

NBER WORKING PAPER SERIES

BORROWING HIGH VS. BORROWING HIGHER:
SOURCES AND CONSEQUENCES OF DISPERSION IN INDIVIDUAL BORROWING COSTS

Victor Stango
Jonathan Zinman

Working Paper 19069
<http://www.nber.org/papers/w19069>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2013

Thanks to Anna Lunn and Gordon Vermeer for research assistance, and to John Driscoll, Charles Sprenger, David Silberman, Justin Sydnor, Jeremy Tobacman, and conference/seminar participants at the Consumer Financial Protection Bureau, 2012 American Economic Association Meetings, Berkeley (Haas), Boston University School of Management, Cornell (Dyson School), the Federal Reserve Bank of Philadelphia, the Federal Reserve Board/George Washington Financial Literacy Seminar Series, ITAM, the NBER (Behavioral Economics, Household Finance), the UC-Boulder Conference on Consumer Financial Decision Making, UC-Davis, and Wharton (Decision Processes) for comments. Thanks to Filene Research Institute, the Federal Bank of Chicago, the Rockefeller Center at Dartmouth College and the Graduate School of Management at UC-Davis for research support. Special thanks to Bob Hunt and Steph Wilshusen, of the Payments Cards Center of the Federal Reserve Bank of Philadelphia, and to Geng Li of the Federal Reserve Board, for sharing data and expertise. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Victor Stango and Jonathan Zinman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Borrowing High vs. Borrowing Higher: Sources and Consequences of Dispersion in Individual Borrowing Costs

Victor Stango and Jonathan Zinman

NBER Working Paper No. 19069

May 2013

JEL No. D14,D22,D4,D83,G21,G23

ABSTRACT

We document cross-individual variation in U.S. credit card borrowing costs (APRs) that is large enough to explain substantial differences in household saving rates. Borrower default risk and card characteristics explain roughly 40% of APRs. The remaining dispersion exists because a borrower can receive offers and hold cards with wide-ranging APRs, as different issuers price the same observable risk metrics quite differently. Borrower debt (mis)allocation across cards explains little dispersion. But self-reported borrower search/shopping (along with instruments for shopping implied by Fair Lending law) can explain APR differences comparable to moving someone from the worst credit score decile to the best.

Victor Stango
Graduate School of Management
University of California, Davis
One Shields Avenue
Davis, CA 95616
vstango@ucdavis.edu

Jonathan Zinman
Department of Economics
Dartmouth College
314 Rockefeller Hall
Hanover, NH 03755
and NBER
jzinman@dartmouth.edu

I. Introduction

Asset-side research on financial decision-making and wealth accumulation asks both whether individuals “under-save,” and whether savers maximize asset returns.¹ In contrast, liability-side research on decision-making and wealth often focuses on whether households “over-borrow,” but gives short shrift to the returns question: do borrowers minimize costs of debt, conditional on how much they borrow?² Borrowing costs merit scrutiny because, for many households, liabilities are a more important determinant of wealth accumulation than financial assets.³

Despite the potential importance of links between household decisions, borrowing costs and household savings rates, research on the topic is underdeveloped and piecemeal.⁴ We address several important gaps in the context of the \$800 billion U.S. credit card market. Rich transaction-level administrative data, credit bureau data, and survey data grant us a uniquely comprehensive view

¹ See, e.g., Barber and Odean (2011) and Campbell (2006) for reviews.

² For discussions of over-borrowing see, e.g., Campbell (2006) and Benton, Meier and Sprenger (2007).

³ Aggregate U.S. household debt holdings are almost as large as aggregate financial asset holdings, and that was true even before the 2000s boom (Dynan 2009). Also, by most measures, more households participate in debt markets than in financial asset markets.

⁴ Several papers suggest that search and switching costs prevent people from shopping perfectly for loans but do not link such contract choices to the cross-section of borrowing costs. See, e.g., Agarwal et al (2006); Berlin and Mester (2004); Bucks et al (2009, A12), Calem et al (2006); Charles et al (2008); FINRA (2009); Kerr and Dunn (2008); Stango (2002); Woodward and Hall (2012). Other papers suggest that debt misallocation – holding higher-cost debt when lower-cost debt or liquid assets are available—may be important, but that work also does not link misallocation to the cross section of “returns” on the liability side. See, e.g., Agarwal, Skiba, and Tobacman (2009); Amar et al (2011); Ponce et al (2012); 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 2013). There is also work on (mis)allocation across (liquid) assets and liabilities, although heterogeneity in the non-price characteristics of the different products complicates that analysis; see, e.g., Gross and Souleles (2002b); Telyukova (2011); Zinman (2007).

of consumers' credit card choice sets, card portfolio holdings, card attributes, and borrowing/allocation decisions. The data allow us to address questions including: is individual-level heterogeneity in borrowing costs substantial? If so, does that heterogeneity simply reflect risk-based pricing and product differentiation by lenders? Or is heterogeneity in consumer decision-making, conditional on choice sets, an important driver of borrowing costs? And if consumer decisions matter, which ones are most important in explaining the cross-section of borrowing costs?

On the first question, we find striking cross-sectional variation in borrowing costs across U.S. credit cardholders. Even after discarding introductory “teaser rates” and “transactors” who never borrow, the balance-weighted interquartile range of APRs paid by credit card borrowers is about 800 basis points. Moving a typical consumer from the 75th to 25th percentile of borrowing costs, holding all else constant, would increase that consumer’s annual household savings rate by more than 1%. (For comparison, the national annual household saving rate during our sample period was about 3.5%.) In short, the cross-section of borrowing costs could explain substantial cross-sectional dispersion in savings rates. This stands in stark contrast to the *lack* of APR dispersion in the 1980s and 1990s⁵ that motivated much of the prior literature on credit card pricing and competition.⁶

Second, observable borrower default risk (commonly priced by issuers) and product differentiation together explain less than half of cross-sectional borrowing cost dispersion. Default risk – measured by credit scores, and in-sample risk metrics such as late payments, borrowing, credit limits, utilization and so on – explains roughly 40 percent of cross-sectional variation in borrowing costs. Product differentiation (card features such as rewards, fees, fixed/variable rate pricing, and so on) and demographics (age, income and education) add only

⁵ See Appendix Figure 1 for an illustration from 1983 and a comparison to 2007.

⁶ See, e.g., Ausubel (1991), Stango (2000; 2002), Knittel and Stango (2003).

slightly to the fit. We find little evidence of a substantial tradeoff between APRs and rewards or fees. The bottom line is that similarly risky borrowers, holding cards with similar characteristics and debt levels, pay substantially different rates. This finding holds both across issuers and within issuers.

Third, we examine cross-issuer differences and find important heterogeneity in the pricing of commonly observed individual-level risk metrics. This translates to consumers facing substantial APR dispersion in their choice sets, both “outside the wallet” and “inside the wallet.” Outside the wallet, data from direct mail solicitations show that APRs offered by different issuers to the *same individual during the same month* often differ by several hundred basis points. We also build pricing models that allow for issuer-specific differences in pricing credit risk, other consumer characteristics, and product characteristics. These models confirm that different issuers systematically price the same factors – even credit scores – differently enough to yield dispersion of several hundred basis points in the APRs faced by a given individual. Inside the wallet, we show that many individuals simultaneously hold multiple credit cards with APRs that vary widely.

Finally, we decompose demand-side sources of borrowing cost dispersion by examining the importance of consumer choices in and out of the wallet.

Inside the wallet, we find little misallocation or heterogeneity: most people allocate debt to the cheapest card(s) they hold, subject to credit limits. The median cardholder leaves close to zero dollars on the table *annually* via misallocation, and fewer than 10% of our credit card users leave more than \$100 on the table annually. “True” misallocation is probably lower when one considers pecuniary and non-pecuniary costs of moving balances from one card to another.

In contrast, heterogeneity outside the wallet – in shopping for credit cards – can explain substantial cross-sectional differences in APRs, even for similarly risky individuals. For a subset of our sample, we observe a self-reported measure of whether a borrower “keeps an eye out for better credit card offers,” allowing us

to estimate the relationship between APRs and search intensity, conditional on all other observable borrower and card characteristics. In estimating that relationship we use several instruments for search intensity, exploiting the fact that Fair Lending law prohibits card issuers from considering marital status or gender when setting APRs (there is a long marketing literature documenting strong relationships between gender or marital status and shopping behavior). Our empirical results are robust to the instrument set, and suggest that “super-shoppers” pay borrowing costs several hundred basis points lower than do non-shoppers – a difference comparable to that between individuals in the best vs. worst decile of credit score. Many consumers behave as though they face large search costs in the credit card market, a finding in accord with related work on mortgages (Campbell 2006; Woodward and Hall 2012).⁷

Our results inform several strands of research, starting with the literatures on credit card pricing, debt shopping and debt (mis)allocation discussed above. Beyond that, our findings represent a liability-side analogue to prior work on the asset side of the balance sheet, where several papers 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). Similarly, 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. Our findings also paint a more complex picture of interactions between consumers and card issuers than one might expect (similar to the idea in Taylor (2003)). That issuers can view the same customer so differently – and price accordingly – is somewhat surprising.

⁷See also <http://www.fanniemae.com/portal/research-and-analysis/housing-quarterly.html>.

Our results also highlight borrowing cost dispersion as a potential driver of cross-sectional dispersion in savings rates and wealth accumulation. While wealth accumulation depends on choices about both borrowing levels (the quantity channel) and APRs (the price channel), research and modeling often focus on the quantity channel, under-emphasizing the price channel as a first-order source of heterogeneity. Our work adds the price channel as a potential explanation for borrowing costs that seem puzzlingly high even by the standards of behavioral models.⁸ We cannot infer whether people are “over-paying” for credit card debt in a behavioral sense, or simply making optimal tradeoffs between shopping costs and benefits, but we hope that our paper will provoke inquiry along these lines.

On a practical level, our results highlight the potential value of innovations—e.g., in technology, standards, and/or disclosure—that lower the cost of loan shopping. We speculate on implications for policy and practice in the Conclusion.

II. Data

A. Data Contents and Sample Characteristics

Our data come from Lightspeed Research (formerly Forrester Research). Individuals in our sample are members of the “Ultimate Consumer Panel,” which is one of many such panels maintained by Forrester/Lightspeed.⁹

⁸ E.g., even with quasi-hyperbolic discounting, the calibrated life-cycle model in Angeletos et al (2001) substantially underpredicts credit card borrowing. Other work on links between borrowing behavior and behavioral biases in preferences, expectations, and/or price perceptions includes 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; 2011).

⁹ We also use the data in Stango and Zinman (2009a; 2013). Other Forrester/Lightspeed panels track consumer behaviors 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), Prince (2008).

The credit card data collected by Lightspeed have four main components. The first component is transaction-level and comes from monthly credit card account statements. The set of transactions includes all credits (payments, refunds) and debits (purchases, fees, interest charges, etc.) on the account. The second component is account-month level and contains data on account terms: APR, cash advance APR, bill date, due date, ending balance on bill date, summaries of credits and debits during the month, and so on.

In addition, we observe credit report data from one of the major bureaus, “pulled” at around the time of the panelist’s registration. The credit report data include data on “trades” (current and past loans of all kinds), delinquency, loan balances, and a credit score on the standard 850-point scale.

Finally, Lightspeed solicits and collects survey data from panelists. All panelists complete a registration survey in which they report demographics and financial characteristics. Lightspeed also periodically invites panelists to take online surveys. The data we use later in the paper regarding credit card shopping come from one of those periodic surveys.

Table 1 summarizes the data. Our data span 2006-2008, and in this paper our main sample consists of the 4,312 panelists who enroll at least one credit card account and for whom we observe credit bureau data. We stratify panelists by their quartile of average “revolving” (i.e., interest-accruing) debt to facilitate analysis that conditions on debt levels, to understand how heavy and light borrowers differ, and because our research questions are most salient for heavier borrowers. Within-panelist revolving debt levels are quite persistent, with a month-to-month serial correlation of 0.96.

Seventy percent of panelists enroll one or two accounts, and the remaining thirty percent enroll three or more. Roughly half of our sample enrolls a “complete” set of credit cards, meaning that the number of accounts enrolled

matches the number of “active credit card lines” on the panelist’s credit report. Appendix Table 1 shows that complete-set panelists look quite similar to the full sample (compare to Table 1). This alleviates selection/measurement concerns, and suggests panelists with “incomplete” sets register the cards that they use regularly.

The 1st quartile of revolving debt contains many “transactors” who essentially never revolve balances but use their cards for purchases. Consumers in the 3rd and 4th quartiles are heavier “revolvers” who consistently carry balances. For these revolvers interest charges are 81% and 92% of total borrowing costs.¹⁰

Interestingly, we see many similarities between individuals in the highest and lowest quartiles of revolving debt. Purchase volume, credit scores, and education are each U-shaped with respect to revolving balances. Income increases with revolving debt. We also see the expected life-cycle pattern, with those in the middle of the age distribution carrying more debt.

Table 1 also shows that credit card interest paid can be substantial relative to income. Interest costs for the median *individual* in the heaviest-borrowing quartile are 2.4% of annual *household* income, and exceed 1% for one-quarter of our full sample.¹¹ As context, the national average annual household savings rate was about 3.5% both during our sample period and over the 10 years prior. So any money left on the table via higher borrowing costs could materially affect household savings rates.

¹⁰ The remainder of borrowing costs comes from annual, late, over-limit, cash advance, balance transfer and other fees. See Stango and Zinman (2009a) for further detail on fees in these data.

¹¹ Restricting the sample to “single, never married” panelists leaves the results unchanged. Only 7% of respondents report registering a card belonging to someone else.

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

The Data Appendix provides many additional details on panel construction, variable definitions, and sample characteristics.

B. Representativeness and External Validity

Our credit data benchmarks reasonably well against various other data sources (the Data Appendix provides detailed comparisons). Our sample is similar to the U.S. population in terms of cardholding, purchases, creditworthiness, APR distribution, and interest costs relative to total borrowing costs. The one key difference is that our cardholders have outstanding balances that are about half the national average. Given that our analysis focuses on identifying borrowing cost dispersion *conditional on debt amounts*, any “missing debt” will lead us to understate the potential impact of borrowing cost dispersion on saving rates.

In terms of demographics, our panelists are younger, more educated, and higher income (conditional on age) than national averages.

The online nature of the panel might affect inferences about the broader population of cardholders, if “being online” is correlated with shopping or allocating debt efficiently. To the extent that our sample is more homogeneous than the population by dint of being younger, “online” and willing to participate in the panel, our results could easily *understate* the level of diversity in borrowing costs, shopping behavior and debt allocation in the population.¹²

¹² Panel participants are also relatively willing to share financial information, raising questions about whether they might be unrepresentative in other, unobservable but

The time period under consideration here, 2006-2008, is also noteworthy. We do not observe a decline in borrowing cost dispersion in the early stages of the financial crisis. Nor do we know of any reason to expect that our results—which are mostly about dispersion—would differ in calmer times, but this is clearly something worth exploring in future research.

III. Cross-sectional Variation in Borrowing Costs

A. Measuring Credit Card Borrowing Costs, With and Without Float and Teasers

We measure borrowing costs for each panelist as the average balance-weighted annual percentage rate (APR) over our sample period. Balances accrue interest charges if they are “revolving”: not fully repaid after the due date of the bill. We focus on APRs because they constitute >80% of borrowing costs for heavier borrowers (Table 1).

The first rows of Table 2 show APR dispersion over revolving and non-revolving (zero-APR) balances. Our primary focus is on revolving APRs, so the next rows exclude the 627 panelists (15% of the sample) who never revolve balances during our sample period. APR dispersion is substantial within every borrowing quartile and similar across the top three, with interquartile ranges of 800-900 basis points (bp), and 10th/90th percentile ranges of 1600-1700bp.

The next rows, and most of the analysis below, discard the account-months we classify as paying “teaser” (introductory) rates (see the Data Appendix for details). Dropping teaser rates from the data has little effect on dispersion.¹³

critical ways. But the same could be said about any data source—including household surveys-- that relies on a clear opt-in from subjects.

¹³ Teaser rates have a negligible effect on cross-individual dispersion for three reasons: 1) teaser rates typically last only six months or so, and represent only a small proportion of

The subsequent rows show that dispersion in the sub-sample of panelists who enroll all of their cards is nearly identical to that in the full sample.

The last rows motivate our focus on the cross-section rather than the time series of APRs. We first present, as an illustration, data from a single month of our data (January 2006). In this month – and every other month – dispersion is virtually identical to dispersion averaged across all months. Second, the last row shows that regressing panelist-month-level APRs on just a set of panelist fixed effects yields a fit of nearly eighty percent. Most APR variation is in the cross-section of panelists rather than within-panelist over time.

B. Scaling the Magnitude of APR Dispersion

APR dispersion could matter a lot economically. Take a borrower at the medians for income, interest costs, and revolving debt in our top revolving debt quartile (keeping in mind that outstanding balances for such an individual in our data are equal to the outstanding balances of the *median* cardholder in the U.S.). That individual's savings rate could rise by 1.2 percentage points if borrowing costs fell from the 75th percentile of APRs to the 25th, or by 1.8pp if borrowing costs fell from the 90th percentile to the 10th. Alternatively, the same individual could hold total interest costs constant with \$4,000 (\$10,000) in additional debt or consumption, moving from the 75th to the 25th (90th to the 10th) percentile. Our APR dispersion seems representative, so these magnitudes should be relevant for U.S. cardholders more broadly (see the Data Appendix Section E for details).

account-months; 2) people have multiple cards, and a mix of teaser and non-teaser rates, at any point in time; i.e., even though one can sort *account-months* into teaser vs. non-teaser, the extent to which this sorting aggregates to the individual level is muted; 3) we do not actually find a significant tradeoff between introductory and post-teaser APRs. Anecdotally, most teaser rate offers recoup their lower APRs via 2-5% balance transfer fees rather than higher post-teaser APRs.

The potential savings-rate implications for heavy borrowers here are slightly smaller than losses incurred by individual investors due to excessive trading in Barber et al. (2009); they are larger than the 75th percentile of losses from investment mistakes among asset holders in Campbell, Calvet and Sodini (2007); they are similar to losses from sub-optimal 401(k) account contributions in Choi et al (2011); and they are larger than (amortized) losses from insufficient mortgage shopping in Woodward and Hall (2012).

IV. Borrowing Cost Dispersion: Explanations and Empirical Strategies

A. Broad Explanations and Data Requirements

What might explain the substantial cross-sectional APR dispersion in Table 2?

Consider two borrowers, Gretchen and Mary. Assume each has two credit card accounts with a total credit limit of \$10,000, and each revolves an average of \$6,000 across those two cards. If we find that Gretchen pays 22% APR on average and Mary pays 14%, what might explain the difference?

Broadly, there are two classes of explanation. One class holds that Gretchen and Mary face different prices from the market because they are *differently risky*, or use cards that are *differentiated products*; if so, their choice sets cannot be compared apples-to-apples. The other broad class of explanation is that Gretchen and Mary in fact face similar prices from the market (or similar distributions of prices), but *make different choices* given the same choice set.

Disentangling those explanations, and assessing their relative importance, requires rich data. One must observe what issuers observe re: customer default risk and card characteristics, and be able to infer how those things are related to the APRs that customers face in the market (and hold in their wallets). The analysis would be enriched with observations of consumer choices with respect to shopping for and using cards.

Our data are up to the task. Our ability to measure risk with credit bureau, usage, and repayment data compares favorably to issuers' abilities; for example, issuers observe only partial usage and repayment data for their cardholders' other accounts, while we observe those data in detail and in closer to real-time. We also observe characteristics associated with product differentiation: issuer identities, and details on card attributes such as credit limits, rewards, fees, and "fixed" vs. variable APRs. Collectively, these data allow us to control quite comprehensively for default risk and product differentiation.

We also observe detailed information about repayments (allocation), and about some borrowers' self-reported shopping for credit cards (contract choice). Card repayment and usage data allow us to examine the importance of decisions "within the wallet": does Mary pay lower APRs because she allocates debt to her low-rate cards better than does Gretchen? Shopping data allow us to assess the importance of behavior "outside the wallet": does Mary pay lower APRs because she has a keener eye for better outside offers from issuers?¹⁴

B. Empirical Strategies

Our empirical approach to exploring the (relative) importance of these different explanations proceeds as follows:

1. First, we use our default risk and card characteristic data to estimate how much those factors contribute to cross-sectional APR variation (Section V). The key statistics are r-squareds revealing how much variation can be explained by observable borrower default risk, unobserved issuer-specific factors, and observable card characteristics;

¹⁴ Most customers still apply for new cards in response to direct mail solicitations, although the online channel is growing. Issuers mailed over five billion solicitations during our sample period, and most of our panelists report having acquired their card(s) in response to direct mail.

2. Next, we assess a threshold question regarding whether borrower behavior might affect borrowing costs: does a given borrower face different APRs from different issuers, even at the same point in time (Section VI)? That must be true either for borrowers' shopping to matter or, strictly speaking, for misallocation "in the wallet" to be possible. Otherwise, the Law of One Price would dictate that borrower behavior is irrelevant in determining APRs held (and paid);
3. Last, we examine the extent to which misallocation in the wallet and shopping can explain the cross-section of APRs paid (Section VII). We also ask how holding more cards in the wallet affects both the APR cost from misallocation and the lowest rate in the wallet.

V. Borrowing Cost Dispersion Conditional on Risk, Product Heterogeneity

A. Specifications: Models Explaining the Cross-Section of Borrowing Costs

The most natural explanations for cross-sectional variation in APRs are default risk and card-level product differentiation.

Our data include much, if not all, of the information used by issuers when setting and adjusting APRs, as well as significant detail about card characteristics. We observe credit score, supplementary credit bureau data (e.g., the number of current and past "lines" of credit of varying kinds), purchase volume and revolving balances, in-sample late/missed payments, credit limits and utilization, demographics (age, income and education categories), fees (annual, balance transfer, cash back, others), rewards and affinity links, and fixed/variable rate pricing. The Data Appendix provides additional details on variable construction.

To assess how well these covariates explain APRs we estimate a series of panelist- and account-month-level models with APRs as the dependent variable, using all of our available data regarding risk factors and product characteristics as

flexibly parameterized covariates. In panelist-level regressions we include panelist-level aggregates, as well as characteristics of panelist's primary card by average revolving balances. The panelist-level models include balance-weighted issuer fixed effects, accounting for average APR differences across issuers stemming from omitted card characteristics, systematic differences in pricing customer risk, and other unobservables. The account-month models include issuer and month-year fixed effects. All models include indicators for panelists' first and last months in the data, accounting for variation in APRs generated by systematic time-varying APRs and panelists' different sample entry/exit dates.

B. Results and Robustness

Table 3 reports the fit of the APR regressions. The broad takeaway is that observable risk and card/issuer characteristics explain 30-40% of cross-sectional variation in borrowing costs. Credit scores alone explain 5-20% of cross-sectional variation in APRs. Including in-sample risk measures adds substantially to the explanatory power of the model, in most cases allowing the model to explain 25-40% of cross-sectional variation. This compares favorably to analogous work predicting credit card delinquency (Gross and Souleles 2002a; Allen, DeLong, and Saunders 2004). Card characteristics and demographics add very little to the explanatory power of the model. Reading across columns, the models do a better job fitting APRs for heavier borrowers than for "transactors." And the panelist-level models generally have better fit than the account-month models.

Appendix Table 2 shows regression coefficients from our best-fitting panelist-level specification (the fourth r-squared row in the last column of Table 3). Because our focus is on improving fit rather than parsimony, we include many sets of covariates that are highly collinear. For example, the model includes revolving balances and credit lines – which together are very highly correlated with utilization – and also includes utilization as well. So the results on

many individual variables do not have clear interpretations. Nevertheless, it bears noting that we do see the expected strong results on credit score and late fees.

Returning to Table 3, the bottom rows compare model fit early in the sample vs. late, highlighting the tradeoffs between informativeness of our credit bureau data (observed at the beginning of the sample period) and informativeness of our in-sample risk metrics (which grow more comprehensively backward-looking by the end of the sample period). Credit scores explain APRs no better overall when timely than when “stale.” This probably reflects the stylized fact that credit scores are very stable within-person, over time. The in-sample risk measures are more informative by the end of the sample, though not dramatically so. Overall, the adjusted fit increases from 0.28 to 0.33 between the first six months of the sample and the last six, again suggesting that panelist-level behavior, and hence risk, is strongly auto-correlated (also recall the within-panelist serial correlation in revolving balances of 0.96). Yet another symptom of this is that the *time-invariant* panelist-level in-sample risk variables explain nearly as much variation in APRs in the “first six months” model (adj. r-squared=0.25) as do the same variables in the “last six months” model (adj. r-squared=0.31), even though the former are almost purely derived from *as-yet-unobserved future* panelist behavior, while the latter are based completely on 3 years of *directly observed recent past* behavior.

Even our richest model leaves more than half of APR dispersion unexplained. Figure 1 illustrates this, showing both the raw (de-meaned) variation in borrowing costs and the residual variation. The inter-quartile range in residual variation is 500 basis points, and the 10th/90th range is 1000bp.

A natural concern is that some of the residual dispersion is driven by omitted variables that are commonly priced by all issuers (any variable that is differently priced across issuers makes consumer decision making about contract choice and allocation important, as we explore in the next sections). We consider this possibility somewhat remote because we do observe the market-wide standard

measure of risk (the credit score), which at panel entry provides a summary measure of all pre-panel signals about risk. We also observe detailed in-sample data on granular behaviors (late/missed payments, utilization and over-limit instances) that are the other primary pricing factors. In all we may observe *more* account-level information than do issuers; e.g., a given card issuer may only observe details on its own accounts, while we observe details for all accounts held by a particular panelist. We also observe data at higher frequency than issuers do, because of reporting lags between issuers and the bureau. Conversations with bankers and industry experts offer reassurance that our model captures nearly all of the key risk metrics and product characteristics that affect pricing.¹⁵

Another concern is that our functional form might not capture the true relationship between these variables and APRs. But our model is extremely flexible – we parameterize nearly every right-hand side variable into deciles or similarly flexible categories. We have also estimated even more flexible specifications, with interactions, to the point of over-fitting, in that these models reduce the adjusted r-squared substantially.¹⁶

Another contributor to unexplained variation in APRs could be randomization by issuers (see, e.g., Day 2003; Shui and Ausubel 2004). We cannot empirically distinguish between randomized pricing and the omitted credit risk story, but

¹⁵ We observe account ages (years since opening) for a subset of panelists, and in that sub-sample do not find account age to be significant correlated with APR. This is unsurprising given that: 1) many panelists have a mix of older and newer accounts; 2) we do observe panelist age, which is correlated with account age at the panelist level; 3) issuers can reprice accounts over time (implying that we wouldn't expect to see APRs that were initially low due to macro conditions "stick" over time); 4) consumers can close any sticky-high APR accounts over time, or move balances out of them (implying that the effects of any sticky-high accounts would be muted).

¹⁶ Another possible issue might be cards shared across individuals, but restricting the sample to "single, never married" panelists leaves the results unchanged. Only 7% of respondents report registering a card belonging to someone else.

intuition argues against randomized prices as a *primary* driver of borrowing cost dispersion, given the considerable resources that issuers expend in developing proprietary internal risk models.

It is also possible that “relationship banking” – benefits granted to cardholders because they also hold, e.g., a deposit account or mortgage with the same bank – could affect APRs (Sumit Agarwal et al. 2009). But in survey responses only 3% of our panelists report paying a lower credit card APR due to relationship banking.

Overall, our finding of substantial cross-sectional dispersion in borrowing costs seems robust to various ways of controlling for credit risk and product differentiation. Nevertheless, we grant that our models fitting the cross-section of borrowing costs might be imperfectly specified. We therefore pursue a complementary approach, one focusing directly on the possibility that different issuers price the same risk characteristics differently, leading similarly risky borrowers to face different prices in the market and hold cards at different APRs, and making borrower behavior in and out of the wallet – allocating debt across cards, and card shopping – an important determinant of dispersion in borrowing costs.

VI. Choice Sets: APR Dispersion in the Market and in the Wallet

A. Offer APR Dispersion

Our first evidence of within-individual offer dispersion comes from a separate dataset on the terms of credit card mailers from Mintel Comperemedia.¹⁷ The Mintel data allow us to measure dispersion in offers received for a particular

¹⁷ We are extremely grateful to Mintel, and to Geng Li at the Federal Reserve Board of Governors, for allowing us to share summary statistics from these data. A paper by Li and coauthors (Han, Keys, and Li 2011) contains more detail about these data.

individual *in a specific month*; that is a lower bound on dispersion measured over a longer time period. Looking at within-month offer APR dispersion eliminates any confounding effect of time-varying credit risk at the individual level. We focus on January 2007 in particular: January because it is a peak month for mail solicitations by credit card issuers, and 2007 because it sits in the middle of our Lightspeed sample period. We condition on having received more than one credit card offer during January 2007, dropping roughly 25% of individuals and leaving us with 1,211 people who received a mean (median) of 4 (3) credit card offers.

To illustrate within-individual dispersion in offers, Table 4a shows the distribution of within-individual differences between the highest and lowest APR offers, calculated two ways.¹⁸ The first APR is the contract or “goto” APR – the APR after any teaser period expires (column 3). The second APR is an estimated “net-of-teaser” APR, which is the 24-month weighted average of the teaser and goto APRs (column 6).¹⁹ The median within-individual and within-month high-low goto (net-of-teaser) rate spread is 434 (750) basis points, and the seventieth percentiles are 725 (986) basis points. These measures of dispersion must, mechanically, be at least weakly larger over longer time periods – longer time periods that are still short enough such that within-person variation in creditworthiness is trivial for nearly all consumers. In short, it is common for an individual to receive credit card offers at very different APRs.

¹⁸ The distribution of APRs shown here lies below that in our data, because these are initial offers and do not reflect the upward shift in APRs that occurs in the group of cardholders who are repriced or incur a penalty rate after accepting the initial offer.

¹⁹ If, for example, the teaser APR is zero for six months and the goto rate is 2000 basis points, the net-of-teaser APR equals $(6/24)*(0)+(18/24)*2000=1500\text{bp}$.

B. Cross-Issuer Heterogeneity in Risk-Based Pricing

Our second analysis of within-individual dispersion in choice sets uses the Lightspeed data to estimate cross-issuer heterogeneity in risk-based pricing. Relatively little is known about such heterogeneity, and whether it leads to significantly different APR offers for a particular individual, in part because issuers invest considerable resources in their internal modeling and view their models as valuable trade secrets. In some sense, of course, the fact that issuers expend significant resources is *prima facie* evidence that different internal models yield different “optimal” APRs for a given individual; otherwise, why invest in the models? Nonetheless, we know of no academic work documenting or estimating the magnitude of this heterogeneity.

Appendix Figure 2 illustrates cross-issuer heterogeneity by plotting distributions of the credit score/APR relationship for each of five large (anonymized) issuers in our data, and also for a sixth “all other issuers” group.²⁰ The plots illustrate three sorts of heterogeneity across issuers, all of which are substantial. The first is that, even within a credit score decile, different issuers can have APR *levels* that differ by several hundred basis points (e.g., compare the horizontal lines denoting the median rate across issuers for the same credit score decile). Another type of cross-issuer heterogeneity is in the credit score-APR *gradient*: the decline in APRs from the worst to best decile. The third type of cross-issuer heterogeneity is in the extent of APR *variance* within credit score deciles. At the least, these types of differences indicate differential emphasis on credit scores vs. other information (such as late payments) in pricing risk.

More formally, in order to quantify the potential impact of these differences on cross-panelist borrowing cost dispersion, we take the simplest or richest

²⁰ Credit scores are defined based on the entire distribution of APRs, so “decile 1” for different issuers captures exactly the same range of scores.

account-month model in Table 3 and allow for issuer-specific coefficients on risk factors for each of the largest six issuers in our sample (which collectively make up 85% of cards in our sample and 75% nationally).

We use the coefficients from these issuer-level pricing models to predict implied APRs for every panelist, in each month, from each of the six issuers. The hypothetical is “what would the set of APRs from these six issuers be, given panelist X’s characteristics and the month-year of the data?” We then calculate the gap between the highest and lowest of these implied APRs, for every panelist in every month. The thought experiment is to ask how much within-individual variation in APRs can be generated simply via systematic cross-issuer differences in risk-based pricing. Note that because our models include month-year effects and estimate within-month high-low differences, time series variation in issuers’ pricing does not contribute to our estimate of within-individual price dispersion.

A useful feature of this approach is that it is quite conservative. It treats all smaller issuers as pricing identically, and we actually exclude the “all other issuers” category from our dispersion calculations below. It treats each larger issuer as applying a single pricing model, when in fact many large issuers employ different models, even internally, for a variety of reasons, one being legacy effects from acquisitions of other issuers with different models. Finally, it is possible that our specification is less flexible than that actually employed by a given issuer, which makes our fitted APRs less dispersed than the ones an issuer would actually set.

Table 4b shows the key results of this exercise: the implied high and low APRs, and the high-low spread. We show data from January 2007 to facilitate comparison with the Mintel data in Table 4a. Dispersion from our predictions (Table 4b) is even greater than that in the Mintel data (Table 4a), perhaps because heterogeneity in ex-post repricing compounds heterogeneity in ex-ante pricing. In any case, the central takeaway is that both prediction model specifications in

Table 4b – “all covariates” and “credit score decile only” – imply substantial price dispersion based simply on differential treatment of identical customer characteristics by the largest six credit card issuers. Even the 10th percentile of the high-low difference is an estimated 500 or 600 basis points. The 90th percentile is estimated at about 1300 basis points in both specifications.

C. Within-Wallet Dispersion

Table 4c shows actual within-wallet APR dispersion for our panelists, measured as the difference between the highest and lowest APRs held for each panelist in January 2007. The first three columns describe dispersion for the entire set of panelists. This is useful as a reference point, but necessarily includes many zeros – dispersion must be zero for any panelist with only one card. Columns 4-6 show in-wallet dispersion for the subset of panelists with more than one card.

These data show substantial within-wallet dispersion for many panelists. Among those with >1 card, the median is 400 basis points. Twenty-five percent of panelists with >1 card hold cards differing by more than 766 basis points.

In all, the evidence in Tables 4a-4c strongly suggests that any given individual receives offers at very different APRs from different issuers, and that many individuals actually hold cards with very different APRs. We reiterate that these estimates are probably conservative.

Price dispersion in choice sets creates the possibility that consumer behavior contributes to conditional dispersion in marginal borrowing costs. If borrowers shop differentially for cards, then better shoppers will obtain lower rates from the distribution of possible APRs. Similarly, if a given borrower holds cards in the wallet with different APRs, effectively allocating debt to lower-APR cards will reduce borrowing costs. We now examine these issues.

VII. Borrower Behavior and Borrowing Costs: Allocation and Shopping

A. Potential Savings from Re-allocating Revolving Debt

Table 5 examines the within-panelist allocation of total revolving balances across different cards. We first calculate the panelist-level “best weighted APR” that would apply if all debt were always allocated to the lowest-rate card in the wallet, up to the credit limit of each card, including teaser rates and fixing the levels of revolving and non-revolving debt. For example, if someone is floating \$1,000 on one card, revolving \$2,000 on another card, and has a zero balance on a third card, we calculate the cheapest APR the panelist *could* pay to revolve that \$2,000, across the APRs on all 3 cards on that day, subject to the credit limit on each card. Subtracting the “best weighted APR” from the “actual weighted APR” yields “APR misallocation,” which we scale 3 different ways: as an APR, as annualized dollars, and as a share of annual interest costs.

Nearly all panelists allocate debt quite efficiently (Table 5). Among all panelists including those with only one card, the 50th percentile of misallocation is zero. The 75th percentile of APR (dollar) misallocation is 48bp (\$8/year), and the 90th is 245bp (\$84/year). Heavier borrowers have greater misallocation in APR and dollar terms, while the share of total interest costs incurred via misallocation is higher among the lightest borrowers.

Limiting the analysis to “complete cards” panelists, for whom we observe all possible misallocation, reveals somewhat more but still modest misallocation. The median is 2bp and the 75th percentile is 114bp (\$20/year). Even in the top borrowing quartile the median is only \$35/year. Economically meaningful misallocation exists at the highest percentiles of the heaviest borrowing quartile, but is pretty rare overall.

Even these estimates of “misallocation” costs, modest as they are, are probably *upper bounds* on money left on the table even in a fairly strictly pecuniary sense. Our measured misallocation ignores credit card rewards (miles,

points, cash) that might render a card more expensive in APR terms but less expensive net of rewards (a pattern supported in some unreported analysis). Further, re-allocation can incur balance transfer fees, in which case some measured “misallocation” could be optimal conditional on balance transfer fees.

We also, in unreported results, have examined allocations of “excess repayments”: payments greater than the monthly minimum (Ponce et al 2012). 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, and efficiency increases with payment size.

B. Shopping/Search Behavior and Borrowing Costs: Descriptive Data

We now examine the link between card shopping and borrowing costs. For a subset of panelists (n=603), we observe agreement (on a 10-point scale) with the statement “I always keep an eye out for better credit card offers.” Panelists supplied responses via one of the periodic and voluntary surveys emailed to panelists by Lightspeed; the survey was administered in the first quarter of 2007.

Table 6 summarizes shopping responses grouped into four categories. The bottom row shows that 34% of panelists report 1-3 on the 10-point scale (“non-shoppers”), 30% report 4-6 (“medium shoppers”), 26% report 7-9 (“high shoppers”), and 10% report 10, the strongest agreement (“super-shoppers”).

The top rows of Table 6 show that self-reported shopping intensity correlates sensibly with other variables that might reflect shopping: current cards held, previous (now closed) accounts, and recent credit card applications. For example, only 15% of non-shoppers hold 5 or more cards, while 33% of super-shoppers

hold 5+ cards; 16% of non-shoppers have 15+ past cards, while 40% of super-shoppers have 15+ past cards; and 5% of non-shoppers report having applied for 2+ cards recently, while 26% of super-shoppers report the same thing.

The next sets of rows provide descriptive evidence previewing our instrumental variables results below: shoppers pay lower APRs, conditional on credit characteristics.

The last columns compare survey respondents to non-respondents. Non-respondents have fewer current/past cards and recent applications, are less creditworthy, borrow more and pay higher APRs. These differences caution against extrapolating our results below from respondents to non-respondents.

C. Shopping Behavior and Borrowing Costs: Regressions

Can individual-level differences in shopping behavior explain meaningful differences in borrowing costs? We examine this question by adding the 10-point shopping intensity variable, in linear form, to our main panelist-level specification from Table 3 and Appendix Table 2.²¹

The key identification issue is that shopping may be endogenous; a high APR “shock” (in the form of APRs in the wallet that a panelist views as higher than his/her risk warrants) might increase shopping effort, and thereby upward-bias the estimate of the relationship between shopping and borrowing costs.

To deal with endogeneity we instrument for the shopping variable with up to three panelist characteristics: marital status, gender and survey-taking behavior (panelists get invited by Lightspeed to take short surveys, about once per quarter).

²¹ We have experimented with other functional forms (e.g., fewer ordinal categories, dummies for “shopper” vs. “non-shopper” at different thresholds) with similar results.

We choose marital status and gender as instruments because they satisfy the exclusion restriction by law: the Equal Credit Opportunity Act (ECOA) prohibits lenders from (price) discriminating based on marital status or gender, regardless of intent. E.g., Fair Lending examiners monitor compliance by testing lenders for “disparate impact”: conditional correlations between protected characteristics and credit outcomes. Lenders have strong incentives to pass these tests; i.e., to ensure they satisfy our exclusion restrictions.²² We choose survey-taking because it too plausibly enters the model only via its indication of search behavior – lenders do not observe survey-taking, and we cannot think of any omitted risk factor for which it might proxy. Exogeneity also requires that APR shocks do not change the IVs themselves. This almost certainly holds for gender and marital status.

Appendix Table 3 sheds light on the first stage for each of the three instruments, in the raw data: single panelists and male panelists search substantially more, and the number of periodic Lighstpeed surveys taken is correlated with shopping intensity (negatively). There are many possible explanations for these correlations; we simply note that gender and marital differences in shopping have long been observed in marketing research,²³ and that one can imagine that online survey-taking could be correlated with online search. Below we treat the first stage more formally and carefully, by systematically varying the instrument set and reporting coefficients and confidence intervals robust to the weak instrument problem.

Table 7 presents our estimates of the effect of shopping on borrowing costs. The first column reports OLS results, while columns 2-6 present IV results for permutations of the instrument set. For each IV specification we report the

²² See http://www.federalreserve.gov/boarddocs/supmanual/cch/fair_lend_over.pdf for Fair Lending guidance pertinent to our sample period (and today).

²³ See, e.g., Laroche et al (2000) and references therein.

standard IV point estimate and standard error, the p-value of the CLR/AR test for whether the test of the null is robust to weak instruments, and the associated corrected-confidence interval for the coefficient on the endogenous regressor.²⁴ The three over-identified models pass Sargan exogeneity tests (see bottom row).

The OLS results reveal no strong relationship between shopping and APRs, but the IV specifications suggest a large, negative effect of search intensity on APRs that is robust to the instrument set and to weak instruments. The point estimates suggest a roughly 100bp reduction in borrowing costs per “point” of shopping intensity. The IQR of shopping intensity in the sub-sample is 5 points – from 2 to 7 – implying a reduction in borrowing costs of 500bp by moving from the 25th to 75th percentiles of shopping intensity. This effect is comparable to that generated by cross-sectional variation in observable and commonly priced default risk; e.g., moving from the 25th to 75th percentile of credit score is also correlated with a 500-600bp reduction in borrowing costs. Note also that the cross-sectional variation implied by differences in shopping dwarfs that deriving from misallocation. The interquartile range of misallocation is 48 basis points in our data: roughly one-tenth the amount implied by differences in self-reported shopping behavior.

So why doesn't everyone shop more? Framed differently, why do many people behave as if they have very large search costs? Answering this question convincingly will require much additional research, but our results suggest a partial and reduced-form explanation: borrowers may face tradeoffs between shopping efficiently and allocating efficiently. Appendix Table 4 explores this by focusing on the number of cards held, since active shoppers hold more cards (Table 6). Holding more credit cards is associated with better APRs in the wallet

²⁴ See Finlay and Magnusson (2009) for a discussion of the weak instrument problem and the Stata routine we use to deal with the issue.

(middle panel of Appendix Table 4), but also with greater misallocation in the wallet. On net, these correlations almost offset: the conditional correlations between cards held and panelist-level borrowing costs are small and weak statistically (Appendix Table 2).

VIII. Conclusion

We document and decompose cross-consumer dispersion in credit card borrowing costs that is large enough to materially affect household saving rates, even after controlling for debt levels, credit risk, and product characteristics. Our results suggest that dispersion is generated by the intersection of heterogeneity in issuer pricing and heterogeneity in consumer contract choice: different issuers offer different APRs to the same individual, and differences in consumer shopping behavior lead otherwise identical consumers to choose contracts at widely differing rates. Little of the cross-sectional dispersion is due to heterogeneity in how consumers allocate debt across their portfolio of cards, even though there is in fact substantial “within-wallet” APR variation.

Our estimates of borrowing cost levels and dispersion, and hence of the potential impact of borrowing cost heterogeneity on the distribution of household savings rates, probably err on the conservative side. Our sample seems to revolve substantially less credit card debt than the broader population, and it may be more efficient in its financial decision making by dint of being online, more-educated, and higher-income. Our sample may also be relatively homogeneous, by dint of everyone being online, and consenting to participate in the panel that generates our data. Moreover, we observe only credit card borrowing costs, and not costs in other, even larger debt markets: mortgages, auto loans, and student loans.

Our results inform interventions designed to help improve credit market outcomes. If credit shopping is more malleable than creditworthiness (credit

scores are quite sticky), then helping people shop for cards may be a relatively effective focus for interventions.²⁵ This is not to say that our results support any particular policy, programmatic, or business tack: they are silent, for example, on how or how cost-effectively one could affect search behavior, and on what the general equilibrium effects of any such innovation would be.

We close with five closely related directions for further research. One is on “borrowing higher” vs. “over-borrowing”; e.g., on the relative importance of, and relationships between, the price and quantity channels in explaining the cross-section of savings rates and wealth accumulation. Second is unpacking what drives search behavior: heterogeneity in standard preferences for leisure, in one or more behavioral factors, in skills/endowments, etc.? Optimal policy and practice may depend on those primitives, and one important task going forward is reconciling the substantial heterogeneity and inefficiency we find on the contract choice margin with the substantial homogeneity and efficiency we find on the allocation margin. Third is identifying how issuers respond to search behavior, and how issuers and consumers interact in equilibrium. Fourth is examining how consumers allocate attention, including shopping effort, across multiple domains in household finance: is attention to different areas positively correlated in the cross-section, or do people substitute attention in one area for attention in another? Our results here suggest tradeoffs between allocation margins and contract choice margins. Fifth is building sharper links between consumer credit market outcomes and wealth accumulation. Much work remains to characterize the nature and sources of financial product price dispersion, and its implications for incidence and efficiency.

²⁵ This harks back to early work on consumer protection in debt markets, which typically focused on improving comparison shopping (National Commission on Consumer Finance 1972). The importance of consumer-specific pricing in the credit card market suggests that “Smart Disclosure” could be useful (e.g., Kamenica, Mullainathan, and Thaler 2011).

REFERENCES

- Agarwal, S., S. Chomsisengphet, C. Liu, and N. Souleles. 2006. "Do Consumers Choose the Right Credit Contracts?"
- Agarwal, Sumit, Souphala Chomsisengphet, Chunlin Liu, and Nicholas Souleles. 2009. "Benefits of Relationship Banking: Evidence from Consumer Credit Markets". Working paper.
- 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 and Credit 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* 48 (Special Issue on Consumer Financial Decision Making): S38–S50.
- 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 M., Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean. 2009. "Just How Much Do Individual Investors Lose by Trading?" *Review of Financial Studies* 22 (2) (February 1): 609–632.
- Barber, Brad, and Terrence Odean. 2011. "The Behavior of Individual Investors."
- Benton, Marques, Stephan Meier, and Charles Sprenger. 2007. "Overborrowing and Undersaving: Lessons and Policy Implications from Research in Behavioral Economics." *Federal Reserve Bank of Boston Community Affairs Discussion Paper* 07 (4).
- Berlin, Mitchell, and Loretta Mester. 2004. "Credit Card Rates and Consumer Search." *Review of Financial Economics* 13 (1-2): 179–198.
- 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* LXI (4): 1553–1604.
- 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. 2011. "\$100 Bills on the Sidewalk: Suboptimal Saving in 401(k) Plans." *Review of Economics and Statistics* 93 (3): 748–763.
- Choi, James, David Laibson, and Brigitte Madrian. 2010. "Why Does the Law of One Price Fail? An Experiment on Index Mutual Funds." *Review of Financial Studies* 23 (4): 1405–1432.
- Day, G.S. 2003. "Creating a Superior Customer-Relating Capability." *MIT Sloan Management Review* 44 (3): 77–82.
- Dynan, Karen. 2009. "Changing Household Financial Opportunities and Economic Security." *Journal of Economic Perspectives* 23 (4): 49–68.
- Finlay, Keith, and Leandro Magnusson. 2009. "Implementing Weak-instrument Robust Tests for a General Class of Instrumental-variables Models." *The Stata Journal* 9 (3): 1–24.
- FINRA Investor Education Foundation. 2009. "Financial Capability in the United States: National Survey – Executive Summary."
- 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, David, and Nicholas Souleles. 2002a. "An Empirical Analysis of Personal Bankruptcy and Delinquency." *The Review of Financial Studies* 15 (1): 319–347.
- . 2002b. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics* 117 (1) (February): 149–85.
- Han, Song, Benjamin Keys, and Geng Li. 2011. "Credit Supply to Bankrupt Consumers: Evidence from Credit Card Mailings."

- 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* 119 (2): 403–456.
- Kamenica, Emir, Sendhil Mullainathan, and Richard Thaler. 2011. "Helping Consumers Know Themselves." *American Economic Review* 101 (3) (May): 417–422. doi:10.1257/aer.101.3.417.
- Kerr, Sougata, and Lucia Dunn. 2008. "Consumer Search Behavior in the Changing Credit Card Market." *Journal of Business and Economic Statistics* 26 (3): 345–353.
- Knittel, C., and V. Stango. 2003. "Price Ceilings, Focal Points, and Tacit Collusion: Evidence from Credit Cards." *American Economic Review* 93 (5) (December): 1703–1729.
- 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.
- Laroche, Michel, Gad Saad, Mark Cleveland, and Elizabeth Browne. 2000. "Gender Differences in Information Search Strategies for a Christmas Gift." *Journal of Consumer Marketing* 17 (6): 500–522.
- Meier, Stephan, and Charles Sprenger. 2010. "Present-biased Preferences and Credit Card Borrowing." *American Economic Journal: Applied Economics* 2 (1): 193–210.
- 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. 2012. "Do Consumers Borrow on Their Cheapest Credit Card? Evidence from Mexico."
- Prince, Jeffrey. 2008. "Repeat Purchase Amid Rapid Quality Improvement: Structural Estimation 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* 53 (5): 1589–1622.
- 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. 2000. "Competition and Pricing in the Credit Card Market." *Review of Economics and Statistics* 82 (3): 499–508.
- . 2002. "Pricing with Consumer Switching Costs: Evidence from the Credit Card Market." *Journal of Industrial Economics* 50 (4): 475–492.
- 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. "Fuzzy Math, Disclosure Regulation, and Credit Market Outcomes: Evidence from Truth-in-Lending Reform." *Review of Financial Studies* 24 (2): 506–534.
- . 2013. "Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Bank Overdraft Fees."
- Taylor, Curtis. 2003. "Supplier Surfing: Competition and Consumer Behavior in Subscription Markets." *The RAND Journal of Economics* 34 (2): 223–246.
- Telyukova, Irina. 2011. "Household Need for Liquidity and the Credit Card Debt Puzzle."
- Woodward, Susan E., and Robert E Hall. 2012. "Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence." *American Economic Review* 102 (7): 3249–3276.
- Zinman, Jonathan. 2007. "Household Borrowing High and Lending Low Under No-arbitrage."

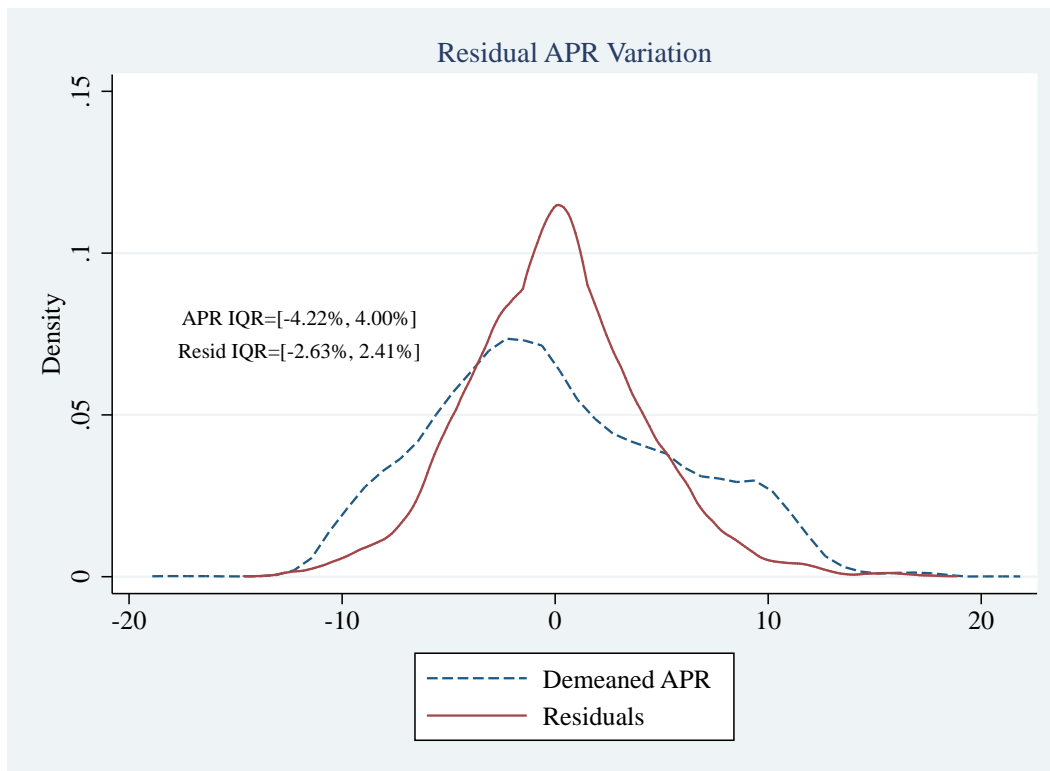


Figure 1. Raw and residual variation in weighted APRs.

Notes: “Demeaned APR” shows the distribution (kernel density) of panelist-level average weighted APRs on all revolving balances during the sample period, demeaned so that they are centered on zero. “Residuals” shows the distribution (kernel density) of residuals from the fullest specification of the panelist-level “above plus demographics” regressions described in Table 3. Fitted values and residuals are calculated using the quartile-specific coefficients in the first four columns of Table 3 (fifth row down).

Table 1. Panelist-Level Summary Statistics

| | Revolving Balance Quartile | | | | | All |
|---------------------------------------------------------|----------------------------|-------------|--------------|---------------|------------|-----|
| | 1 | 2 | 3 | 4 | | |
| Quartiles [revolving balances, \$] | [0, 499] | [499, 1534] | [1534, 4586] | [4586, 62515] | [0, 62515] | |
| Cards held (share of panelists, N= 4312) | | | | | | |
| 1 | 0.50 | 0.51 | 0.39 | 0.30 | 0.42 | |
| 2 | 0.26 | 0.29 | 0.31 | 0.27 | 0.28 | |
| 3 | 0.12 | 0.10 | 0.15 | 0.16 | 0.13 | |
| 4 | 0.06 | 0.05 | 0.07 | 0.09 | 0.07 | |
| 5+ | 0.06 | 0.05 | 0.08 | 0.17 | 0.09 | |
| Average purchases per month (\$, N =4312) | | | | | | |
| 25th | 36 | 12 | 27 | 71 | 28 | |
| 50th | 291 | 57 | 148 | 305 | 173 | |
| 75th | 1025 | 314 | 491 | 911 | 688 | |
| 90th | 2090 | 1362 | 1497 | 1934 | 1722 | |
| Average revolving balances (\$, N = 4312) | | | | | | |
| 25th | 0 | 363 | 1447 | 5773 | 187 | |
| 50th | 0 | 528 | 2029 | 8471 | 1061 | |
| 75th | 42 | 769 | 2866 | 13544 | 3965 | |
| 90th | 125 | 934 | 3536 | 20981 | 10102 | |
| Annualized interest costs (\$, N =4312) | | | | | | |
| 25th | 0 | 63 | 264 | 970 | 30 | |
| 50th | 0 | 102 | 363 | 1487 | 197 | |
| 75th | 7 | 152 | 515 | 2332 | 695 | |
| 90th | 21 | 203 | 716 | 3778 | 1741 | |
| Interest costs/total borrowing costs, average (N =4312) | 0.48 | 0.66 | 0.81 | 0.92 | 0.75 | |
| Annualized interest costs/annual income (N= 4312) | | | | | | |
| 25th | 0.000 | 0.001 | 0.004 | 0.014 | 0.001 | |
| 50th | 0.000 | 0.002 | 0.007 | 0.024 | 0.004 | |
| 75th | 0.000 | 0.003 | 0.011 | 0.042 | 0.012 | |
| 90th | 0.000 | 0.005 | 0.018 | 0.068 | 0.030 | |
| Credit score [N= 4312] | | | | | | |
| 25th | 696 | 562 | 602 | 651 | 616 | |
| 50th | 767 | 631 | 669 | 699 | 694 | |
| 75th | 803 | 728 | 741 | 749 | 768 | |
| 90th | 819 | 796 | 792 | 789 | 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=4312] | | | | | | |
| HS or less | 0.08 | 0.12 | 0.10 | 0.08 | 0.10 | |
| Some college | 0.23 | 0.34 | 0.31 | 0.28 | 0.29 | |
| College degree + | 0.69 | 0.53 | 0.59 | 0.64 | 0.61 | |
| 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 | 2,079 | 1,987 | 2,247 | 2,994 | 9,307 | |
| Panelist-months | 18,561 | 19,761 | 21,030 | 21,960 | 81,312 | |
| Account-months | 29,438 | 29,681 | 35,117 | 47,851 | 142,087 | |

Notes: All variables measured at panelist level. Panelist-level averages are across all panelist-days in the sample. Income statistics have smaller sample sizes due to item-nonresponse on registration survey. "Cards held" is the maximum number of distinct cards (accounts) observed on any one day in the Lightspeed data, at the panelist level. Interest costs are calculated using daily balances and APRs for all card/days in the sample, and annualized. "Total borrowing costs" include interest costs, annual fees, late and over-limit fees, cash advance fees and balance transfer fees. Credit score is from one of the three major bureaus, observed upon entry into the panel. Income, education and age are self-reported upon entry into the panel. Revolving balance quartiles are calculated using panelist-level average daily revolving balances.

Table 2. Borrowing Costs in the Cross-Section of Panelists

| | Revolving Balance Quartile | | | | Total |
|-----------------------------------------------------------------------------------------------------|----------------------------|-------------|--------------|---------------|------------|
| | 1 | 2 | 3 | 4 | |
| Quartile cutoffs (revolving balances) | [0, 499] | [499, 1534] | [1534, 4586] | [4586, 62515] | [0, 62515] |
| Panelist-level weighted APR, all balances, all panelists (N=4312) | | | | | |
| 10th | 0.00 | 3.04 | 6.38 | 8.80 | 0.00 |
| 25th | 0.00 | 8.21 | 11.21 | 11.91 | 3.45 |
| 50th | 0.00 | 15.96 | 16.18 | 16.13 | 13.17 |
| 75th | 1.08 | 21.11 | 21.68 | 20.77 | 19.53 |
| 90th | 7.57 | 25.14 | 25.90 | 25.42 | 24.38 |
| Panelist-level weighted APR, revolving balances (N=3685) | | | | | |
| 10th | 11.99 | 11.99 | 10.56 | 10.07 | 10.99 |
| 25th | 14.90 | 15.26 | 14.90 | 13.40 | 14.79 |
| 50th | 17.59 | 19.34 | 18.47 | 17.28 | 18.21 |
| 75th | 20.92 | 23.94 | 23.47 | 21.92 | 23.05 |
| 90th | 26.26 | 28.24 | 28.04 | 26.53 | 27.63 |
| Panelist-level weighted APR, revolving balances, no teaser rates (N=3629) | | | | | |
| 10th | 12.24 | 12.90 | 11.90 | 11.51 | 11.96 |
| 25th | 14.99 | 15.74 | 15.24 | 14.01 | 14.99 |
| 50th | 17.80 | 19.46 | 18.90 | 17.78 | 18.36 |
| 75th | 21.07 | 24.03 | 23.78 | 22.31 | 23.21 |
| 90th | 26.32 | 28.29 | 28.15 | 26.83 | 27.84 |
| Panelist-level weighted APR, revolving balances, no teaser rate, complete cards sub-sample (N=1742) | | | | | |
| 10th | 11.70 | 12.74 | 11.34 | 11.24 | 11.58 |
| 25th | 14.90 | 15.43 | 14.99 | 13.34 | 14.73 |
| 50th | 17.36 | 19.46 | 17.95 | 17.00 | 17.99 |
| 75th | 21.22 | 23.99 | 22.43 | 21.61 | 22.78 |
| 90th | 25.84 | 28.24 | 27.66 | 25.99 | 27.20 |
| Panelist-level weighted APR, revolving balances, no teaser rates, January 2006 (N=2495) | | | | | |
| 10th | 12.23 | 12.07 | 12.23 | 11.16 | 11.77 |
| 25th | 14.90 | 14.90 | 15.14 | 14.05 | 14.90 |
| 50th | 17.40 | 17.68 | 17.93 | 17.23 | 17.45 |
| 75th | 20.07 | 21.32 | 22.91 | 22.07 | 22.08 |
| 90th | 28.15 | 28.05 | 28.24 | 27.24 | 28.08 |
| R-sq.: monthly borrowing costs on panelist FEs | 0.78 | 0.76 | 0.78 | 0.76 | 0.77 |

Notes: Weighted APR is at panelist level across all card/days (or card/days without teaser APRs) in sample, weighted by total balances or only revolving balances. Balances that are non-revolving have an APR of zero. "Teaser rates" are defined by the authors as any APR below 7.99%. Complete cards sub-sample is defined as in Table A1. January 2006 sub-sample summarizes panelist-month-level weighted APRs for the month of January 2006. R-squared is from a regression of panelist-month-level weighted APRs on revolving balances on a set of panelist fixed effects; the r-squared therefore identifies the share of variation in panelist-month-level APRs that is identified by time-invariant differences in APRs across panelists (i.e., the cross-section).

Table 3. Explaining Borrowing Cost Cross-Sections Using Observable Risk, Card Characteristics/Effects, Demographics, and Issuer/Time Fixed Effects

| | Revolving Balance Quartile | | | | Total |
|-------------------------------------------------------------------|-----------------------------------------|-------------|-------------|-------------|--------------|
| | 1 | 2 | 3 | 4 | |
| Panelist-level models: | | | | | |
| | R-squared (unadjusted R-squared) | | | | |
| Credit score decile | 0.06 (0.07) | 0.14 (0.15) | 0.21 (0.20) | 0.13 (0.12) | 0.15 (0.15) |
| Above plus in-sample risk | 0.09 (0.25) | 0.24 (0.30) | 0.34 (0.39) | 0.30 (0.35) | 0.27 (0.29) |
| Above plus "issuer effects" | 0.13 (0.38) | 0.39 (0.46) | 0.38 (0.47) | 0.35 (0.45) | 0.34 (0.37) |
| Above plus card fees/characteristics | 0.17 (0.42) | 0.39 (0.47) | 0.39 (0.48) | 0.36 (0.46) | 0.35 (0.38)* |
| Above plus demographics | 0.16 (0.46) | 0.38 (0.48) | 0.39 (0.50) | 0.37 (0.49) | 0.35 (0.39) |
| N | 448 | 1062 | 1061 | 10588 | 3629 |
| Account-month-level models: | | | | | |
| | R-squared (unadjusted R-squared) | | | | |
| Credit score decile | 0.07 (0.07) | 0.16 (0.16) | 0.17 (0.17) | 0.06 (0.06) | 0.12 (0.12) |
| Time-invariant panelist-level variables | 0.27 (0.27) | 0.38 (0.38) | 0.32 (0.32) | 0.26 (0.26) | 0.27 (0.27) |
| Time-varying usage and risk, issuer and time effects | 0.24 (0.24) | 0.36 (0.36) | 0.29 (0.29) | 0.21 (0.21) | 0.25 (0.25) |
| Time-varying usage and risk, issuer and time effects, card chars. | 0.24 (0.25) | 0.36 (0.37) | 0.30 (0.30) | 0.21 (0.22) | 0.26 (0.26) |
| All covariates above | 0.32 (0.33) | 0.41 (0.42) | 0.34 (0.35) | 0.49 (0.38) | 0.30 (0.30) |
| N | 28375 | 28708 | 33480 | 44804 | 135367 |
| Account-month-level models, first six months of sample: | | | | | |
| | R-squared (unadjusted R-squared) | | | | |
| Credit score decile | 0.10 (0.10) | 0.15 (0.15) | 0.16 (0.16) | 0.07 (0.07) | 0.13 (0.13) |
| Time-invariant panelist-level variables | 0.27 (0.28) | 0.35 (0.36) | 0.31 (0.32) | 0.23 (0.24) | 0.25 (0.25) |
| Time-varying usage and risk, issuer and time effects, card chars. | 0.21 (0.22) | 0.33 (0.34) | 0.27 (0.27) | 0.20 (0.20) | 0.24 (0.24) |
| All covariates above | 0.30 (0.31) | 0.38 (0.39) | 0.32 (0.33) | 0.27 (0.28) | 0.28 (0.28) |
| N | 7591 | 8070 | 9224 | 11786 | 36671 |
| Account-month-level models, last six months of sample : | | | | | |
| | R-squared (unadjusted R-squared) | | | | |
| Credit score decile | 0.03 (0.03) | 0.20 (0.20) | 0.23 (0.23) | 0.06 (0.06) | 0.13 (0.13) |
| Time-invariant panelist-level variables | 0.23 (0.26) | 0.50 (0.53) | 0.42 (0.45) | 0.35 (0.38) | 0.31 (0.32) |
| Time-varying usage and risk, issuer and time effects, card chars. | 0.27 (0.30) | 0.46 (0.49) | 0.41 (0.44) | 0.28 (0.30) | 0.30 (0.31) |
| All covariates above | 0.33 (0.37) | 0.52 (0.55) | 0.44 (0.48) | 0.35 (0.38) | 0.33 (0.34) |
| N | 2123 | 1783 | 1961 | 3122 | 8989 |

Notes: Each cell reports the r-squared (unadjusted r-squared) from a regression of APRs on the set of listed covariates. Panelist-level models use as the dependent variable the panelist-level APR paid on all revolving balances, excluding teaser rates, weighted by balances across all accounts and days in sample period. Account-month-level models use the account-month-level APR as the dependent variable. Covariates are listed below and described in fuller detail in the Data Appendix. Full results from asterisked specification * are shown in Appendix Table 3.

"**Credit score decile**" is a full set of indicator variables for the panelist-level credit score. Base model also includes indicators for sample entry/exit timing.

"**In-sample risk**" (or "**time-invariant panelist-level variables**") include the number of cards held (indicators up to 5+), panelist-level average daily total credit line across all cards (decile indicators), panelist-level indicators for quintile of total late fees in-sample and quartile of total over-limit fees in-sample, panelist-level credit utilization decile indicators, average monthly purchase volume quartile indicators, and average monthly revolving balance quartile indicators.

"**Issuer effects**" in panelist-level regressions are a vector measuring for each panelist the average shares of revolving balances allocated to each distinct issuer in the data.

"**Card fees/characteristics**" include average fees paid per year (annual, balance transfer and cash advance) across all cards, and indicators for whether the panelist's primary card (the one with the highest level of revolving balances, on average) has an annual fee, has a variable rate, and is a rewards card.

"**Demographics**" include indicators for age category, income category and education category (see Table 1).

"**Time-varying usage and risk**" include panelist-month level indicators for utilization decile, credit line decile, total late fees to date in sample and total over-limit fees to date in sample.

"**Issuer and time effects**" in account-month models include issuer fixed effects and month-year fixed effects.

"**Card characteristics**" in account-month models include card-level indicators for whether the card has an annual fee, is a rewards card or has a variable rate, and interactions between the variable rate indicator and month-year fixed effects.

Table 4. APR Dispersion in Choice Sets, in the Market and in the Wallet

| Percentile | Goto APR | | | Net-of-teaser APR | | |
|------------|----------|-------|------------|-------------------|-------|------------|
| | High | Low | Difference | High | Low | Difference |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| 10th | 10.99 | 8.99 | 0.00 | 9.90 | 3.75 | 0.46 |
| 30th | 13.99 | 9.90 | 2.25 | 12.24 | 4.95 | 4.25 |
| 50th | 16.15 | 9.99 | 4.34 | 14.24 | 6.45 | 7.50 |
| 70th | 18.24 | 10.99 | 7.25 | 18.24 | 9.31 | 9.86 |
| 90th | 19.50 | 14.90 | 9.25 | 18.99 | 13.39 | 13.95 |

Notes: Data from Mintel Comperemedia Inc. (<http://www.comperemedia.com/>). Sample covers all reported credit card direct mail offers for 1211 individuals in the Comperemedia sample, from January 2007. "Goto" APR is the rate at which balances incur interest charges after expiration of the introductory "teaser" period (if any). "Net-of-teaser" APR is the average of the teaser and goto APRs over the first 24 months of the offer.

Table 4b. Estimated within-panelist (within-month) high APR "offer," low APR "offer" and high-low APR "offer" spread, January 2007.

| Percentile | All covariates | | | Credit score decile only | | |
|------------|----------------|-------|------------|--------------------------|-------|------------|
| | High | Low | Difference | High | Low | Difference |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| 10th | 16.71 | 9.25 | 5.05 | 17.73 | 9.26 | 5.89 |
| 25th | 18.52 | 10.91 | 6.42 | 18.74 | 10.67 | 7.11 |
| 50th | 21.38 | 12.72 | 8.19 | 21.73 | 12.42 | 8.73 |
| 75th | 24.12 | 14.75 | 10.43 | 24.10 | 14.55 | 10.84 |
| 90th | 26.67 | 16.79 | 13.20 | 26.62 | 16.08 | 13.50 |

Notes: Estimated APR "offers" are calculated using our Lightspeed data. We first estimate OLS APR regressions for each of the largest six issuers, letting the relationship between panelist characteristics and APR differ by issuer. Each model includes a full set of panelist-month-level and card-month-level covariates described in Section V and the Data Appendix, (Columns 1-3 above), or just credit score decile and month-year fixed effects (Columns 4-6 above). We use the coefficients from each model to predict six fitted APRs for each panelist in each month - a hypothetical set of "offers" from the largest six issuers. This allows us to estimate a high APR, low APR, and high-low spread for each panelist.

Table 4c. Actual within-panelist within-month APR differences "in the wallet," January 2007.

| Percentile | All panelists | | | Panelists with >1 card | | |
|------------|---------------|-------|------------|------------------------|-------|------------|
| | High | Low | Difference | High | Low | Difference |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| 10th | 14.99 | 8.99 | 0 | 16.99 | 8.90 | 0 |
| 25th | 18.24 | 10.99 | 0 | 18.55 | 9.97 | 0.75 |
| 50th | 21.65 | 14.79 | 0 | 23.24 | 12.99 | 4.00 |
| 75th | 29.24 | 17.49 | 3.00 | 29.99 | 15.90 | 7.66 |
| 90th | 30.49 | 21.65 | 8.01 | 32.24 | 18.24 | 12.00 |

Notes: Table shows actual APRs on cards held by panelists in our data during January 2007 (we choose this month-year to facilitate comparison with 4a and 4b above). "High" column shows the distribution of APRs across panelists using only each panelist's highest APR "in the wallet." "Low" column shows the distribution of APRs across panelists using only each panelist's lowest APR in the wallet. "Difference" shows the distribution of the max-min spread in the wallet, again during only January 2007. First column shows distributions for all panelists in that month (N=2808), and second column shows distributions for the subset of panelists with more than one card in that month (N=1197).

Table 5. Is Borrowing Cost Heterogeneity Driven By "Misallocation", Very Strictly Defined?

| | Revolving Balance Quartile | | | | Total |
|-----------------------------------------------------------------|----------------------------|-------|-------|--------|--------|
| | 1 | 2 | 3 | 4 | |
| All panelists with revolving balances | | | | | |
| Average within-wallet APR max-min difference | | | | | |
| 25th | 0 | 0 | 0 | 0 | 0 |
| 50th | 0 | 0 | 0.01 | 1.66 | 0.0 |
| 75th | 2.11 | 2.00 | 4.51 | 6.38 | 4.08 |
| 90th | 5.09 | 5.66 | 8.34 | 10.63 | 8.30 |
| Average APR misallocation | | | | | |
| 25th | 0 | 0 | 0 | 0 | 0 |
| 50th | 0 | 0 | 0 | 0.05 | 0 |
| 75th | 0.09 | 0.09 | 0.46 | 0.93 | 0.48 |
| 90th | 4.00 | 1.73 | 2.05 | 2.62 | 2.45 |
| Annualized dollars lost from misallocation | | | | | |
| 25th | 0 | 0 | 0 | 0 | 0 |
| 50th | 0 | 0 | 0 | 4.96 | 0 |
| 75th | 0.02 | 0.55 | 9.20 | 100.23 | 8.22 |
| 90th | 2.81 | 9.08 | 45.14 | 321.19 | 84.13 |
| Misallocation costs as percentage of annual interest costs | | | | | |
| 25th | 0 | 0 | 0 | 0 | 0 |
| 50th | 0 | 0 | 0 | 0 | 0 |
| 75th | 0.01 | 0 | 0.02 | 0.06 | 0.03 |
| 90th | 0.23 | 0.11 | 0.13 | 0.17 | 0.16 |
| 95th | 0.45 | 0.25 | 0.30 | 0.26 | 0.28 |
| N | 451 | 1078 | 1078 | 1078 | 3685 |
| Panelists with revolving balances, complete cards sub-sample | | | | | |
| Average within-wallet APR max-min difference | | | | | |
| 25th | 0 | 0 | 0 | 0 | 0 |
| 50th | 0.46 | 0 | 1.56 | 3.72 | 1.21 |
| 75th | 3.65 | 3.48 | 5.91 | 7.89 | 5.23 |
| 90th | 5.86 | 7.06 | 8.66 | 11.95 | 9.14 |
| Average APR misallocation | | | | | |
| 25th | 0 | 0 | 0 | 0 | 0 |
| 50th | 0 | 0 | 0.06 | 0.39 | 0.02 |
| 75th | 1.15 | 0.50 | 0.85 | 1.71 | 1.14 |
| 90th | 4.76 | 3.28 | 3.49 | 3.88 | 3.78 |
| Annualized dollars lost from misallocation | | | | | |
| 25th | 0 | 0 | 0 | 0 | 0 |
| 50th | 0 | 0 | 0.99 | 35.21 | 0.27 |
| 75th | 0.45 | 2.86 | 21.13 | 180.01 | 19.80 |
| 90th | 3.24 | 15.53 | 75.15 | 435.23 | 133.24 |
| Misallocation costs as percentage of annual interest costs | | | | | |
| 25th | 0 | 0 | 0 | 0 | 0 |
| 50th | 0 | 0 | 0 | 0.02 | 0 |
| 75th | 0.07 | 0.03 | 0.06 | 0.11 | 0.07 |
| 90th | 0.30 | 0.20 | 0.22 | 0.24 | 0.24 |
| n | 252 | 522 | 511 | 484 | 1769 |

Notes: "Average APR misallocation" is the average balance-weighted daily reduction in APR that would be obtained by *costlessly* transferring all balances to lowest-rate cards, conditional on current credit limits. "Dollars lost" is the average annualized dollar savings the panelist would enjoy by reallocating perfectly throughout the sample period. "Misallocation costs as percentage..." compares the annualized dollar savings from reallocation to the annualized dollar interest costs from Table 1, at the panelist level. Complete cards sub-sample includes panelists for whom the number of card accounts in our data matches "active bankcard lines" from the credit report.

Table 6. Self-Reported Search Intensity, Other Shopping-Related Variables, and Credit Characteristics

| Variable | 10-point scale (10 highest), "I always keep an eye out for better credit card offers" | | | | All respondents | Non-respondents |
|-------------------------------------------------|---------------------------------------------------------------------------------------|---------|---------|--------|-----------------|-----------------|
| | [1, 3] | [4, 6] | [7, 9] | 10 | | |
| Current (in-sample) credit card accounts: | | | | | | |
| 1 | 0.27 | 0.22 | 0.21 | 0.28 | 0.24 | 0.45 |
| 2 | 0.25 | 0.25 | 0.21 | 0.20 | 0.24 | 0.29 |
| 3 | 0.22 | 0.19 | 0.18 | 0.12 | 0.19 | 0.12 |
| 4 | 0.11 | 0.13 | 0.10 | 0.07 | 0.11 | 0.07 |
| 5+ | 0.15 | 0.21 | 0.29 | 0.33 | 0.22 | 0.07 |
| Previous (closed) credit card accounts | | | | | | |
| 0-2 | 0.23 | 0.22 | 0.22 | 0.15 | 0.22 | 0.25 |
| 3-5 | 0.23 | 0.15 | 0.15 | 0.10 | 0.17 | 0.22 |
| 6-9 | 0.21 | 0.24 | 0.20 | 0.23 | 0.22 | 0.23 |
| 10-14 | 0.17 | 0.17 | 0.19 | 0.12 | 0.17 | 0.15 |
| 15+ | 0.16 | 0.21 | 0.24 | 0.40 | 0.22 | 0.14 |
| Self-reported recent credit card applications | | | | | | |
| 0 | 0.70 | 0.57 | 0.41 | 0.35 | 0.55 | 0.57 |
| 1 | 0.25 | 0.32 | 0.42 | 0.39 | 0.33 | 0.36 |
| 2 | 0.05 | 0.10 | 0.14 | 0.20 | 0.10 | 0.06 |
| 3 | 0.00 | 0.01 | 0.04 | 0.06 | 0.02 | 0.01 |
| Credit score (median) | 753 | 728 | 705 | 703 | 723 | 690 |
| Average weighted APR, no misallocation (median) | 16.33 | 15.43 | 16.12 | 14.90 | 15.88 | 17.94 |
| Revolving credit card balances (\$, median): | 553.91 | 1088.83 | 1313.87 | 839.65 | 890.00 | 1080.00 |
| N (panelists) | 205 | 181 | 157 | 60 | 603 | 3710 |

Notes: Search intensity is self-reported agreement with the statement "I always keep an eye out for better credit card offers," on a 10 point scale with 1 being "Does not describe me at all" and 10 being "Describes me perfectly." Survey question was part of a larger email survey sent to all panelists in the first quarter of 2007. Survey content was not announced prior to the decision to take the survey. In-sample credit card accounts is defined as in Table 1. Previous accounts is the number of previously held but closed credit card accounts from the panelist's credit bureau file. "Applications" are the sum of affirmative responses to survey questions asking "Have you applied for any new credit cards in the last 12 months?" Surveys were emailed to panelists in 2004Q4, 2005Q1 and 2006Q1. Only those panelists taking each survey (751 for the first, 972 for the second, and 1354 for the last) could have provided an affirmative response. Non-respondent column shows data for panelists who did not take the survey containing the search question.

Table 7. Search Behavior and Borrowing Costs

| | | Dependent variable: weighted best APR (mean=16.35) | | | | | |
|-------------------------------------------------|--------------------------------------------------|----------------------------------------------------|--------------------|--------------------|-------------------|---------------------|----------------------|
| | | (1) | (2) | (3) | (4) | (5) | (6) |
| Self-reported search intensity (10-point scale) | | -0.083 (0.078) | -1.119* (0.577) | -1.065* (0.583) | -0.950 (0.620) | -1.101** (0.430) | -1.116*** (0.360) |
| | N | 497 | 497 | 476 | 497 | 476 | 476 |
| | r-squared | 0.59 | 0.41 | 0.45 | 0.55 | 0.44 | 0.50 |
| | full set of control variables? | yes | yes | yes | yes | yes | yes |
| | Instruments: None (OLS) | | | | | | |
| | Marital status | | yes | | | yes | yes |
| | Gender | | | yes | | yes | yes |
| | Lightspeed surveys taken | | | | yes | | yes |
| | CLR/AR test robust to weak instruments (p-value) | n/a | 0.03 | 0.03 | 0.08 | 0.00 | 0.00 |
| | 95% CI, robust to weak instruments | n/a | [-3.38, -0.14] | [-3.26, -0.12] | [-3.38, 0.06] | [-2.51, -0.37] | [-2.12, -0.37] |
| | Sargan test (p-value) | n/a | 0.83 | n/a | n/a | 0.35 | 0.98 |

Notes: All models are estimated at the panelist level. "Weighted best APR" is the lowest APR the panelist could pay in that month, averaged across days, if balances were allocated to lowest-rate cards conditional on credit limits. Search intensity is self-reported agreement on a scale of 1-10 with the statement "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." All models include the full set of regressors described in Tables 3 and A3. Marital status is a set of four indicator variables for "single, never married," "married," "divorced/widowed/separated" and "other." Lightspeed surveys taken is a panelist-level total of periodic surveys taken by panelists; see Stango and Zinman (2013) for details. CLR/AR test is for significance of the endogenous regressor (search intensity), given that the instruments may be weak, and 95% CI is calculated using standard errors robust to the presence of weak instruments. Sargan test is for exogeneity of instruments (where rejection of the null indicates endogeneity), and is only applicable when the model is over-identified as in columns (2), (5) and (6). See <http://econ.tulane.edu/kfinlay/pdf/FinlayMagnusson2009.pdf> for a discussion of the weak instrument problem and the Stata routine we use to deal with the issue.

Data Appendix

A. Panel Construction and Maintenance

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. Our sample size falls over time, however, because later panelists tend not to have matched credit report data. Appendix Table 5 shows some data on how the number of panelists and their characteristics evolve over time. Because we focus on cross-panelist differences and generally employ panelist-level *time-invariant* variables in the analysis, those dynamics are not a focus of the analysis. Where appropriate, we do account for panelists' sample entry/exit dates in the empirics. We also check that our results are robust to using only individual months, or the first six months, of data.

B. Measuring Credit Risk

Our data include much, if not all, of the information used by issuers when setting and adjusting APRs:

1. *Credit scores:* A credit score from one of the major three bureaus is probably the single best summary source of information about credit history and risk. We observe one credit score for each panelist at entry into the sample, which is generally in January 2006, but occasionally later. The score, on the standard 850-point scale, summarizes risk by 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.
2. *Supplementary credit bureau data:* We also observe other information from the report including total debt, the number of active credit cards, total credit available, the number of

active auto and mortgage “lines” (loans), the total number of past (closed) credit card accounts, and a few other variables.¹

3. *The number of credit cards held:* For each panelist on each day, we observe the number of registered credit card accounts. We define for each panelist the number of cards held as the maximum number of cards held on any one day. We have defined the variable other ways without any difference, because the number of cards held is very stable for a panelist over time.
4. *Purchase volume and revolving balances:* For each panelist we calculate average monthly purchase volume and average monthly revolving balances (these can be very different depending on whether the panelist revolves). We then bin each panelist into one of four quartiles based on each variable.
5. *In-sample late/missed payments:* A late or missed payment can trigger a “default” APR on the account in question, and is also in many cases reported to the credit bureau, leading other issuers to incorporate the late/missed payment history into APRs on new offers or existing cards. The credit score mentioned above should capture information about late/missed payments leading up to the panelist’s enrollment in Lightspeed, and once the panelist enters our data we directly observe late/missed payments. We measure running late payment counts for each account, a running count of late payments at the panelist level across all accounts, and several panelist-level and time-invariant measures of “total late fees,” “average late fees per month” and “any late fee in-sample.”
6. *Limits and utilization:* Issuers generally consider utilization (the ratio of balances to available credit) as a signal about risk. Cardholders may face higher APRs or offers either by having what an issuer considers “high” utilization, or by exceeding their credit line (going “over-limit”) on one or more cards. Again, the credit score we observe at panel enrollment incorporates all available information about utilization as of enrollment; after enrollment we

¹ Beyond the credit score itself, issuers may also incorporate this disaggregated information from the credit report into risk modeling for new account offers. In practice, adding such information non-parametrically to our models has little effect on the fit. This is partly because we use rich, disaggregated data on within-sample account performance, as described below. Customers may also self-report income, education, and other demographics on their applications, but an issuer generally does not directly observe those things. We include such demographics in our models and find that they do not improve the fit.

observe utilization levels (including both credit limits and card balances) and “over-limit” instances directly, at the card and panelist level. As with late/missed payments we here we calculate running utilization levels and over-limit instance counts, and also construct panelist-level time-invariant “Over-limit fees per month” and “Any over-limit fee in-sample” measures.

7. *Demographics*: We observe from the registration survey categorical variables measuring age, income and education. These may not be directly observed by issuers, but may proxy for variables (such as time since opening first credit card) that issuers incorporate into pricing.

Collectively, these variables are quite comprehensive. They likely compare favorably to the data observed by issuers on their own cards, although individual issuers may of course employ those data differently. They may dominate data observed by issuers on other cards (i.e., on accounts issued by other issuers).

C. Measuring Non-APR Account Attributes

We also observe a variety of card- and issuer-level characteristics:

1. *Annual fees*: For each card in the data, we observe annual fees incurred. We measure annual fees both as a cardinal number – the average annual fees paid per year, either at the card or panelist level – and using an indicator for “any annual fee incurred during the sample period,” again either at the panelist or account level.
2. *Other fees (balance transfer, cash advance, etc.)*: We observe balance transfer fees, cash advance fees, late fees, and over-limit fees as they are incurred, and include them as annual dollar costs per account or panelist. This is imperfect because we only observe fees that are incurred, rather than the contingent price that might be incurred. We have experimented with a variety of alternative approaches to this issue – inferring fees even when they are never incurred from data on actual fees paid by other panelists with the same card, for example – with little effect on the results.
3. *Rewards*: We observe for every card in the data its “card name” as a text string, which is the issuer’s name for the card. An example would be “MBNA CREDIT CARD.” The card name often reveals information about rewards or “affinity” links (e.g., “AMERICA WEST FLIGHTFUND CREDIT CARD,” “GREEN BAY PACKERS VISA”). We also observe its “account name,” which is an issuer- or panelist-defined name for the account and also

contains information about affinity/reward links (e.g., “NATIONAL WILDLIFE FEDERATION PLATINUM PLUS MASTERCARD,” “PLATINUM DELTA SKYMILES”). We do not directly observe rewards, but in practice the dollar value of rewards does not vary by much across cards. We have experimented with separate variables for rewards and affinity status, or a single combined indicator for the presence of either.

4. *Fixed/variable rate pricing:* 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.² We construct an indicator measuring whether the rate is fixed or variable.
5. *Unobserved issuer-specific and state-specific effects:* We also observe the issuer (e.g., Bank of America, Capital One, Citi, etc.) for each credit card in the data. This allows us to construct a set of “issuer effects” measuring average APR differences across issuers, which might come from omitted card characteristics, or from systematic differences in pricing customer risk.³ Because a given panelist may have balances allocated across multiple cards from different issuers, our panelist-level regressions measure the average share of revolving balances held on cards of each issuer. In card-level regressions we simply include a fixed effect for the card issuer. (We’ve also estimated specifications with fixed effects for the card name, since, e.g., MBNA cards may be remain branded as “MBNA”, even after MBNA gets acquired by Bank of America. This alternative definition of issuer does not affect the results.) We also observe the panelist’s state of residence and in unreported specifications include fixed effects for state of residence; those effects might capture any number of omitted influences on state-level supply or demand for credit.
6. *Sample entry/exit dates:* Because panelists may be in the data for less than the entire sample period, we include a set of indicators for the panelist’s first and last months in the data. This corrects for variation in APRs generated by systematic time-varying APRs, combined with differential entry/exit dates by panelists.

² See Stango (2000) for a detailed discussion of fixed and variable rate pricing in credit cards.

³ Issuer effects are de-identified when we report the results, per confidentiality provisions of our data licensing agreement with Lightspeed.

D. Classifying “Teaser” Rates

Our data do not identify teaser rates as such, but they are fairly easy to classify empirically because they are significantly lower than even the lowest contract rate offered to the best credit risks during our sample period. We classify any APR below 7.99% as a teaser rate (source: tabulations from the Mintel data discussed in Section VI-A). This discards 5% of account-months, and 1% of panelists who always pay teasers in-sample.

E. Representativeness

Starting at the top of Table 1, our cardholding distribution matches up well with data from the 2007 Survey of Consumer Finances (SCF), particularly when one uses our “complete cards” sub-sample (Appendix Table 2) as the benchmark.⁴ Purchases also match up well with the SCF, in which the comparable 25th, 50th, and 75th percentiles of the weighted data are \$20, \$250, and \$1000.⁵ 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).⁶ This suggests that our data may understate the *level* of debt and total interest costs in the broader population. For our panelists the share of total credit card costs from interest (vs. fees) is 74%, as compared to an estimate of 80% from 2007 issuer-side data (source: Cards&Payments).

Data from other sources on APR distributions is limited, but comparing our data 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, even if one restricts our data to the subsample of panelists’ “primary cards.” The dispersion we observe also looks similar to that in more recent administrative data from the OCC.⁷ How the central tendency in our data compares to other data is murkier. The APRs we observe are higher on average than the self-reported APRs in the 2007 SCF, but are similar to those in the OCC

⁴ Zinman (2009) shows that cardholding in the SCF matches up well with issuer-side data.

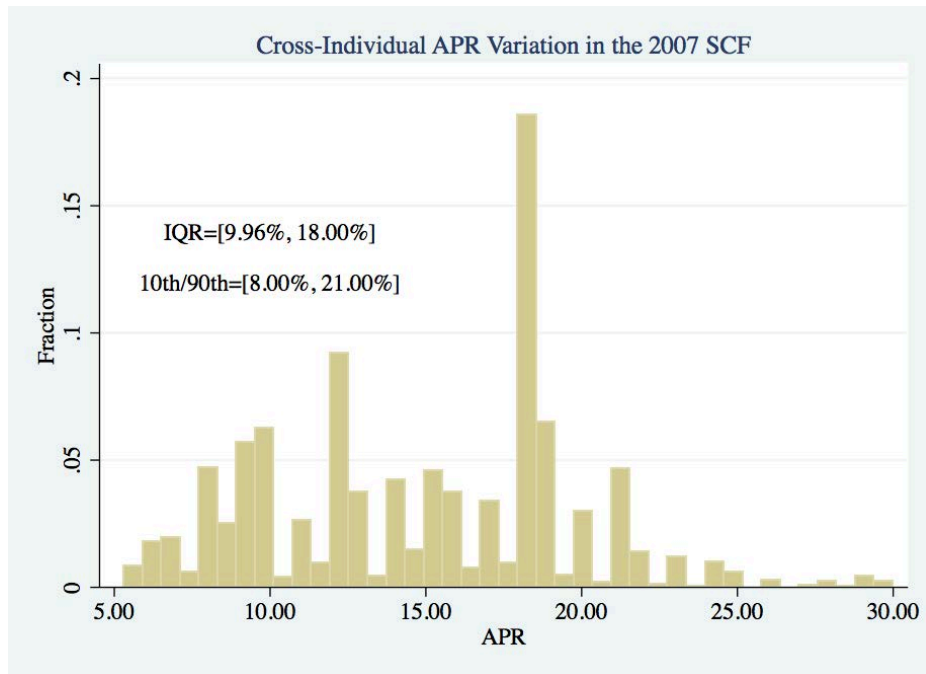
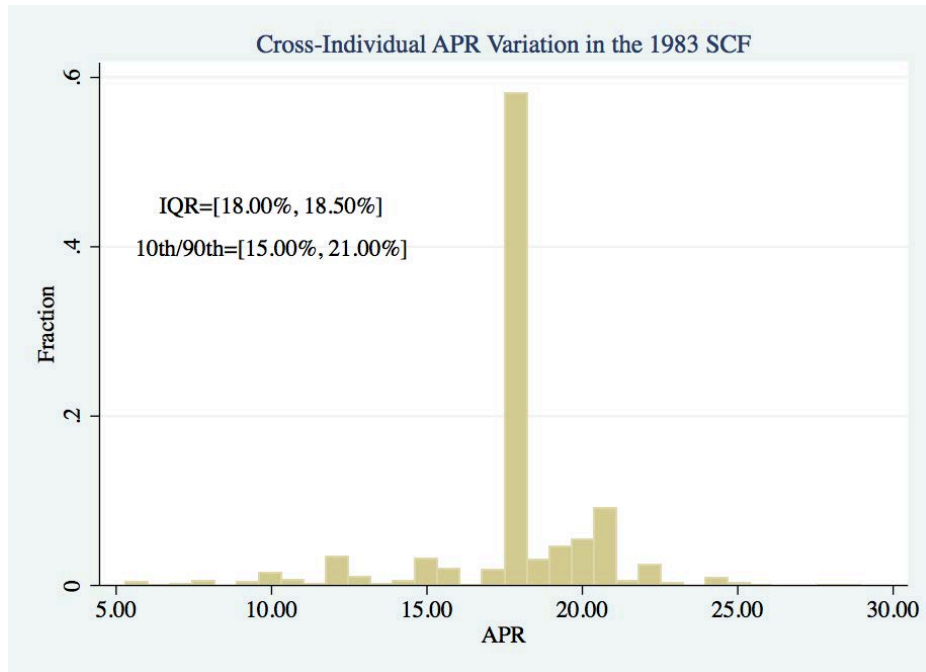
⁵ Zinman (2009) shows that card purchases in the SCF match up well with issuer-side data.

⁶ 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.

⁷ Source: <http://files.consumerfinance.gov/f/2011/03/OCC-Presentation.pdf> .

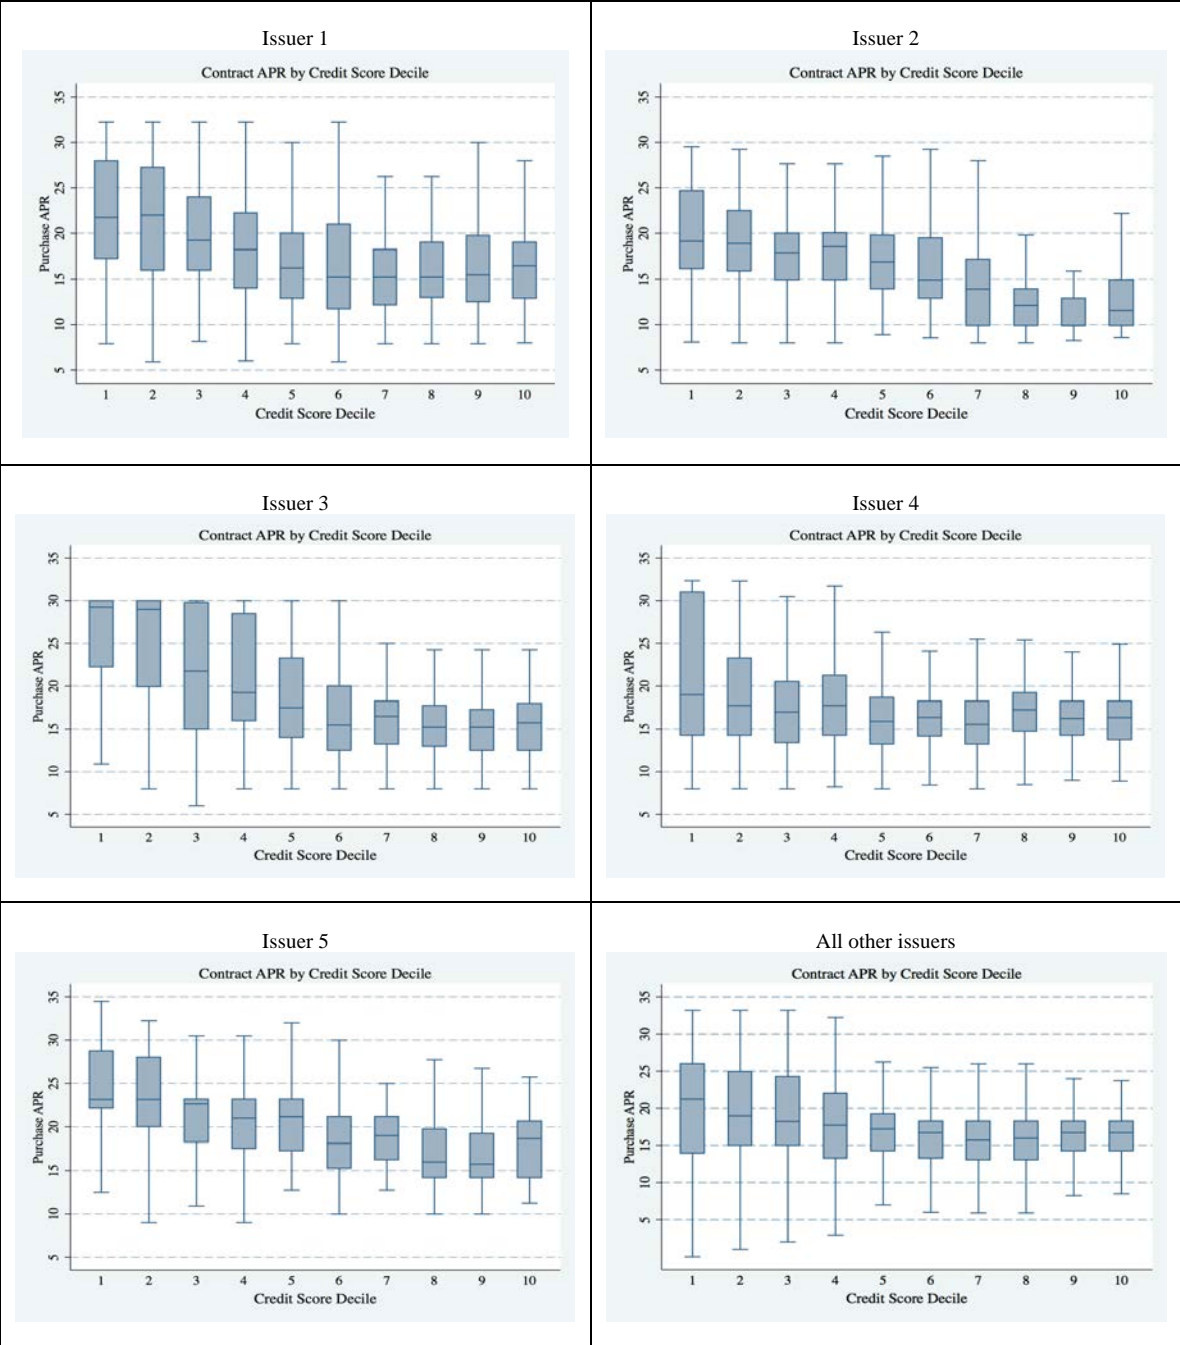
administrative data from 2009. In short, we see little reason to believe that the dispersion we observe is uncharacteristic of the national population of U.S. credit cardholders.

In terms of demographics, our panelists are younger, more educated, and higher income (conditional on age) than national averages. The overall credit score distribution looks representative, conditional on demographics (source: tabulations from the Payment Cards Center of the Federal Reserve Bank of Philadelphia).



Appendix Figure 1. APR variation in the 1983 and 2007 Surveys of Consumer Finances.

Notes: Figures show distributions of answers to SCF open-ended questions regarding credit card interest rates in the 1983 and 2007 Surveys of Consumer Finances. In 1983 interviewers asked for “Respondent’s best guess as to average interest rate he/she pays (annualized) if the full bill is not paid on the bank or storecard he/she uses the most often.” In 2007 interviewers asked “What interest rate do you pay on the card where you have the largest balance?”



Appendix Figure 2. Credit scores and APRs for five large issuers and all other issuers.

Notes: Each pane shows the relationship between credit score decile and the within-decile distribution of contract APRs, for five largest issuers in sample and the remaining smaller issuers (the latter appearing in the bottom right pane). Each box-and-whisker plot shows the median APR as a solid horizontal line within the box, the 25th/75th percentiles as the top and bottom of the box, and the 5th/95th percentiles as the whiskers.

Appendix Table 1. Panelist-Level Summary Statistics, Complete Cards Sub-sample (Compare to Table 1)

| | | Revolving Balance Quartile | | | | All |
|---------------------------------------------------------|------------------------------------|----------------------------|-------------|--------------|---------------|------------|
| | | 1 | 2 | 3 | 4 | |
| | Quartiles (revolving balances, \$) | [0, 499] | [499, 1534] | [1534, 4586] | [4586, 62515] | [0, 62515] |
| Cards held (share of panelists, N= 2134) | | | | | | |
| | 1 | 0.35 | 0.38 | 0.25 | 0.15 | 0.29 |
| | 2 | 0.29 | 0.31 | 0.32 | 0.26 | 0.29 |
| | 3 | 0.17 | 0.15 | 0.19 | 0.19 | 0.17 |
| | 4 | 0.10 | 0.08 | 0.10 | 0.11 | 0.10 |
| | 5+ | 0.10 | 0.09 | 0.14 | 0.29 | 0.15 |
| Average purchases per month (\$, N =2134) | | | | | | |
| | 25th | 52 | 21 | 53 | 126 | 45 |
| | 50th | 445 | 89 | 218 | 450 | 264 |
| | 75th | 1197 | 565 | 682 | 1224 | 986 |
| | 90th | 2220 | 1821 | 1837 | 2335 | 2070 |
| Average revolving balances (\$, N = 2134) | | | | | | |
| | 25th | 0 | 355 | 1423 | 5630 | 98 |
| | 50th | 0 | 525 | 2005 | 8637 | 877 |
| | 75th | 38 | 771 | 2894 | 13666 | 3535 |
| | 90th | 125 | 916 | 3528 | 21692 | 9529 |
| Annualized interest costs (\$, N =2134) | | | | | | |
| | 25th | 0 | 62 | 240 | 942 | 16 |
| | 50th | 0 | 100 | 353 | 1450 | 157 |
| | 75th | 6 | 144 | 497 | 2239 | 578 |
| | 90th | 20 | 200 | 667 | 3707 | 1591 |
| Interest costs/total borrowing costs, average (N =2134) | | | | | | |
| | | 0.46 | 0.66 | 0.81 | 0.93 | 0.74 |
| Annualized interest costs/annual income (N= 2134) | | | | | | |
| | 25th | 0.000 | 0.001 | 0.004 | 0.013 | 0.000 |
| | 50th | 0.000 | 0.002 | 0.007 | 0.024 | 0.003 |
| | 75th | 0.000 | 0.003 | 0.011 | 0.041 | 0.011 |
| | 90th | 0.000 | 0.005 | 0.018 | 0.073 | 0.029 |
| Credit score [N= 2134] | | | | | | |
| | 25th | 695 | 553 | 600 | 649 | 619 |
| | 50th | 763 | 634 | 673 | 702.5 | 703 |
| | 75th | 798 | 743 | 749 | 755 | 773 |
| | 90th | 817 | 802 | 797 | 791 | 807 |
| Income [N=2031] | | | | | | |
| | under \$25,000 | 0.14 | 0.19 | 0.20 | 0.11 | 0.16 |
| | \$25k-\$45k | 0.17 | 0.25 | 0.20 | 0.19 | 0.20 |
| | \$45k-\$87.5 | 0.48 | 0.41 | 0.44 | 0.49 | 0.45 |
| | \$87.5-\$125k | 0.11 | 0.08 | 0.09 | 0.12 | 0.10 |
| | \$125k+ | 0.11 | 0.06 | 0.07 | 0.10 | 0.09 |
| Education [N=2134] | | | | | | |
| | HS or less | 0.09 | 0.14 | 0.10 | 0.10 | 0.11 |
| | Some college | 0.21 | 0.37 | 0.36 | 0.30 | 0.31 |
| | College degree + | 0.70 | 0.49 | 0.54 | 0.60 | 0.59 |
| Age [N=2134] | | | | | | |
| | Under 30 | 0.25 | 0.33 | 0.30 | 0.25 | 0.28 |
| | 30-39 | 0.28 | 0.27 | 0.30 | 0.34 | 0.30 |
| | 40-49 | 0.20 | 0.19 | 0.22 | 0.19 | 0.20 |
| | 50-59 | 0.16 | 0.15 | 0.12 | 0.16 | 0.14 |
| | 60+ | 0.11 | 0.07 | 0.07 | 0.07 | 0.08 |
| | Panelists | 617 | 522 | 511 | 484 | 2,134 |
| | Accounts | 1,441 | 1,195 | 1,280 | 1,734 | 5,650 |
| | Panelist-months | 11,379 | 10,216 | 10,601 | 10,718 | 42,914 |
| | Account-months | 20,424 | 17,724 | 19,937 | 27,945 | 86,030 |

Notes: All variables are as described in Table 1. "Complete cards" sub-sample includes panelists for whom the number of cards held in the Lightspeed data is at least as great as the number of "active credit card lines" reported in the credit bureau data.

Appendix Table 2. More Results from a Panelist-Level Model (asterisked specification in Table 3, row 4, last column).

| Dependent variable: panelist weighted average APR on revolving balances, no teaser rates: mean (LHS)=17.48 | | | | | | |
|------------------------------------------------------------------------------------------------------------|-------------|----------------|------------------------------------------------|-------------|----------------|--|
| Variable | Coefficient | Standard error | Variable | Coefficient | Standard error | |
| Credit score: decile 2 | -0.439 | 0.295 | Total late fees >0: quintile 1 | 1.073*** | -0.387 | |
| Credit score: decile 3 | -0.522 | 0.320 | Total late fees >0: quintile 2 | 1.900*** | -0.42 | |
| Credit score: decile 4 | -1.263*** | 0.355 | Total late fees >0: quintile 3 | 2.586*** | -0.413 | |
| Credit score: decile 5 | -1.705*** | 0.380 | Total late fees >0: quintile 4 | 4.018*** | -0.431 | |
| Credit score: decile 6 | -1.797*** | 0.392 | Total late fees > 0: quintile 5 | 5.830*** | -0.475 | |
| Credit score: decile 7 | -2.169*** | 0.417 | Total over-limit fees > 0: tertile 1 | -0.13 | -0.367 | |
| Credit score: decile 8 | -2.530*** | 0.449 | Total over-limit fees >0: tertile 2 | -0.095 | -0.412 | |
| Credit score: decile 9 | -2.393*** | 0.473 | Total over-limit fees > 0: tertile 3 | 0.109 | -0.433 | |
| Credit score: decile 10 | -2.167*** | 0.484 | Average utilization: decile 2 | -0.105 | 0.457 | |
| Two cards held | -0.361* | 0.204 | Average utilization: decile 3 | -0.034 | 0.472 | |
| Three cards held | -0.245 | 0.271 | Average utilization: decile 4 | 0.225 | 0.473 | |
| Four cards held | -0.259 | 0.351 | Average utilization: decile 5 | -0.287 | 0.501 | |
| 5+ cards held | -0.567 | 0.366 | Average utilization: decile 6 | -0.479 | 0.519 | |
| Total credit line: decile 2 | -0.108 | 0.335 | Average utilization: decile 7 | -0.046 | 0.537 | |
| Total credit line: decile 3 | -0.834** | 0.411 | Average utilization: decile 8 | 0.702 | 0.563 | |
| Total credit line: decile 4 | -1.575*** | 0.474 | Average utilization: decile 9 | 0.805 | 0.579 | |
| Total credit line: decile 5 | -2.249*** | 0.510 | Average utilization: decile 10 | 1.382** | 0.589 | |
| Total credit line: decile 6 | -1.769*** | 0.555 | Average monthly purchase volume: quartile 2 | -0.07 | 0.222 | |
| Total credit line: decile 7 | -2.056*** | 0.587 | Average monthly purchase volume: quartile 3 | -0.097 | 0.288 | |
| Total credit line: decile 8 | -1.902*** | 0.636 | Average monthly purchase volume: quartile 4 | 0.438 | 0.368 | |
| Total credit line: decile 9 | -2.431*** | 0.686 | Average monthly revolving balances: quartile 2 | -0.890** | 0.352 | |
| Total credit line: decile 10 | -2.422*** | 0.779 | Average monthly revolving balances: quartile 3 | -1.048** | 0.449 | |
| Average annual fees paid/year, all cards | 0.035** | 0.017 | Average monthly revolving balances: quartile 4 | -1.148** | 0.584 | |
| Average balance transfer fees paid/year, all cards | -0.004 | 0.035 | Constant | 21.475*** | 1.058 | |
| Average cash advance fees paid/year, all cards | 0.029 | 0.049 | | | | |
| Panelist's primary card: variable rate? | -0.416* | 0.251 | | | | |
| Panelist's primary card: annual fee? | 1.029*** | 0.184 | | | | |
| Panelist's primary card: rewards? | 0.677*** | 0.225 | | | | |
| Panelist-level "issuer effects": p-value | | 0.000 | | | | |
| Indicators for last month in sample: p-value | | 0.000 | | | | |
| N | | | 3629 | | | |
| r-squared (unadjusted) | | | 0.35 (0.38) | | | |

Notes: Coefficients are from an OLS regression at the panelist level. Dependent variable is the panelist-level weighted average APR on revolving balances during the sample period, excluding balances with "teaser rates." Sample begins with 4312 panelists from Table 1, dropping 627 who never revolve balances and 56 who borrow on teaser rates for entire time in sample. "Cards held" is the maximum number of different accounts open on any one day during the sample period. "Total credit line" quintile is measured using the average daily credit line on all cards. "Average utilization" is the average across all days in the sample of daily balances (revolving or not) divided by total credit line, across all cards. "Primary card" is the card on which a majority of balances are held during the sample period, across all days. Panelist-level "issuer effects" are a vector measuring for each panelist the average shares of revolving balances allocated to each distinct issuer in the data.

Appendix Table 3. Self-reported search intensity and instruments for search intensity

| Variable | 10-point scale (10 highest), "I always keep an eye out for better credit card offers" | | | | All respondents | Non-respondents |
|---------------------------------|---------------------------------------------------------------------------------------|--------|--------|------|-----------------|-----------------|
| | [1, 3] | [4, 6] | [7, 9] | 10 | | |
| Marital status: | | | | | | |
| Single, never married | 0.15 | 0.21 | 0.20 | 0.22 | 0.19 | 0.16 |
| Married | 0.66 | 0.59 | 0.61 | 0.53 | 0.61 | 0.67 |
| Divorced, separated or widowed | 0.15 | 0.19 | 0.15 | 0.18 | 0.17 | 0.11 |
| Other | 0.03 | 0.01 | 0.05 | 0.07 | 0.03 | 0.05 |
| Female | 0.69 | 0.66 | 0.59 | 0.47 | 0.63 | 0.69 |
| Lightspeed surveys taken (mean) | 5.43 | 5.30 | 4.97 | 4.70 | 5.20 | 0.70 |
| N (panelists) | 205 | 181 | 157 | 60 | 603 | 3710 |

Notes: Search intensity is self-reported agreement with the statement "I always keep an eye out for better credit card offers," on a 10 point scale with 1 being "Does not describe me at all" and 10 being "Describes me perfectly." Survey question was part of a larger email survey sent to all panelists in the first quarter of 2007. Survey content was not announced prior to the decision to take the survey. The final row shows the total number of Lightspeed surveys voluntarily taken by panelists; the maximum possible is 18. Non-respondent column shows data for panelists who did not take the survey containing the search question.

Appendix Table 4. Costs and Benefits of Holding Multiple Cards

| | Revolving balance quartile | | | | |
|-----------------------------------------------|----------------------------------------------------------------|----------------------|----------------------|---------------------|----------------------|
| | 1 | 2 | 3 | 4 | All |
| | Dependent variable: average max-min difference "in the wallet" | | | | |
| Two cards held | 1.355*** (0.150) | 1.907*** (0.163) | 2.220*** (0.233) | 2.404*** (0.288) | 1.964*** (0.103) |
| Three cards held | 2.349*** (0.204) | 3.065*** (0.248) | 4.104*** (0.294) | 3.754*** (0.354) | 3.388*** (0.136) |
| Four cards held | 3.479*** (0.271) | 4.275*** (0.342) | 4.499*** (0.399) | 4.929*** (0.429) | 4.346*** (0.178) |
| 5+ cards held | 4.630*** (0.308) | 5.294*** (0.400) | 6.194*** (0.429) | 7.258*** (0.409) | 6.240*** (0.184) |
| Dependent variable: weighted best APR | | | | | |
| Two cards held | -2.197*** (0.663) | -0.526 (0.352) | -1.116*** (0.416) | -0.975** (0.445) | -0.947*** (0.212) |
| Three cards held | -1.983** (0.882) | -1.653*** (0.537) | -2.099*** (0.526) | -0.720 (0.548) | -1.548*** (0.279) |
| Four cards held | -1.497 (1.120) | -2.003*** (0.741) | -2.583*** (0.716) | -1.344** (0.668) | -1.844*** (0.363) |
| 5+ cards held | -4.413*** (1.278) | -4.199*** (0.868) | -4.333*** (0.771) | -1.460** (0.637) | -2.877*** (0.373) |
| Dependent variable: APR cost of misallocation | | | | | |
| Two cards held | 1.659*** (0.435) | 0.383*** (0.131) | 0.442*** (0.135) | 0.280** (0.114) | 0.522*** (0.082) |
| Three cards held | 2.625*** (0.579) | 0.937*** (0.200) | 0.928*** (0.171) | 0.637*** (0.141) | 1.089*** (0.108) |
| Four cards held | 1.224* (0.735) | 1.523*** (0.276) | 1.049*** (0.233) | 1.049*** (0.171) | 1.384*** (0.141) |
| 5+ cards held | 2.751*** (0.838) | 2.554*** (0.323) | 2.264*** (0.251) | 1.317*** (0.163) | 2.038*** (0.145) |

Notes: All models are OLS, estimated at the panelist level. "Average max-min difference" is the unweighted average across all days of the difference between the highest and lowest APRs simultaneously held by the panelist (see Table 5). "Weighted best APR" is balance-weighted APR the panelist would have paid over the sample period if all balances were always allocated to lowest-rate cards (conditional on credit limits). "APR cost of misallocation" is the balance-weighted reduction in APR that would have been obtained if all balances had been costlessly transferred to lowest-rate cards conditional on credit limits throughout the sample period. Models include all covariates from the fullest models described in Table 3. See Appendix Table 2 for conditional correlations on net: between cards held and average borrowing costs.

Appendix Table 5. Observations and panelist characteristics by month/year

| | Panelists | Accounts | Panelist-level median | | |
|--------|-----------|----------|-----------------------|--------------|--------------|
| | | | Revolving balances | Weighted APR | Credit Score |
| Jan-06 | 3,656 | 6,448 | 1154 | 18.23 | 694 |
| Feb-06 | 3,688 | 6,532 | 1169 | 18.25 | 694 |
| Mar-06 | 3,673 | 6,572 | 1168 | 18.24 | 696 |
| Apr-06 | 3,631 | 6,461 | 1193 | 18.24 | 697 |
| May-06 | 3,589 | 6,285 | 1201 | 18.24 | 697 |
| Jun-06 | 3,512 | 6,124 | 1208 | 18.24 | 697 |
| Jul-06 | 3,462 | 5,979 | 1226 | 18.21 | 697 |
| Aug-06 | 3,418 | 5,867 | 1227 | 18.19 | 698 |
| Sep-06 | 3,302 | 5,642 | 1242 | 18.14 | 699 |
| Oct-06 | 3,204 | 5,490 | 1247 | 18.13 | 700 |
| Nov-06 | 3,118 | 5,253 | 1248 | 18.12 | 700 |
| Dec-06 | 3,009 | 4,962 | 1260 | 18.10 | 699 |
| Jan-07 | 2,889 | 4,591 | 1258 | 18.11 | 699 |
| Feb-07 | 2,818 | 4,460 | 1265 | 18.09 | 700 |
| Mar-07 | 2,829 | 4,654 | 1260 | 18.05 | 700 |
| Apr-07 | 2,808 | 4,704 | 1260 | 18.02 | 700 |
| May-07 | 2,719 | 4,571 | 1248 | 18.01 | 700 |
| Jun-07 | 2,726 | 4,619 | 1275 | 17.96 | 702 |
| Jul-07 | 2,755 | 4,727 | 1268 | 17.93 | 703 |
| Aug-07 | 2,679 | 4,654 | 1286 | 17.90 | 704 |
| Sep-07 | 2,233 | 4,020 | 1337 | 17.73 | 709 |
| Oct-07 | 1,243 | 2,553 | 1128 | 17.00 | 719 |
| Nov-07 | 1,211 | 2,502 | 1132 | 16.99 | 720 |
| Dec-07 | 1,185 | 2,288 | 1224 | 16.90 | 721 |
| Jan-08 | 1,169 | 2,209 | 1239 | 16.86 | 722 |
| Feb-08 | 1,161 | 2,180 | 1239 | 16.88 | 722 |
| Mar-08 | 1,142 | 2,142 | 1245 | 16.87 | 722 |
| Apr-08 | 1,115 | 2,092 | 1255 | 16.82 | 723 |
| May-08 | 1,096 | 2,033 | 1233 | 16.76 | 724 |
| Jun-08 | 1,070 | 1,974 | 1255 | 16.73 | 723 |
| Jul-08 | 1,031 | 1,919 | 1249 | 16.69 | 724 |
| Aug-08 | 1,008 | 1,872 | 1233 | 16.69 | 724 |
| Sep-08 | 981 | 1,834 | 1255 | 16.68 | 724 |
| Oct-08 | 951 | 1,766 | 1285 | 16.63 | 724 |
| Nov-08 | 884 | 1,609 | 1286 | 16.42 | 728 |
| Dec-08 | 347 | 499 | 1750 | 16.12 | 738 |

Notes: number of panelists/accounts reflects both entry of new panelists and attrition. For a given panelist the set of cards is registered upon entry and not updated (i.e., we do not observe new accounts for panelists who remain in the sample). The attrition beginning in September 2007 coincides with the Forrester/Lightspeed spinoff (which required panelists to renew their consent for participation).