrowers have similar learning advantages relative to older borrowers.

In summary, our findings imply that a high rate of knowledge depreciation offsets learning. Nevertheless, learning dominates knowledge depreciation. On average, fees fall over the life of the credit card. These learning dynamics are most advantageous for high-income and middle-aged borrowers.

We organize our paper as follows, Section 2 summarizes our data and presents our basic evidence for learning and backsliding/forgetting. Section 3 analyzes various alternative (non-learning) explanations for our findings. Section 4 discusses extensions to our analysis on learning and forgetting, including results on the demographics of learning. In Section 5, we draw some conclusions.

2 Two Patterns in Fee Payment

In this section, we describe the data. We then show that fee payments decline sharply with account tenure. We also show that the learning dynamics exhibit a recency effect: a late payment charge from the *previous* month is strongly associated with fee avoidance this month, and this elasticity sharply declines as the time gap increases between the previous fee payment and the current period.

2.1 Data

We use a proprietary panel dataset from a large U.S. bank that issues credit cards nationally. The dataset contains a representative random sample of about 128,000 credit card accounts followed monthly over a 36 month period (from January 2002 through December 2004). The bulk of the data consists of the main billing information listed on each account's monthly statement, including previous payment, purchases, credit limit, balance, debt, amount due, purchase APR, cash advance APR, date of previous payment, and fees incurred. At a quarterly frequency, we observe each customer's credit bureau rating (FICO score) and a proprietary (internal) credit 'behavior' score. We have credit bureau data for the number of other credit cards held by the account holder, total credit card balances, and mortgage balances. We have data on the age, gender and income of the account holder, collected at the time of account opening. Further details on the data, including summary statistics and variable definitions, are available in the appendix.

We focus on three important types of fees, described below: late fees, over limit fees, and cash advance fees.[8]

1. **Late Fee**: A *direct* late fee of $30 or $35 is assessed if the borrower makes a payment beyond the due date on the credit card statement. If the borrower is late by more than 60 days once, or by more than 30 days twice within a year, the bank may also impose *indirect* late fees by raising the APR to over 24 percent.[9] Such indirect fees are referred to as 'penalty pricing.' The bank may also choose to report late payments to credit bureaus, adversely affecting consumers' FICO scores. Our analysis measures *only* direct late fees (and therefore excludes consumer costs associated with penalty pricing).

2. **Over Limit Fee**: A direct over limit fee, also of $30 or $35, is assessed the first time

[8] Other types of fees include annual, balance transfer, foreign transactions, and pay by phone. All of these fees are relatively less important to both the bank and the borrower. Fewer issuers (the most notable exception being American Express) continue to charge annual fees, largely as a result of increased competition for new borrowers (Agarwal et al., 2006). The cards in our data do not have annual fees. A balance transfer fee of 2-3% of the amount transferred is assessed on borrowers who shift debt from one card to another. Since few consumers repeatedly transfer balances, borrower response to this fee will not allow us to study learning about fee payment. The foreign transaction fees and pay by phone fees together comprise less than three percent of the total fees collected by banks.

[9] If the borrower does not make a late payment during the six months after the last late payment, the APR will revert to its normal (though not its promotional) level.

the borrower exceeds his or her credit limit in a given month. Penalty pricing also results from over limit transactions. As above, our analysis measures *only* direct over limit fees.

3. **Cash Advance Fee**: A direct cash advance fee of 3 percent of the amount advanced or $5 (whichever is greater) is levied for each cash advance on the credit card. Unlike the first two types of fees, a cash advance fee can be assessed many times per month. Cash advances do not invoke penalty pricing. However, the APR on cash advances is typically greater than that on purchases, and is usually 16 percent or more. Our analysis measures only direct cash advance fees (and not subsequent interest charges).

2.2 Fee payment by account tenure

Figure 1 reports the frequency of each fee type as a function of account tenure. The regression — like all those that follow — controls for time effects, account fixed effects, and time-varying attributes of the borrower (e.g., variables that capture card utilization each pay cycle). The data plotted in Figure 1 is generated by estimating,

$$
\begin{aligned}
(1) \qquad f^{j}_{i,t} &= \alpha + \phi_i + \psi_{time} + Spline(Tenure_{i,t}) \\
&\quad + \eta_1 Purchase_{i,t} + \eta_2 Active_{i,t} + \eta_3 BillExist_{i,t-1} \\
&\quad + \gamma_1 Util_{i,t-1} + \epsilon_{i,t}.
\end{aligned}
$$

The dependent variable $f^{j}_{i,t}$ is a dummy variable that takes the value 1 if a fee of type j is paid by account i at tenure t. Note that t indexes account tenure – not calendar time. When we refer to calendar time we use subscript $time$. Fee categories, j, include late payment fees — $f^{Late}_{i,t}$ — over limit fees — $f^{Over}_{i,t}$ — and cash advance fees — $f^{Advance}_{i,t}$. Parameter α is a constant; ϕ_i is an account fixed effect; ψ_{time} is a time fixed-effect; $Spline(Tenure_{i,t})$ is a spline[10] that takes account tenure (time since account was opened) as its argument; $Purchase_{i,t}$ is the total quantity of purchases in the current month; $Active_{i,t}$ is a dummy variable that re-

[10] The spline has knots every 12 months through month 72. We have also tried replacing the spline with dummies for each month of account tenure; the results are quantitatively similar. We report results with a spline as our baseline specification because the reduction in the number of parameters meaningfully decreases the computation time required to estimate the model; each regression has 4,000,000 observations of credit card statements.

flects the existence of any account activity in the current month; $BillExist_{i,t-1}$ is a dummy variable that reflects the existence of a bill with a non-zero balance in the previous balance; $Util_{i,t}$, for utilization, is debt divided by the credit limit; $\epsilon_{i,t}$ is an error term. Table 1 provides the regression results.

Figure 1 plots the expected frequency of fees as a function of account tenure, holding the other control variables fixed at their means.[11] This analysis shows that fee payments are fairly common when accounts are initially opened, but that the frequency of fee payments declines rapidly as account tenure increases. In the first four years of account tenure, the monthly frequency of cash advance fees drops from 57% of all accounts to 13% of all accounts. The frequency of late payment fees drops from 36% to 8%. Finally, the frequency of over limit fees drops from 17% to 5%.

To establish that the estimated pattern is robust to alternative specifications, we estimate several variants. We estimate equation 1 as a conditional logit; the results are qualitatively similar. We also repeat the analysis controlling for behavior and FICO scores, both lagged by three months to reflect the fact that they are only computed quarterly. The coefficients on the spline of tenure are almost unchanged. Finally, to eliminate the possibility that attrition from the sample has distorted the results, we have tried restricting the sample to only those account present for all 36 months; the results are little changed.

It is also of interest to know how the amount of fees paid vary by account tenure. For the late fee and over limit fees, the fee amount is constant, while for the cash advance fee, the fee amount can vary over time. Figure 2 reports the average value of each fee type as a function of account tenure, conditional on other factors that might affect fee payment. The data plotted in Figure 2 is generated by estimating,

$$\begin{aligned} (2) \qquad V^j_{i,,t} &= \alpha + \phi_i + \psi_{time} + Spline(Tenure_{i,t}) \\ &\quad + \eta_1 Purchase_{i,t} + \eta_2 Active_{i,t} + \eta_3 BillExist_{i,t-1} \\ &\quad + \gamma_1 Util_{i,t-1} + \epsilon_{i,t}. \end{aligned}$$

The dependent variable $V^j_{i,t}$ is the *value* of fees of type j paid by account i at tenure t. All other variables are as before. Table 2 reports the regression results.

[11] Tenure in all figures starts at month two since borrowers cannot, by definition, pay late fees or over limit fees in the first month their accounts are open.

Figure 2 shows that, when an account is opened, the card holder pays \$6.65 per month in cash advance fees, \$5.63 per month in late fees, and \$2.46 per month in over limit fees. These numbers understate the total cost incurred by fee payments, as these numbers do not include interest payments on the cash advances, the effects of penalty pricing (i.e., higher interest rates), or the adverse effects of higher credit scores on other credit card fee structures. Like Figure 1, Figure 2 shows that the average value of fee payments declines rapidly with account tenure.

Figures 1 and 2 imply that fee payments fall substantially with experience. We next turn to a second pattern in our data.

2.3 The impact of past fee payment on current fee payment

There are four reasons to expect fee payments of agent i to be correlated (positively or negatively) at two arbitrary dates t and $t+k$.

First, the (cross-sectional) type of the card holder (e.g. forgetful) may influence fee paying behavior. If person i pays a fee in period t then person i is more likely to be of the type that pays fees in general, implying that person i has a higher likelihood of paying a fee in period $t+k$ relative to other subjects in our sample. As long as the cross-sectional 'type' of the account holder is a fixed characteristic, this first source of intertemporal linkage could be modeled as a fixed effect. If the *true* fixed effect is added to the model, the residual variation is no longer correlated across time. However, as we discuss below, Nickell bias (1981) introduces a wrinkle when the fixed effect needs to be estimated.

Second, transitory shocks that persist over more than one month (for instance, an unemployment spell) may influence fee paying behavior, causing fees to be positively autocorrelated.

Third, transitory shocks that are negatively correlated across months (for instance, an annual summer vacation) will cause fees to be negatively autocorrelated.

Fourth, fee payments may engender learning, causing fees to be negatively autocorrelated.

One natural approach to estimating the force of these four effects would be to estimate an autoregressive model, in which current fee payment would be allowed to depend on lagged fee payment, controlling for account fixed effects and time-varying characteristics. However, Nickell (1981) showed that including fixed effects in dynamic panel data models causes the autoregressive coefficients to exhibit a bias of order $-1/T$, where T is the number of time series observations. We thus adopt two approaches to deal with this problem. The first is to calculate the following statistic:

$$
\begin{aligned}
L_{t,k} &\equiv \frac{E\left[f_t \mid f_{t-k}=1\right]}{E\left[f_t\right]} \\
(3) \qquad &= \frac{\text{Probability of paying a fee at tenure } t \text{ given the agent paid a fee } k \text{ periods ago}}{\text{Probability of paying a fee at tenure } t}
\end{aligned}
$$

without conditioning on the RHS variables from the previous subsection (most importantly, we do not use person fixed effects to calculate $L_{t,k}$). Specifically, $E\left[f_t\right]$ is just the average frequency of fee payments in period t.[12] Likewise, $E\left[f_t \mid f_{t-k}=1\right]$ is the average frequency of fee payments in period t among the account holders who paid a fee at time $t-k$.

Conditioning on this sparse information, a consumer who paid a fee k periods ago has a probability of paying a fee equal to the baseline probability, $E\left[f_t\right]$, multiplied by $L_{t,k}$. A value of 1 for $L_{t,k}$ indicates that having paid a fee k periods does not change the expected probability of paying a fee this period; a value less than one indicates lagged fee payment is associated with a reduction in the expected probability, and a value greater than one indicates lagged fee payment are associated with an increase in the expected probability. For example, if $L_{t,1}=0.6$, a consumer who paid a fee last month has a probability of paying a fee this month that is 40% below the baseline probability.

We report averages of $L_{t,k}$:

$$L_k \equiv \frac{1}{T}\sum_{t=1}^{T} L_{t,k}.$$

Hence, L_k is the average relative likelihood of paying a fee, if the account holder paid a fee k periods ago. The L_k statistic illustrates some important time series properties in our data while avoiding econometric problems associated with estimating probit or logit models with fixed effects for a large N dataset.

Figure 3 plots L_k for all three types of credit card fees for values of k ranging from 1 to 35. All three lines start below 1, indicating that a fee payment last month is associated with a less than average likelihood of making a fee payment this month. For both cash advance and late fees, having paid a fee one month ago is associated with a 40 percent reduction in the likelihood of paying a fee in the current month. For over limit fees, having paid a fee

[12] Observations used for the calculation of L_k are from subjects who are in our sample at *both* date t and date $t-k$.

last month is associated with a 50 percent reduction in the likelihood of paying a fee in the current month.

The L_k plots rise with k, indicating that as a given fee payment recedes into the past, the negative association between this fixed lagged fee payment and current fee payments is diminished. By the time one year has passed, the association between the lagged fee payment and current fee payment has almost vanished.

For large values of k, all three graphs rise above 1. This asymptotic property reflects the variation in fee payments that is driven by cross-sectional variation in the (persistent) type of the borrower. Some individuals have a relatively high long-run likelihood of paying fees. For instance, imagine that 20 percent of consumers never pay a fee (maybe because they are very disciplined), while the others have a long-run monthly probability $b > 0$ of paying a fee. Then, the long run L_k is $1/0.8 = 1.25$.[13]

2.3.1 Dynamic panel models with fixed effects

Our second approach to estimating the impact of past fee payments on current fee payments is to estimate dynamic panel data models with fixed effects for each fee, while using Monte Carlo simulations to bound the size of the bias. As noted above, Nickell (1981) analytically derived that the size of the bias for an AR(1) with fixed effects is of order $-1/T$. Since in our estimation, $T = 36$, the implied bias is approximately $-1/36 = -.028$. This amount is much too small to explain the large reduction in fee frequency this month associated with a fee payment last month. However, we are interested in the impact of fee payments at longer lags than one month on current fee payment; no analytical results exist for determining the size of the bias for higher-order autoregressive models. We thus use Monte Carlo simulations to determine the size of the bias in such cases.

Specifically, we first draw 10,000 times from a uniform (0,1) distribution; call each draw α_i, where $i = 1, ..., 10,000$. For each i, we then draw 36 times from a uniform (0,1), recoding the results as 1 if the draw is less than α_i and 0 otherwise. This algorithm simulates 36 i.i.d. draws from a binomial distribution with probability of success α_i (for 10,000 households). On these N =10,000 and $T = 36$ observations, we then estimate an AR(1), AR(12) and AR(18) with fixed effects. We repeat the whole process 5,000 times, and report mean and standard

[13] The numerator of (3) is b, while the denominator is $0.8b$. So $L_k = b/(0.8b) = 1/0.8$.

deviations. Note that, if there were no Nickell bias, all of the autoregressive coefficients should be zero.

Table 3 reports the mean values of those coefficients for all three cases. The AR(1) case shows a bias of about $-.028$, confirming the analytical results of Nickell (1981). The AR(12) and AR(18) cases show somewhat larger biases – for the first lag, about -0.06 and -0.10, respectively – that decline in absolute value with the lag.[14] Although this declining pattern does mimic in a qualitative fashion the recency effects we see above, the magnitude of the bias is again not nearly sufficient to explain the size of the recency effects that we measure.

We show this in Table 4, which estimates:

$$
\begin{aligned}
f_{i,t}^{j} &= \alpha + \phi_i + \psi_{time} + A(L) f_{i,t-1}^{j} \\
&\quad + \eta_1 Purchase_{i,t} + \eta_2 Active_{i,t} + \eta_3 BillExist_{i,t-1} \\
&\quad + \gamma_1 Util_{i,t-1} + \epsilon_{i,t}.
\end{aligned} \tag{4}
$$

where $A(L)$ is a twelfth-order lag polynomial for each of the three kinds of fees. As with the L_k measure, all three sets of results show a strong recency effect, far larger than can be explained by Nickell bias. The coefficients decrease in absolute value with lag, reaching zero by the eleventh or twelfth lag. For late fees, the coefficient on the first lag shows that having paid a fee one month ago is associated with nearly a 50 percent reduction in the frequency of current fee payments – or about a 45 percent reduction once the bias resulting from the Monte Carlo simulations is subtracted. For the over limit and cash advance fees, the reductions are slightly over 40 percent and 30 percent, respectively (again controlling for the bias). These large reductions are all within ten percentage points of those associated with the L_k approach estimated above.

We note that a third way of estimating the recency effect would be to use the instrumental variables approach of Arellano and Bond (1991) or the conditional logit estimators derived by Chamberlain (1993) and Honoré and Kyriazidou (2000), which allow for the presence of lagged endogenous variables. We do not follow these approaches because they all limit the extent to which the disturbance terms may be serially correlated We think it likely that agents do face autocorrelated shocks that may affect their likelihood of fee payment. Honoré and Kyriazidou (2000) also require that there not be time effects; we in turn think that agents

[14] We are grateful to Devin Pope for pointing out this pattern of bias in higher-order autoregressive models.

may face different shocks at different times–for example, due to changing macroeconomic conditions.

2.4 Summary

The data exhibit a robust time-series pattern. Paying a fee last month is associated with a sharply reduced likelihood of paying a fee this month (relative to other members of your account holding cohort). Paying a fee a year ago has little relationship to the likelihood of paying a fee now (relative to other members of your account holding cohort). Paying a fee two years ago is associated with a 20% elevation in the likelihood of paying a fee now (relative to other members of your account holding cohort).

Our findings imply that the there must be a mechanism that produces the short-run negative association (or recency effect). Moreover, this mechanism must be strong enough to temporarily overwhelm the positive long-run association in fee payments driven by type variation.[15]

3 Alternative Explanations

The patterns of fee payments that we document is explained by a model in which consumers learn to avoid fees by first experiencing them. The learning dynamics are complicated by partial backsliding or forgetting. In our view, this is like the effect of getting a speeding ticket; a driver may slow down for a few weeks, but will partially revert to type and speed again.[16] In our credit card analysis, the net effect of learning and backsliding appears to be positive, since fee payments do fall on average with tenure.

On the other hand, some or all of the data patterns that we observe may be explained by mechanisms other than learning and forgetting. In this section, we first discuss a few alternative explanations; we find that the available evidence does not support these alternatives.

[15] The short-run drop in L_k would be even bigger if it were not offset by the positive autocorrelation in fees produced by both variation in types and transitory (multi-month) variation in fee-paying propensities.

[16] We are grateful to Devin Pope for suggesting this analogy.

3.1 Potential correlation between financial distress and credit card tenure.

The tendency to observe declining fees may reflect a tendency for new account holders to experience more financial/personal distress than account holders with high tenure. To test this hypothesis, we determined whether FICO scores and behavior scores (two inverse[17] measures of financial distress) correlate with account tenure.

Figure 4 plots FICO scores and behavior scores by account tenure, demeaned and normalized. To calculate the FICO variable, a single FICO mean is calculated for all accounts over all periods in our sample. This mean is used for the demeaning. A single FICO standard deviation is calculated for all accounts over all periods in our sample. This standard deviation is used for the normalization. An analogous method is used for the behavior score.

No time trend is apparent in the normalized data. To more formally measure the FICO-tenure relationship, we predict FICO with an account-tenure spline using annual knots (controlling for account and time fixed effects). The estimated tenure spline exhibits slopes that bounce around in sign and are all very small in magnitude. For example, at a horizon of 5 years, the spline predicts a total (accumulated) change in the FICO score of 18 units since the account was opened. At a horizon of 10 years the spline predicts a total (accumulated) change in the FICO score of -0.04 units since the account was opened. Recall that the mean FICO score is 732 and the standard deviation of the FICO score is 81. Hence, financial distress does not show significant variation with account tenure.

3.2 Movers

Moving to a new home could potentially cause both patterns of fee payment that we see. Disruptions associated with the move and additional needs for cash could lead to payment of late, over limit, and cash advance fees. Over time, consumers would revert to their normal pattern of fee payment.

To test for this possibility, since we know when account holders move, we examine the difference between the frequency of fee payment during the move (defined as payment from two months before to two months after the move date) and fee payment in all other months.

[17] A high FICO or behavior score implies that the individual is a reliable creditor. A behavior score is a proprietary measure of credit risk calculated by the card-issuing institution.

We find that account holders are only about 1 percent more likely to pay late or over limit fees during the move, and 2 percent more likely to pay cash advance fees. These small differences are not enough to account for the large reductions in fee payment by tenure or the recency effects.

3.3 Potential correlation between purchasing patterns and credit card tenure.

The tendency to observe declining fees may reflect a tendency for new account holders to spend more than account holders with high tenure. To test this hypothesis, we determined if purchases correlate with account tenure. Figure 4, which plots the demeaned and normalized level of purchases, again shows no economically significant time trend.

3.4 Non-utilization of the card.

The fee dynamics that we observe could be driven by consumers who temporarily or permanently stop using the card after paying a fee on that card. We look for these effects by estimating a regression model in which the outcome of "no purchase in the current month" is predicted by dummies for past fee payments and control variables, including account and time fixed effects as well FICO, Behavior, and Utilization. We find very small effects of past fee payments on subsequent card use. For example, (controlling for account fixed effects) somebody who paid a fee every month for the past six months is predicted to be only 2% less likely to use their card in the next month relative to somebody with no fee payments in the last six months. Such small effects cannot explain our learning dynamics, which are over an order of magnitude larger. Figure 5 also plots the absolute level of utilization (demeaned and normalized), which exhibits no time-series pattern.

3.5 Time-varying financial service needs.

Time-varying financial service needs may also play an important role in driving fee dynamics. To illustrate this idea, let ν_t represent a time-varying cost of time, so that

$$\Pr\left(f_t = 1\right) = \nu_t, \tag{5}$$

where ν_t is an exogenous process, that causes fee use, but is not caused by it. To explain our recency effect, one needs ν_t to be negatively autocorrelated at a monthly frequency. To see this, consider the regression,

$$f_t = \theta f_{t-1} + \text{controls}. \tag{6}$$

If (5) holds, then the regression coefficient is $\theta = cov\left(\nu_t, \nu_{t-1}\right)/var\left(f_{t-1}\right)$.

We run this regression, including all of our usual control variables, that is, time- and account-fixed effects, a tenure spline, *Purchase*, *Active*, *BillExist*, and *Util*. We also include *Behavior* and *FICO*.[18]

$$\begin{aligned} f_{i,t}^{j} &= \theta f_{i,t-1}^{j} + \alpha + \phi_i + \psi_{time} + Spline(Tenure_{i,t}) \\ &\quad + \eta_1 Purchase_{i,t} + \eta_2 Active_{i,t} + \eta_3 BillExist_{i,t-1} \\ &\quad + \eta_4 FICO_{i,t-3} + \eta_5 Behave_{i,t-3} + \eta_6 Util_{i,t} + \epsilon_{i,t}. \end{aligned}$$

Results for the three types of fees are given in Table 5. We find that θ is -0.75 for the late fee, -0.52 for the over limit fee, and -0.28 for the cash advance fee. We call this the "recency effect," since the payment of a fee last month greatly reduces the probability that a fee will be paid this month.[19]

The empirical finding of $\theta < 0$ implies $corr\left(\nu_t, \nu_{t-1}\right) < 0$. Hence, to explain the "recency effect" with time-varying financial needs, it would need to be the case that ν_t is negatively autocorrelated. The autocorrelation of ν_t would need to be not only negative, but also greater than 0.75 (in the case of the late fee) in absolute value: $corr\left(\nu_t, v_{t-1}\right) \leq \theta = -0.75$.[20]

We think that such a very strong negative autocorrelation of monthly needs is unlikely.[21]

[18] The results do not differ if we instead begin the regressions in month 2 and exclude the behavior and FICO scores.

[19] There is a potential small sample bias (Nickell 1981), to which we thank Peter Fishman and Devin Pope for drawing our attention. To see how large it is, we note that if f_t is i.i.d., then in the regression $f_t = \theta f_{t-1}$+constant, done over a T periods, the expected value of θ is $-1/T$. With $T = 24$, the bias is -0.05. We conclude that, in our study, the small sample bias is very small compared to the large negative θ that we find.

[20] It is easy to see that under (5), $cov\left(f_t, f_{t-1}\right) = cov\left(\nu_t, \nu_{t-1}\right)$, and $var\left(f_t\right) = E\left[\nu_t\right]\left(1 - E\left[\nu_t\right]\right) \geq E\left[\nu_t^2\right] - E\left[\nu_t\right]^2 = var\left(\nu_t\right)$, as $\nu_t \in [0,1]$. So, $\theta = cov\left(f_t, f_{t-1}\right)/var\left(f_{t-1}\right)$ satisfies $|\theta| \leq \left|cov\left(\nu_t, \nu_{t-1}\right)\right|/var\left(\nu_t\right) = \left|corr\left(\nu_t, \nu_{t-1}\right)\right|$, and θ and $corr\left(\nu_t, \nu_{t-1}\right)$ have the same sign.

[21] The least implausible type of negatively autocorrelated process in economics is a "periodic spike" process, which take a value of a every K periods, and $b \neq a$ otherwise. It has an autocorrelation of $-1/\left(K-1\right)$. We fail to find evidence for such a pattern in credit card use other than fees. For instance, expenses across time

First, since the regression results include time fixed effects, such autocorrelations could not occur from events that happen at regular intervals during the year — e.g., from summer vacations. We separately verify that, in any case, fee payments are not seasonal (although spending certainly is). Second, the presence of highly negative autocorrelations at a monthly level would rule out events that last more than one month. For example, a personal crisis that raised the opportunity cost of time for two months would create a positive autocorrelation in time needs and fee payments over the two months, not a negative one. Third, the time-varying needs would have to produce a higher than average fee payment in one month followed by a lower than average fee payment in the following month. This would rule out episodes of high opportunity cost of time for one month followed by a return to the status quo.

For most plausible processes, needs are likely to be positively autocorrelated. For example, the available evidence implies that income processes are positively autocorrelated (e.g. Guvenen (2007)). While we cannot rule out the "negatively autocorrelated needs" story, existing microeconomic evidence suggests it is highly unlikely to be the right explanation for the empirical patterns that we observe. We conclude that the finding of $\theta < 0$ in (6) is most plausibly explained by a recency effect – consumers become temporarily vigilant about fee avoidance immediately after paying a fee.

3.6 Medical expenses

Negative autocorrelation in fee payments could also be induced by a one-time medical emergency that is not repeated in subsequent months. In our sample, less than 3% of account-holder spending is in medical-related categories. Moreover, spending in such categories does not increase during periods in which a fee is paid.

4 Extensions

In this section, we present ancillary evidence that extends our analysis of learning and backsliding. Teasing out some of the determinants of learning is challenging, since we do not observe many of the underlying factors that influence learning dynamics. For example, we

are *positively* autocorrelated.

do not see the printed format of the bill that was used during our sample period, and can not tell how salient the fees were (although we suspect that the credit card company likely did not go out of its way to call attention to them).

Below, we evaluate several factors that we think might influence the rate of learning and backsliding in fee payment, including differences across demographic groups.

4.1 Late payments and timeliness of subsequent payments

One possible learning mechanism is variation in salience: paying a fee brings the existence of that fee to the account holder's attention, leading her to change subsequent behavior to reduce the frequency of fee payment. We do not directly observe many of the factors that contribute to fee salience–for example, how many times or how carefully the account holder reviews her monthly statement. However, one consequence of a shift in salience of the late payment fee is that the account holder will begin to pay on an earlier day in the billing cycle. Moreover, we hypothesize that the shift to earlier payments will be particularly pronounced in the cycle immediately after a late fee payment has been incurred.

Figure 5 presents a histogram of the days before (negative numbers) and days after (positive numbers) the due date the bill is paid for consumers with either 1 month or 36 months of account tenure. As expected, the distribution shifts to the left over account tenure–reflecting the net reduction in late fee payment observed over tenure. The mass of the distribution largely lies within a two-week period before the due date and a one-week period afterwards.

Figure 6 presents histograms of the days before (negative numbers) or after (positive numbers) the due date the bill is paid for consumers with account tenures of 1-12 months or 25-36 months. for the billing cycle following the incursion of a late fee. The distribution of dates for high-tenure account holders has shifted to the left–about 42 percent of high-tenure account holders pay their bills more than two weeks early, while only about 27 percent of low-tenure account holders do so–a 15 percentage point increase. The daily frequency of fee payment within two weeks of either side of the due date do not show similarly-sized changes. For both account tenures, the frequency of payment more than two weeks early is much higher than the frequency of payment for that range for all account holders of those tenures.[22]

These figures are consistent with the idea that paying a fee brings the fee payment to the account holder's attention, leading him or her to try to be extra careful the next month by paying the account earlier in the cycle. The frequency of these 'early' payments rises with account tenure.

4.2 Borrower Characteristics

In this subsection, we characterize how fee payment varies by borrower characteristic– marital status, gender, age, income, and FICO (credit) score Note that most of these characteristics – e.g. FICO score – are endogenously determined jointly with fee payment, so our interpretations are not causal.[23]

Table 6 presents the correlation matrix for the characteristics. The correlations are all small and positive, ranging between 1 and 15 percent. Gender and marital status have the lowest correlations with the other characteristics, with correlations ranging between 1 and 7 percent for each category. Age has correlations between 2.5 and 9 percent for all but one category, but a correlation of 15 percent with income.[24] The financial variables show somewhat higher pairwise correlations, although still small in absolute value.

To measure the impact of particular characteristics, for each characteristic we first divide the sample, as appropriate, into two sub-categories (in the case of gender and marital status) or three sub-categories (bottom, middle, and top third for age, income, FICO, credit limit, and card utilization). For each characteristic, and for each sub-category, we estimate the same regressions as in section 2–specifically, equations 1 and 4.[25]

Estimating a regression for each characteristic sub-category for all three types of fees

[22] The distribution is much flatter than the distribution of payment dates presented in the previous figure because this figure averages over account tenures, and the peak day of payment is moving earlier over time.

[23] We have also looked at the association with credit limits and card utilization, with somewhat similar results.

[24] The pattern of income with age is an inverse U-shape, as has been found in other samples.

[25] In principle, we would like to know the marginal effect of varying one borrower characteristic while holding the other characteristics fixed. In practice, two difficulties complicate such analysis. First, given the very large number of observations (4,000,000 account statements), the six-fold expansion in parameters required to control for all seven demographic characteristics makes estimation computationally difficult. Second, although this problem could be ameliorated either by throwing out observations or controlling for smaller sets of characteristics, some parts of the characteristic space are sparse. For example, there are not many elderly, low-income men in the sample.

leads to a total of 57 account-tenure regressions and 57 fixed-effect autoregressions.[26] We summarize these results as follows. For the tenure regressions (which trace out the decline in fee payment by account tenure), we regress the estimated coefficients on the spline for tenure on an exponential function, specifically estimating

$$Spline_k = a + be^{-\phi(k-2)},$$

where $Spline_k$ is the coefficient for tenure of month k. At the initial account tenure, $k = 2$, estimated fee payment is given by $a + b$. As k increases, estimated fee payment declines to a. Thus, the total percent reduction in fee payment by tenure is given by $b/(a + b)$. The 'half-life' of fee payment–that is, the number of periods after which fee payment will be cut in half–is given by $\ln 2/\phi$; thus ϕ can be thought of as a measure of the speed of learning (where larger ϕ implies faster learning).

For the fixed-effects autoregression on fees, if we denote the coefficient on lag k as β_k, then we estimate

$$\beta_k = a + be^{-\phi(k-1)}.$$

Recall that these regressions (and the related L_k calculations) showed a recency effect–so that at small lags, $\beta_k \ll 0$, indicating that having paid a fee recently is associated with a large reduction in current fee payment frequency. For $k = 1$, $\beta_k = a + b$, while as k gets larger, β_k approaches a. So the total value of the recency effect is equal to b, and the percent decline associated with the recency effect is $b/(a + b)$. As before, $\ln 2/\phi$ gives the number of periods after which the recency effect has declined by half–a rough estimate of the speed of forgetting.

Both sets of regressions show that the net reductions in fee payment and recency effect seen in the whole dataset are not attributable to any single borrower characteristic; that is, it is not the case that, for example, the reduction in fee payment is associated with just high income borrowers, while borrowers with other incomes see little change in fee payment by account tenure. Rather, across and within all borrower characteristics, we see net reductions in fee payments and a recency effect.

[26] Two characteristics with two sub-categories each added to five characteristics with three sub-categories equals 19 total sub-categories. Estimating one regression for three fee types for each of these sub-categories yields $19 \times 3 = 57$ total regressions.

From the tenure regressions, broadly, two points emerge:

1. **The net reduction in fee payment over time is roughly about 70 percent, regardless of fee type or borrower characteristic**.There is some economically significant variation in initial fee payment across different categories of some borrower characteristics; for example, borrowers with low credit limits have about a 39 percent frequency of paying a late fee, while those with high limits have a 29 percent frequency. However, even in this case, for both categories of credit limit, the net reduction in fee payment over time is about 70 percent.

2. **However, the 'speed of learning' can vary by a factor of two or more within borrower characteristic categories.** For example, for the late fee the estimated values of ϕ for young, middle-aged, and older adults imply half-lives of fee frequency payment of one year, six months, and eight months, respectively.[27] Singles learn more than twice as fast as couples; the speed of learning is increasing with income, so that high-income borrowers learn more than twice as fast as low-income borrowers. Similar comparisons apply to credit limit and card utilization. There is no appreciable difference by gender. Interestingly, FICO score shows a U-shaped pattern, as fee reductions occur most quickly for account holders with the middle third of FICO scores.

These differences can be seen in Figure 7, which plots the tenure effects for the five categories Table A2 in the appendix provides the regression results for a, b, and ϕ.

The autoregressions broadly show two points:

1. **The size of the recency effect varies strongly with demographics**. For example, middle-aged borrowers see a 62 percent reduction in late fees, while younger borrowers and older borrowers see reductions of 58 percent and 41 percent in the same fee category, respectively. High-income borrowers see a reduction of about 65 percent in late fee payment–nearly double that of low-income borrowers.

2. **The speed of forgetting also varies strongly with demographics–by more than a factor of two in some cases**. For example, the estimated values of ϕ

[27] This disadvantage of older adults confirms, with a different task, the qualitative findings of Agarwal et al. (2009).

implies that it only takes about two and one-half months for half of the recency effect to disappear for older account holders, while for middle-aged holders it takes nearly six months. Similarly, for high income borrowers, it takes nearly six months for half of the recency effect to disappear, while for low-income borrowers it takes a little over two months.

Figure 8 plots the recency effect for several of the categories, and table A3 in the appendix provides the regression results.

Jointly, both sets of results show that a slower pace of net learning is associated with a smaller recency effect and a faster level of forgetting. For example, as noted above the rate of tenure-based declines in fees is about half as fast for lower-income borrowers as for higher-income borrowers; the size of the recency effect is also about half as great for lower-income borrowers; and the speed of forgetting is more than twice as great for lower-income borrowers. The same is true for older borrowers, as compared with middle-aged borrowers and for couples vs. singles.

4.3 Learning from other types of fees

Payment of one kind of fee may alert account holders to the possibility that they may have to pay other types of fees on their credit card, thereby increasing their general vigilance. To evaluate this possibility, we compute versions of L_1 in which we calculate the conditional probability of paying a fee of type j after having paid a fee of a different type, normalized by the unconditional probability of paying a fee of type j at time t. As before, a value of 1 will imply no impact of lagged fee payment on current fee payment, while values much less than one imply a substantial impact.

Table 7 presents the results. We find the impact of paying other fees to be small, but not negligible. For the six possible cases, values of L range from 0.90, in the case of the association between a late fee payment last month and a cash advance fee this month, to 0.96, in the case of a cash advance fee payment last month to a limit fee payment this month–that is, the measured reductions in fee payment likelihood range between 4 and 10 percent. These figures are much smaller than the reductions of over 40 percent attributable to payment of the own type of fee.

4.4 Cash advance fees and the availability of other sources of liquidity

Cash advance fee payment may be affected by different factors than payment of other kinds of fees. Cash advances are an expensive substitute for other kinds of liquidity, and borrowers may choose to take out cash advances when currency is needed for payment, but other potential sources of currency are not available. Thus, the availability of ATMs may affect the frequency of cash advances; account holders in areas with a large number of ATMs would be more likely to have lower absolute levels of cash advance fee payment and to show a faster rate of reduction in cash advance fee payment (as consumers successfully seek out ATMs as an alternative).

We do not directly observe the number of ATM machines in an account holder's home market. However, we do observe population density by zip code of the account holder's address, which may be a good proxy for ATM coverage.[28]

Estimating the same set of regressions as in the previous subsection provides some support for this hypothesis. The reduction in initial cash advance fee payment generated by going from a medium- to a high-density area is about 4 percentage points. The 'half-life' of fee payment frequency is about 19 months in a medium-density area, as opposed to 13 months in a high-density one. The recency effect is about 28 percentage points higher in the higher-density area, with a half-life of about seven months, as opposed to five months in the medium-density area.

4.5 Other tests

Any learning effect might be stronger for a consumer's first credit card. While we do not know the cardholders in our dataset for whom this may be the first card, we do observe two types of consumers for whom this is more likely to be true: young consumers and consumers for whom this is their only credit card. Limiting our sample to these consumers, we reestimate the relationships between (i) tenure and fee payments and (ii) past fee payment and current fee payment. We find no notable differences between those households with only one card and those who have more than one card.

[28]This relationship may not hold for poorer urban areas, which have traditionally been underserved by banks. However, few of the borrowers in our sample live in such areas.

We have also analyzed fee payments of borrowers who exhibit two other types of behavior. First, we study borrowers who spend on gambling-related institutions such as casinos. We find no significant difference for fee payment between this group and the total population of cardholders. Second we study 'transactors'–those account holders who pay their balance in full by the due date, thus enjoying the benefit of the float and avoiding paying interest charges; this group comprises 27 percent of our sample. We again see no significant differences in behavior between this group and those account-holders who carry a balance ('revolvers').

5 Conclusion

Credit card users sometimes learn about add-on fees by paying them. With years of experience, credit card customers substantially reduce these fee payments. We document this process using a three-year panel dataset representing 120,000 accounts.

In our data, new accounts generate direct fee payments of $15 per month. The data implies that negative feedback — i.e., paying fees — teaches consumers to avoid triggering fees in the future. Controlling for account fixed effects, monthly fee payments fall by 75% during the first four years of account life.

We also find that learning is not monotonic. We estimate that knowledge depreciates 10% per month. As previous fee-paying lessons recede into the past, consumers tend to backslide. However, on net, knowledge accumulation dominates knowledge depreciation. Over time, fee payments drastically fall.

Our findings support the view that some consumers have limited attention, and that they do not initially understand all the terms of their financial contracts. We have also found that the passage of time and the accumulation of experience significantly improves performance. Our results suggest that new customers are far more profitable than experienced ones, since new customers pay far more add-on fees. All of these considerations may help to explain the marketing decisions that firms make and the nature of the complex contracts that they offer prospective customers.

Future work should explore the endogenous design of contracts, of marketing campaigns, and of the ongoing feedback that firms give their customers. Our research implies that profit-maximizing firms should try to write contracts with novel, shrouded fees and that marketing will be targeted at potential customers who are the most likely to stumble into these fees. Our results also highlight the *potential* social value of clear disclosure of fees before a financial relationship begins and of salient ongoing disclosures during the life of the relationship. For example, occasionally reminding consumers about the *aggregate* fees they have paid since they opened their account, might make fees more salient. To understand consumer behavior and firm behavior, and to design and evaluate regulations, economists and other regulators will need to understand the mechanisms that we have begun to explore in this paper.

References

Agarwal, Sumit, Souphala Chomsisengphet, Chunlin Liu, and Nicholas S. Souleles. 2006. "Do Consumers Choose the Right Credit Contracts?" University of Pennsylvania Working Paper.

Agarwal, Sumit, John C. Driscoll, Xavier Gabaix, and David I. Laibson. 2009. "The Age of Reason: Financial Mistakes Over the Lifecycle." *Brookings Papers on Economic Activity,* 2009(2): 51-117.

Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles. 2007. "The Reaction of Consumer Spending and Debt to Tax Rebates– Evidence from Consumer Credit Data." *Journal of Political Economy*, 115(6), 986-1019

Arellano, Manuel and Stephen R. Bond. 1991. "Some Specification Tests for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Review of Economic Studies*, 58(1): 277-298.

Argote, Linda, Sara Beckman, and Dennis Epple. 1990. "The Persistence and Transfer of Learning in Industrial Settings." *Management Science*, 36(1): 141-15.

Aumann, Robert, Sergiu Hart, and Motty Perry. 1997. "The Absent-Minded Driver." *Games and Economic Behavior*, 20(1): 102-116.

Ausubel, Lawrence M. 1991. "The Failure of Competition in the Credit Card Market." *American Economic Review*, 81(1): 50-81.

Ausubel, Lawrence M. 1999. "Adverse Selection in the Credit Card Market." University of Maryland Working Paper.

Bahk, Byong Hong, and Michael Gort. 1993. "Decomposing Learning by Doing in New Plants." *Journal of Political Economy*, 101(4): 561-583.

Benkard, C. Lanier. 2000. "Learning and Forgetting: The Dynamics of Aircraft Production." *American Economic Review*, 90(6): 1034-1054.

Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman. 2010. "What's Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment." *Quarterly Journal of Economics,* 125(1): 263-306.

Besanko, David, Ulrich Doraszelski, Yaroslav Kryukov, Mark Satterthwaite. 2010. "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics." *Econometrica*, 78(2): 453-508.

Camerer, Colin F. 2003. *Behavioral Game Theory.* Princeton, NJ: Princeton University Press.

Chamberlain, Gary. 1993. "Feedback in Panel Data Models." Harvard University Working Paper.

Crawford, Vincent P. 1995. "Adaptive Dynamics in Coordination Games." *Econometrica*, 63(1): 103-143.

DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature,* 47(2): 315-372.

DellaVigna, Stefano, and Ulrike Malmendier. 2004. "Contract Theory and Self-Control: Theory and Evidence." *Quarterly Journal of Economics*, 119(2): 353-402.

Fishman, Peter, and Devin G. Pope. 2006. "Specific Deterrence: Evidence and Implications." UC Berkeley Working Paper.

Gabaix, Xavier, and David I. Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics*, 121(2): 505-540.

Grubb, Michael. 2009. "Selling to Overconfident Consumers." *American Economic Review,* 99(5): 1770-1807.

Guvenen, Fatih. 2007. "Learning Your Earnings: Are Labor Income Shocks Really Very Persistent?" *American Economic Review*, 97(3): 687-712.

Ho, Teck-Hua, and Juin-Kuan Chong. 2003. "A Parsimonious Model of Stockkeeping-Unit Choice." *Journal of Marketing Research*, 40(3): 351-365.

Honoré, Bo E., and Ekaterini Kyriazidou, 2000. "Panel Data Discret Choice Models with Lagged Dependent Variables" *Econometrica*, 68(4): 839-874.

Jin, Ginger, and Philip Leslie. 2003. "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards." *Quarterly Journal of Economics*, 118(2): 409-451.

Kerr, Sougata, and Lucia Dunn. 2008. "Consumer Search Behavior in the Changing Credit Card Market." *Journal of Business and Economic Statistics,* 26(3): 345-353.

Koszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics*, 121(4): 1133-1166.

Lehrer, Ehud. 1988. "Repeated Games with Stationary Bounded Recall Strategies." *Journal of Economic Theory*, 46(1): 130-144.

Lemieux, Thomas, and W. Bentley MacLeod. 2000. "Supply Side Hysteresis: The Case of the Canadian Unemployment Insurance System." *Journal of Public Economics*, 78(1): 139-170.

Malmendier, Ulrike, and Devin Shanthikumar. 2007. "Are Small Investors naïve about Incentives?" *Journal of Financial Economics*, 85(2): 457-489.

Marimon, Ramon and Shyam Sunder. 1994. "Expectations and Learning under Alternative Monetary Regimes: An Experimental Approach." *Economic Theory,* 4(1): 131-162.

Miravete, Eugenio J. 2003. "Choosing the Wrong Calling Plan? Ignorance and Learning." *American Economic Review*, 93(2): 297-310.

Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The Market for News," *American Economic Review,* 95(4): 1031-1053.

Nickell, Stephen. 1981. "Biases in Dynamic Models with Fixed Effects." *Econometrica*, 49(6): 1417-1426.

Odean, Terrance, Michal Strahilevitz and Brad M. Barber. 2010. "Once Burned, Twice Shy: naïve Learning, Counterfactuals,and the Repurchase of Stocks Previously Sold." UC Davis Working Paper.

Oster, Sharon M., and Fiona M. Scott Morton. 2005. "Behavioral Biases Meet the Market: The Case of Magazine Subscription Prices." *Advances in Economic Analysis and Policy*, 5(1): Article 1.

Roth, Alvin E., and Ido Erev. 1995. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior*, 8(2): 164-212.

Shui, Haiyan, and Lawrence M. Ausubel. 2004. "Time Inconsistency in the Credit Card Market." University of Maryland Working Paper.

Spiegler, Rani. 2011. *Bounded Rationality and Industrial Organization.* Oxford: Oxford University Press.

Thornton, Rebecca A., and Peter Thompson. 2001. "Learning from Experience and Learning from Others: An Exploration of Learning and Spillovers in Wartime Shipbuidling." *American Economic Review*, 91(6): 1350-1368.

Van Huyck, John B., Joseph P. Cook, and Raymond C. Battalio. 1994. "Selection Dynamics, Asymptotic Stability, and Adaptive Behaviour." *Journal of Political Economy*, 102(6): 975-2005.

Wixted, John. 2004. "The Psychology and Neuroscience of Forgetting." *Annual Review of Psychology*, 55: 235-269.

Appendix A: Data Description

The total sample consists of 125,384 accounts open as of January 2002 and 22,392 opened between January and December of 2002 observed through December 2004. These accounts were randomly sampled from several million accounts held by the bank. From this sample of 147,776, we drop accounts that were stolen, lost, or frozen (due to fraud). We also exclude accounts that do not have any activity (purchases and payments) over the entire period. This leaves 128,142 accounts. Finally, we also remove account observations subsequent to default or bankruptcy, as borrowers do not have the opportunity to pay fees in such instances. This leaves us with an unbalanced panel with 3.9 million account observations.

Table A1 provides summary statistics for variables related to the accounts, including account characteristics, card usage, fee payment, and account holder characteristics. The second column notes whether the variable is observed monthly ('M'), quarterly ('Q'), or at account origination ('O'), the third column reports variable means, and the fourth column variable standard deviations. Note that the monthly averages for the 'Fee Payment' variables imply annual average total fees paid of \$141 (=\$11.75×12), with about 7.52 fee payments per year. Higher interest payments induced by paying fees (which raise the interest rate on purchases and cash advances) average about \$226 per year.

The accounts also differ by how long they have been open. Over 31 percent of the accounts are less than 12 months old, 20 percent are between 12 and 24 months old, 18 percent are between 24 and 36 months old, 13 percent are between 36 and 48 months old, 10 percent are between 48 and 60 months old, and 8 percent are more than 60 months old.

Table A1: Variable Descriptions and Summary Statistics

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

Notes: The "Credit Bureau Risk Score" is provided by Fair, Isaac and Company (hence 'FICO'). The greater the score, the less risky the consumer is. The "Payment Behavior Score" is a proprietary score based on the consumer's past payment history and debt burden, among other variables. It is created by the bank to capture determinants of consumer payment behavior not accounted for by the FICO score. "Q" indicates the variable is observed quarterly, "M" monthly, and "O" only at account origination.

Table A2: Tenure Results By Characteristic

	Late Fee						Over Limit Fee					
	a	s.e.	b	s.e.	phi	s.e.	a	s.e.	b	s.e.	phi	s.e.
A) Marital Status												
couples	0.101**	0.024	0.237**	0.035	0.052**	0.009	0.172**	0.040	0.399**	0.093	0.087**	0.014
singles	0.096**	0.025	0.192**	0.036	0.119**	0.009	0.163**	0.037	0.326**	0.097	0.199**	0.012
B) Gender												
males	0.095**	0.026	0.220**	0.032	0.084**	0.009	0.173**	0.035	0.357**	0.088	0.121**	0.011
females	0.100**	0.023	0.213**	0.036	0.094**	0.009	0.176**	0.040	0.343**	0.087	0.150**	0.014
C) Age												
young	0.101**	0.018	0.231**	0.030	0.061**	0.009	0.179**	0.030	0.371**	0.050	0.097**	0.015
middle age	0.093**	0.029	0.188**	0.040	0.114**	0.009	0.158**	0.048	0.317**	0.067	0.190**	0.015
old	0.123**	0.019	0.232**	0.046	0.086**	0.009	0.205**	0.031	0.392**	0.078	0.146**	0.015
D) Income												
low	0.113**	0.023	0.248**	0.035	0.051**	0.009	0.203**	0.035	0.426**	0.062	0.087**	0.015
medium	0.114**	0.025	0.223**	0.038	0.092**	0.009	0.190**	0.037	0.404**	0.063	0.153**	0.015
high	0.094**	0.027	0.191**	0.033	0.114**	0.009	0.157**	0.035	0.342**	0.067	0.190**	0.016
E) FICO												
low	0.104**	0.028	0.241**	0.037	0.053**	0.009	0.202**	0.037	0.406**	0.062	0.110**	0.015
medium	0.095**	0.024	0.208**	0.042	0.115**	0.009	0.182**	0.033	0.346**	0.070	0.189**	0.016
high	0.118**	0.024	0.222**	0.044	0.085**	0.009	0.194**	0.041	0.397**	0.062	0.155**	0.015

	Cash Advance Fee					
A) Marital Status						
couples	0.064**	0.014	0.158**	0.0270	0.032**	0.006
singles	0.063**	0.013	0.123**	0.0270	0.074**	0.005
B) Gender						
males	0.067**	0.017	0.148**	0.0231	0.054**	0.006
females	0.066**	0.015	0.132**	0.0241	0.063**	0.006
C) Age						
young	0.067**	0.011	0.143**	0.0188	0.034**	0.006
middle age	0.059**	0.018	0.115**	0.0264	0.072**	0.006
old	0.081**	0.012	0.146**	0.0305	0.056**	0.006
D) Income						
low	0.086**	0.015	0.169**	0.0197	0.032**	0.006
medium	0.072**	0.017	0.150**	0.0202	0.059**	0.006
high	0.058**	0.015	0.121**	0.0240	0.069**	0.006
E) FICO						
low	0.089**	0.012	0.146**	0.0209	0.054**	0.006
medium	0.052**	0.012	0.149**	0.0211	0.080**	0.006
high	0.061**	0.015	0.129**	0.0271	0.062**	0.006

*: significant at 95% level. **: significant at 99% level.

Table A3: Effects of Past Fee Payment By Characteristic

	Late Fee						Over Limit Fee					
	a	s.e.	b	s.e.	phi	s.e.	a	s.e.	b	s.e.	phi	s.e.
A) Marital Status												
couples	0.008*	0.004	-0.624**	0.066	0.179**	0.064	0.002	0.003	-0.812**	0.068	-0.164**	0.054
singles	0.007	0.006	-0.426**	0.074	0.358**	0.045	0.005**	0.001	-0.405**	0.074	0.306**	0.064
B) Gender												
males	0.010	0.008	-0.529**	0.039	0.164*	0.070	0.009**	0.001	-0.549**	0.062	0.240**	0.079
females	0.009	0.005	-0.493**	0.051	0.284**	0.100	0.001	0.001	-0.453**	0.044	0.215*	0.085
C) Age												
young	0.003	0.005	-0.575**	0.039	0.143*	0.061	0.004**	0.000	-0.780**	0.031	0.209**	0.052
middle age	0.009	0.009	-0.616**	0.058	0.129*	0.060	0.010	0.009	-0.687**	0.020	0.132	0.091
old	0.009**	0.002	-0.411**	0.057	0.268**	0.073	0.007**	0.001	-0.440**	0.052	0.274**	0.090
D) Income												
low	0.006	0.004	-0.352**	0.064	0.275		0.002	0.002	-0.360**	0.036	0.238**	0.082
medium	0.008	0.007	-0.506**	0.034	0.192**	0.069	0.010	0.009	-0.478**	0.040	-0.171**	0.064
high	0.009	0.005	-0.649**	0.044	0.121*	0.052	0.004	0.006	-0.586**	0.072	0.126	0.093
E) FICO												
low	0.010	0.006	-0.385**	0.070	0.251**	0.041	0.002	0.007	-0.238**	0.087	0.242**	0.070
medium	0.002	0.010	-0.533**	0.041	0.208**	0.062	0.008	0.007	-0.518**	0.097	0.231**	0.060
high	0.008	0.005	-0.710**	0.026	0.163*	0.078	0.003	0.005	-0.601**	0.058	0.158*	0.069
	Cash Advance Fee											
A) Marital Status												
couples	0.002	0.002	-0.516**	-0.516	0.162**	0.055						
singles	0.007	0.007	-0.413**	-0.413	0.285**	0.049						
B) Gender												
males	0.002	0.002	-0.510**	-0.510	0.151**	0.054						
females	0.003	0.003	-0.542**	-0.542	0.287**	0.069						
C) Age												
young	0.005	0.005	-0.506**	-0.506	0.143*	0.060						
middle age	0.007	0.007	-0.573**	0.573	-0.133	0.082						
old	0.005	0.005	-0.445**	-0.445	0.231**	0.054						
D) Income												
low	0.009	0.009	-0.350**	-0.350	0.276**	0.071						
medium	0.005	0.005	-0.483**	-0.483	0.247**	0.063						
high	0.000	0.000	-0.644**	-0.644	0.144*	0.071						
E) FICO												
low	0.007*	0.007	-0.373**	-0.373	0.294**	0.088						
medium	0.006	0.006	-0.503**	-0.503	0.231*	0.100						
high	0.007**	0.007	-0.621**	-0.621	0.181*	0.079						

*: significant at 95% level. **: significant at 99% level.

Table 1: Regressions of a Dummy for Fee Payment on Account Tenure and Controls

	Late Fee		Over Limit Fee		Cash Advance Fee	
	Coefficient	Std. Err.	Coefficient	Std. Err.	Coefficient	Std. Err.
Intercept	0.228**	0.079	0.139**	0.022	0.421*	0.186
Account Age ≤ 12	–0.012**	0.003	–0.008**	0.002	–0.022**	0.007
Account Age 12 to 24	–0.005*	0.002	–0.003**	0.001	–0.012**	0.004
Account Age 24 to 36	–0.005*	0.002	–0.0004*	0.0002	–0.005*	0.002
Account Age 36 to 48	–0.003	0.002	–0.0001	0.00009	–0.002	0.002
Account Age 48 to 60	–0.0004	0.002	0.000	0.000	–0.0002	0.007
Account Age ≥ 60	–0.0001	0.001	0.000	0.00002	0.00008	0.0008
Bill Existence Dummy$_{t-1}$	0.061*	0.024	0.015	0.023	0.099**	0.034
Bill Activity Dummy	0.007	0.005	0.003**	0.001	0.010	0.005
Purchases	0.00006	0.00004	0.00002**	0.000	0.00008	0.00005
Utilization (Debt/Limit)	0.068	0.007	0.023	0.006	0.069	0.015
Account Fixed Effects	Yes		Yes		Yes	
Time Fixed Effects	Yes		Yes		Yes	
Number of Observations	4.13 Million		4.13 Million		4.13 Million	
Adj. R^2	0.043		0.041		0.042	

This table reports fixed-effects panel regressions of dummies for fee payment on a spline for account tenure (with yearly knot points), a (lagged) dummy variable for existence of a bill, a dummy variable for account activity, the dollar value of purchases (in hundreds), and utilization (as measured by debt/limit). See Appendix Table A1 for variable summary statistics. Standard errors are Hubert/White/Sandwich. * denotes statistical significance at a 95 percent confidence level, and ** at a 99 percent confidence level.

Table 2: Regressions of Fee Dollar Amount on Account Tenure and Controls

	Late Fee		Over Limit Fee		Cash Advance Fee	
	Coefficient	Std. Err.	Coefficient	Std. Err.	Coefficient	Std. Err.
Intercept	4.724**	0.888	1.974**	0.363	5.674**	1.268
Account Age ≤ 12	–0.128**	0.032	–0.062**	0.015	–0.137*	0.058
Account Age 12 to 24	–0.093**	0.031	–0.050**	0.013	–0.099*	0.049
Account Age 24 to 36	–0.064*	0.029	–0.030**	0.011	–0.077*	0.039
Account Age 36 to 48	–0.050	0.028	–0.014	0.010	–0.059	0.038
Account Age 48 to 60	–0.021	0.021	–0.009	0.009	–0.029	0.030
Account Age ≥ 60	–0.001	0.021	–0.002	0.008	–0.002	0.024
Bill Existence Dummy$_{t-1}$	0.763**	0.265	0.218	0.137	0.826	0.562
Bill Activity Dummy	0.064	0.047	0.037**	0.011	0.099**	0.037
Purchases	0.00007	0.00005	0.00004**	0.00001	0.00008*	0.00003
Utilization (Debt/Limit)	0.267	0.059	0.278	0.066	0.257	0.087
Account Fixed Effects	Yes		Yes		Yes	
Time Fixed Effects	Yes		Yes		Yes	
Number of Observations	4.13 Million		4.13 Million		4.13 Million	
Adj. R^2	0.041		0.045		0.043	

This table repeats the regressions reported in Table 1, but replacing the dependent variables with the dollar values of fees paid. * denotes statistical significance at a 95 percent confidence level, and ** at a 99 percent confidence level.

Table 3: Monte Carlo Simulations of Fixed Effects Panel Autoregressions

	One Lag		Twelve Lags		Eighteen Lags	
Lag	Coefficient	Std. Dev.	Coefficient	Std. Dev.	Coefficient	Std. Dev.
1	–0.029	<.001	–0.060	<.001	–0.101	<.001
2			–0.060	<.001	–0.101	<.001
3			–0.059	<.001	–0.102	<.001
4			–0.057	<.001	–0.101	<.001
5			–0.056	<.001	–0.100	<.001
6			–0.054	<.001	–0.098	<.001
7			–0.052	<.001	–0.096	<.001
8			–0.050	<.001	–0.093	<.001
9			–0.048	<.001	–0.089	<.001
10			–0.046	<.001	–0.085	<.001
11			–0.043	<.001	–0.081	<.001
12			–0.040	<.001	–0.076	<.001
13					–0.070	<.001
14					–0.064	<.001
15					–0.058	<.001
16					–0.051	<.001
17					–0.044	<.001
18					–0.037	<.001

This table reports the results of Monte Carlo simulations of fixed effects panel autoregressions of a dummy variable drawn from an i.i.d. uniform distribution.

Table 4: Fixed Effects Panel Autoregressions

	Late Fee		Over Limit Fee		Cash Advance Fee	
Lag	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff	Std. Err.
1	–0.512**	0.004	–0.467**	0.007	–0.377**	0.005
2	–0.475**	0.004	–0.383**	0.006	–0.328**	0.004
3	–0.422**	0.004	–0.326**	0.005	–0.251**	0.005
4	–0.362**	0.004	–0.295**	0.004	–0.198**	0.005
5	–0.305**	0.004	–0.250**	0.006	–0.152**	0.001
6	–0.250**	0.003	–0.146**	0.006	–0.115**	0.006
7	–0.198**	0.003	–0.096**	0.006	–0.094**	0.006
8	–0.151**	0.003	–0.072**	0.006	–0.064**	0.006
9	–0.100**	0.003	–0.016*	0.006	–0.025**	0.006
10	–0.060**	0.003	0.006	0.007	–0.007	0.007
11	–0.031**	0.003	0.035**	0.007	0.003	0.004
12	0.007*	0.003	0.092**	0.007	0.002	0.005
Fixed Effects	Yes		Yes		Yes	
Adj. R^2	0.05		0.04		0.04	
Obs	2.1 Million		2.1 Million		2.1 Million	

This table reports fixed effects panel autoregressions of dummy variables for three kinds of fee payments. The regressions also include the same control variables as in Tables 1 and 2, but those results have been suppressed for simplicity. See Appendix Table A1 for variable summary statistics. Standard errors are Hubert/White/Sandwich. * denotes statistical significance at a 95 percent confidence level, and ** at a 99 percent confidence level.

Table 5: Time Varying Needs

	Late Fee		Over Limit Fee		Cash Advance Fee	
	Coefficient	Std. Err.	Coefficient	Std. Err.	Coefficient	Std. Err.
Intercept	0.2487*	–0.08236	0.1267**	–0.0284	0.4263**	–0.1547
$f^j_{i,t-1}$	–0.7483**	–0.1403	–0.5248**	–0.1076	–0.2784**	–0.064
Account Age $\leq$ 12	–0.0103*	–0.0037	–0.0077	–0.0024	–0.0237**	–0.0068
Account Age 12 to 24	–0.0059*	–0.0023	–0.0025**	–0.001	–0.0127**	–0.0048
Account Age 24 to 36	–0.0044*	–0.0019	–0.0004	–0.0001	–0.0057*	–0.0028
Account Age 36 to 48	–0.0021	–0.0015	–0.0001	–0.0001	–0.0021	–0.0018
Account Age 48 to 60	–0.0003	–0.0016	–0.0001	–0.0001	–0.0002	–0.0068
Account Age $\geq$ 60	–0.0001	–0.0014	–0.0001	–0.0001	–0.0001	–0.0008
Purchases/$100_{i,t}$	0.0052	–0.0034	0.0021**	–0.0003	0.0073	–0.0053
$\text{Active}_{i,t}$	0.0071	–0.0048	0.0026**	–0.0009	0.0093	–0.0058
Bill Existence Dummy_{t-1}	0.0618**	–0.0257	0.0179*	–0.0084	0.0964**	–0.0389
$\text{Behave}_{i,t-3}$	–0.0035**	–0.0008	–0.0028**	–0.0007	–0.0053*	–0.0025
$\text{FICO}_{i,t-3}$	–0.0027**	–0.0005	–0.0014**	–0.0006	–0.0046*	–0.0021
Utilization (Debt/Limit)	0.0506**	–0.0074	0.0283**	–0.008	0.0693**	–0.0182
Adj. R^2	0.0416		0.0484		0.0497	
Number of Observations	3.9 Million		3.9 Million		3.9 Million	

This table reports regressions of dummies for fee payment on a constant, a lagged dummy for fee payment, a spline for account tenure (with yearly knot points), a (lagged) dummy variable for existence of a bill, a dummy variable for account activity, the dollar value of purchases (in hundreds), and utilization (as measured by debt/limit). See Appendix Table A1 for variable summary statistics. Standard errors are Hubert/White/Sandwich. * denotes statistical significance at a 95 percent confidence level, and ** at a 99 percent confidence level.

Table 6: Correlation Matrix

	Marital Status	Gender	Age	Income	FICO
Marital Status	1.00				
Gender	0.02	1.00			
Age	0.05	0.03	1.00		
Income	0.06	0.07	0.15	1.00	
FICO	0.07	0.01	0.05	0.09	1.00

This table reports cross-correlations of account-holder characteristics.

Table 7: Impact of Other Fees

Cross Correlations						
Month (t and $t-1$)	$\frac{\mathrm{PF}_t \mid \mathrm{PF}_{t-1}}{\mathrm{PF}_t}$	$\frac{\mathrm{PF}_t \mid \mathrm{NPF}_{t-1}}{\mathrm{PF}_t}$	$\mathrm{PF}_t \mid \mathrm{PF}_{t-1}$	$\mathrm{PF}_t \mid \mathrm{NPF}_{t-1}$	PF_t	NPF_t
$\mathrm{LateFee}_t$ & $\mathrm{LimitFee}_{t-1}$	0.941	1.016	0.2045	0.2210	0.217	0.783
$\mathrm{LateFee}_t$ & $\mathrm{CashAdvFee}_{t-1}$	0.929	1.020	0.2020	0.2217	0.217	0.783
$\mathrm{CashAdvFee}_t$ & Limt Fee_{t-1}	0.958	1.022	0.3238	0.3454	0.338	0.662
$\mathrm{CashAdvFee}_t$ & $\mathrm{LateFee}_{t-1}$	0.903	1.050	0.3054	0.3548	0.338	0.662
$\mathrm{LimitFee}_t$ & $\mathrm{LateFee}_{t-1}$	0.922	1.010	0.1016	0.1113	0.110	0.890
$\mathrm{LimitFee}_t$ & $\mathrm{CashAdvFee}_{t-1}$	0.960	1.005	0.1058	0.1108	0.110	0.890

The first column of this table reports the probability of paying a fee of one type in period t conditional on having paid a fee of another type in period $t-1$, normalized by the unconditional probability of paying a fee of the first type in period t. The second column reports the (normalized) probability of paying a fee of type t conditional on not having paid a fee of another type in period $t-1$. The next two columns report the non-normalized versions of these two quantities, and the last two columns the unconditional probabilities of having paid and not having paid a fee, respectively.

Figure 1: Fee Frequency and Account Tenure

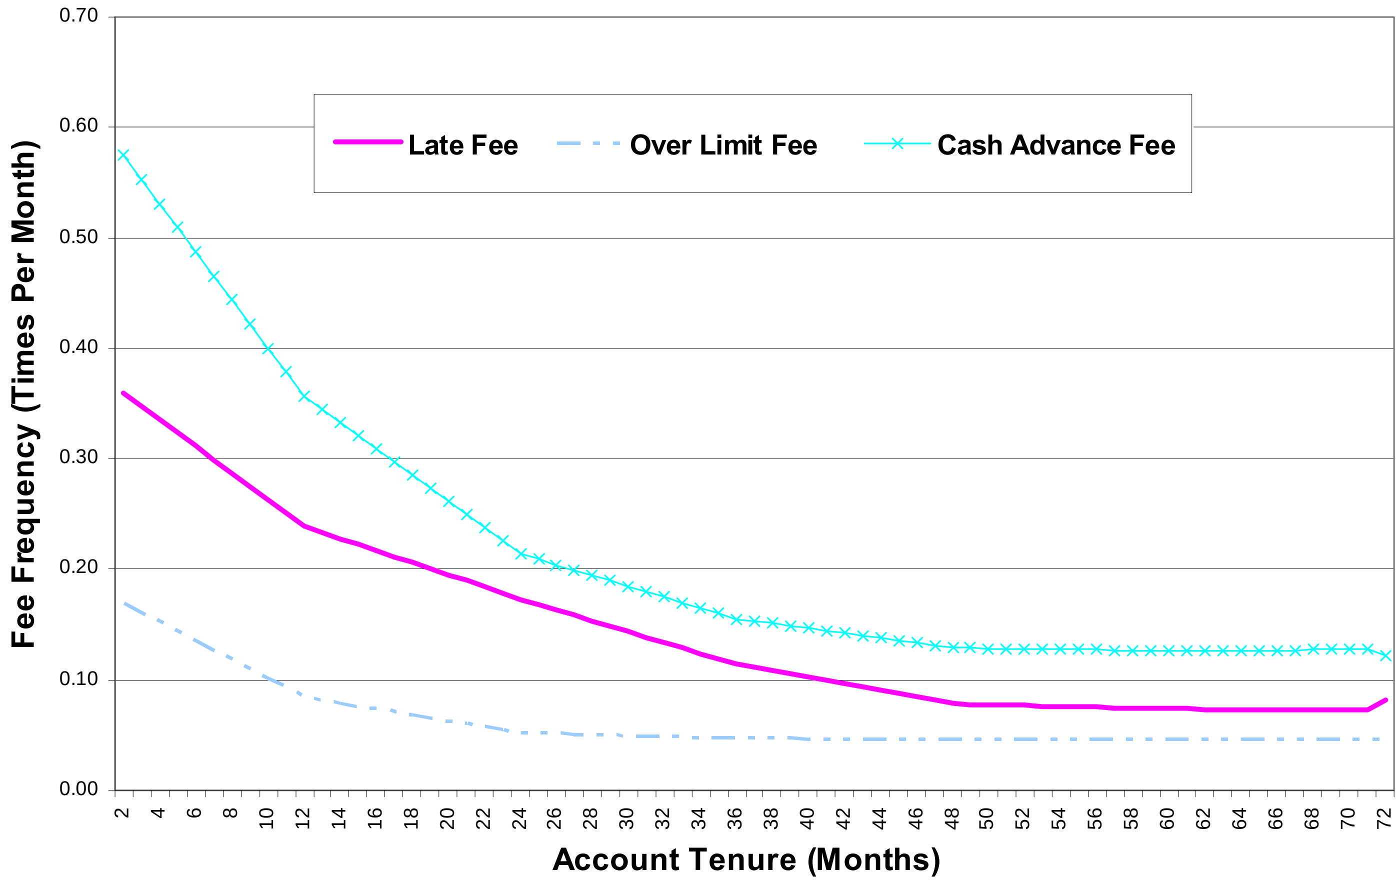

Notes: This figure plots the fitted values of regressions of fee frequency (times per month fees are paid) on a continuous piecewise linear function of account tenure (the function is a spline, with knots every twelve months, on the time since the account was opened), a constant, account- and time-fixed effects, and control variables (utilization (debit/limit), purchase amount, and dummy variables for any account activity this month and the existence of a bill last month). The intercept is computed by summing the constant with the product of the estimated coefficients on the control variables and their average values (the account and time-fixed effects sum to zero by construction). Tenure starts at the second month because account holders are, by definition, unable to pay late or over limit fees in their first month of account tenure.

Figure 2: Fee Value and Account Tenure

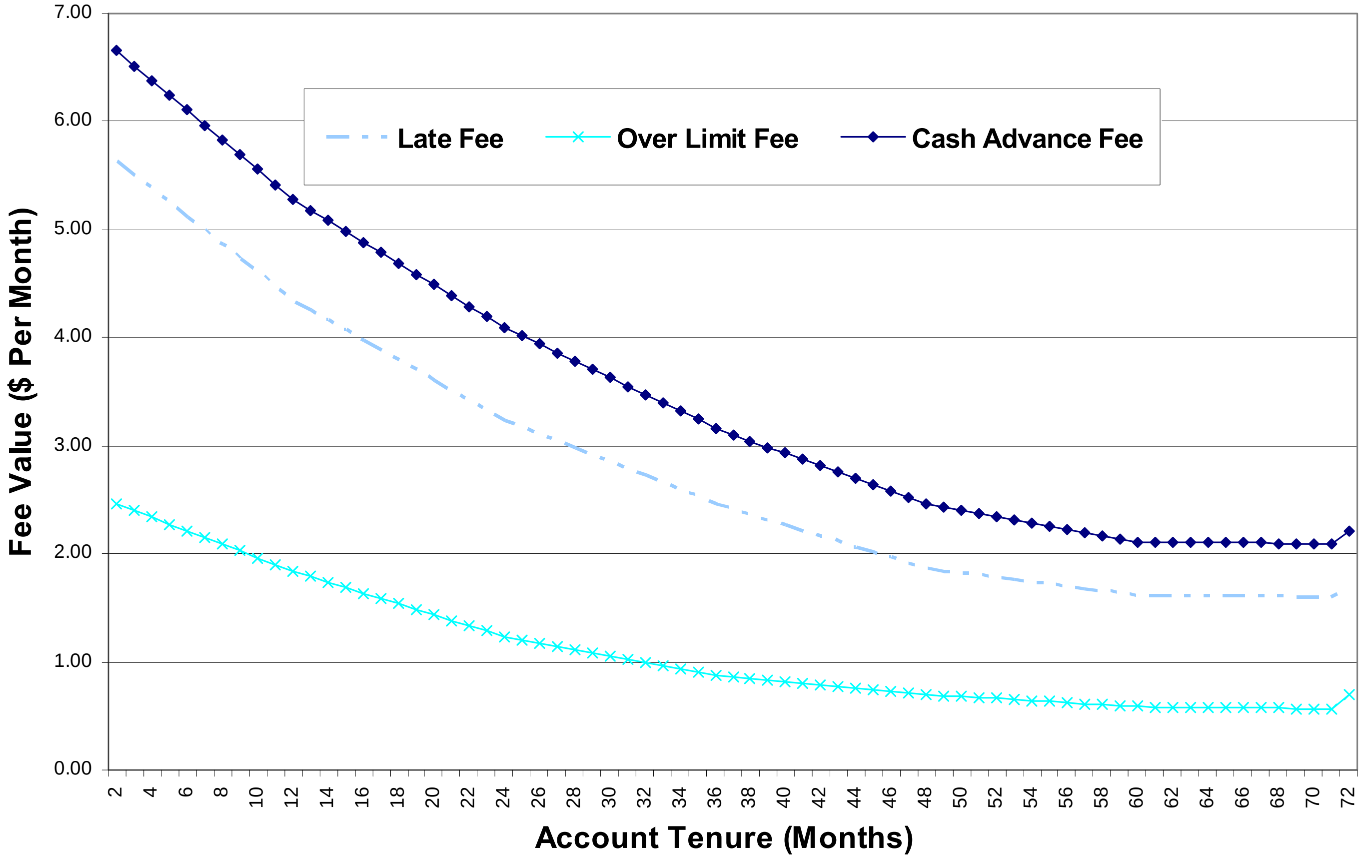

Notes: This figure plots the fitted values of regressions of fee value (dollars per month in fees paid) on a continuous piecewise linear function of account tenure (the function is a spline, with knots every twelve months, on the time since the account was opened), a constant, account- and time-fixed effects, and control variables (utilization (debit/limit), purchase amount, and dummy variables for any account activity this month and the existence of a bill last month). The intercept is computed by summing the constant with the product of the estimated coefficients on the control variables and their average values (the account and time-fixed effects sum to zero by construction). Tenure starts at the second month because account holders are, by definition, unable to pay late or over limit fees in their first month of account tenure.

Figure 3: Impact of Fees Paid k Months Ago on Fees Paid Now

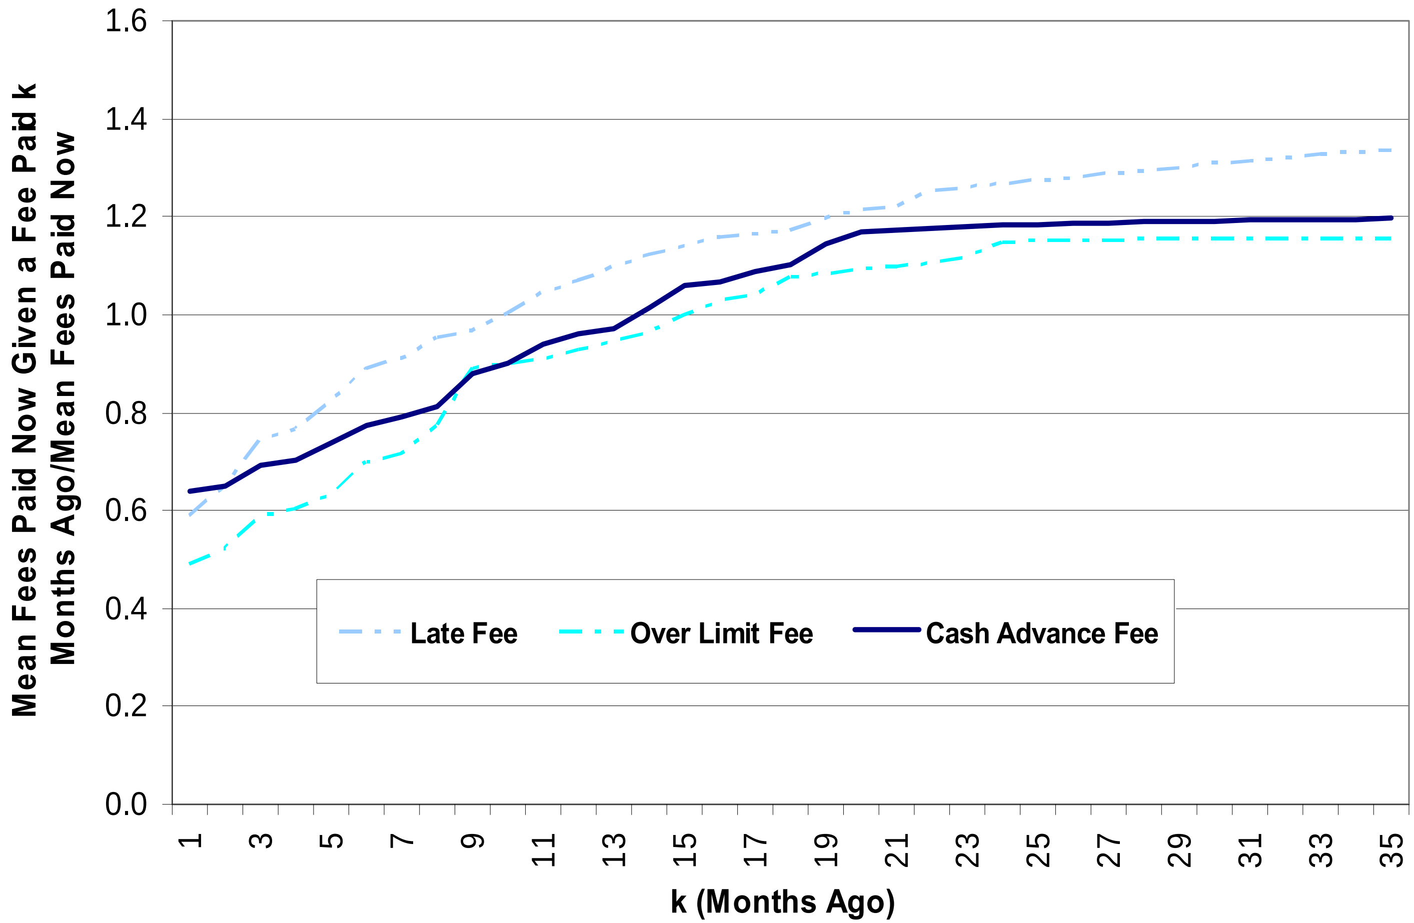

Notes: This figure plots $L_k=E(f_t|f_{t-k}=1)/E(f_t)$, the ratio of the conditional mean of fees f_t paid now given a fee was paid k months ago to the mean of fees paid now. If this value is 1, having paid a fee k months ago has no effect on current fee payment; if it is less than one, having paid a fee k months ago reduces current fee payment; if it is greater than one, it increases fee payment.

Figure 4: Demeaned and Normalized FICO Scores, Behavior Score, Purchases, and Utilization Rates

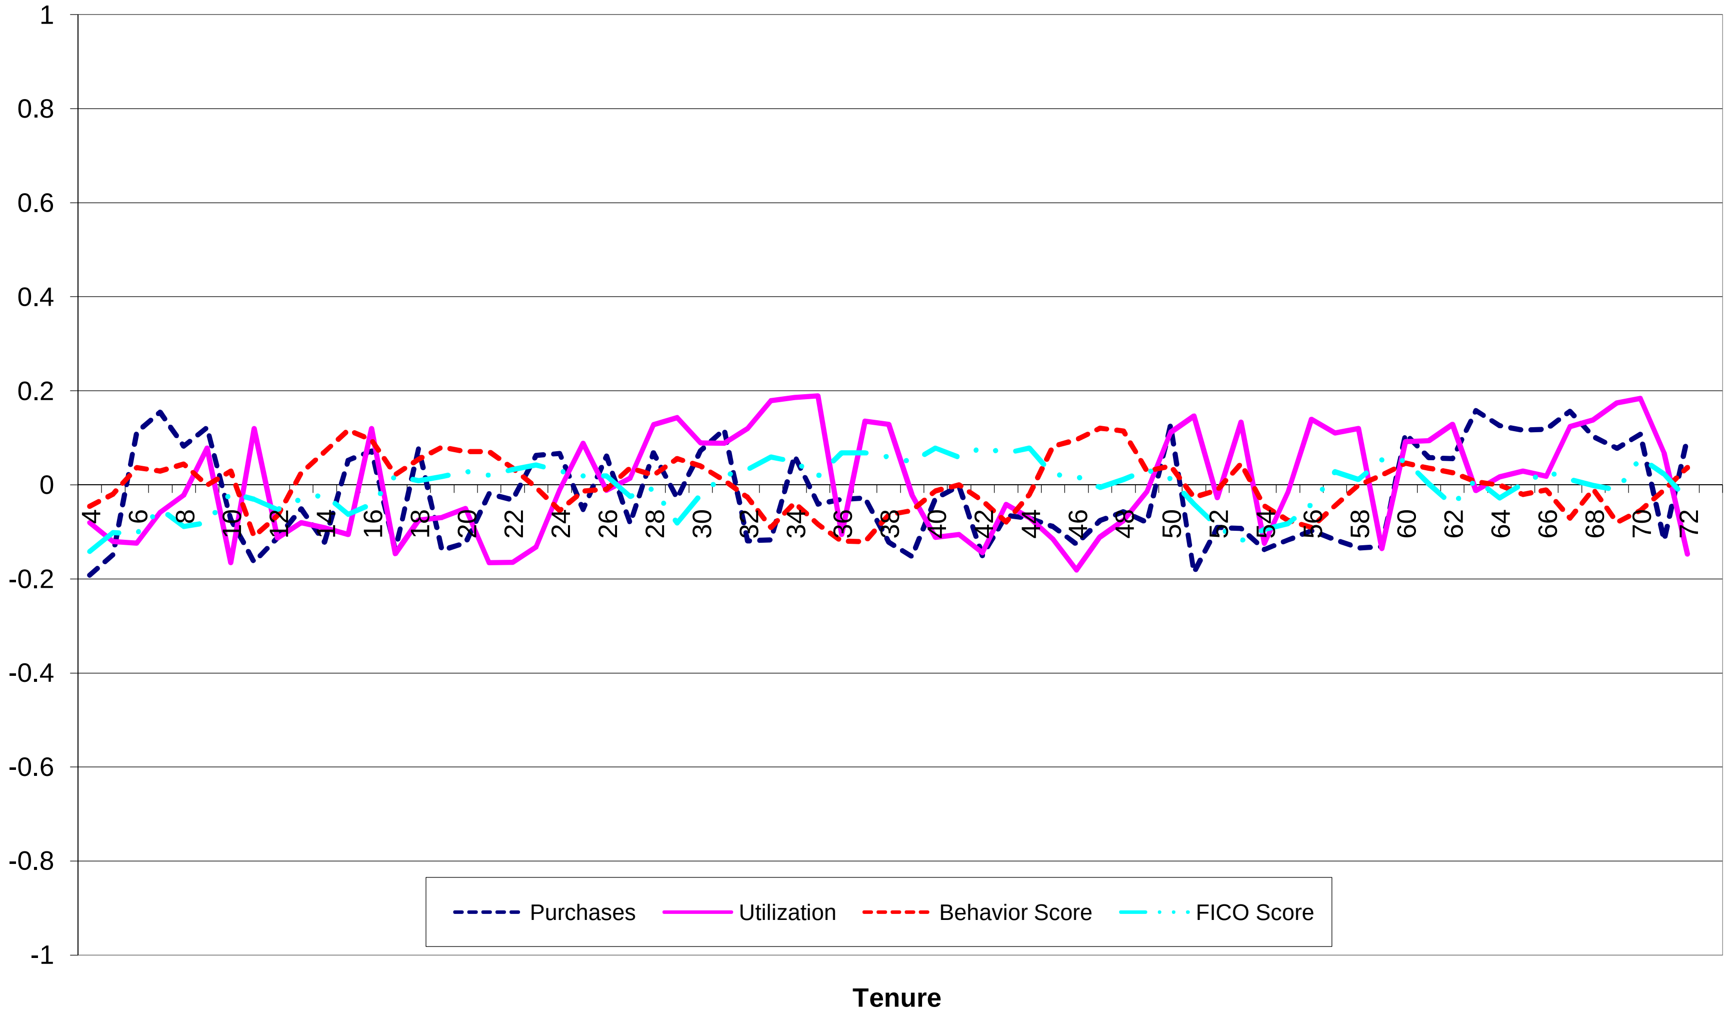

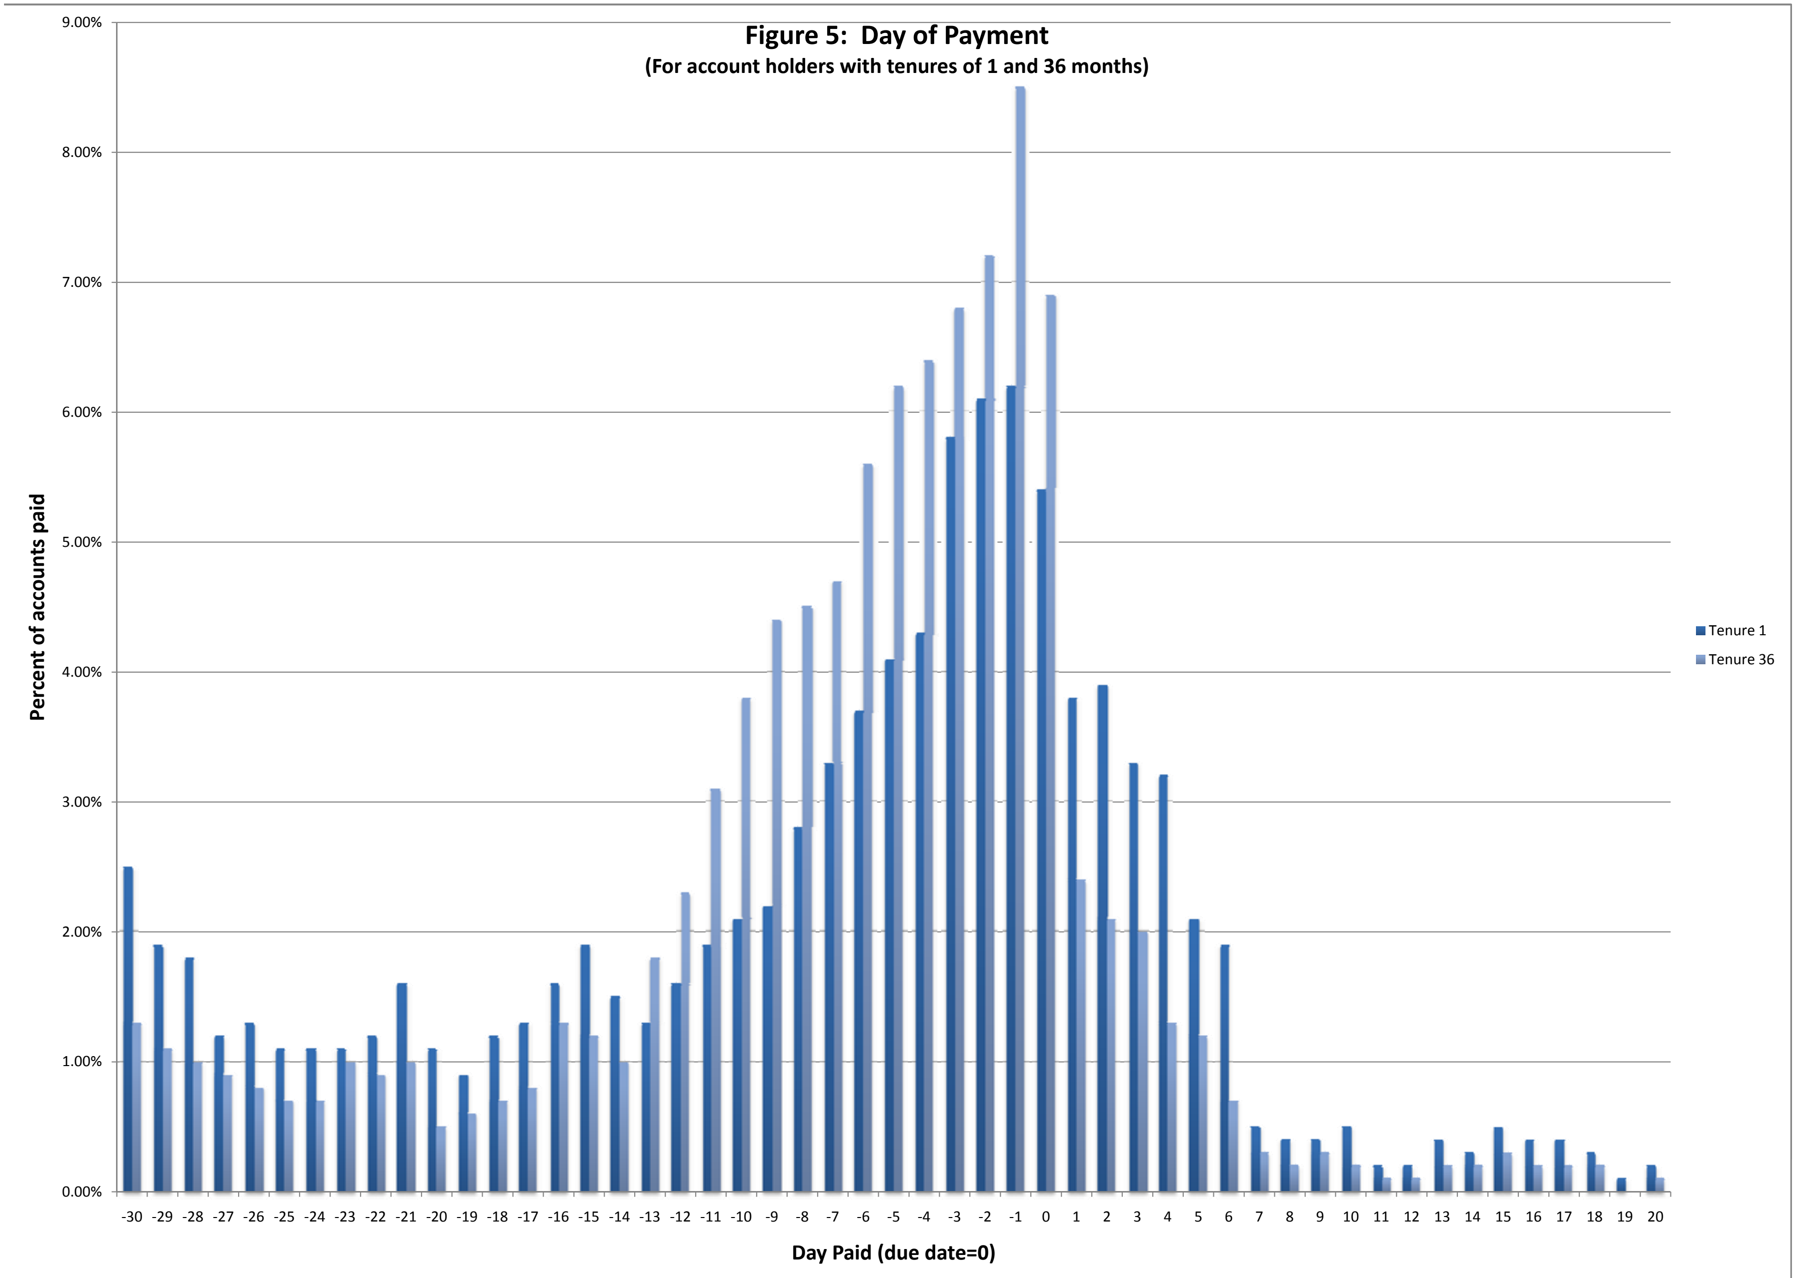

Notes: This figure plots the distribution of payment days. The due date is day zero; thus dates below or at zero reflect on-time payments, and dates above zero late payments.

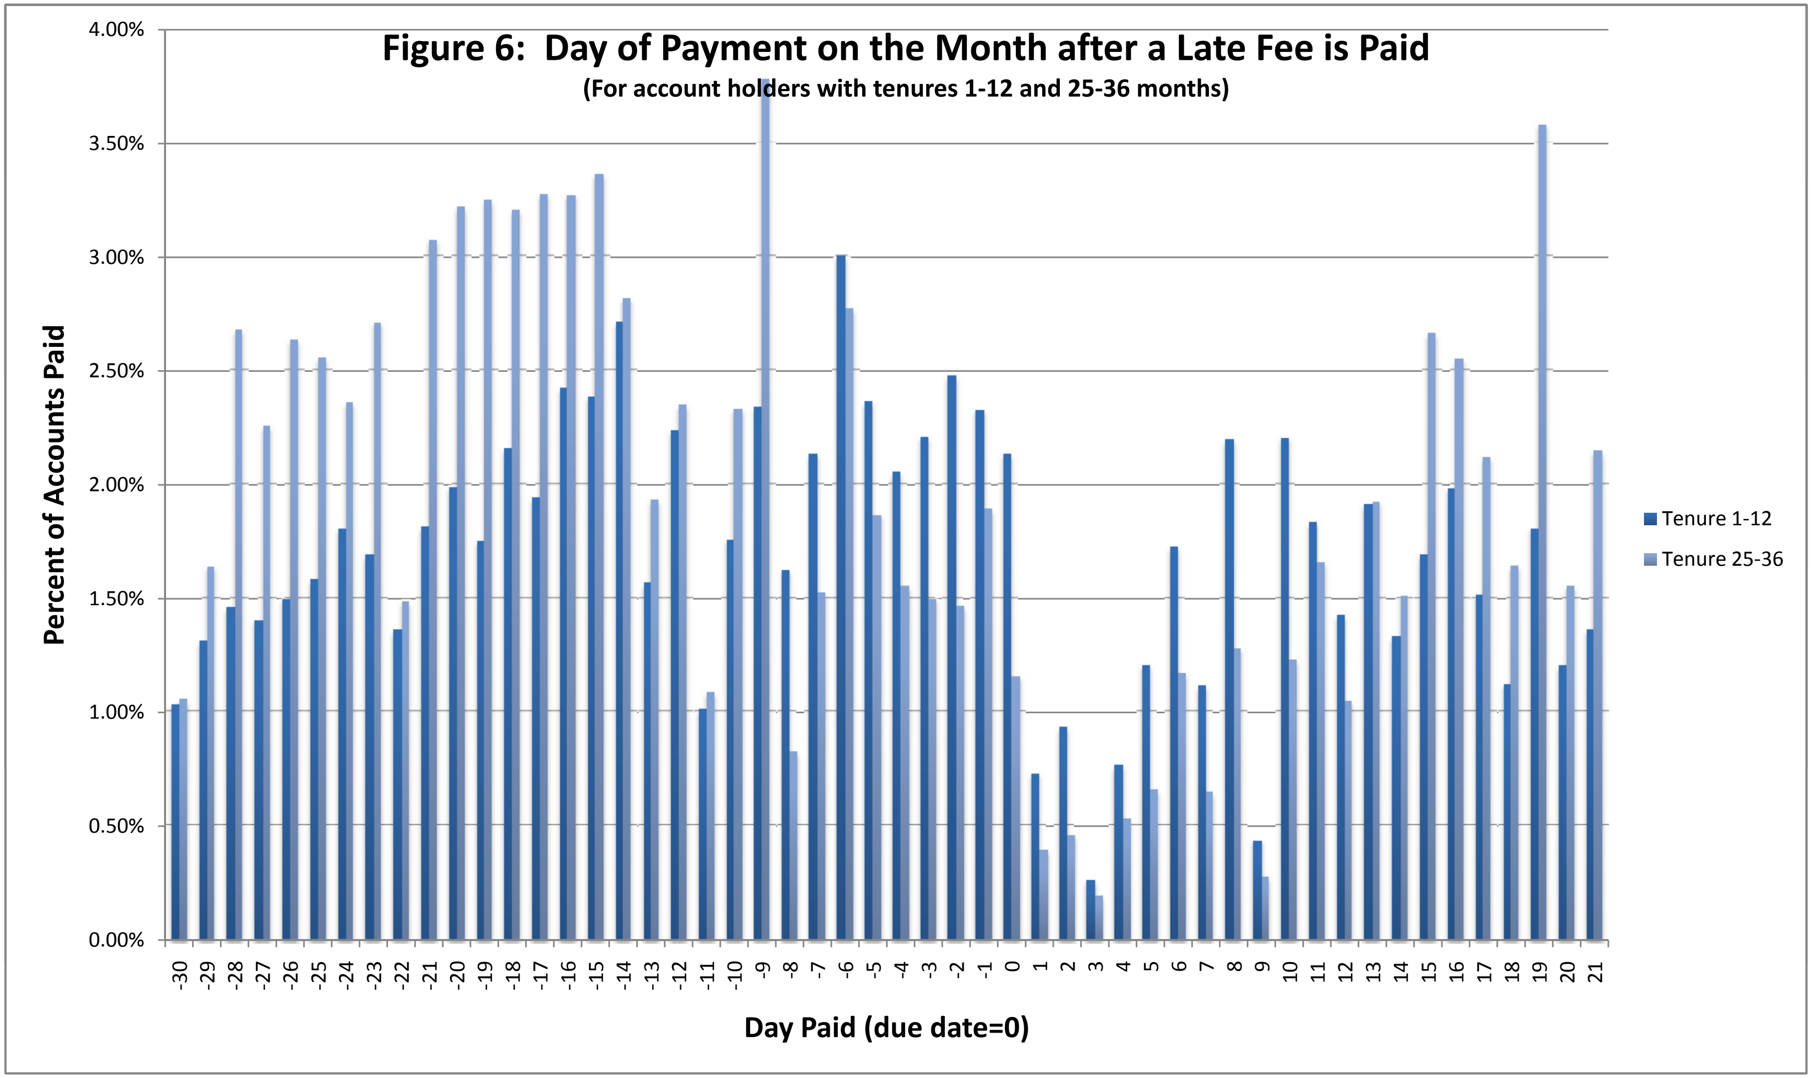

Notes: This figure plots the day of statement payment on the month following a month during which the account holder paid a late fee.
The due date is day zero; thus dates below or at zero reflect on-time payments, and dates above zero late payments.

Figure 7: Late Fee Frequency and Account Tenure, by Borrower Characteristic

Age Group

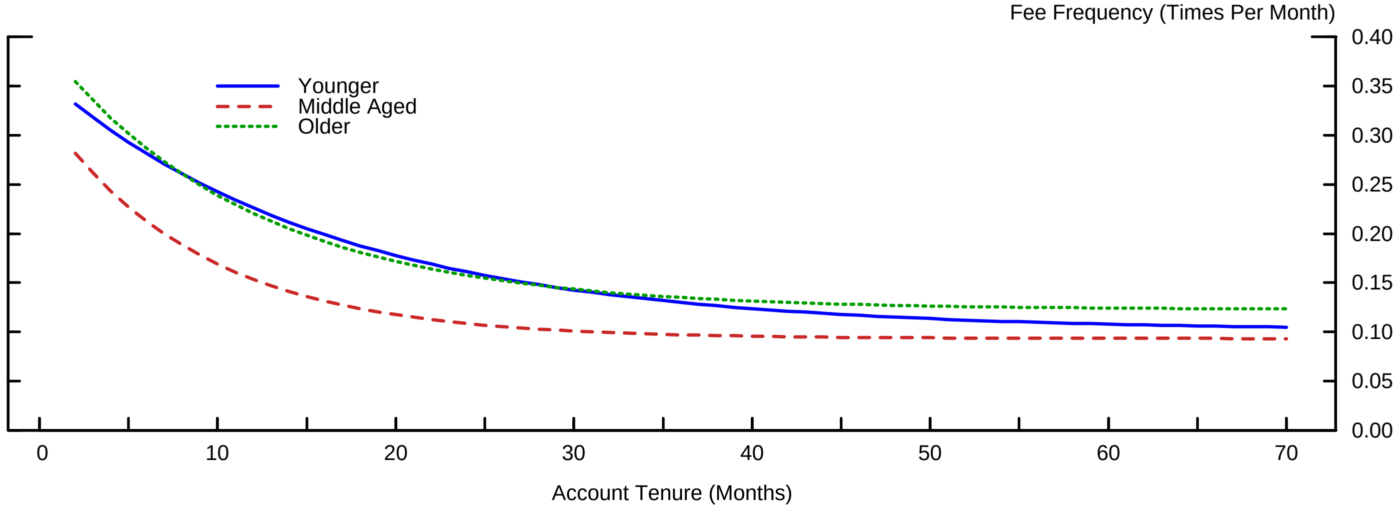

Gender

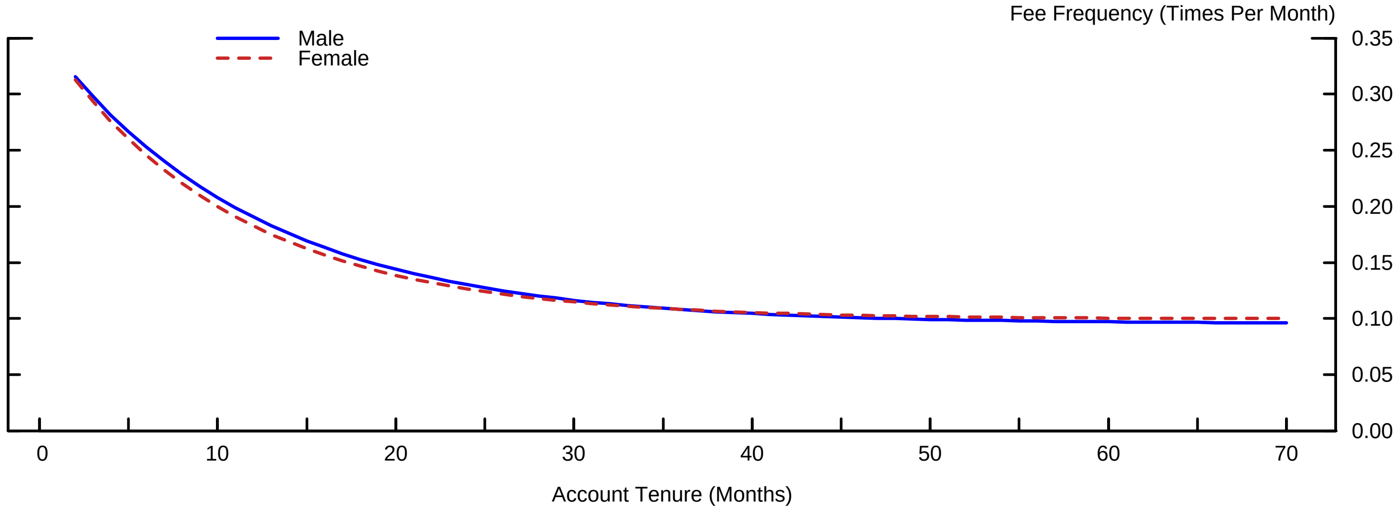

Martial Status

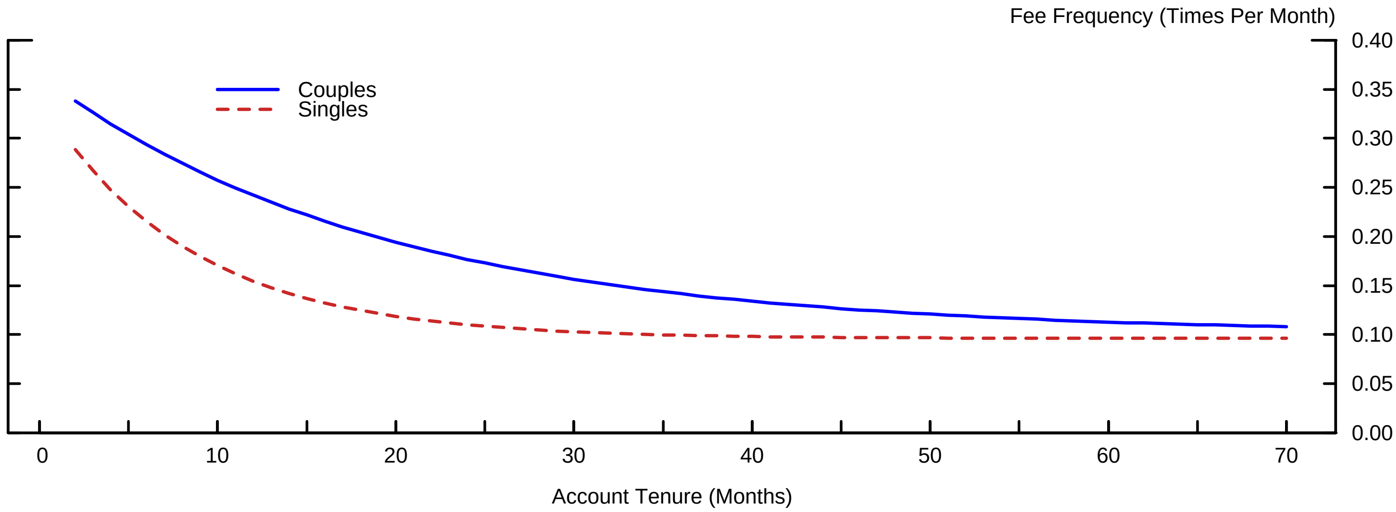

Figure 7: Late Fee Frequency and Account Tenure, by Borrower Characteristic

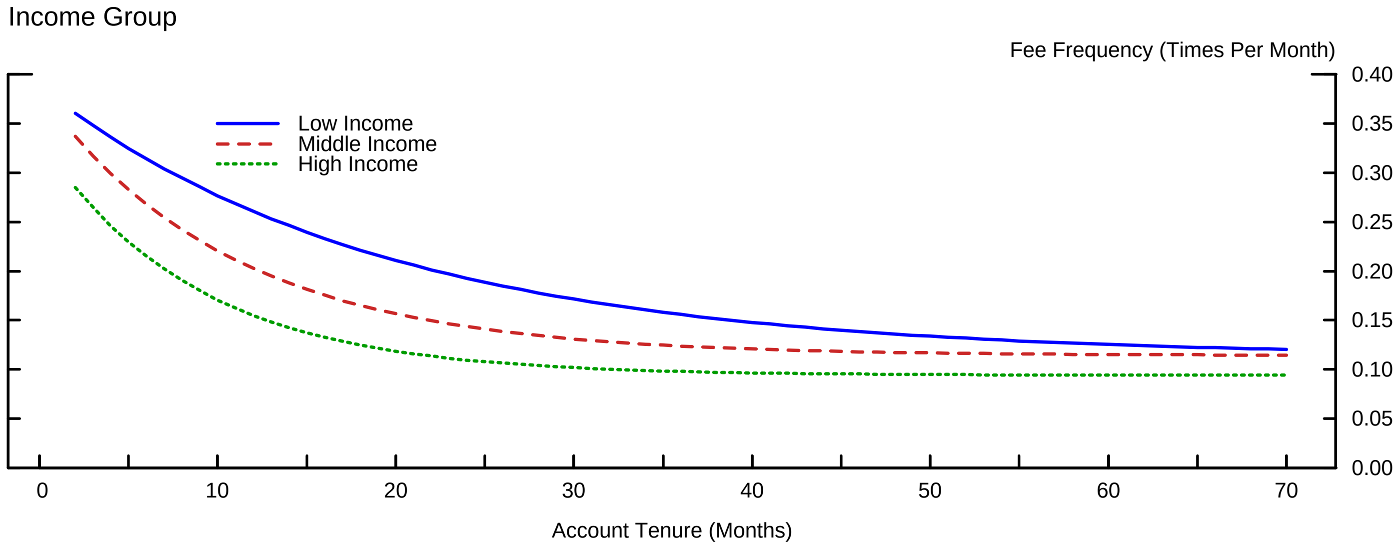

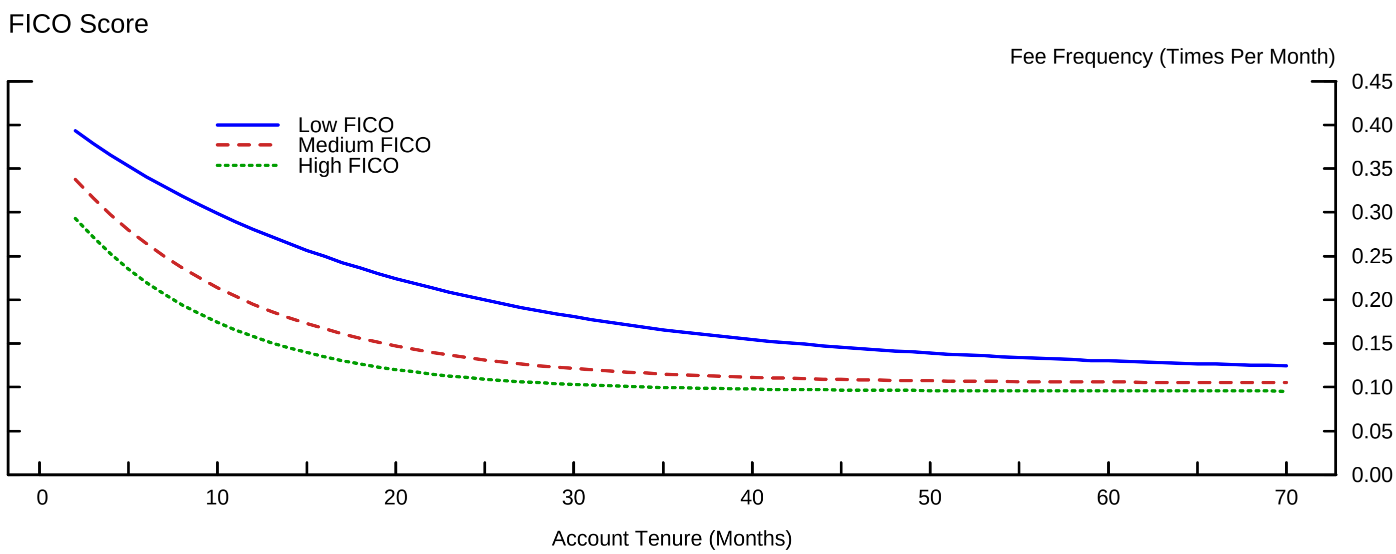

Figure 8: Impact of Late Fees Paid k Months Ago on Late Fees Paid Now by Borrower Characteristic

Age Group

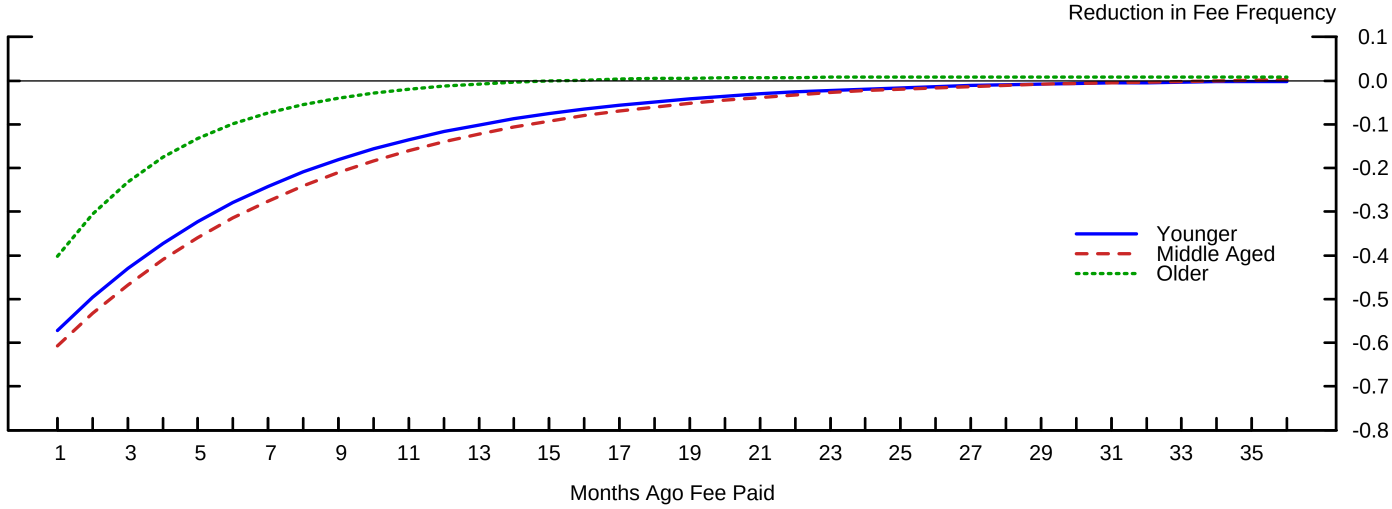

Gender

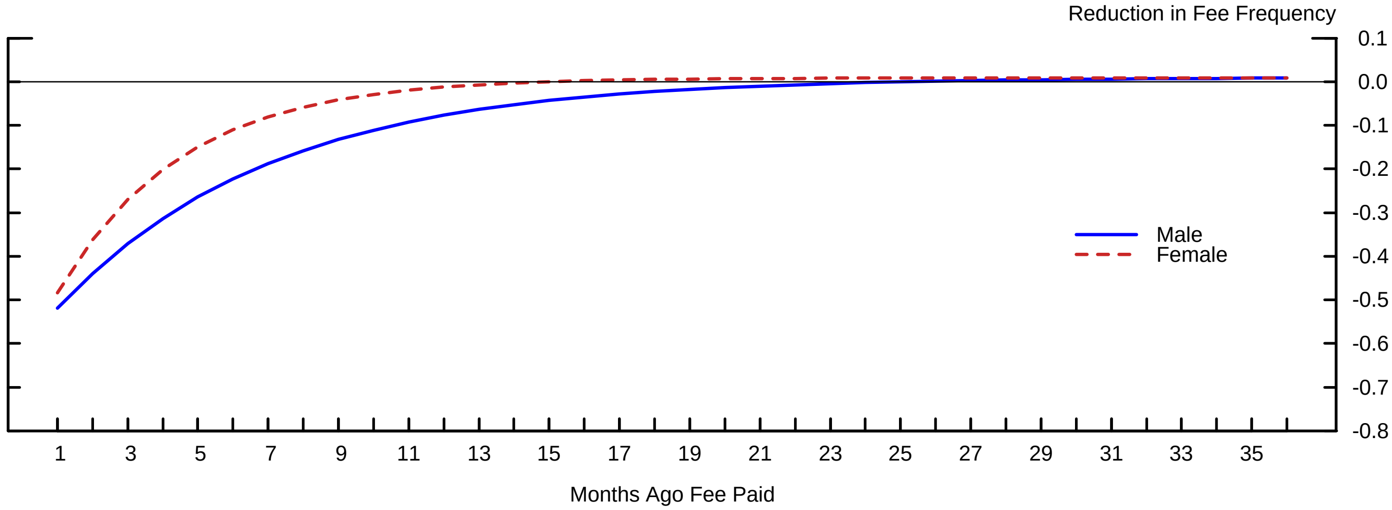

Martial Status

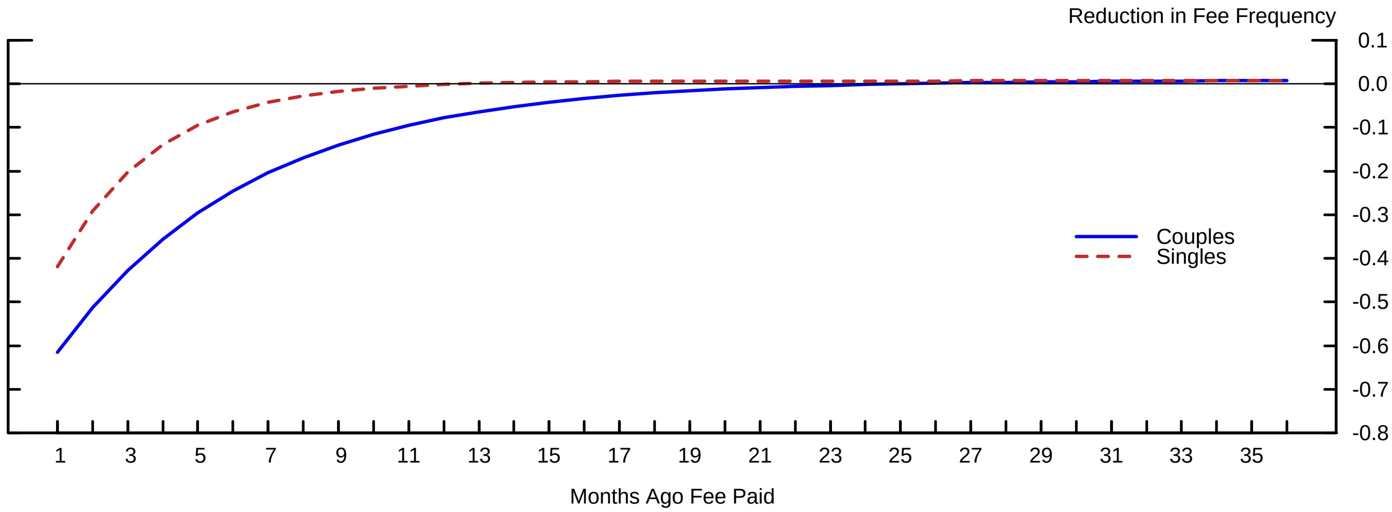

Figure 8: Impact of Late Fees Paid k Months Ago on Late Fees Paid Now by Borrower Characteristic

Income Group

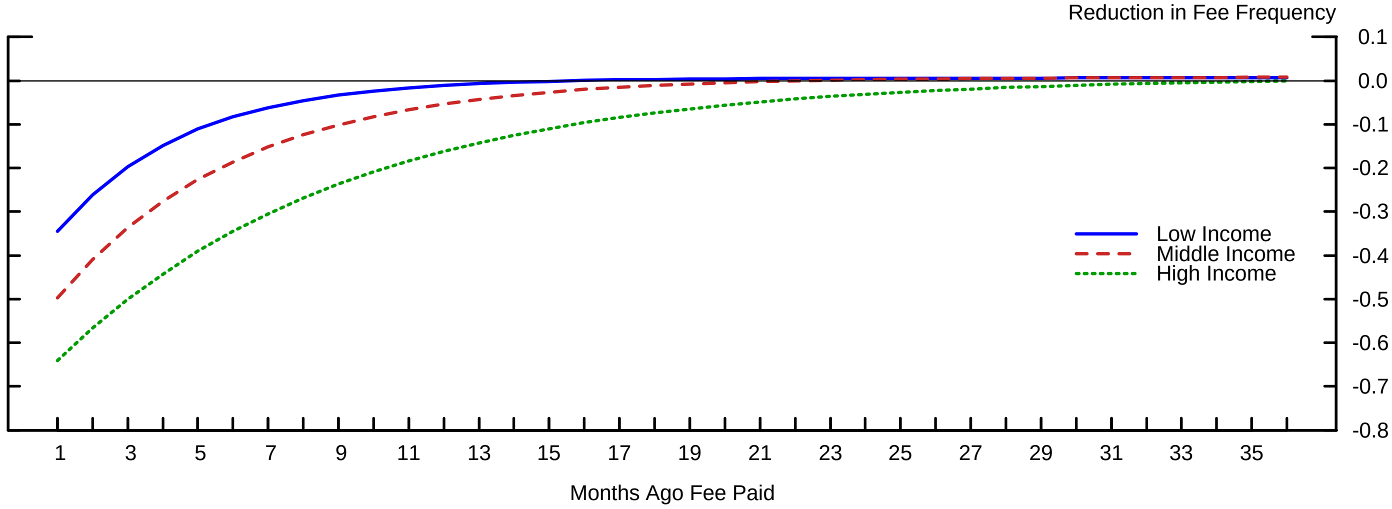

FICO Score