Borrowing High vs. Borrowing Higher:

Sources and Consequences of Dispersion in Individual Borrowing Costs[*]

Victor Stango   Jonathan Zinman

Graduate School of Management   Department of Economics

University of California, Davis   Dartmouth College and NBER

Preliminary and Incomplete: Do Not Cite

November 2011

For many households, paying lower borrowing costs is the surest, fastest way to increase net worth. Using administrative, credit bureau and survey data on U.S. credit cards, we find pervasive and systematic cross-individual variation in borrowing costs. Credit risk and product differentiation explain about one-third of that variation. The remaining risk-adjusted dispersion can materially affect wealth accumulation: moving heavy borrowers from the 75[th] to the 25[th] percentile of risk-adjusted borrowing costs increases their savings rates by more than a percentage point. Debt (mis)-allocation conditional on cards held could matter in principle, but appears to matter very little in practice, because most people allocate debt to their lowest-rate cards. Rather, similarly risky borrowers often hold cards with very different contract APRs. Heterogeneity in consumer search behavior appears to be an important factor in explaining that contract APR variation – a factor nearly as important as credit risk in explaining the cross section of borrowing costs.

I. Introduction

Research on household balance sheets, financial decision making and wealth accumulation often focuses on the asset side.[1] That work identifies substantial individual-level heterogeneity in risk-adjusted net returns – heterogeneity that materially affects how quickly people build net worth, even conditional on savings. It also identifies several sources of that heterogeneity, including portfolio inefficiency and differences in search/shopping or trading behavior.

For many households, however, borrowing costs and not asset yields are the more important determinant of wealth accumulation. Many more households hold debt than hold financial assets. Aggregate debt holdings are almost as large as aggregate financial asset holdings, and that was true even before the 2000s boom (Dynan 2009). Loan interest rates meet or exceed even risky asset returns; for example, historical credit card and auto interest rates exceed historical equity returns. So the marginal borrowing cost is the true opportunity cost of capital for many households, and efficiently minimizing borrowing costs is a critical determinant of wealth accumulation.

Despite the importance of borrowing costs for wealth accumulation, research on the topic is underdeveloped. Data limitations have made it difficult to answer the threshold question: does the individual-level heterogeneity observed in risk-adjusted asset returns also exist in borrowing costs? Beyond that, we also know relatively little about the determinants of individual-level variation in borrowing costs. There is some evidence showing that search and switching costs prevent customers for shopping perfectly,[2] but relatively little work mapping those costs into individual-level differences in borrowing costs. There is a growing literature on (mis)allocation across debt instruments but no consensus on how widespread or important misallocation is, in large part because most datasets do not cover enough of the liability side in the requisite detail.[3]

---

[1] See, e.g., Campbell (2006) and Barber and Odean (2011) for reviews.
[2] See, e.g., Agarwal et al (2006); Berlin and Mester (2004); Calem et al (2006); Charles et al (2008); Crook (2002); Kerr and Dunn (2008); Shui and Ausubel (2004); Stango (2002); Woodward and Hall (2010).
[3] See, e.g., Agarwal, Skiba, and Tobacman (2009); Amar et al (2011); Ponce et al (2009); Stango and Zinman (2009a).There is related work on fee avoidance in credit cards and bank accounts, suggesting that many of those borrowing costs are incurred due to limited information, memory, and/or attention (Sumit Agarwal et al. 2011; Stango and Zinman 2011). There is also work on (mis)allocation across (liquid) assets and liabilities, although heterogeneity in the non-price

We contribute to the debt-side literature on wealth accumulation by assembling a more comprehensive picture of borrowing costs, individual heterogeneity and sources of that heterogeneity in the $800 billion credit card market. We document that credit-risk-adjusted individual-level borrowing costs vary enough to generate substantive cross-sectional differences in household wealth accumulation, conditional on amounts borrowed. We then decompose that variation, to identify what consumer behavior is most important in explaining variation in borrowing costs. Our data span multiple accounts per individual and come from credit card statements, credit reports, and survey responses.

Our first finding is that borrowing costs vary substantially across households, even for the relatively homogeneous product (credit cards) we examine. The raw difference in borrowing costs between the 25[th] and 75[th] percentiles of credit card borrowers is 700-800 basis points. We measure borrowing costs as an individual-level balance-weighted average annual percentage rate (APR), across all credit cards and days with revolving balances. We focus on APRs and largely ignore fees, despite the recent surge in research and policy interest in fees, because fees account for only 6% of borrowing costs in our data.[4] We find balance-weighted APR dispersion *within* every part of the credit card debt distribution: not much price dispersion comes from differences between light and heavy borrowers. Among heavier borrowers, dispersion can explain differences in household saving rates of 1-2 percentage points.[5]

Substantial cross-sectional variation in borrowing costs remains when we adjust for observable individual-level and time-varying credit risk,[6] a variety of measurement issues, and product differentiation across cards. A credit score from one of the major bureaus explains less than one quarter of cross-sectional variation in weighted APRs. Adding measures of time-varying risk - within-sample late fees, over-limit fees, and credit utilization rates - accounts for another 10% or so of weighted APR dispersion. Product differentiation in credit card contracts (low introductory rates, fee/APR tradeoffs,

---

characteristics of the different securities complicates that analysis; see, e.g., Gross and Souleles (2002); Telyukova (2011); Zinman (2007).

[4] [benchmark to national data]

[5] An analogous exercise on the asset side of the balance sheet would be to measure risk-unadjusted yields relative to income, taking the amount of assets as exogenous.

[6] The average within-person, year-to-year correlation in credit scores is close to one [Karlan and Zinman in progress; Payment Cards Center 2011].

card rewards, etc.) contributes modestly to cross-household price dispersion but leaves the qualitative finding of large variation unaffected. Under our most conservative set of assumptions we still find interquartile ranges of 500+ basis points in weighted risk-adjusted APRs between observably similar borrowers. Among similarly risky borrowers, someone at the 75[th] percentile of borrowing costs could reduce interest costs by roughly 33% simply by paying an APR at the 25[th] percentile.

What leads card users to pay substantively different prices for an essentially homogeneous product? In the spirit of prior work on both assets and liabilities, we consider two explanations. One is is heterogeneity in how individuals allocate debt across cards, conditional on the contract APRs held. The other is heterogeneous shopping behavior that leads to different contract APRs across similar people. Both explanations are plausible given that we observe substantial contract APR variation both across and within individuals, and given prior research on how people use credit cards.

Differences due to debt allocation explain relatively little cross-sectional variation in borrowing costs. We do find some evidence of behavior consistent with the "debt puzzles" identified by other work – using more expensive cards when cheaper cards are in the wallet, or repaying cheaper debt before costlier debt – but such behavior is far from the norm. Even under an extremely generous set of assumptions—no switching costs (e.g., no balance transfer fees), and no non-price differences across-cards-within-household (e.g., from rewards)—the median cardholder leaves zero money on the table *annually* via misallocation, and fewer than 15% of our credit card users leave more than $100 on the table annually. Our results on debt misallocation dovetail with those of Campbell, Calvet and Sodini (2007), who find that asset allocation mistakes have substantive costs for only a few investors. They stand in contrast to the findings of Ponce et al. (2009), who find more misallocation in a sample of credit card users from Mexico.

We then examine the role of search/shopping behavior in driving cross-sectional dispersion in borrowing costs, using several complementary approaches. First, for a sub-sample of panelists (fewer than 500 of 4300), we observe a direct self-reported measure of how intensively a panelist "keeps an eye out for better credit card offers." In a cross-sectional regression controlling for credit score, time-varying risk, credit utilization, debt demand, demographics and other household characteristics, we confirm that shopping

behavior explains meaningful variation in borrowing costs – this despite the fact that the small sub-sample size and the likely measurement error in search behavior would attenuate our ability to make such a statement. We find the same relationships in card-month- or panelist-month-level regressions that also control for issuer, month/year effects, card features (rewards, fees, fixed/variable pricing, etc.) and card-level default behavior.

We then develop a complementary but more broadly measured variable capturing shopping behavior: the number of credit cards held, conditional on all other observable characteristics including total credit limit, credit score, time-varying risk, credit utilization, debt demand, demographics and other household characteristics. This is potentially a more precise measure, because it measures actual search/switching behavior rather than a self-reported preference. And, it could be more impactful in equilibrium, because it is directly observable by issuers: there is anecdotal evidence that issuers use actual search/switching behavior to price discriminate, offering lower APRs to cardholders who shop more intensively. For an issuer, how many cards the cardholder has "in the wallet" is a good signal of such shopping behavior.

To address the concern that cards held might be correlated with omitted variables unrelated to search behavior, we do two things. First, we show that cards held are significantly correlated with APRs (weighted and contract) in the sub-sample of panelists for whom we observe the self-reported search behavior variable. Second, we show that when one instruments for cards held with the directly measured search variable, in the subsample where we measure search directly, the estimated effect is *larger* in point terms. That result suggests that an OLS estimate does not overstate the true search-related relationship between cards held and APRs. And again, we measure cards held conditional on many other variables, most notably total credit line and utilization ratio.

We then estimate the OLS relationship between cards held and APR in the full sample. The results suggest a strong link between search behavior and APRs paid. The average weighted APR falls substantially (an estimated 50-100 bp) with each additional card held, conditional on our rich set of other borrower characteristics, usage, timing, and other account terms. To put the estimated importance of search in perspective, otherwise similar individuals who differ only in holding one vs. five cards pay borrowing costs

more different than individuals in the worst vs. best credit score decile. In all, this suggests that search behavior can matter nearly as much as observable credit risk in explaining how borrowing costs differ across individuals.

Further support for the importance of search behavior comes from the link between APRs held/paid and other panelist characteristics. Income is positively correlated with APRs, opposite to the prediction of a credit-scoring model but in line with what would see in a search model where income is positively correlated with the opportunity cost of time. We see a similar pattern with age, which has an inverse u-shaped relationship with APRs (i.e., APRs are highest for those in their 40s and 50s). Both purchase volume and revolving balances are *negatively* correlated with APRs, as one would expect if those variables were correlated with the benefits and/or costs of comparison shopping, and contrary to what a default risk story would suggest. While there are certainly confounding factors in interpreting these other coefficients, the overall pattern is consistent with an important role for search behavior.

Our finding that search behavior strongly correlates with outcomes adds to the literature on credit card search, which to date has been limited to using aggregate or survey data, which has not attempted to tie individual-level search behavior to dispersion in borrowing costs, and which has found mixed evidence regarding the importance of search behavior (see cites in footnote 2). We also add to the broader literature on credit shopping.[7]

These results help to frame policy discussions, particularly to the extent that consumer protection policy focuses, due to limited resources, on *either* helping people manage the accounts they have, or helping people to choose accounts in the first place. Our results suggest that the latter tack may be more impactful; this is in line with some related work by Agarwal et al. [2006] on credit card contract choices, where consumers are trading higher APRs for lower annual fees. Our results also have implications for policy intended to help individuals improve their creditworthiness: if credit shopping is more malleable (or remediable) than creditworthiness (which, as we discuss below, is

---

[7] Recent nationally representative survey evidence also finds substantial heterogeneity in "shopping for borrowing" (Bucks et al. 2009, A12), and for credit cards (FINRA Investor Education Foundation 2009). About half of credit card users report that they don't comparison shop for cards at all. See also cites in footnote 2.

quite sticky), then focusing on helping people shop for cards may be more socially cost-effective. This harks back to early work on consumer protection in debt markets, which typically focused on improving comparison shopping.[8]

Our results also inform the literature on "over-borrowing" and "under-saving." Calibrated life-cycle models struggle to make sense of high credit card debt levels, whether discounting is specified as exponential (Carroll 2001) or quasi-hyperbolic (Angeletos et al. 2001). Explanations for over-borrowing tend to emphasize the underestimation of borrowing likelihood and costs, and/or or the temptation of consuming out of liquidity.[9] Our results suggest a complementary explanation: people (over-)consume leisure rather than shop for credit card offers. Differences in wealth accumulation due to "over-paying" via higher-than-necessary borrowing costs may be as substantial as differences in wealth accumulation from "over-borrowing" per se.

Our findings generally dovetail with prior work on the asset side of the balance sheet. Several papers on the asset side find that search/shopping costs are substantial, and heterogeneous, enough to generate large differences in net asset returns (Choi, Laibson, and Madrian 2010; Hortascu and Syverson 2004; Sirri and Tufano 1998).

We highlight some caveats, related to external validity, before proceeding to the meat of the paper. Our data come from an online panel that is not necessarily representative on some dimensions: our sample is younger, higher-income, and better-educated than average [(although our distribution of credit scores and credit card usage is similar to the U.S. overall)]. But there is little reason to believe that our panelists have higher or more diffuse search, switch, or transaction costs than many populations of interest. Our sample period, 2006-2008, is also unusual in historical perspective. It is not clear how the phenomena we document would be different during other times before or since.

---

[8] See e.g., National Commission on Consumer Finance (1972).
[9] See, e.g., Ausubel (1991) Gabaix and Laibson (2006); Heidhues and Koszegi (2010); Laibson (1997); Meier and Sprenger (2010); Shui and Ausubel (2004); Soll et al (2011); Stango and Zinman (2009b).

II. Data and Sample Characteristics

A. Our Data and Sample Characteristics

Our data come from Lightspeed Research (formerly Forrester Research). Panelists in our sample are members of the "Ultimate Consumer Panel," which is one of many such panels maintained by Forrester/Lightspeed.[10]

Panelists enter the Ultimate sample by providing Lightspeed with access to at least two online accounts (checking, credit card, savings, loan or time deposit) held within the household. Panelists have typically participated in other Forrester/Lightspeed panels; the incremental payment for enrolling in the Ultimate panel averages $20. After initial enrollment, panelists need take no action to maintain membership in the panel, and a panelist may request to leave the panel at any time. Enrollment of new panelists occurs consistently throughout our sample period, as Lightspeed attempts to keep panel size constant by balancing enrollment against attrition.

The credit card data collected by Lightspeed have two main components. The first component is transaction-level and comes from monthly credit card account statements. The statements contain every accounting debit and credit to the account: purchases, transfers, fees, interest charges, payments, and so on. The second component is account-month level and contains data on terms: APR, cash advance APR, bill date, due date, ending balance on bill date, summaries of credits and debits in the month, and so on.[11]

In addition to the transaction/account data, Lightspeed collects survey data from panelists. All panelists complete a short online registration survey when they sign-up for the panel; this gives us some baseline information on demographics, financial characteristics, and respondent-assessed financial literacy. Once in the panel, panelists are then invited to take online surveys that are offered periodically. Finally, for a subset of panelists, Forrester procured credit report data from one of the major bureaus, around

---

[10] Other Forrester/Lightspeed panels track consumer behavior of interest to market researchers, such as the use and purchases of new technology. Those panels are widely used by industry researchers and academics; see, e..g, Goolsbee (2000; 2001), Kolko (2010), and Prince (2008).

[11] Some accounts may have more than one APR at once due to balance transfers, which can have different APRs than new purchases. We observe only one purchase APR per account; it is not clear whether this is the only APR.

the time of the panelist's registration. The credit report data include the number of "trades" (current and past loans of all kinds), as well as a credit score on the standard 850-point scale.

Table 1 summarizes the data. Our data span 2006-2008, and in this paper we use data from the 4,312 panelists who enroll at least one credit card account and for whom we observe a credit report. The mean number of enrolled accounts is two (top and bottom sets of rows in Table 1), with greater borrowing associated with holding more cards. Comparing accounts (cards held) to a baseline survey question and to credit bureau data, we estimate that we have the complete set of credit card accounts (at least initially) for roughly half of our panelists.[12] Panelists remain in the data for about 20 months on average.

The second and third sets of rows in Table 1 summarize monthly spending and borrowing figures. We measure these as panelist-level averages for the sample period, across all days on which the panelist is present in the data. We stratify panelists by their quartile of average revolving debt throughout the paper, in part to facilitate analysis that conditions on demand for debt, and in part to help identify the incidence of the impacts of dispersion in household borrowing costs. Our sample has a relative dearth of pure "convenience users": fewer than 5% of the sample always "float" and never revolve debt, although many panelists revolve balances only infrequently.

The next rows show annualized interest costs, which are substantial in the groups with heavier credit card borrowing. The 3rd quartile of average revolving balances has a median annualized interest cost of $436, and the interquartile range is [$321, $627]. In the 4th quartile median interest costs are $1633 per year (roughly $133 per month). The 75th percentile is $2480 per year, and the 90th percentile of interest costs is $4077 per year or $340 per month. Relative to all fees, interest costs are by far the dominant component, 94%, of overall account costs, as the next set of rows show; see Table A1 for summary

_____

[12] The credit bureau report contains two measures of credit cardholding: the number of "open bankcard trades," and the number of "active bankcard trades." For 30% of our panelists the number of accounts in our data matches (or in rare cases exceeds) the number of open trades, and for 60% the number matches/exceeds the number of active trades.

data on the individual fee components (annual, late payment, over-limit, cash advance and balance transfer).

To place these figures in the context of wealth accumulation, the next rows in Table 1 show how interest costs compare to annual income. This ratio is essentially the contribution of credit card interest costs to net household savings rates. Income is self-reported and categorical, and reported pre-tax rather than after-tax, so it is best to view these figures as approximate (and a bit on the conservative side, because we take the upper end of each category as the value of income). Nonetheless one can see quite clearly that credit card interest costs can substantially affect household net savings rates, for heavy and even moderate levels of debt. To frame the numbers, consider that over the period 1990-2010 household savings rates in the United States averaged roughly 4.5%; during our sample period (2006-2008) they were much lower, although they have risen substantially since then.

The last rows show credit scores and simple demographics. Compared to national averages, our panelists are younger, more educated, and higher income (especially conditional on age). The overall credit score distribution looks representative, conditional on our observable demographics.[13]

Credit scores, income and education follow a somewhat U-shaped pattern with respect to revolving balances; the heaviest borrowers are as well-educated and earn income comparable to the lightest borrowers, and if one accounts for the fact that debt *per se* reduces credit scores, those in the fourth quartile are similarly creditworthy to those in the first quartile. We discuss this later, but it is also the case that credit utilization (total balances divided by total limit across all cards) varies across groups: generally speaking those who borrow more have higher utilization. However, even heavy borrowers have substantial access to liquidity – the 25[th] percentile of "minimum available credit on any one day during the sample" is over $4500 for borrowers in quartile 4, and the median is over $7000. Most borrowers in the sample have plenty of available credit, although we we show below, APRs are strongly correlated with credit utilization rates in the cross section.

---

[13] [Benchmark using PCC data].

Perhaps the most noteworthy overall pattern in Table 1 is the substantial heterogeneity, both within and across revolving quartiles, in every variable. "Who borrows?" and "At what cost?" are not easily explained by observable individual characteristics.

B. Representativeness? Comparisons to Other Data Sources

Here we discuss how some of our key variables (in Table 1, and the top distribution of APR contract rates in Table 3) compare to other data sources, principally the Survey of Consumer Finances, nationally representative samples of credit bureau data (as analyzed by the Federal Reserve Banks of New York (Brown et al. 2011) and Philadelphia (tabulations provided to the authors by the Payments Card Center), and a public comment from the industry that reports on issuers covering an estimated 70% of outstanding balances (Morrison & Forrester LLP 2008). We discuss rather than report tabs from these other data sources to conserve space in our tables, but the other tabulations are available upon request.

Starting at the top of Table 1, our cardholding distribution matches up well with all other available sources: the SCF (which for this variable matches up well with issuer-side data (Zinman 2009)), and the Philadelphia Fed. We see slightly significantly more mass at one card (not surprising since our panelists can enroll a checking and credit card account instead of two credit card accounts), a bit more mass in the right tail, and hence a bit less mass in the middle (2-3 cards). Our purchases distribution also matches up well with the SCF. Zinman [2009] shows that this SCF variable matches up well with issuer-side data. The comparable 25th, 50th, and 75th percentiles of the weighted SCF are $20, $250, and $1000.

Comparisons of revolving debt are more problematic, given substantial underreporting in the SCF (Zinman 2009, Brown et al 2011), and the lack of distinction between revolving and transaction balances in credit bureau data (and in the data that issuers report to regulators). But if we look simply at outstanding balances, we see about 50% less in our data than in the bureau (Brown et al Appendix Table 1). This may be explained in part by the life-cycle u-shaped pattern of credit card debt (Brown et al Figure 4), coupled with the fact that our sample is relatively young.

Other data on APR distributions is limited, but comparing our top rows in Table 3 to the SCF (which asks about a single APR, on the card used most often), we find similar dispersion; the interquartile range in the SCF is 900 basis points, which is comparable to what we observe (Table 3). Morrison & Forrester (2008) reports an average contract APR of 17.2% on revolving accounts over April 2006-February 2008; the comparable mean in our data is 18.08%.

Finally, our sample is higher-income, higher-educated, and younger than national averages.

In all, it seems that our sample is more representative in its card terms and usage (albeit with less borrowing than typical) than in its demographics. As we noted at the outset, the implications of these patterns for the external validity of our key findings, on the magnitude and sources of dispersion in cross-individual borrowing costs, are unclear.

III. Cross-sectional Variation in Risk-adjusted Borrowing Costs

In this section we detail the key facts motivating our analysis: 1) credit card borrowing costs vary dramatically across individuals, even after controlling for credit risk and other (non-APR) attributes of credit card contracts, 2) the cross-sectional variation in borrowing costs that remains is large enough to generate substantial differences in household (dis)savings rates and hence in wealth accumulation.

A. Measuring Credit Card Borrowing Costs

To measure individual-level borrowing costs, we calculate for each panelist the average balance-weighted annual percentage rate (APR) over our sample period.[14] This individual-level and time-invariant measure in fact accounts for the lion's share of overall variation in borrowing costs; there is relatively little within-individual variation over time in borrowing costs, even during this somewhat volatile period.[15] Variation in the individual-level measure indicates cross-sectional variation in average borrowing costs.

---

[14] The average weights APRs by balances held each day, and averages (equally-weighted) across days. Weighting each day by balances held yields nearly identical results.

[15] If one regresses daily weighted APRs (across all cards) on panelist fixed effects, the r-squared is roughly 0.75.

We can then explore how much cross-sectional variation comes from sources such as credit scores, panelist-specific default and/or credit utilization behavior, and non-price terms.

The top rows of Table 2 show weighted APRs that include both revolving and non-revolving balances. Including non-revolving balances (which have an APR of zero) shifts the distribution of our APR measure leftward, and more so for borrowers in the lightest-borrowing quartiles.

Our primary focus is on variation in borrowing costs conditional on the level of (revolving) debt, so for our purposes the next set of rows is more informative; those rows show variation across panelists in borrowing costs only for revolving balances. For the sample as a whole, the interquartile range (IQR) of borrowing costs ("Panelist-level weighted APR, revolving balances") is roughly 750 basis points, and the 10^{th}/90^{th} percentile range is 1500 basis points. Dispersion is substantial at all borrowing levels. Dispersion is greatest (IQR = [13.16%-21.45%) among the heaviest borrowers, and lowest (IQR = [14.90%-20.79%) among the lightest borrowers.

To again frame the magnitudes in terms of wealth accumulation, consider a typical "heavy borrower" household paying $1600 per year in interest costs with disposable income of $45,000 (both of these numbers are the medians for someone in the top quartile of revolving balances). If that household currently pays an APR in the 75^{th} percentile of borrowing costs, moving that household to the 25^{th} percentile, holding principal balances and everything else constant, would increase its household savings rate by 1.3 percentage points. For purposes of comparison, the average household savings rate as calculated by the Bureau of Economic Analysis (BEA) was 4% over 1990-2010, and was roughly 2.5% during our sample period (2006-2008).

B. Discarding Teaser Rates

Our measurement of borrowing costs is ideally something that reflects *contract* (steady-state) APRs. One measurement issue is created by low introductory "teaser" rates, which revert to a contract rate after 6-12 months. These are fairly easy to identify, simply because the range of teaser rates is significantly below even the lowest contract APR offered to the best credit risks. Here, we assume that any APR lower than 7.99 is a

teaser rate; we base this on independent data on credit card offers, in which 7.99% is the lowest "Regular Contract APR" observed in the 2006-2008 period.[16] This classification labels roughly 5% of account-month APRs as teaser rates.

One might wonder whether post-teaser contract APRs exceed those that did not begin with a teaser rate, but we have seen little evidence (in our data or elsewhere) suggesting that post-teaser rates are much higher than initial non-teaser contract rates. Most teaser rates seem to be targeted to those with lower switching costs, and recoup their initial losses via 2-5% balance transfer fees rather than higher ex post APRs.[17]

The impacts of discarding teaser rates can be seen in the third panel of Table 2. We still find a cross-individual interquartile range of 700-800bp, with little difference across categories of borrowing intensity.

C. Risk-adjusted Borrowing Costs: Credit Scores and Within-sample Risk Measures

A more important explanation for different borrowing costs across individuals is default risk. Card issuers expend considerable resources to measure that risk and tailor credit card offers based on risk (Allen, DeLong, and Saunders 2004; Edelberg 2006; D. Gross and Souleles 2002), so it seems plausible that risk would explain cross-sectional variation in borrowing costs. We can assess the importance of risk in several ways during our sample period.

Our primary measure of risk is each panelist's credit score (observed from one of the three major credit bureaus). Credit scores are measured when the panelist enters our data, which is generally in January 2006, but occasionally later. The credit report also contains detailed information on "lines" or accounts. We observe lines by category of borrowing (e.g., "mortgage lines," "finance company installment lines," etc.) and the total dollar balance on each category of line at the time of the credit score "pull." We also observe, for each panelist, basic demographic data from Lightspeed's registration survey, including income, education, and age. If credit score imperfectly captures the relationship

---

[16] We are grateful to Steph Wilshusen at FRB Philadephia for sharing these data with us.
[17] Most teaser rate offers charge a one-time fee on transferred balances, ranging from 2-5% of the transferred amount. We do not typically observe these particular fees, incurred on brand-new accounts, because the panel construction enrolls existing accounts.

between these variables and risk, or if the scoring agency does not observe these variables but the issuer does, then these might be correlated with pricing.

We augment the credit score data in several ways, which can capture dimensions of risk unmeasured by the score, and time-varying risk within the sample period. To place our approach in context, it is useful to discuss how issuers price risk on both current and prospective customers, and also to understand how the construction of our data affects what risk is likely to be observable or unobservable.

For issuers generating a new offer to a non-customer, a credit score is the primary source of information about risk. The credit score is a summary measure of that risk, incorporating information about total debt, debt utilization, default history ranging several years into the past, and the number of "pulls" or applications for new credit. The three major credit bureaus each have a proprietary algorithm for summarizing that information in a single three-digit score taking a maximum of 850. Generally, scores from the three bureaus are extremely highly correlated, and they rarely differ substantively for a particular individual.

Beyond the credit score itself, issuers may also incorporate disaggregate information from the credit report into their proprietary risk modeling. The key pieces of information an issuer might use are information about late payments on credit cards and other loans, information about credit utilization – utilization being the ratio of balances to the borrower's credit line – and information about recent applications for new credit. All of this information may affect the initial contract rate offered to a customer. Customers also self-report income and things such as education on their applications, but an issuer generally does not directly observe those things.

Issuers also can modify APRs for existing cardholders, based on assessments of changes in cardholder risk over time. Changes in credit bureau data often form the basis for such changes. But for existing cardholders an issuer can directly observe behavior on the card – late payments and credit utilization being the most important determinants of APR changes. Issuers can also observe such variables for other cards in the holder's portfolio, albeit less precisely.

In our case, we observe the cardholder's credit bureau data and credit score when the panelist enters the sample. This means that our credit score is likely an accurate measure

of what issuers would see given the APRs that prevail for that panelist at the beginning of the sample. Over time the credit score might change, introducing measurement error, but credit scores actually tend to be very stable within-individual, over-time. Conversations with regulators who have access to repeated observations on individual scores over time indicate that they find year-to-year within-individual correlations of quite close to one, even during our sample period. Our data corroborate this view : our credit scores do not lose predictive power for APRs as the individual "ages" through our sample period. The correlation between borrowing costs and credit score is 0.33 January 2006, and 0.34 in December 2008.

Our conversations with industry experts indicate that the second type of information – account- and panelist-level information about credit card usage, utilization and repayment on existing cards, and the incorporation of that information by a cardholder's current issuers is at least as important as credit score once a cardholder acquires a panelist. With that in mind, we construct a large and detailed set of variables measuring risk from cardholders' actual behavior:

- *Late payment fees:* Most issuers impose a fee when the cardholder fails to make the minimum monthly payment. For each account, we calculate a running total of late fees incurred during the sample. We use this to construct panelist-level running totals of late fees across all accounts, as well as time-invariant "Late fees per month" and "Any late fee in-sample" measures.

- *Over-limit fees:* Issuers generally allow cardholders to exceed their issuer-defined account-level credit limits, but then impose a monthly fee as long as the account balance remains over the limit. Again, here we calculate running totals of in-sample over-limit fees incurred during the sample, at both the account and panelist levels, and use these to code time-invariant "Over-limit fees per month" and "Any over-limit fee in-sample" measures.

- *Credit lines and utilization*: Issuers view approaching the limit of one's credit as a strong signal of default risk. The standard metric of that risk is the utilization rate, measured as a ratio of balances to total limit. We measure that utilization rate at the account level and across all cards. We also allow for the possibility that a given ratio may mean different things depending on the total

credit limit (85% of $10,000 may present risk different from 85% of $50,000), so we measure the total credit limit across all cards.

Collectively, these variables are quite comprehensive. They likely compare favorably to the data observed by any one issuer, although individual issuers may of course employ those data differently. And while we do not observe these variables until panelists enter the sample, most of the variables display extremely strong serial correlation. For example, only 17% of panelists with no late payments in 2007-08 had one or more late payments in 2006. Credit lines and balances are extremely stable over time within panelist. We also measure both time-invariant metrics of risk ("any late fee in sample") that are presumably correlated with that time-invariant behavior leading up to the sample period, and time-varying metrics on the same dimensions ("number of late fees up until this month in the sample") that capture variation in a person's risk.

As a first pass at understanding the links between observable risk and borrowing costs, we construct three categories of risk ("low," "medium" and "high"). "Low risk" are cardholders with no in-sample late or over-limit fees, average credit utilization below 0.50 and a credit score exceeding 720.[18] "High risk" are cardholders in 3rd-5th quintile by total in-sample late fees, with utilization >0.70 and with a credit score below 720. "Medium" risk are the remainder.

The bottom three sets of rows in Table 2 show how these categories explain cross-sectional variation in borrowing costs, and the variation that remains within-category. Even these fairly coarse risk categories explain large *mean* shifts in average borrowing costs: APRs in the "low risk" bin are roughly 800bp lower than those in the "high risk" bin. But even *within* each risk bin substantial variation remains, and if anything it is larger than cross-category variation for heavier borrowers. The IQR of borrowing costs is over 500bp in the low-cost/heavy borrower group, and over 700bp in the high-cost/heavy borrower group. Figure 2 shows box-and-whisker plots to illustrate overlap in the distributions by risk category (grouping all borrowing categories together).

To get a more precise fix on how well observable risk explains the cross section of borrowing costs, we show at the bottom of Table 2 the fit of a series of regressions. Each

---

[18] Under the confidentiality terms of our contract with Lightspeed, we can not present descriptive statistics or tabulations by credit score bin alone.

regression uses a panelist as the level of observation, with that panelist's weighted APR as the dependent variable. The independent variables consist of fifteen measures, decomposed into 128 discrete and continuous variables:

- Credit score decile
- Indicators for the panelist's first and last months in the sample (to account for variation over time in market APRs)
- Decile of total credit line (average over the sample)
- Decile of average credit utilization
- Indicators for "incurred any late fee in-sample" and "incurred any over-limit fee in-sample"
- Quintiles of total late and over-limit fees incurred in-sample
- The average number of late fees incurred per month
- Quartile of average monthly purchase volume
- Quartile of average monthly revolving balances
- Average annual fees paid
- Age, income and education categories as described in Table 1

The bottom rows of the table show the fit from these regressions, and contributions to the fit of sub-groups of the control variables.

Including only credit score and the sample entry/exit controls explains roughly one-fifth of cross-sectional variation in borrowing costs, with the lion's share of explanatory power coming from the credit score deciles. Adding the in-sample risk measures increases the fit to 0.34, and including the demographics adds very little beyond that. We have experimented with a variety of more flexible functional forms, but have been unable to improve the fit of the model.

Figure 2 illustrates the variation in the raw data, and compares it to variation in the residuals from our most heavily parameterized model. The "Demeaned APR" density centers the distribution of weighted APRs (omitting teaser rates) on zero. The "Residuals" density shows a kernel density estimate of the distribution of unexplained variation in APRs. The biggest effect of controlling for risk is to shift down APRs at the high end of the distribution, but the IQR of the residuals is nearly 600 basis points, and

the $10^{th}/90^{th}$ percentile gap is over 1100 basis points. A good deal of variation in APRs remains to be explained, after controlling for our measures of risk.

It is possible that we omit some useful (and observable to issuers) measure of risk, but we do observe the market-wide standard measure of risk (the credit score). And, we observe the individual-specific measures of behavior (late payments and utilization) widely regarded as most informative about risk by industry experts. In fact, because we observe data for multiple accounts, we may observe *more* account-level information than would any one issuer (because, for example, issuers report to the bureaus with a bit of a lag).

One possibility is that our credit scores may not fully reflect risk of those who opened the account, because people share cards. However, our results are unchanged if we restrict the analysis to those who are "single, never married." Sharing of cards between people who are not married is rare.

It is also surely true that there are issuer-specific and card-specific features that might affect APRs. We estimate richer models below that incorporate such characteristics. For now, we simply note that the substantial cross-sectional dispersion in borrowing costs is not fully explained by (our measures of) observable credit risk. In the next section we document that much of this variation is due to large cross- and within-individual differences in contract APRs held, even among similarly risky panelists. We then shed further light on the individual behavior that explains why similar individuals pay different borrowing costs.

D. Other Sources of Product Heterogeneity and Pricing Tradeoffs?

We also explore whether these results are driven by other types of product heterogeneity and/or pricing tradeoffs: rewards and affinity links, fixed vs. variable rate pricing, and annual fees. They are not. We do not show results in tables here to conserve space (although we do control directly for such card characteristics and terms in our later regressions), but there are several reasons why these factors do not contribute materially to dispersion. One is lack of strong tradeoffs, even conditional on credit risk; e.g., we do not see a significant fixed vs. variable rate tradeoff, and annual fee cards have -higher- APRs than no-fee cards (see, e.g., Table 7). A second reason is that we are interested in

cross-individual dispersion, and many individuals hold a mix of cards. A third reason is that the (potential) economic magnitudes involved are small relative to the dispersion we find; e.g., a standard valuation of rewards is 100 basis points, and most rewards go unclaimed.

IV. Contract APRs, Debt Misallocation and Search Behavior

Cross-sectional dispersion in borrowing costs has two plausible proximate sources. One is *contracts held*: if similar borrowers hold cards with different APRs. The second is *contract usage*: if similar borrowers allocate debt across their cards differently, conditional on contracts held.

A. Contract APR Distributions: Cross- and Within-Individual Variation

Table 3 suggests that both contracts held and contract usage are potential sources of cross-sectional dispersion in borrowing costs: there is substantial variation in both cross-sectional contract APRs (opening the possibility that contracts held matters), and in within-individual contract APRs (opening the possibility that contract usage matters).

The first rows show total variation at the card-month (or card-billing date, to be precise) level. Sample-wide, the interquartile range of contract APRs is 675 basis points, and again we see greater dispersion among heavier borrowers.

The next rows break the sample into our coarse low/medium/high risk categories, and illustrate variation in each panelist's average lowest contract APR across all days in the sample. This reveals how much variation in borrowing costs we might see purely from cross-individual dispersion in contract APRs, if individuals always use the lowest-rate card in the wallet. The pattern here echoes that in Table 2: we still see substantial variation across individuals, and more variation among riskier individuals.

However, we also see substantial APR variation *within* individuals. The next sets of rows show the average high-low APR spread within each panelist's wallet, averaged over the sample period.[19] We discard individuals with one card, who represent roughly one-third of heavier borrowers and two-fifths of the sample overall (see the bottom rows).

---

[19] Each day, we calculate the spread between the highest and lowest APR in the panelist's wallet. We average that spread over all days on which the panelist is in the sample

For panelists with more than one card, the median high-low spread is roughly 300bp. Again, we observe more scope for variation among heavier borrowers; in that group, even the low-risk panelists have a median high-low spread of 382bp, and the 75[th] percentile is 560bp. Among high-risk borrowers the spreads are greater: for heavy high-risk borrowers the median (75[th]) is 596bp (1006bp).

The upshot of this table is that cross-sectional variation in borrowing costs has two plausible explanations. First, cardholders holding identical portfolios of APRs might use their cards differently. The hypothesis that borrowing costs might differ substantially based on how cardholders allocate debt across their different cards has found some support in prior work, although the prevalence and underlying causes of such behavior is unsettled. Amar et al (2011) find that people sometimes ignore prices when repaying debt in incentivized laboratory simulations, preferring to repay low-balance (but lower-APR) cards before high-balance (but higher-APR) cards. Ponce et al (2009) presents evidence, from administrative data in Mexico, that borrowers often borrow on more expensive cards while holding cheaper cards, or repay low-rate before high-rate debt. Ponce et al argues that such behavior is consistent with mental accounting or other psychological drivers of behavior. There is also evidence that some people are unfamiliar with the terms of their cards [FINRA, Lusardi and Tufano].

Alternatively, similarly risky borrowers could hold cards with different contract APRs due to search/shopping behavior. Prior work on search in the credit card market is somewhat mixed, although it has generally faced data limitations, and has not focused on our particular questions. Calem and Mester (1995) establish a negative correlation between outstanding balances and individuals' propensity to search for low rates, and argue that the correlation indicates that people who dislike search also prefer current to future consumption. Crook (2002) finds that the correlation disappears in more recent data, and also notes that high balances should increase the benefits of search. Berlin and Mester (2004) examine cross-issuer price dispersion during the 1980s, and argue that such dispersion is poorly fit by models of search.[20] Ausubel's (1991) influential piece on

---

[20] Whether or not the models considered by Berlin and Mester were an apt description of credit card pricing during the 1980s, they are clearly less so now, when intra-issuer dispersion as well as individual-specific pricing based on shopping behavior are far more common. Also, the model they consider has other unusual features, such as a perfect division of consumers into "searchers"

credit card rate stickiness argues that search/shopping behavior might be an influence on market outcomes.

B. Cards Held and Debt Misallocation

Table 4 examines the allocation of credit card debt by our panelists. The top row reproduces the panelist-level weighted APR on revolving balances (from the second set of rows, Table 2). The next row shows statistics for the "best" weighted APR on revolving balances, averaged across all days. We calculate the best weighted APR by re-allocating balances to lowest-rate cards first, up to the credit limit of each card. We also calculate in the next set of rows the APR savings from re-allocating balances to lower-rate cards, averaged at the panelist level across all days with revolving balances, the panelist-level annualized dollar savings from re-allocation, and the share of total panelist-level interest costs that could be avoided via reallocation.[21]

These results suggest that panelists generally are quite effective in allocating balances to their lowest-rate cards. At every percentile actual and best rates are quite similar. For the sample overall, the median average APR savings from re-allocation is 0 basis points, meaning that all balances are allocated to the lowest-rate cards for the entire sample period. The seventy-fifth percentile of average savings is 51 basis points, and the 90[th] percentile is 222 basis points. Misallocation is slightly more common among heavier borrowers, who are more likely to have more cards.

The annualized dollar savings from re-allocating balances to the lowest-rate cards are similarly small: zero for the median panelist, $15 per year at the 75[th] percentile, and $110 per year at the 90[th]. Even in the heaviest-borrowing quartile, annual savings at the 75[th]/90[th] percentiles are $143/$431. As a share of total interest costs, costs due to misallocation are very small – the 75[th] percentile is 0.04.

Some of these findings are mechanical, because this table includes panelists with one card (in order to facilitate comparison with the prior tables). But the results are not much

---

and "non-searchers" and a uniform and fixed maximum willingness to pay for credit. Without that latter feature the primary empirical prediction tested by the authors need not hold.

[21] We include balances on teaser and penalty rates in the actual APR, and allow for re-allocation to/from teaser and penalty APRs in the best rate calculation. This allows for greater savings than would excluding these cards, because it allows panelists to re-allocate away from very high penalty rates and toward very low teaser rates.

different if we restrict the analysis to those panelists with more than one card, and for whom we observe the full set of cards (Table A2). Even among those panelists, the median share of interest costs due to misallocation is 0.03.

It is possible that these findings, modest as they are, even overstate the true degree of misallocation. They do not control for credit card rewards (miles, points, cash), which might lead borrowers to use a card that seems more expensive in APR terms but is actually less expensive net of rewards. In unreported results, we find some evidence that transactions made with the lowest-rate card are less likely to carry rewards than transactions using a "wrong" card, which does suggest that rewards play a role. It also appears that rewards cards are more heavily used by panelists who do not revolve (borrow) on their cards. Finally, we implicitly assume above that panelists could move balances costlessly to lower-rate cards. In practice people face issuer-imposed switch costs, in the form of balance transfer fees.

We also, in unreported results, have examined allocations of "excess repayments": payments greater than the monthly minimum . This is a somewhat cleaner test, because although rewards might affect purchase choice, once rewards have been obtained a borrower should always allocate excess repayments to the highest-APR card. Again, we see that nearly all repayments are allocated efficiently; sample-wide, all excess repayments go to highest-rate cards in 80% of panelist-months. Large repayments tend to be allocated more efficiently than small repayments. The latter finding suggests that ignorance or misperception of APRs cannot fully explain misallocation, because both of those should be independent of the dollar value of excess repayments.

The overall picture painted by our debt/repayment misallocation analysis is this: we find evidence that most borrowers allocate debt and repayments to their lowest-rate card(s). Overall, misallocation explains very little of the cross-sectional variation in borrowing costs.

C. Contract APR Variation and Search Behavior

If misallocation conditional on cards held explains only little of cross-individual variation in borrowing costs, then it must be that cross-panelist variation in contract APRs explains differences in borrowing costs. What then drives cross-individual

variation in contract APRs? One explanation offered is heterogeneity in consumer search/shopping behavior.

The notion that consumers do not always obtain credit at the lowest possible price due to information frictions is long-held. Work from the 1960s, around the advent of the original Truth-in-Lending Act, frequently highlighted the difficulty that many borrowers face in comparison shopping for cards and other debt products. Theoretical and empirical models of search behavior more generally have a long history in economics and marketing. [cites]

Search seems to be particularly likely to matter in credit cards versus other types of debt, for several reasons. Many thousands of issuers offer credit cards to customers. Some internet sites aggregate offers, but this is an imperfect shopping solution because the shopping engines do not (yet) tailor offers based on an individual's credit risk. A customer generally only learns terms of an issuer-specific offer after completing as least some, and perhaps most, of the application process, making a full comparison of prices at all issuers extremely costly in terms of time. The other primary mechanism by which customers see offered prices is via the mailout solicitations sent by issuers. Even today, mailout solicitations are the dominant form of credit card marketing and new account acquisition [cite]. Issuers collectively mail roughly [cite] billion solicitations per year.

On the supply side, there also appears to be substantial price dispersion conditional on observable risk, because issuers' internal credit scoring models differ and hence two issuers might view a particular cardholder in different ways. A second source of price dispersion is the more classic type that itself arises from the presence of search behavior. Search behavior can generate ex post price dispersion, even when buyers and sellers are ex ante similar, or when buyers differ only in their opportunity cost of search. In that case, issuers might pursue different pricing strategies, trading higher prices for fewer customers along a zero-profit frontier. Issuers might even randomize their offers to observably identical customers, trading higher prices conditional on acquisition against lower acquisition rates. Indeed, there is evidence that for experimental/learning reasons issuers introduce randomization into their pricing.

A third source of price dispersion in the credit card market occurs ex post, after issuers learn about a customers' search/shopping behavior. At that point, issuers can price

discriminate based on search behavior (or switching costs, which are an input into equilibrium search). Discussions of such behavior have a long history in the credit card market. Customers are often advised by financial advice columnists that simply calling their existing issuer can induce the issuer to cut the interest rate on the card. A cardholder's current issuers can also observe card switches, because the credit report contains a list of current and past "lines" (accounts), along with balances on each card for all open lines. Issuers often respond to switches away from their cards by sending a "counter-offer" in the form of a lower teaser or contract rate, to the customer switching balances away from the card.

Such strategic behavior by issuers can itself induce strategic behavior by cardholders. The practice of "surfing" from one card to another in order to obtain better terms is advocated by consumer advocates, as is directly negotiating better terms. A forward-looking customer has an incentive to signal low search/switching costs, by contacting current issuers, by holding multiple cards, and by switching balances between cards (Taylor 2003).

What this means for the empirical distribution of credit card interest rates is that it will reflect at least two differences from the empirical distribution of offers. First, different search behavior across customers will generate a distribution based on the joint density of offers and search behavior. Second, issuer-specific responses to customer-specific signals about search behavior will lead some customers to face choice sets that differ systematically from those faced by others with identical observable risk.

D. Measuring Search/Shopping Behavior

We have both direct and indirect measures of search behavior in our sample. For a small subsample (N=505) of panelists, we observe a 10-point scale response indicating agreement with the statement "I always keep an eye out for the best credit card deals." This is the most direct measure of search that we observe, but the fact that we observe it for so few people is problematic for statistical reasons. In what follows below, to gain precision we collapse the 10-point scale into a single "search-intensive customer"

variable, which takes on a value of one if the panelists reports 5-10 on the ten-point scale.[22]

A second measure, which we observe for all panelists, is the number of cards held. For an active credit card shopper, there is little downside to keeping old cards even when a new low-APR card is found. Indeed, for the strategic reasons mentioned above, there may be considerable benefit holding a portfolio of cards, even if they are rarely/never used, to induce competition between issuers "in the wallet."[23] So, this is likely to be a good measure of how active a customer has been in acquiring new cards. It is also, in contrast to the self-reported search variable, something that is directly observable by issuers. If issuers price discriminate based on observed search behavior, this variable is likely to matter more than any inherent (and unobservable to the issuer) preference for shopping. And finally, this variable is more concrete relative to the self-reported, qualitative and perhaps measured-with-error search intensity variable.

Of course, card count could measure things other than shopping behavior, and unconditionally that is undoubtedly true. We emphasize that our analysis of its importance conditions on all other characteristics considered thus far, including most notably total credit limit, monthly purchases/balances, and credit utilization. We discuss what else "cards held" could measure, and show IV results that should allay some concerns about omitted variable bias, below.

E. Sub-Sample Results: the Direct Search Measure and Borrowing Costs

We first correlate our direct search measure, in the subsample of panelists for which it is observed, with a variety of APR measures (Table 5). The first APR measure is the panelist-level weighted best APR, as measured in Table 3. The "best APR" considers both contract APRs and credit limits on each card; it is possible that shopping for credit limits is also important, and differences in credit limits will affect this measure. The second measure is the lowest APR on any card held by the panelist during the sample

---

[22] The modal responses are 1, 5 and 10 (each with ~15% of responses). We base our definition of "search-intensive" on the point estimates of relationships between search intensity and APRs. The data do not reject the coefficient restriction.

[23] Inducing such competition is only valuable in the presence of search-related friction; else, competition in the wallet is no different from competition from firms outside the wallet.

period; one could view this as a lower bound on the APR the panelist could pay. The third and fourth measures use account-month level APR data. The third measure uses the account-month as the unit of observation, and therefore may include multiple observations for a panelist in a particular month; this measures the relationship between any panelist-level covariate and *average* APRs across all cards. The last measure is the lowest APR "in the wallet" measured at the panelist-month level. This allows us, as does the third measure, to control variation over time in panelist, card and market characteristics, in addition to time-invariant panelist characteristics.

The control variables in the first two models are identical to those described in the weighted APR regressions earlier (re: the r-squareds reported in Table 2): credit score decile, indicators for the panelist's first and last months in the sample (to account for variation over time in market APRs), decile of total credit line (average over the sample), decile of average credit utilization, indicators for "incurred any late fee in-sample" and "incurred any over-limit fee in-sample," quintiles of total late and over-limit fees incurred in-sample, the average number of late fees incurred per month, quartile of average monthly purchase volume, quartile of average monthly revolving balances, average annual fees paid, and the age, income and education categories as described in Table 1. For the purposes of exposition in Tables 5 and 7, we simplify the functional form of many variables. As we mentioned above, we use a binary "search-intensive customer" indicator rather than the full 10-point scale for search behavior. We also collapse credit score into three rather than ten categories; we include the cardinal value (1-10) of credit line decile, linearly, rather than the full vector of dummies; we use a "high utilization" dummy (>0.70=1) rather than a full vector of dummies; and so on. None of these restrictions are rejected in the data, and more important, none materially affects either the fit of the model or the coefficients of interest. [full sets of results to be placed in appendix].

The latter two columns in Tables 5 and 7 measure APRs at the monthly level, so we include controls for time-varying changes in all APRs (month/year fixed effects). We

also construct panelist-level running totals, in the sample period, of late and over-limit fees.[24]

For the account-level specifications (columns 3 and 4 of Tables 5 and 7) we can and do also include card characteristics: card issuer fixed effects, whether the card has an annual fee, whether it carries rewards or has an affinity link (to e.g., a sports team, university, etc.), and whether it has a "variable rate." A credit card APR may be "fixed," meaning not pegged to another market rate, or "variable," meaning that the APR moves monthly or quarterly with some market interest rate. Issuers may change either the fixed rate or the variable margin at any time, however, making the use of the terms "fixed" and "variable" non-standard relative to, say, terminology regarding mortgages.[25] Nonetheless we classify any APR that changes at least quarterly in the sample period as variable.

The coefficient of interest is that on "search-intensive customer." It is negative in every model, and statistically significant in the second and fourth (we cluster standard errors at the panelist level in the account-level specifications). The point estimate suggests that intensive search reduces rates paid by 80bp. This is not terribly large in economic terms, and we offer a few thoughts about that. First, the qualitative nature of the responses probably introduces measurement error; one person's intensive search may be another's casual search. That measurement error might bias the results toward zero. Second, we are forced due to the small sample to impose a coefficient restriction that might mask sharper effects among the most aggressive searchers.

The small sample size and the large number of explanatory variables renders many of the other coefficient estimates imprecise; we defer discussion of those variables until we show the models below that estimate relationships in the full sample, because we have more precision there.

F. OLS and IV Estimates Using Cards Held

In Table 6 we re-estimate the models in Table 5, substituting cards held for the search intensity variable (OLS row), but using only observations for which we also observe the search intensity variable. We then instrument for cards held (IV row) using the search

---

[24] Running totals for other things like utilization at the card level are not significant in any specification – credit lines and utilization are strongly serially correlated.
[25] See Stango (2000) for a detailed discussion of fixed and variable rate pricing in credit cards.

intensity variable – for flexibility, the instrument vector includes indicators for each value on the full 10-point scale, and also interacts that vector with panelist age categories (because all else equal, older panelists should have acquired more cards).

We discuss the magnitudes in detail below, but the key point illustrated by this exercise is that the IV estimates are no different, statistically, from the OLS estimates. In fact, in every specification the IV point estimate is *larger* than the OLS point estimate.

What we take from the results is that while it is certainly possible that "cards held" captures supply- or demand-side phenomena other than search intensity, it does not appear that this biases the estimates downward (toward more negative values). If anything, the bias appears to work the other way: OLS understates the strength of the negative relationship between cards held and APRs. And again, it is useful to remember that these coefficients come from models that also control for credit risk, debt demand, spending levels, and most importantly total credit lines and utilization. The comparison here is between people with identical observable risk, identical observable debt demand, similar card/account features, and most importantly identical (and perfectly observed) credit lines/utilization.[26]

Table A3 shows the first stage of the IV estimates. Search is significant in explaining cards held, conditional on all of the other covariates, and also significant in explaining cards held from the credit report. As a (weak) falsification test, we also show that search is not correlated other debt "lines" on the credit report.

G. Full Sample Results

With the sub-sample results in hand we now turn to the full sample results. Table 7 shows results from specifications that are identical to those in Table 5 except for the substitution of cards held for the survey-based search measure.

The central finding, in the first, second and fourth models, is a large and significant negative relationship between cards held and APRs. In the first (second) model the point estimate is 75 (100) bp per card held, suggesting that all else equal, moving from one to 5+ cards held is associated with an APR 375 (500) bp lower. The effect is not significant

---

[26] The coefficients are more negative, but not significantly, if we drop the credit line and utilization variables.

in the third model (although the IV estimate is significant for this model in Table 6). The coefficient in the fourth model (lowest APR in the wallet, month by month) is also economically meaningful, suggesting a 200+ bp reduction from going to 5+ from one card.

To put the magnitudes on the cardholding variable into perspective, they compare to what we would see from fairly large changes in observable risk. Moving from one to three cards more than completely offsets moving from the worst to the best credit score bin. That holds everything else constant, of course, and people in different bins also differ in utilization and late payments. But in the fourth model moving from one to four cards more than offsets the joint effect of moving from the worst to the best credit score bin, moving utilization below 0.70, and reducing at least some late fees incurred. That difference is more than the effects of moving from "medium" to "low" risk in Table 2, and slightly smaller than the effects of moving from "high" to "medium."

The OLS coefficients here are similar to those in Table 6, suggesting that this sample overall does not differ markedly from the direct search sub-sample.

Why are the magnitudes on the cards held variable (Table 7) so much larger than those on the survey-based search variable (Table 5)? To reiterate, it may be that (1) the search variable is measured with error, and that the measurement error attenuates the estimated effect in that model, and/or that (2) cards held is itself a direct signal of both actual shopping behavior, and a signal about search/shopping/switching behavior that is observable to issuers, and induces them to cut APRs in order to price discriminate.

We can now discuss the results for the other variables. Most have intuitive effects, with variables measuring higher risk associated with higher APRs. The in-sample risk variables (late fees and "high utilization," in particular) have quite dramatic effects, particularly relative to credit scores. These variables collectively can explain large variation in APRs across individuals. They also, for that matter, explain credit scores; regressing credit scores on our constructed risk measures yields a fit of 0.65.

The fourth model shows significant negative relationships between both purchase volume/revolving balances and APRs (in the other three models, the six coefficients on these variables are insignificant but all negative). This may indicate that those who benefit more from search (because of high balances), or who have more frequent

interactions with credit card issuers (because of high purchase volume), engage in more search. A default risk interpretation of these variables would predict coefficients of the opposite sign. This consistency with a model of search fits with the pattern of results on age and income described below.

The card-level characteristics are less important. We see no significant effects of rewards or variable rate pricing. The correlation between APRs and annual fees is positive, which is consistent with some other work (Stango 2002).

The fixed effects for month/year of sample entry/exit, for month/year of the data and for issuer are all jointly significant at very low p-values.

The "Table 7 continued" pane focuses on the correlation between demographics and APRs, because the pattern here is also consistent with search behavior. APRs are highest among the middle-aged, which is consistent with a model where search is a function of the opportunity cost of time. They are interesting in the context of the results in Agarwal et al (2009), who show that middle-aged individuals make fewer financial mistakes. The difference does not necessarily represent a contradiction, because search behavior might operate differently from mistake-making conditional on search.

Similarly, there is a strong *positive* correlation between income and APRs. This stands in contrast to what one would expect if income measured credit risk, but is consistent with what one would expect if income is positively correlated with the opportunity cost of time searching, or applying for new cards.[27]

V. Conclusion

We find that similarly risky borrowers often pay dramatically different borrowing costs. Conditional on risk and principal owed, moving someone from the 75[th] to the 25[th] percentile of borrowing costs would increase their household savings rate by more than a percentage point. We find that these results are at least partially explained by heterogeneity in borrower search behavior for good deals (contracts), and are essentially

---

[27] We find statistically weaker correlations, and smaller point estimates, on education.

not at all explained by heterogeneity in debt allocation conditional on contracts held. In fact we find very little meaningful misallocation at all.

The programmatic and policy implications of the results are somewhat "old-school": search behavior can have a big impact on market outcomes and incidence. More specifically our results suggest that helping people choose good contracts in the first place may have greater impacts than helping people manage the contracts they have. This is not to say that our results support any particular policy, programmatic, or business tack: our results are silent, for example, on how or how cost-effectively one could reduce affect search behavior, and on what the general equilibrium effects of any such intervention or innovation would be.

We close with three closely related directions for further research. One is on the relationship between, and the relative importance of, over-paying vs. over-spending. Our results suggest that over-paying, conditional on the amount borrowed/spent, could be as important as over-spending per se. This is noteworthy because both neoclassical and behavioral models have struggled to explain the high levels of U.S. credit card debt. Incorporating borrowing cost heterogeneity and/or its determinants may improve model performance. Second is what drives search heterogeneity. Is search behavior driven by heterogeneity in standard preferences for leisure, in one or more key behavioral factors, in skills/endowments, etc.? Optimal policy and practice may depend on those primitives. Third is how issuers respond to search behavior, and how issuers and consumers interact in equilibrium.

REFERENCES

Agarwal, S., S. Chomsisengphet, C. Liu, and N. Souleles. 2006. Do Consumers Choose the Right Credit Contracts?

Agarwal, S., J. Driscoll, X. Gabaix, and D. Laibson. 2009. "The Age of Reason: Financial Decisions over the Lifecycle with Implications for Regulation." *Brookings Papers on Economic Activity* Fall: 51-117.

Agarwal, Sumit, John Driscoll, Xavier Gabaix, and David Laibson. 2011. Learning in the Credit Card Market.

Agarwal, Sumit, Paige Skiba, and Jeremy Tobacman. 2009. "Payday Loans and Credit Cards: New Liquidity andCredit Scoring Puzzles?" *American Economic Review Papers and Proceedings* 99 (2): 412-17.

Allen, Linda, Gayle DeLong, and Anthony Saunders. 2004. "Issues of Credit Risk Modeling in Retail Markets." *Journal of Banking and Finance* 28: 727-752.

Amar, Moty, Dan Ariely, Shahar Ayal, Cynthia Cryder, and Scott Rick. 2011. "Winning the Battle but Losing the War: The Psychology of Debt Management." *Journal of Marketing Research*.

Angeletos, George-Marios, David Laibson, Jeremy Tobacman, Andrea Repetto, and Stephen Weinberg. 2001. "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives* 15 (3): 47-68.

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

Barber, Brad, and Terrence Odean. 2011. The Behavior of Individual Investors.

Berlin, Mitchell, and Loretta Mester. 2004. "Credit Card Rates and Consumer Search." *Review of Financial Economics*.

Brown, Meta, Andrew Haughwout, Donghoon Lee, and Wilbert van der Klaauw. 2011. "Do we know what we owe? A comparison of borrower- and lender-reported debt." *Federal Reserve Bank of New York Staff Report* (523).

Bucks, Brian, Arthur Kennickell, Traci Mach, and Kevin Moore. 2009. "Changes in U.S. Family Finances from 2004 to 2007: Evidence from the Survey of Consumer Finances." *Federal Reserve Bulletin* February: A1-A56.

Calem, Paul, Michael Gordy, and Loretta Mester. 2006. "Switching costs and adverse selection in the market for credit cards: New evidence." *Journal of Banking and Finance* 30 (6): 1653-1685.

Calvet, Laurent, John Y. Campbell, and Paolo Sodini. 2007. "Down or out: Assessing the welfare costs of household investment mistakes." *Journal of Political Economy* 115 (5): 707-747.

Campbell, John Y. 2006. "Household Finance." *Journal of Finance*.

Carroll, C.D. 2001. "A Theory of the Consumption Function, With and Without Liquidity Constraints." *Journal of Economic Perspectives* 15 (3): 23-45.

Charles, Kerwin, Erik Hurst, and Melvin Stephens. 2008. "Rates for Vehicle Loans: Race and Loan Source." *American Economic Review Papers and Proceedings* 98 (2): 315-320.

Choi, James, David Laibson, and Brigitte Madrian. 2010. "Why Does the Law of One Price Fail? AnExperiment on Index Mutual Funds." *Review of Financial Studies*.

Crook, J. 2002. "Adverse Selection and Search in the Bank Credit Card Market." *Working Paper, Credit Research Centre, University of Edinburgh*.

Dynan, Karen. 2009. "Changing Household Financial Opportunities and Economic Security." *Journal of Economic Perspectives*.

Edelberg, Wendy. 2006. "Risk-Based Pricing of Interest Rates for Consumer Loans." *Journal of Monetary Economics* 53 (8): 2283-2298.

FINRA Investor Education Foundation. 2009. *Financial Capability in the United States: National Survey – Executive Summary*. http://www.finrafoundation.org/web/groups/foundation/@foundation/documents/foundation/p120535.pdf.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded attributes, consumer myopia, and information suppression in competitive markets." *Quarterly Journal of Economics* 121 (2): 505-540.

Goolsbee, Austan. 2000. "In a World without Borders: The Impact of Taxes on Internet Commerce." *The Quarterly Journal of Economics* 115 (2) (May): 561=576.

———. 2001. "Competition in the Computer Industry: Online versus Retail." *Journal of Industrial Economics* 49 (4 (Symposium on E-Commerce)) (December): 487-499.

Gross, D., and N. Souleles. 2002. "An Empirical Analysis of Personal Bankruptcy and Delinquency." *The Review of Financial Studies* 15 (1): 319-347.

Gross, David B, and Nicholas S Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *The Quarterly Journal of Economics* 117 (1): 149-85.

Heidhues, Paul, and Botond Koszegi. 2010. "Exploiting Naïvete about Self-Control in the Credit Market." *American Economic Review* 100 (5): 2279-2303.

Hortascu, Ali, and Chad Syverson. 2004. "Product Differentiation, Search Costs, and the Welfare Effects of Entry: A Case Study of S&P 500 Index Funds." *Quarterly Journal of Economics*.

Kerr, Sougata, and Lucia Dunn. 2008. "Consumer Search Behavior in the Changing Credit Card Market." *Journal of Business and Economic Statistics*.

Kolko, Jed. 2010. "A new measure of US residential broadband availability." *Telecommunications Policy* 34 (3) (April): 132-143.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2): 443-477.

Meier, Stephan, and Charles Sprenger. 2010. "Present-biased preferences and credit card borrowing." *American Economic Journal: Applied Economics* 2 (1): 193-210.

Morrison & Forrester LLP. 2008. *Letter on Proposed Rulemakings*.

National Commission on Consumer Finance. 1972. *Consumer credit in the United States*. Washington, DC: U.S. Government Printing Office.

Ponce, Alejandro, Enrique Seira, and Guillermo Zamarripa. 2009. Do consumers borrow on their cheapest credit card? Evidence from Mexico.

Prince, Jeffrey. 2008. "Repeat Purchase amid Rapid Quality Improvement: StructuralEstimation of Demand for Personal Computers." *Journal of Economics and Management Strategy* 17 (1): 1-33.

Shui, Haiyan, and Lawrence Ausubel. 2004. Time Inconsistency in the Credit Card Market.

Sirri, Erik, and Peter Tufano. 1998. "Costly Search and Mutual Fund Flows." *Journal of Finance*.

Soll, Jack B., Ralph L. Keeney, and Richard P. Larrick. 2011. Consumer Misunderstanding of Credit Card Use, Payments, and Debt: Causes and Solutions.

Stango, Victor. 2002. "Pricing with Consumer Switching Costs: Evidence from the Credit CardMarket." *Journal of Industrial Economics*.

Stango, Victor, and Jonathan Zinman. 2009a. "What do consumers really pay on their checking and credit card accounts? Explicit, implicit, and avoidable costs." *American Economic Review Papers and Proceedings* 99 (2): 424-29.

———. 2009b. "Exponential growth bias and household finance." *Journal of Finance* 64 (6): 2807-2849.

———. 2011. Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Bank Overdraft Fees.

Telyukova, Irina. 2011. Household need for liquidity and the credit card debt puzzle.

Woodward, Susan E., and Robert E. Hall. 2010. "Consumer Confusion in the Mortgage Market: Evidence of Less Than a Perfectly Transparent and Competitive Market." *American Economic Review Papers and Proceedings*.

Zinman, Jonathan. 2007. Household borrowing high and lending low under no-arbitrage.

———. 2009. "Where is the missing credit card debt? Clues and implications." *Review of Income and Wealth* 55 (2): 249-265.

Table 1. Summary Statistics for the Sample.

		Revolving Balance Quartile				
		1	2	3	4	All
Quartile cutoffs (revolving balances)		[0, 499]	[499, 1534]	[1534, 4586]	[4586, 62515]	[0, 62515]
Cards (share of panelists)						
	1	0.61	0.44	0.36	0.28	0.42
	2-3	0.33	0.44	0.46	0.43	0.42
	4+	0.06	0.12	0.17	0.28	0.16
Average purchases per month ($)						
	25th	12	22	44	77	28
	50th	65	120	252	361	173
	75th	260	715	1022	1083	688
	90th	623	1441	2435	2514	1722
Average revolving balances ($)						
	25th	45	695	1959	6411	499
	50th	198	909	2542	9357	1534
	75th	357	1187	3469	14161	4587
	90th	450	1399	4105	22040	10841
Panelist-level annualized interest costs ($)						
	25th	7	120	321	1061	88
	50th	33	170	436	1633	267
	75th	63	227	627	2480	794
	90th	87	293	825	4077	1871
Interest costs as share of total account costs		0.87	0.92	0.98	1.00	0.94
Interest costs as share of income						
	25th	0.000	0.002	0.004	0.014	0.002
	50th	0.001	0.003	0.007	0.025	0.005
	75th	0.001	0.005	0.013	0.043	0.014
	90th	0.002	0.008	0.020	0.071	0.032
Credit score						
	25th	599	581	619	657	616
	50th	707	660	695	704	694
	75th	786	759	770	753	768
	90th	812	804	806	791	805
Income [n=4106]						
	under $25,000	0.20	0.18	0.15	0.09	0.16
	$25k-$45k	0.22	0.22	0.18	0.17	0.20
	$45k-$87.5	0.42	0.46	0.45	0.51	0.46
	$87.5-$125k	0.09	0.08	0.12	0.12	0.10
	$125k+	0.07	0.06	0.10	0.11	0.08
Education [n=3892]						
	HS or less	0.12	0.12	0.10	0.08	0.11
	Some college	0.34	0.35	0.32	0.28	0.32
	College degree +	0.54	0.53	0.58	0.63	0.57
Age [n=4312]						
	Under 30	0.27	0.27	0.26	0.21	0.25
	30-39	0.25	0.29	0.28	0.31	0.28
	40-49	0.21	0.20	0.22	0.23	0.21
	50-59	0.18	0.16	0.15	0.18	0.16
	60+	0.09	0.08	0.09	0.08	0.08
Panelists		1,078	1,078	1,078	1,078	4,312
Accounts		1,655	2,108	2,390	3,154	9,307
Panelist-months		16,295	21,069	21,767	22,181	81,312
Purchase transactions		277,542	196,549	206,120	259,939	940,150

Notes: All variables measured at panelist level. Panelist-level averages are across all panelist-days in the sample. Note

Table 2. The Cross Section of Borrowing Costs.

		Revolving Balance Quartile				
		1	2	3	4	Total
Quartile cutoffs (revolving balances)		[0, 499]	[499, 1534]	[1534, 4586]	[4586, 62515]	[0, 62515]
Panelist-level weighted APR, all balances						
	10th	0.00	7.67	8.59	9.96	5.97
	25th	4.22	11.06	12.02	12.85	10.09
	50th	9.90	16.89	16.11	16.71	15.34
	75th	16.86	22.35	21.49	21.17	20.72
	90th	22.55	26.26	26.32	26.04	25.79
Panelist-level weighted APR, revolving balances						
	10th	11.89	11.63	10.14	10.08	10.90
	25th	14.90	15.46	14.14	13.16	14.51
	50th	17.59	18.60	17.83	17.06	17.80
	75th	20.79	23.28	22.05	21.45	22.16
	90th	25.73	27.75	27.24	26.33	26.74
Panelist-level weighted APR, revolving balances, no teaser rates						
	10th	12.87	12.99	11.72	11.41	11.90
	25th	14.90	15.90	14.90	13.79	14.90
	50th	17.80	18.81	17.96	17.30	17.98
	75th	20.98	23.42	22.30	22.02	22.38
	90th	25.90	28.03	27.37	26.47	26.93
Panelist-level weighted APR, revolving balances, no teasers, "low risk" (N=1007)						
	10th	11.84	10.99	10.11	9.97	10.80
	25th	14.01	13.97	13.08	11.84	13.24
	50th	16.45	16.28	16.09	14.00	16.03
	75th	17.99	18.05	18.15	17.16	17.99
	90th	19.24	19.80	20.09	19.61	19.79
Panelist-level weighted APR, revolving balances, no teasers, "medium risk" (N=2404)						
	10th	12.90	13.00	11.66	11.41	11.90
	25th	15.09	15.83	14.77	13.74	14.90
	50th	18.31	18.80	17.82	17.06	17.97
	75th	22.59	23.05	21.58	20.61	22.07
	90th	26.87	26.78	26.78	25.87	26.41
Panelist-level weighted APR, revolving balances, no teasers, "high risk" (N=716)						
	10th	15.67	18.20	15.85	15.78	16.80
	25th	19.94	20.78	19.80	18.87	19.93
	50th	23.12	24.21	23.90	22.91	23.58
	75th	26.57	27.71	27.46	26.79	27.26
	90th	28.52	29.16	29.39	29.22	29.09
R-squared (adjusted): credit score decile		0.25 (0.19)	0.34 (0.30)	0.28 (0.23)	0.22 (0.17)	0.23 (0.22)
R^2 (adj.): above plus time-varying risk, util.		0.38 (0.30)	0.45 (0.39)	0.40 (0.32)	0.40 (0.32)	0.34 (0.32)
R^2 (adj.): above plus demographics		0.40 (0.29)	0.48 (0.38)	0.43 (0.31)	0.44 (0.33)	0.35 (0.32)

Notes: Weighted APR is at panelist level across all cards (or cards without teaser/penalty APRs) and days in sample, weighted by balances or revolving balances. "Teaser rates" are those below 7.99%. "Low risk" are cardholders with no in-sample late or over-limit fees, average credit utilization below 0.50 and a credit score exceeding 720. "High risk" are cardholders in 3rd-5th quintile by total in-sample late fees, with utilization >0.70 and with a credit score below 720. "Medium" risk are the remainder.

R-squared notes: Rows show fit from panelist-level OLS regressions with no-teaser weighted APR as the dependent variable. First row shows fit from model with credit score decile and controls for sample entry/exit dates on the RHS. Second row also includes average annual fees, quintiles of late/over-limit fees, purchase volume and revolving balance quartiles, and deciles for total credit line and credit utilization. Third row also includes categorical variables measuring age, income and education.

Table 3. Possible Sources of APR Dispersion: Cross-individual and within-individual APR variation

		Revolving Balance Quartile				
		1	2	3	4	Total
Contract APR, no teaser rates						
	10th	11.65	12.12	10.99	10.99	10.99
	25th	14.49	14.90	14.24	13.90	14.24
	50th	17.24	17.80	17.49	17.24	17.32
	75th	19.80	21.90	21.24	21.24	20.99
	90th	23.90	28.15	28.15	28.99	28.15
Lowest contract APR in-sample, no teaser, "low risk"						
	10th	10.16	10.30	9.91	9.82	9.99
	25th	13.19	12.95	12.30	10.99	12.50
	50th	15.75	15.27	14.89	13.07	14.99
	75th	17.54	17.19	17.56	16.08	17.34
	90th	18.99	18.74	19.15	18.48	18.97
Lowest contract APR in-sample, no teaser, "medium risk"						
	10th	12.89	12.22	10.75	10.12	11.04
	25th	14.99	14.90	13.35	12.35	13.79
	50th	18.24	17.91	16.86	15.63	17.14
	75th	22.87	22.38	19.88	18.78	20.65
	90th	26.96	26.34	25.79	24.64	25.90
Lowest contract APR in-sample, no teaser, "high risk"						
	10th	14.90	14.92	13.64	12.40	14.07
	25th	19.74	19.07	17.04	16.56	17.57
	50th	22.81	22.94	21.26	19.92	21.53
	75th	25.90	27.03	25.90	24.93	25.96
	90th	28.53	28.96	28.80	28.08	28.87
High-low contract APR difference (>1 cards), "low risk"						
	10th	0.00	0.00	0.00	0.00	0.00
	25th	0.00	0.77	0.09	1.24	0.23
	50th	1.22	2.10	2.03	3.82	2.08
	75th	3.60	4.29	4.50	5.60	4.54
	90th	6.11	6.06	7.90	8.01	7.05
High-low contract APR difference (>1 cards), "medium risk"						
	10th	0.00	0.00	0.00	0.00	0.00
	25th	0.00	0.08	0.45	0.77	0.16
	50th	0.59	2.19	3.26	4.34	3.00
	75th	2.97	5.16	6.88	8.42	6.47
	90th	5.67	7.97	9.41	12.18	10.22
High-low contract APR difference (>1 cards), "high risk"						
	10th	0.00	0.00	0.00	0.35	0.00
	25th	0.00	0.13	1.64	3.05	1.03
	50th	0.06	2.80	4.20	5.96	3.93
	75th	1.48	5.65	7.72	10.06	7.96
	90th	3.74	10.25	12.17	12.67	11.65
Cards (share of panelists)						
	1	0.61	0.44	0.36	0.28	0.42
	2-3	0.33	0.44	0.46	0.43	0.42
	4+	0.06	0.12	0.17	0.28	0.16

Notes: Unit of observation in first set of rows is card-month; all others are panelist-level averages. "Lowest" contract APRs are panelist-level averages across all days in the sample. "High-low contract APR difference" is the average across all days of the daily panelist-level difference between the highest and lowest contract APRs "in the wallet." "Teaser rates" are those below 7.99%. "Low risk" are cardholders with no in-sample late or over-limit fees, average credit utilization below 0.50 and a credit score exceeding 720. "High risk" are cardholders in 3rd-5th quintile by total in-sample late fees, with utilization >0.70 and with a credit score below 720. "Medium" risk are the remainder. "Cards" is measured at panelist level as maximum number of distinct accounts on any day in the sample period.

Table 4. (Mis)allocation of balances.

| | | \multicolumn{5}{c}{Revolving Balance Quartile} | | | | |
		1	2	3	4	Total
Weighted average actual APR						
	25th	14.90	15.46	14.14	13.16	14.51
	50th	17.59	18.60	17.83	17.06	17.80
	75th	20.79	23.28	22.05	21.45	22.16
	90th	25.73	27.75	27.24	26.33	26.74
Weighted average best APR						
	25th	14.09	14.66	12.90	11.97	13.36
	50th	17.26	18.15	17.13	16.17	17.23
	75th	20.20	23.06	21.50	20.76	21.64
	90th	25.36	27.56	27.04	25.90	26.41
Average APR savings from re-allocation						
	25th	0.00	0.00	0.00	0.00	0.00
	50th	0.00	0.00	0.00	0.08	0.00
	75th	0.00	0.27	0.70	1.01	0.51
	90th	1.22	1.75	2.66	2.80	2.22
Annualized dollar savings from re-allocation						
	25th	0	0	0	0	0
	50th	0	0	0	10	0
	75th	0	4	24	134	15
	90th	5	25	98	431	110
Savings as percentage of annual interest costs						
	25th	0.00	0.00	0.00	0.00	0.00
	50th	0.00	0.00	0.00	0.01	0.00
	75th	0.00	0.02	0.06	0.08	0.04
	90th	0.19	0.18	0.25	0.22	0.21

Notes: "Actual APR" is weighted APR as in Table 2. "Best APR" re-allocates balances to lowest rate cards, conditional on available credit limits."Average APR savings" is the average daily reduction in APR obtained by transferring all balances to lowest-rate cards, conditional on available credit. "Dollar savings" multiplies the average APR savings by average daily revolving balances. "Savings as percentage..." compares dollar savings from reallocation to the dollar interest costs from

Table 5. Search Behavior and Borrowing Costs: Direct Search Measure.

	Dependent variable:			
	Weighted best APR (panelist)	Lowest APR (panelist)	Account APR (account-month)	Lowest account APR (panelist-month)
Search-intensive customer (5-10 out of 10)	-0.458 (0.419)	**-0.816**** (0.372)	-0.399 (0.257)	**-0.848***** (0.314)
Credit Score 600-719	-0.235 (0.717)	-0.907 (0.649)	**-1.011*** (0.570)	-0.812 (0.632)
Credit score 720+	0.375 (0.802)	-0.577 (0.738)	-1.037 (0.648)	-0.669 (0.719)
Panelist-level credit line decile	**-0.608***** (0.119)	**-0.590***** (0.106)	-0.127 (0.077)	**-0.285***** (0.087)
High credit utilization (>0.70)	**2.905***** (0.692)	**1.357**** (0.613)	**1.798***** (0.547)	**1.832***** (0.528)
Any in-sample late fee	0.507 (0.470)	0.573 (0.418)	-0.057 (0.265)	0.173 (0.324)
Any in-sample over-limit fee	1.125 (0.697)	-0.100 (0.616)	**1.254***** (0.472)	0.244 (0.507)
Average late fees/month	**2.700***** (1.033)	**2.236**** (0.909)	**1.168*** (0.598)	-0.227 (0.829)
Total late fee quintile 4-5	**1.752**** (0.829)	0.778 (0.735)	-0.143 (0.475)	0.800 (0.609)
Purchase volume quartile	0.094 (0.266)	-0.054 (0.238)	-0.154 (0.163)	-0.189 (0.188)
Revolving balance quartile	-0.073 (0.231)	-0.221 (0.209)	-0.011 (0.132)	-0.171 (0.159)
Account-level late fees (censored at 10)			**0.370**** (0.153)	**0.561***** (0.189)
Account-level over-limit fees (censored at 10)			0.086 (0.215)	0.005 (0.266)
Average annual fees, panelist level	0.033 (0.044)	-0.031 (0.039)	0.025 (0.018)	0.028 (0.020)
Account level: has annual fee?			0.403 (0.319)	**0.802*** (0.431)
Account-level cash rewards			-1.938 (1.485)	-3.007 (1.838)
Account-level affinity link			2.251 (1.482)	**3.761**** (1.821)
Account level: variable rate			**-0.795***** (0.217)	-0.225 (0.289)
N	493	482	25906	11060
r-squared	0.47	0.45	0.36	0.48
Sample entry/exit indicators	yes (0.01)	yes (0.01)	yes (0.00)	yes (0.00)
Age categories	yes (0.38)	yes (0.46)	yes (0.16)	yes (0.42)
Education categories	yes (0.46)	yes (0.61)	yes (0.97)	yes (0.72)
Income categories	yes (0.13)	yes (0.63)	yes (0.38)	yes (0.84)
Month-year indicators			yes (0.00)	yes (0.00)
Issuer indicators			yes (0.00)	yes (0.00)

Notes: First two models are at panelist level. "Weighted best APR" allocates balances to the lowest-rate cards, conditional on credit limits. "Lowest APR" is simply the lowest APR in the panelist's wallet across the entire sample. Third model is at account-month level, and fourth model is at panelist-month level. Third and fourth models cluster standard errors at panelist level. Coefficients shown are for specifications using linear specifications for selected variables/categories, to conserve space. "Intensive search" assigns a value of one to responses 5-10 on the scale "I always keep an eye out for better credit card offers," with 1 meaning "Does not describe me at all" and 10 meaning "Describes me perfectly." "Sample entry/exit" indicators are for the first and last months of the panelist's presence in the data. Panelist-level fee variables are across all cards. "Panelist late fees per month," "Panelist late fee quintile," and the panelist utilization/revolving balance/purchase volume variables are measured at the panelist level across the entire sample. "Rewards/affinity" are two indicators at the account level summarizing cash rewards and affinity links. Late/over-limit fee indicators are time-varying running sums of account- or panelist-level fees incurred within the sample period.

Table 6. OLS and IV Estimates of the Search/Borrowing Cost Relationship.

	Dependent variable:			
	Weighted best APR (panelist)	Lowest APR (panelist)	Account APR (account-month)	Lowest account APR
Current cards held (number, censored at 5+): OLS	**-0.737***** (0.172)	**-1.028***** (0.157)	-0.192 (0.136)	**-0.536***** (0.158)
Current cards held (number, censored at 5+): IV	**-0.802*** (0.444)	**-1.354***** (0.410)	**-0.520*** (0.288)	**-0.757**** (0.331)
N	480	482	25906	11060

Notes: Models are identical to those in Table 5, except that "Cards held" replaces "Search-intensive customer." OLS treats "Cards held" as exogenous; IV instruments for "Cards held" using indicators for search intensity (the full vector of responses on the 10-point scale), and interacts that vector with panelist age.

Table 7. Search Behavior and Borrowing Costs.

	Weighted best APR (panelist)	Lowest APR (panelist)	Account APR (account-month)	Lowest account APR (panelist-month)
	\multicolumn Dependent variable:			
Current cards held (number, censored at 5+)	**-0.745***** (0.082)	**-1.003***** (0.075)	-0.039 (0.062)	**-0.447***** (0.071)
Credit Score 600-719	**-0.735***** (0.246)	**-0.765***** (0.225)	**-0.551**** (0.227)	**-0.597**** (0.241)
Credit score 720+	**-1.450***** (0.302)	**-1.563***** (0.275)	**-1.475***** (0.274)	**-1.482***** (0.284)
Panelist-level credit line decile	**-0.414***** (0.058)	**-0.273***** (0.052)	**-0.213***** (0.047)	**-0.276***** (0.049)
High credit utilization (>0.70)	**1.169***** (0.249)	**0.875***** (0.228)	**0.932***** (0.227)	**0.893***** (0.230)
Any in-sample late fee	**1.266***** (0.201)	**0.614***** (0.184)	**0.334**** (0.144)	**0.343**** (0.163)
Any in-sample over-limit fee	**0.810***** (0.248)	-0.175 (0.226)	**0.476**** (0.234)	-0.005 (0.231)
Average late fees/month	**3.000***** (0.398)	**2.004***** (0.358)	**1.837***** (0.287)	1.736*** (0.399)
Total late fee quintile 4-5	**1.658***** (0.307)	-0.082 (0.277)	0.008 (0.253)	0.071 (0.286)
Purchase volume quartile	-0.019 (0.106)	-0.143 (0.097)	**-0.139*** (0.079)	**-0.174**** (0.087)
Revolving balance quartile	-0.067 (0.097)	-0.055 (0.088)	-0.077 (0.072)	**-0.186**** (0.080)
Account-level late fees (censored at 10)			**-0.757***** (0.221)	-0.185 (0.287)
Account-level over-limit fees (censored at 10)			1.216 (0.743)	**3.258***** (0.948)
Average annual fees, panelist level	**0.066***** (0.020)	0.013 (0.018)	**0.059***** (0.015)	**0.054***** (0.014)
Account level: has annual fee?			**0.325**** (0.162)	**0.686***** (0.192)
Account-level cash rewards			-0.674 (0.607)	-0.045 (0.607)
Account-level affinity link			0.790 (0.595)	0.512 (0.584)
Account-level: variable rate			-0.674 (0.607)	-0.045 (0.607)
N	3535	3536	107400	60679
r-squared	0.38	0.33	0.32	0.36
Sample entry/exit indicators	yes (0.00)	yes (0.00)	yes (0.00)	yes (0.01)
Age categories	yes (0.05)	yes (0.40)	yes (020)	yes (0.19)
Education categories	yes (0.04)	yes (0.03)	yes (0.20)	yes (0.35)
Income categories	yes (0.00)	yes (0.06)	yes (0.00)	yes (0.01)
Month-year indicators			yes (0.00)	yes (0.00)
Issuer indicators			yes (0.00)	yes (0.00)

Notes: First two models are at panelist level. "Weighted best APR" allocates balances to the lowest-rate cards, conditional on credit limits. "Lowest APR" is simply the lowest APR in the panelist's wallet across the entire sample. Third model is at account-month level, and fourth model is at panelist-month level. Third and fourth models cluster standard errors at panelist level. Coefficients shown are for specifications using linear specifications for selected variables/categories, to conserve space. "Intensive search" assigns a value of one to responses 5-10 on the scale "I always keep an eye out for better credit card offers," with 1 meaning "Does not describe me at all" and 10 meaning "Describes me perfectly." "Sample entry/exit" indicators are for the first and last months of the panelist's presence in the data. Panelist-level fee variables are across all cards. "Panelist late fees per month," "Panelist late fee quintile," and the panelist utilization/revolving balance/purchase volume variables are measured at the panelist level across the entire sample. "Rewards/affinity" are two indicators at the account level summarizing cash rewards and affinity links. Late/over-limit fee indicators are time-varying running sums of account- or panelist-level fees incurred within the sample period.

Table 7, continued. Search Behavior and Borrowing Costs (demographics).

	Dependent variable:			
	Weighted best APR (panelist)	Lowest APR (panelist)	Account APR (account-month)	Lowest account APR (panelist-
Age: 30-39	0.061	0.009	0.195	0.156
	(0.234)	(0.213)	(0.179)	(0.197)
Age: 40-49	**0.714*****	0.356	**0.378***	**0.442****
	(0.254)	(0.231)	**(0.196)**	**(0.213)**
Age: 50-59	0.282	0.233	**0.544****	**0.553****
	(0.277)	(0.251)	**(0.214)**	**(0.232)**
Age: 60+	0.221	0.370	0.409	0.375
	(0.344)	(0.313)	(0.251)	(0.280)
Income: $25,001-$45,000	0.344	0.064	0.125	0.017
	(0.283)	(0.258)	(0.225)	(0.239)
Income: $45,001-$87,500	**0.876*****	**0.462****	**0.684*****	**0.562*****
	(0.254)	**(0.231)**	**(0.204)**	**(0.218)**
Income: $87,501-$125,000	**0.993*****	**0.552***	**0.921*****	**0.647****
	(0.347)	**(0.316)**	**(0.258)**	**(0.279)**
Income: $125,001+	**1.497*****	**0.818****	**0.761*****	**0.798*****
	(0.380)	**(0.347)**	**(0.284)**	**(0.292)**
Education: Some college	**-0.660****	-0.220	**-0.407***	-0.353
	(0.287)	(0.261)	**(0.233)**	(0.256)
Education: college+	-0.317	0.242	-0.269	-0.200
	(0.284)	(0.258)	(0.230)	(0.251)
N	3535	3536	107400	60679

Notes: Continued from Table 7 on previous page.

Table A1. Average Annualized Fee Costs.

		Revolving Balance Quartile			
	1	2	3	4	Total
Quartile cutoffs (revolving balances)	[0, 499]	[499, 1534]	[1534, 4586]	[4586, 62515]	[0, 62515]
Average yearly late fees					
mean	13	60	70	90	58
50th	0	18	17	14	0
75th	0	85	87	95	60
90th	32	201	226	269	185
Average yearly over limit fees					
mean	3	39	39	44	31
50th	0	0	0	0	0
75th	0	55	37	29	19
90th	0	134	138	154	112
Average yearly annual fees					
mean	10	25	19	16	18
50th	0	0	0	0	0
75th	0	36	22	0	3
90th	34	77	70	57	62
Average yearly balance transfer fees					
mean	1	0	3	11	4
50th	0	0	0	0	0
75th	0	0	0	0	0
90th	0	0	0	35	0
Average yearly cash advance fees					
mean	0	1	2	5	2
50th	0	0	0	0	0
75th	0	0	0	0	0
90th	0	0	0	0	0

Notes: All variables measured at panelist level.

Table A2. Mis-allocation of balances, full/multiple cards sub-sample

				Revolving Balance Quartile		
		1	2	3	4	Total
Weighted average actual APR						
	25th	15.09	15.07	13.45	12.12	13.76
	50th	17.90	18.43	17.76	15.98	17.56
	75th	21.90	22.75	21.55	21.04	21.95
	90th	26.51	26.60	25.90	25.84	26.17
Weighted average best APR						
	25th	13.25	13.51	11.52	11.05	12.07
	50th	17.24	17.14	16.16	14.28	16.26
	75th	20.83	22.02	20.63	20.19	20.88
	90th	25.36	26.29	25.26	24.87	25.58
Average APR savings from re-allocation						
	25th	0.00	0.00	0.05	0.14	0.00
	50th	0.01	0.17	0.44	0.66	0.34
	75th	0.70	1.16	1.42	1.84	1.42
	90th	3.75	3.24	3.92	3.50	3.50
Annualized dollar savings from re-allocation						
	25th	0	0	2	15	0
	50th	0	2	17	80	11
	75th	4	18	57	252	59
	90th	20	45	137	551	216
Savings as percentage of annual interest costs						
	25th	0.00	0.00	0.00	0.01	0.00
	50th	0.00	0.02	0.03	0.05	0.03
	75th	0.07	0.10	0.15	0.14	0.12
	90th	0.43	0.31	0.32	0.30	0.31

Notes: Sub-sample includes panelist with more than one account in Lightspeed data, and where number of cards in our data matches "active bankcard lines" from the credit report. "Actual APR" is weighted APR as in Table 2. "Best APR" re-allocates balances to lowest rate cards, conditional on available credit limits. "Average APR savings" is the average daily reduction in APR obtained by transferring all balances to lowest-rate cards, conditional on available credit. "Dollar savings" multiplies the average APR savings by average daily revolving balances. "Savings as percentage..." compares dollar savings from reallocation to the dollar interest costs from Table 1, at the panelist level.

Table A3. Search intensity and card-holding: first stage of IV.

	Dependent variable:				
Dependent variable	Credit cards (admin. data)	Credit cards (bureau)	Current mortgages	Current bank inst. loans	Current fin. inst. loans
Search-intensive customer (5-10 out of 10)	**0.361*****	**0.681*****	-0.161	0.089	0.073
	(0.109)	**(0.204)**	(0.225)	(0.246)	(0.295)
r-squared	0.49	0.27	0.19	0.11	0.17

Notes: Regression is at panelist level. All models include full set of controls from Tables 6 and 7, and have 495 observations. Coefficients are OLS estimates where the dependent variable is the panelist's number of cards/accounts. "Credit cards admin. data" is the measure of cards (accounts) in the Lightspeed data. "Credit cards (bureau) is the number of active bankcard lines on the credit bureau report. "Current mortgages," "bank installment loans" and "finance company installment loans" are counts taken from credit bureau data.
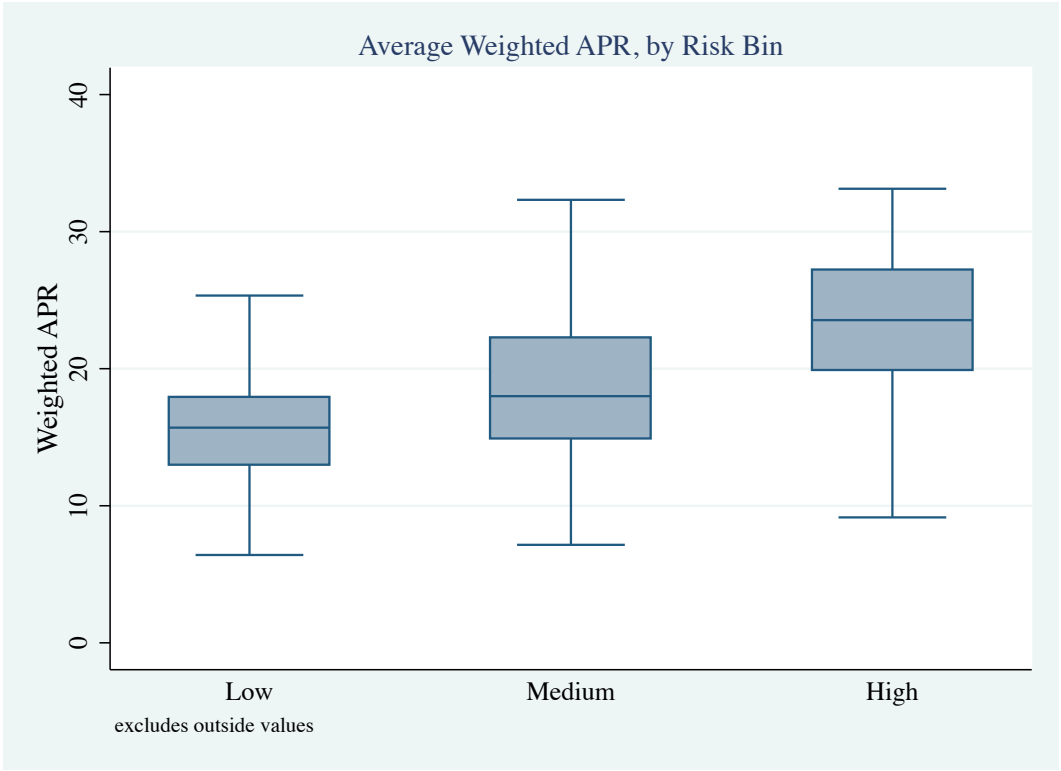
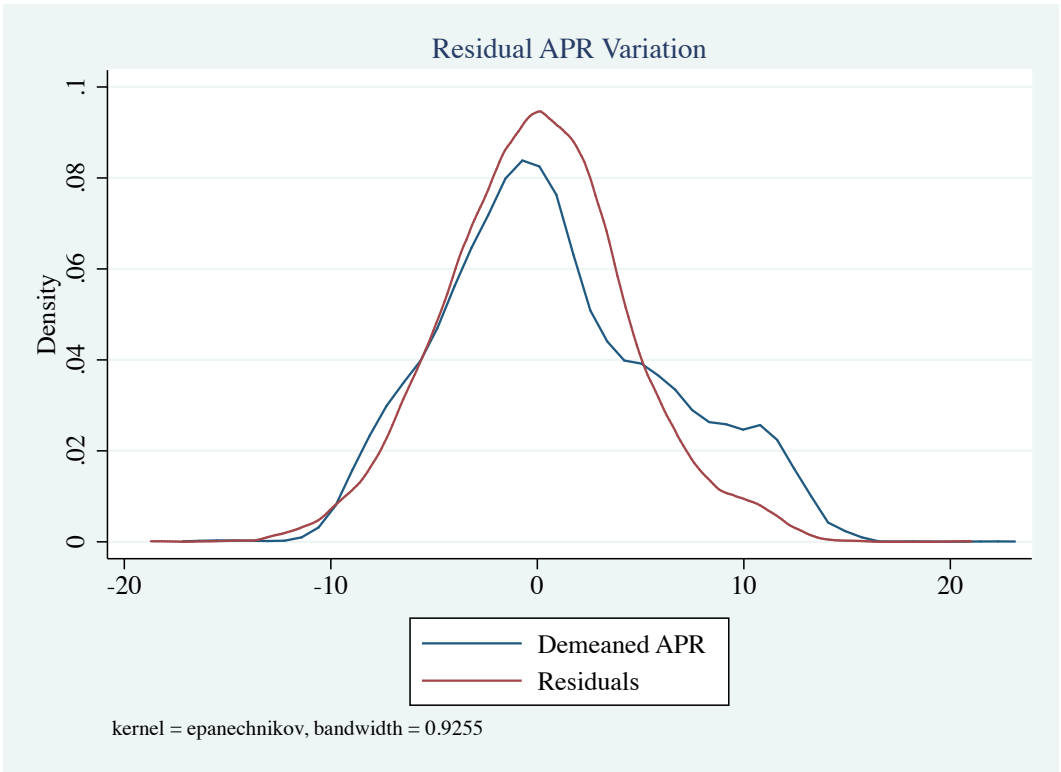Figure 1. Box-and-whisker plot of weighted APR, by risk category.

Figure 2. Raw and residual variation in weighted APRs.