



WORKING PAPERS

RESEARCH DEPARTMENT

WORKING PAPER NO. 11-17
LIMITED AND VARYING CONSUMER ATTENTION:
EVIDENCE FROM SHOCKS TO THE SALIENCE OF
BANK OVERDRAFT FEES

Victor Stango
University of California, Davis

Jonathan Zinman
Dartmouth College
and
Visiting Scholar, Federal Reserve Bank of Philadelphia

April 2011

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/research-and-data/

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

Victor Stango
Graduate School of Management
University of California, Davis

Jonathan Zinman
Department of Economics
Dartmouth College

April 2011

ABSTRACT

We explore the dynamics of limited attention in the \$35 billion market for checking overdrafts, using survey content as shocks to the salience of overdraft fees. Conditional on selection into surveys, individuals who face overdraft-related questions are less likely to incur a fee in the survey month. Taking multiple overdraft surveys builds a “stock” of attention that reduces overdrafts for up to two years. The effects are significant among consumers with lower education and financial literacy. Consumers avoid overdrafts not by increasing balances but by making fewer debit transactions and cancelling automatic recurring withdrawals. The results raise new questions about consumer financial protection policy.

JEL Classification Codes: D14, D18, G13, G21

Keywords: household finance, attention, overdraft fees, intertemporal choice

* vstango@ucdavis.edu, jzinman@dartmouth.edu. Thanks to the Filene Research Institute, the Rockefeller Center at Dartmouth College, and UC-Davis for funding. Thanks to Iwan Barankay, Wilko Bolt, Brianna Cardiff, Stefano DellaVigna, Bob Hunt, Ulrike Malmendier, Paul Oyer, Karen Pence, Julio Rotemberg, Marc Rysman, Scott Schuh, Nick Souleles, and conference/seminar participants at the 2011 AEA meetings, the FTC/Northwestern Micro Conference, Kellogg, NBER Behavioral Economics, NBER Law and Economics, RAND, UC-Boulder, UCLA, and the Federal Reserve Banks of Boston, Chicago, New York, and Philadelphia for helpful comments.

Does limited attention hinder consumers from acquiring and using readily available information when making financial decisions? If yes, how does consumer attention respond to shocks and evolve over time? Answering these questions is critical for understanding how people make financial tradeoffs and for designing sound public policy in retail financial markets. In this paper we explore the role and dynamics of limited consumer attention in the payment of checking account overdraft fees, broadly defining limited attention as incomplete consideration of elements and/or prices in one's choice set.

Checking account overdraft fees are an important dimension of bank pricing, but they have been largely unstudied by economists. U.S. consumers spend an estimated \$35 billion annually on fees charged for initiating transactions that would put their account balance in negative territory, an amount greater than what they spend on fresh fruit or large appliances (Parrish 2009). Over the course of the 2000s, banks have relied increasingly heavily on overdraft fees, with such fees reaching an estimated 74% of revenue from deposit account service charges and 6% of net operating revenue in recent years (Burrhouse et al. 2008).

The growing importance of overdraft fees has generated considerable controversy and policy scrutiny. Consumer groups allege that overdraft fees are poorly disclosed, exploitative, and disproportionately paid by the poor and uneducated.¹ This sentiment has led to lawsuits and new regulations. Recent audits and class-action lawsuits allege violations of disclosure and fair-dealing laws.² The Federal Reserve Board recently issued new rules that require banks to secure consumer opt-in to standard overdraft coverage on debit card and ATM transactions.³ Many banks have modified their overdraft policies in the wake of these events, and some industry observers conjecture that these changes could fundamentally alter how banks price deposit account services.⁴ Understanding *why* people overdraw their accounts is crucial for optimal policy design, and to date, we know very little about that.

To shed light on these issues we examine data from thousands of consumers' checking accounts, covering up to three years of complete account activity. The data are collected by a

¹ See, e.g., <http://www.consumerfed.org/financial-services/credit-and-debt/overdraft-loans>. We discuss relationships between disclosure, other information, and attention below

² On disclosure violations, see, e.g., General Accounting Office (2008). On class action suits, see, e.g., <http://www.bank-overdraft.com/>.

³ See http://www.federalreserve.gov/consumerinfo/wyntk_overdraft.htm for a summary geared to consumers.

⁴ See, e.g., <http://online.wsj.com/article/SB10001424052748703513604575311093932315142.html>.

market research firm that pays panelists for permission to access their accounts. For every account in the data we observe the full set of transactions, account balances, and all fees incurred on the account, including overdraft fees.

We then ask how overdrafting is affected by shocks to customers' attention. Limited attention seems to be a plausible and potentially important explanation for some overdrafts. To cite three motivating patterns (see Section II-D for more), our prior work finds that many overdraft fees are avoidable without forgoing consumption and with small changes in behavior; for example, many consumers could avoid paying overdraft fees by using a credit card instead of a debit card at the point of sale (Stango and Zinman 2009), trading a \$35 overdraft fee for just a few pennies in credit card interest. There is also survey evidence that most people overdraw accounts because they "thought there was enough money in my account." Finally, panelists frequently come close to the margin of overdrafting: among panelists who ever overdrew an account in our sample period, an overdraft occurs in 31% of panelist-months - but available balances fall below \$100 in 83% of all panelist-months. These descriptive statistics are consistent with (though not proof of) an important link between attention to account balances and overdrafting.

Our measure of attention comes from online survey questions answered by panelists.⁵ The market research firm invites all panelists to take surveys periodically. All of the surveys have something to do with household financial services, but specific topics and questions vary from survey to survey. Topics/questions are not announced before a consumer chooses to take a survey. Most surveys do not contain overdraft-related questions (we call these "generic" surveys). But a handful of surveys plausibly draw consumers' attention to overdraft fees by asking questions about overdraft fee payment, use of overdraft protection programs, or (dis)satisfaction with overdraft fees. We label these surveys "overdraft-related." Most of the overdraft-related surveys convey relatively little information in the traditional sense: the questions do not describe specific account terms faced by panelists, and with one exception, overdraft-related questions represent less than 5% of the total survey. The exception is an "overdraft-focused" survey that asks a series of overdraft-related questions and little else, but even that survey provides no consumer-specific information about account terms. All surveys are administered online, so we observe questions exactly as panelists observe them.

⁵ There are large literatures on how surveys and other "primes" can affect behavior, and we discuss these below.

In DellaVigna's taxonomy (2009, p. 349), overdraft-related surveys may increase the *salience* of overdraft fees and may therefore affect consumers' effort to acquire and/or use overdraft-related information. For example, an overdraft-related survey might induce a respondent to monitor account balances more closely, to increase balances with a buffer, and/or to reduce spending.

Conditional on selection into surveys, panelist fixed effects, and time-varying sample-wide shifts in the propensity to incur overdraft charges, we find that taking an overdraft-related survey has a substantial immediate effect on behavior: within the month of the survey, taking an overdraft-related survey reduces the probability of incurring any overdraft fee by an estimated 3.7 percentage points on a base of 30%. There is weak evidence that overdrafting is positively related to selection into survey-taking overall, meaning that selection into surveys may attenuate the effect of taking an overdraft-related survey.

We also find evidence of a strong stock effect: *each* overdraft-related survey taken within a two-year period reduces the probability of paying an overdraft fee by 1.7 percentage points.⁶ Again, this is a within-individual effect and conditional on the timing and number of surveys taken. Selection into survey-taking also attenuates this effect: there is a positive relationship between the number of generic surveys taken and incurring overdraft fees.

Overall, the immediate and stock effects of overdraft-related surveys are strong for those with lower education, lower financial literacy, and lower income, and are not statistically significant for those with higher education or higher financial literacy.⁷ We do not find significant effects on the intensive margin (e.g., total dollar value of fees incurred) of overdrafts, suggesting that attention has a relatively discrete effect on overdrafting.

Additional light on the mechanics and dynamics of attention comes from more detailed exploration of relationships between survey content and overdraft behavior. As we mention above, one of the surveys focuses exclusively on overdrafts (in contrast, overdraft questions never comprise more than 5% of the other surveys). The stock effect of taking the "overdraft-focused" survey is significantly greater than the stock effect of taking the more subtle "overdraft-mentioning" surveys. We also find significant effects on overdrafts of non-overdraft questions

⁶ A model in which stock effects last for two years outperforms, in terms of fit and robustness, a model with one-year stock effects or permanent (non-depreciating) effects (see Section IV-B for details).

⁷ We measure financial literacy using responses to a question from a survey consumers complete when they are registering their accounts in the panel: "I know more than most people when it comes to managing my money and investments [agree/disagree/neutral]."

about spending control, monitoring account balances, and other bank fees. We find no such relationships for a set of “placebo” questions about auto loans, gift cards, or contactless RFID cards.

Panelists who take overdraft-related surveys engineer reductions in overdraft fees by reducing spending transactions. There is some evidence they do this both at high and low frequencies: we see significant reductions in both debit cards and recurring “autodebit” transactions. That pattern suggests that reductions occur on transactions where the balance-at-clearing is difficult to monitor/forecast. There is little evidence that panelists avoid overdrafts by maintaining higher account balances. The full picture of results is consistent with consumers responding to attention shocks by paying greater attention to account balances.

In all, our results suggest that consumers have a limited, time-varying, associative and malleable stock of attention paid to day-to-day household finance. Surveys can operate as shocks to attention, changing behavior and outcomes. Repeated shocks to attention cause sustained changes in behavior that depreciate fairly slowly.

Our results provide the first evidence from a U.S. setting that being surveyed *per se* changes subjects’ subsequent behavior, and are consistent with other work on survey-taking and decision-making.⁸ Zwane et al. (2011) show that being surveyed about health and household finances increases the subsequent purchase of insurance and use of water purification. This follows a large literature showing that asking subjects more directly about their intent or likelihood of engaging in a targeted behavior changes that behavior (Dholakia 2010). A literature on priming and other context effects suggests that small changes in economically irrelevant stimuli can have substantial, and long-lasting, effects on consumer choice, at least in the lab (Chartrand et al. 2008; Sela and Shiv 2009; LeBoeuf, Shafir, and Belyavsky Bayuk 2010).⁹

Our results also add to a growing body of evidence pointing to an important role for limited attention in household finance and other domains (see DellaVigna 2009 for a review).¹⁰ They

⁸ If surveys change behavior, this raises ethical and methodological issues about relying on surveys to measure behavior and elicit behavioral parameters. We distinguish between being surveyed *per se*, about topics *related* to the target behaviors of interest, and being asked more directly about the target behavior of interest, in the sense of questions that are administered by a party with a vested interest in the target behavior (e.g., a firm selling the product under study), and/or questions that ask respondents about their intent or likelihood of engaging in the target behavior. See Zwane et al. (2011) for discussions of the evidence on how surveys affect behavior.

⁹ Persuasive advertising is an important application of priming effects; e.g., Bertrand et al. (2010) find that small changes in uninformative ad content have large effects on the take-up of expensive consumer loans.

¹⁰ See also and Lee and Malmendier (2010) for new theory and empirics on limited attention in auctions, Lacetera et al. (2010) on limited attention in the car market, Alba, Hutchinson, and Lynch (1991), Eliaz and Spiegler (2011),

provide the first clear evidence we know of on the *dynamics* of limited attention.¹¹ Understanding those dynamics is critical, because while we have prior evidence that attention influences decisions, we know very little about how attention accumulates and decays.

Some caveats apply. Our sample is not representative along some dimensions. The external validity of our estimated treatment effects is uncertain. We lack the statistical power to cleanly benchmark our effects to other types of information and/or attention shocks such as a prior overdraft.¹² While we identify significant average treatment effects, it is likely that both attention and the treatment effects vary significantly across individuals. Perhaps most important, we caution against taking welfare implications from our results. We cannot say, for example, whether consumers pay too little or too much attention in the absence of shocks.¹³ We do not observe whether attention to other matters falls when attention to overdrafts rises. In all, it would be imprudent to draw conclusions about the welfare effects of public policy, overdraft-related or other, from our findings.

What we do hope for is further inquiry, because understanding the role of limited attention in financial decisions is a necessary step toward designing sound public policy in household finance. Much of the recent policy discussion regarding overdrafts, and consumer financial protection more generally, revolves around mandated disclosure of account terms that firms might “shroud” absent regulation.¹⁴ The central issue there is whether providing consumers with *more information* helps them to make better decisions. A model of limited attention suggests a different question: can reminders or other shocks to attention, even if they are *uninformative* in the classic sense, help consumers make better decisions?¹⁵ And are ongoing reminders more effective than (or a useful complement to) one-shot, upfront disclosure? Understanding the empirical importance of limited attention can frame and sharpen that policy discussion. We discuss these points further in the Conclusion.

and Koszegi and Szeidl (2011) on consideration sets, and Karlan et al. (2011) on how reminders affect saving behavior.

¹¹ Agarwal et al.’s(2008) evidence on the dynamics of credit card penalty fee payment is similar but may have more to do with the effects of information rather than attention, since the shock that changes behavior in their setup is actually incurring a fee.

¹² Overdrafting is strongly serially correlated in our data; see Section III-C.

¹³ Barber and Odean (2008) describe how news may create buy-side pressure on stocks, leading to losses for investors who buy them.

¹⁴ See Government Accountability Office (2008) for a topical discussion of overdraft disclosure regulation. Gabaix and Laibson (2006) provide an equilibrium model of shrouding in which firms fail to disclose account terms and exploit uninformed customers.

¹⁵ See Grubb (2011) for a model of inattention and a discussion of policy remedies.

Our paper proceeds as follows. Section II describes the overdraft market. Section III describes our data and empirical strategy. Section IV presents our main results. Section V presents results on behavior behind those main results, i.e., on how consumers achieve reductions in overdraft fees. Section VI provides a conceptual framework and some evidence on the cognitive mechanics underlying our results. Section VII concludes.

II. Checking Overdraft Fees: Background and Public Policy

A. Overdrafts and Bank Policies

An overdraft occurs when a customer initiates a transaction that would bring the account balance below zero. Consumers can overdraw an account by making ATM withdrawals, writing checks, or using a debit card. Automated withdrawals can also generate overdrafts.¹⁶

When an overdraft occurs, a bank can pay or decline the transaction according to the overdraft policy in the customer's checking account agreement. The agreement specifies both the manner in which overdrafts are paid and any overdraft or nonsufficient fund (NSF) fees that an overdraft will incur. Banks sometimes call this policy "bounce protection" because by paying the transaction, the bank helps the customer avoid bounced check fees charged by merchants; we will follow more recent usage and refer to this as "standard overdraft service."¹⁷ Standard overdraft service policy generally grants banks complete discretion over whether to pay or decline a particular overdraft. Banks typically will pay overdraft transactions up to an internally specified limit on the account balance and decline any transactions exceeding the limit; the limit protects the bank against default by customers who never bring balances back into positive territory.¹⁸

The most common standard overdraft service pays overdrafts up to the internal limit and charges a fixed per-transaction overdraft fee ranging from \$20-35.¹⁹ Some banks charge NSF or "return item" fees for transactions they decline to pay. Some also charge daily "negative balance

¹⁶ A Federal Deposit Insurance Corporation (FDIC) report (Burrhouse et al. 2008) provides an extensive description of bank overdraft pricing and policies, and we draw heavily on that report here.

¹⁷ The Federal Reserve uses "overdraft service." See <http://www.federalreserve.gov/bankinforeg/reglisting.htm> under Regulation E.

¹⁸ The internal limit varies across banks, across customers within banks, and even for a given customer over time. Most large banks use fraud/default protection software similar to that used by credit card companies; if a particular account displays suspicious behavior, or behavior suggesting a high likelihood that the customer plans to default on the negative balance, the bank will start to decline overdraft transactions.

¹⁹ See reports by the FDIC (Burrhouse et al. 2008), the Government Accountability Office (2008), and others.

fees” of \$2-5 while the account balance remains negative and/or charge interest on the overdrawn amount.

Another type of overdraft policy is “overdraft protection,” which links the customer’s checking account to another account or source of liquidity such as a credit card. When the customer overdraws the account, the bank charges the overdraft amount to the linked account or credit card and also often charges a transfer fee and/or interest on the overdrawn balance. Transfer fees are typically \$5 or \$10, but the fee may depend on the amount overdrawn.

Most banks post overdraft fees and account balances daily and show them as debits on checking account statements. Many banks mail (or e-mail) customers notification about overdrafts; those typically reach customers within a few days after the overdraft. More recently, some banks (and third parties) have begun marketing balance-monitoring and other messaging services that, e.g., send an account holder a text message and/or e-mail when the account balance falls below a pre-specified threshold.

Customers choose an overdraft program when opening an account, and can change their overdraft program at any time, or choose to not have overdrafts paid by the bank at all. Most deposit account agreements “default in” customers to standard overdraft service (Burrhouse et al. 2008; Center for Responsible Lending 2010). National surveys produce a broad range of estimates of the prevalence of accounts with overdraft protection: from 16 percent (Burrhouse et al 2008) to roughly 40 percent (Parrish 2008; Center for Responsible Lending 2009).

B. Overdraft Prevalence and Importance

Overdraft fees have become common in recent years as banks shifted from a pricing model based primarily on monthly fees (see, e.g., Stavins (1999)) to one based primarily on fees for service (ATM fees, overdraft fees). The FDIC (Burrhouse et al. 2008) and industry analysts such as Moebs document this transition and find that overdraft revenue rose to about 75% of explicit deposit account revenue and 6% of *total net* operating revenue in recent years. That explicit revenue has become more important to bank income statements in recent years.

Industry reports estimate that annual overdraft fees paid are roughly \$30-40 billion. Roughly speaking, this averages \$150 per year per checking account. There are limited customer-level data on overdrafting, and most are self-reported (American Bankers Association 2007; Burrhouse et al. 2008; Parrish 2008; Center for Responsible Lending 2009). It is possible that

customers under-report their own overdrafting; such under-reporting has been documented for other types of unsecured borrowing (e.g., see Zinman (2009) on credit card debt). Thus, examining administrative data, as we do, is important for understanding how consumers behave.

C. Controversy and Public Policy Regarding Overdrafts

In recent years bank overdraft practices have become controversial. Consumer groups allege that many banks fail to effectively disclose both the terms of their overdraft programs and the choices that customers face (Government Accountability Office 2008). Complaints also note that many transactions triggering overdraft fees are quite small; if a \$4 cup of coffee generates a \$35 overdraft fee, the argument goes, then the customer has been exploited.²⁰ Consumer advocates also argue that overdraft fees disproportionately affect the poor. Finally, the practice of many large banks to batch-process transactions daily, clearing and settling them in order from highest to lowest dollar amount, has also proven controversial, because it increases the total number of overdraft fees assessed.

The Federal Reserve Board recently required banks to secure affirmative consumer permission (i.e., an active opt-in) to pay overdrafts and charge accompanying fees on most debit card and ATM transactions.²¹ Another recent regulation requires banks to more prominently disclose overdraft fees that a customer has paid in any periodic statements issued to that customer.²² Some recent proposals would require banks to notify customers at the point of sale when an overdraft is about to occur.

We do not take a stance on this controversy, except to note that it highlights the need to understand more about the reasons that consumers overdraw their accounts. The merits of more complete up-front disclosure, for example, are greater if poor up-front disclosure generates overdrafts, but they may be negligible if poor up-front disclosure is not a driving force behind overdrafts. More generally, it is difficult to assess the merits of any particular policy given what we currently know. Our paper is a step toward filling that void.

²⁰See, e.g., <http://www.responsiblelending.org/overdraft-loans/research-analysis/quick-facts-on-overdraft-loans.html>

²¹ See http://www.federalreserve.gov/consumerinfo/wyntk_overdraft.htm for a summary geared toward consumers.

²² 12 CFR Part 230 [Regulation DD; Docket no. R-1315], effective January 1, 2010.

D. Limited Attention and Overdrafting

A variety of economic and psychological models can generate overdrafts.

In a purely neoclassical, full-information model, a checking account holder will pay an overdraft fee due to some combination of motives to ease liquidity constraints and/or economize on transaction costs. It can be perfectly rational to pay a \$34 fee to ensure that a \$20 transaction can be completed, if the marginal utility from that transaction is worth at least \$54 in utility terms, and cheaper sources of liquidity are not available (where “cheap” includes transaction costs, broadly defined). We do not attempt to rule out this sort of traditionally rational explanation for overdrafting; indeed, the overdraft-focused survey in this panel finds that a substantial fraction of customers express a willingness to pay a \$35 overdraft fee even on very small transactions.

Limited attention also might play a role in overdrafting. Our view of what limited attention means is quite general: we take it to mean incomplete consideration of information that would inform choices, whether that information is about account terms or available balances. Our definition encompasses that of Grubb (2011), who takes inattention to mean consumers are unaware of their own past account usage (implying uncertainty about available balances in our setting). In both Grubb’s definition and ours, a consumer with limited attention will be uncertain about the *marginal* price associated with a transaction that might overdraw the account. Occasional overdrafting is a natural consequence of such uncertainty.

Survey evidence suggests that limited attention is a plausible explanation for overdrafting. In one survey question answered by panelists taking the “overdraft-focused” survey studied below (see Section III-E), 60% of overdrafters report overdrafting because they “thought there was enough money in my account.” Most of the remainder report that “the money I deposited was not yet available.” Both reasons are consistent with limited attention to checking account balances. There is also evidence of limited attention to account terms. In that same survey, 24 percent of checking account holders did not know/remember whether the bank described different overdraft coverage options at the time of account opening. In two other nationally representative surveys, 12 and 13% of checking account holders report not knowing whether they *currently have* overdraft protection on their checking account (Parrish 2008; Center for Responsible Lending 2009).

Descriptive evidence from administrative data is also consistent with limited attention as a reason for overdrafts. Consumers could avoid many overdraft fees by tapping readily available sources of liquidity (Stango and Zinman 2009). One conservative measure classifies a fee as avoidable if the overdraft amount is exceeded by the *minimum* available liquidity in another bank account or credit card during the month of that overdraft. Under this measure, among consumers with both a checking account and a credit card, over 50% of overdraft fees are avoidable.

If answering overdraft-related survey questions makes overdraft fees more salient and thereby turns the customer's attention to overdrafting, it could in turn affect the frequency of overdrafting even if the survey does not provide information *per se*. For example, getting asked a question like: "Do you have overdraft protection?" could remind someone about the opportunity to enroll in such a program. Alternatively, that or another overdraft-related question could induce the customer to monitor balances more closely, to make purchases with a credit card when balances are low, to keep a higher buffer stock of balances in the account, to cut back on spending generally, or to do several of these things simultaneously.

In short, limited attention is a plausible explanation for overdrafting, and it is also plausible that surveys could increase attention and affect the frequency with which people overdraw their accounts.

III. Data

A. Overview of Panel Creation and Data Content

Our data come from Lightspeed Research (formerly Forrester Research). Panelists in our sample are members of the "Ultimate Consumer Panel," which is one of many such panels maintained by Forrester/Lightspeed.²³

Panelists enter the Ultimate sample by providing Lightspeed with access to at least two online bank 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

²³ Other Forrester/Lightspeed panels track consumer behavior of interest to market researchers, such as the use and purchases of new technology. Those panels are widely used by industry researchers and academics; see, e.g., Goolsbee (2000; 2001), Kolko (2010), and Prince (2008).

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.

In addition to the account data, Lightspeed collects survey data on Ultimate panelists. All panelists complete a short online *registration survey* when they sign-up for the panel; this gives us some baseline information on demographics, financial characteristics, and respondent-assessed financial literacy. Once in the panel, panelists are then invited to take online surveys that are offered periodically. Survey topics are not preannounced, and we use variation in survey content to identify attention shocks.

B. Details on Account/Transaction Data

The checking account data collected by Lightspeed have two main components. The first component is transaction-level and comes from monthly checking (and credit card) account statements. The statements contain every accounting debit and credit to the account: check deposits and withdrawals, debit card purchases, ATM deposits/withdrawals, automated clearinghouse (ACH) debits and credits such as bill payments and PayPal transactions, transfers to/from other accounts, and so on. Every fee on the account is also recorded as an accounting debit (or credit, if the fee is refunded). For every transaction we observe whether it is a debit or credit, the transaction date, and the transaction amount. Other critical information comes from a set of text strings that identify the bank (“BANK OF AMERICA ALL EXCEPT CA, WA & ID), account name (e.g., “MYACCESS CHECKING5266”) and transaction description. The last string is most important for our purposes, as it nearly always contains information about the payment medium, payee/payer, merchant, location, and so on. We use that text string to identify transaction types, fees in general and overdraft fees in particular. An example of a debit card transaction description is this: “CHECKCARD 0607 QUIZNO'S SUB #6431 Q54 HARRISONVILLEMO 2432...” The string therefore indicates payment method, merchant, store ID (#6341) and location (Harrisonville, MO). That is in addition to transaction amount, bank, account name, and date. A typical overdraft fee transaction description looks like this: “WITHDRAWALS/D OVERDRAFT FEE, 25.00.” A fee associated with a linked account might look like this: “OVERDRAFT XFER FROM CREDIT CARD OR LINE.” We use a text

substring identification routine to classify transaction types and fees. While we can generally distinguish between standard overdraft fees and fees related to overdraft protection, we cannot distinguish between accounts with overdraft protection and accounts with standard overdraft service. We discuss the implication of this limitation below.²⁴

A second component of account data is a near-daily running “available” and “current” account balance.²⁵ Lightspeed obtains these data by logging in and “scraping” the consumer’s account between three and seven times per week. These running balances provide a relatively complete picture of the funds available in the account on a day-to-day basis.

C. Panelists and Overdrafting in Our Sample

The sub-sample of panelists we examine here are those with at least one “active” bank checking account. An active account is one with at least one transaction that could generate an overdraft fee (i.e., an accounting debit) during our sample period. For most of our panelists, the checking account in the Lightspeed data seems to be the only one held by the panelist: roughly 95% of our panelists report holding only one open checking account in their registration survey.

An observation in the raw administrative data is a panelist-transaction, but for this paper we aggregate to the panelist-month. We do this because while we observe transactions and fees on a day-by-day basis, our identifying variation in shocks to attention exists only at the monthly level (we know only the month, not the day/time, that someone took a survey). Table 1 shows summary data on panelists in our sample. The administrative data cover 2006-2008; we report on the month-years Feb-06 to Dec-08, inclusive.²⁶ We observe a total of 7,448 panelists with an active checking account, and 102,334 panelist-months of data; Appendix Table A1 provides more information on the time series of participation in the panel month-by-month.²⁷

²⁴ The central problem is that we observe fees only when they are incurred. Thus, an account with no fees (or a period without fees) may have either standard overdraft service or overdraft protection. It is also difficult to use fees paid to infer what overdraft service a customer has. For example, the majority of accounts in which an overdraft protection fee has been incurred *later* incur a standard overdraft fee – presumably because the linked account has been depleted or hit its limit.

²⁵ The available balance is what the consumer may withdraw before overdrafting. The current balance is the available balance plus holds on the account and netting out pending debits/credits.

²⁶ We observe transactions in January 2006, but we do not use those data in the empirical work because we often condition on lagged monthly information – which is first available as of Feb-06.

²⁷ This table shows that sample composition changes over time. Section IV-F discusses how we control for this.

Of our active *panelists*, roughly half incur at least one overdraft fee during the sample (Table 1). We define an overdraft fee as any explicit pecuniary cost associated with a negative account balance. Thus, this definition includes the \$20-35 fee associated with standard overdraft service. It also includes fees associated with returned items, linked accounts, and so on. In our data, the most common fee by far is the standard overdraft fee, but we include others for completeness.

Across all panelist-*months*, the share with at least one overdraft fee is 16%. That number is naturally higher (31%) among panelists who pay at least one overdraft fee in our data. Overdrafting itself is not uncommon, and being close to the margin of overdrafting is extremely common: 72% of all panelist-months show at least one available balance below \$100, and that figure is higher (83%) among panelists with at least one in-sample overdraft fee. It is worth noting that even panelists who never overdraw in-sample often have low balances: among non-overdrafters, balances fall below \$100 in a full 56% of all panelist-months. This is *prima facie* evidence of systematic variation across customers in attention paid to balances; it is *not* simply the case that non-overdrafters never face the prospect of overdrafting.

Roughly 15% of all panelists, and 28% of those with at least one overdraft, have at least one “snowball” month with five or more overdraft fees. Three percent of all panelist-months are snowball months, and that number rises to 6% among those who have ever overdrawn an account.

There is substantial heterogeneity across panelists and months in the frequency of overdrafting and the dollar amount of fees paid. Figures A1 and A2 illustrate that heterogeneity. While roughly half of all panelists never overdraw in our data, a substantial share of panelists serially overdraw (Figure A1). Two percent of panelists who overdraw do so at least once in *every* month. Dollar amounts of overdraft fees charged are similarly skewed. As Figure A2 shows, there are modes at levels representing integer multiples of common overdraft fees (\$25-30, \$50-60, \$80-100, and so on). But in roughly 8% of cases with at least one overdraft fee, the total dollar value of fees exceeds \$250, and in numerous panelist-months total overdraft fees exceed \$500.

There is also a very strong month-to-month correlation in overdraft fees, within-panelist. The last three rows of Table 1 illustrate this. Fifty-four percent of panelists who incur overdraft fees

in a given month do so again in the next month. Eighty-one percent of panelists who incur 5+ fees incur at least one in the following month.

D. Details on Registration Survey

Along with collecting the administrative data (which is done passively from the panelist's perspective), Lightspeed also actively solicits survey information from its panelists. All panelists complete a "registration survey" at the time of enrollment covering demographics, household financial assets/liabilities (such as stock market participation, the number of credit cards held, etc.), and attitudes in household finance and other domains. In the empirics below we use data on three panelist characteristics collected at registration: education, income, and financial literacy. All are self-reported and categorical. We do not observe these variables for every panelist, because panelists register at different times and registration survey questions change over time. We measure literacy using the response to: "I know more than most people when it comes to managing my money and investments," using "agree" as an indicator of high self-assessed financial literacy, and "neutral" and "disagree" as indicators of medium/low literacy. The correlation between "disagree" (i.e., rating oneself below average) and credit score is -0.15, in the subsample of panelists for whom we also observe credit scores.²⁸ The correlation between rating oneself above average and credit score is 0.17.

Table A2 shows simple relationships between education/literacy/income and overdrafting and also provides information about representativeness along these dimensions. At the panelist level, those with higher education and financial literacy incur overdraft fees less often; the relationship between income and incurring fees is weaker. Our sample is drawn from those with higher income and education on average.

E. Details on Periodic Surveys

Lightspeed also periodically gives panelists opportunities to complete additional surveys on financial product use and satisfaction. Invitations are sent via e-mail to all panelists, and survey respondents are compensated by being entered into a prize lottery. Figure 1 shows the click-through screen that a panelist sees before taking a survey; for our purposes it is important to note

²⁸ We use the self-assessed literacy measure rather than credit score because we observe the latter only for a small subsample of panelists.

that the click-through screen reveals nothing about survey content (overdraft-related vs. non-overdraft-related). Lightspeed classifies any respondent who clicks through as having taken the survey.

We have complete data on survey taking, content and responses for our panelists starting in August 2004 (the date of the first survey) and continuing to the end of 2008. The survey data predate our administrative data, so that even in the first month of administrative data we have comprehensive information for each panelist on past surveys administered and taken. We observe survey content and responses for 21 surveys overall. We observe the month in which the survey is administered, but not the precise date of administration. Table 2 shows each month in which a periodic survey was offered and also shows information about survey content. Response rates are typically in the 20-30% range for the 2006-2008 surveys.²⁹ We cannot measure response rates for the pre-2006 surveys, because we lack the administrative data that would give us the total number of panelists.

Content of the periodic surveys varies and provides the identifying variation in attention that we exploit in our empirical work. Each survey contains roughly 60 multiple choice/response questions. Many of the questions appear on nearly every survey and ask panelists to report the types of bank accounts they hold (credit cards, checking, savings etc.), and the bank(s) and/or credit card companies with which they do business. Many surveys also ask questions about usage of credit cards vs. debit cards. Some surveys collect attitudinal information about borrowing and saving. Some are more narrowly focused (say, on different loans held by the panelist, including amounts and repayment behavior). Some focus on panelists' use of or willingness to adopt new financial products (such as "contactless cards" that use RFID technology).

Most important for our purposes, six of the surveys ask "overdraft-related" questions. Table 2 summarizes which surveys asked overdraft-related questions, and Table A3 panel A shows the overdraft-related questions themselves. Some surveys ask whether panelists have overdraft protection. Some ask whether overdraft fees are a source of "dissatisfaction" with the panelist's current bank. We classify any survey containing at least one question referring to overdraft fees or protection as "overdraft-related."

²⁹ Response rates across surveys vary for several reasons. The set of panelists varies over time and may introduce differences in average panelist-level propensity for taking surveys. Some surveys are left "open" online longer than others. And there are small variations from survey to survey in prizes/compensation for participating.

In five out of the six overdraft surveys, the overdraft-related questions make up a small fraction of total questions on the survey: never more than 5%, and typically closer to 1%. None of the questions on these surveys mentions specific fee amounts or defines overdraft protection (i.e., outside options to standard overdraft services are not described in detail). We label these surveys as “overdraft-mentioning” surveys.

The other overdraft survey is what we call “overdraft-focused.” The October 2006 survey was commissioned by the Center for Responsible Lending (CRL), a consumer advocacy group that lobbies on overdraft policy, and 12 of the 15 questions mention overdrafts. Table A3 panel B lists the questions, which ask both about actual behavior and about consumers’ preferences (e.g., to have banks pay or decline overdrafts by default). While this survey does provide some information on prices and outside options, none of that information is customer-specific.

For estimating the relationship between overdrafting and having taken overdraft-related surveys, the empirical ideal would be random assignment into overdraft-related surveys. Our identifying assumption, given that panelists consciously decide whether to take a survey when confronted with the click-through screen, is that *conditional on a panelist’s decision to take a survey, the survey content is random*. Support for this assumption comes from the click-through screen (Figure 1), which contains no information about survey content. Coupled with the irregular intervals at which surveys mention overdrafts (see Table 2), it is unlikely that panelists are able to consciously select into overdraft-related surveys. We do allow for selection into *surveys overall* in the empirical work and find some support for selection bias along that dimension.

We also find no statistical evidence that consumer-level characteristics are correlated with taking overdraft-specific surveys, once we condition on the total number of surveys taken (Table A4, column 2). Table A4 shows the results of OLS regressions in which the dependent variable is the count of overdraft-related surveys taken. The RHS variables are panelist demographics and variables measuring how long a panelist is in the panel. Without controlling for the total number of surveys taken, the panelist demographics are significant in explaining variation in overdraft surveys (p-value=0.00), but when we condition on the total number of surveys taken, that correlation disappears (p-value=0.69). Thus, we cannot reject the hypothesis that customer characteristics are unrelated to selection into overdraft surveys conditional on the propensity to take any survey.

Table 2 also classifies surveys based on other content that is plausibly related to overdrafting behavior, particularly if attention is driven in part by salience that works in an associative way. Some surveys ask about spending control, monitoring checking balances or other non-overdraft fees. We explore the role of this survey content in the empirical work below.³⁰ The table also classifies surveys based on other content that appears at similar frequencies to overdraft questions but is plausibly *unrelated* to overdrafting behavior: contactless (RFID) cards, gift cards and auto loans. Below we use these variables for a robustness check and a form of placebo test on our main results.

Table 3 tabulates panelist-level data summarizing total and overdraft-related surveys taken. Only 30% of panelists never take any survey. Another 30% take one or two surveys. Twenty-seven percent of panelists take at least one overdraft-related survey, and 16% take more than one. The data therefore display substantial variation across and within panelists in both total and overdraft surveys taken; the within-panelist variation in surveys taken and survey content drives our empirics.

F. Advantages and Limitations of the Data

The Lightspeed data are unique (to our knowledge) in three important respects. First, the account and transaction data span multiple banks. Some recent work in household finance employs administrative transaction-level data, but typically only from a single provider.³¹ Second, our data link administrative measures of actual behavior to a rich set of panelist characteristics collected in the registration survey. Third, Lightspeed matches its administrative data with periodic surveys regarding consumer behavior, preferences and other attitudes.

The main disadvantage of the Lightspeed data is that they are not nationally representative. The requirement that panelists register accounts online selects relatively educated and higher-income people.³² Panelists are necessarily those who are comfortable sharing sensitive financial

³⁰ Example of spending control question: “I like being in control of my spending [agree/disagree scale].” Example of balance-monitoring question: “How often do you log into your [Q1 BANK] banking website specifically to look at your account balance?” Example of other fee question: “Debit cards have too many fees associated with them [agree/disagree scale].”

³¹ The study by Agarwal et al. (2009) is a nice exception, drawing on administrative data from multiple financial institutions.

³² We do not report the results, but our sample is also younger than average, and higher-income and more-educated even conditional on age.

information (in exchange for the compensation they get for participating), although nearly all household surveys on consumer finance face a similar selection issue. Most of our panelists manage other aspects of their finances online. It is also worth noting that the distribution of our self-reported financial literacy measure is roughly centered on average, and even tilted a bit toward below average. This is somewhat surprising given other evidence that people tend to self-assess (too) favorably in household finance (Lusardi and Tufano 2009)³³ and related domains such as stock-picking (Glaser, Weber, and Langer 2010).

These representativeness issues are worth noting, but it is equally important to note that they predict nothing about the central representativeness question here, which relates to possible heterogeneity in comparative statics: do our panelists respond more or less to shocks that might affect attention paid to daily household finance? We are not aware of any evidence, theoretical or empirical, that produces sharp priors.

IV. Empirical Strategy and Results

A. Model and Identification

The research question is whether shocks to attention affect overdrafting, in both the short run and over sustained periods. Our empirical model approaches that question by estimating the relationship between overdrafting and within-panelist variation in overdraft surveys taken. The unit of observation is a panelist-month, because we observe only the month in which a survey was administered.

The primary estimating equation is:³⁴

$$\begin{aligned} AnyOD_{it} = & \beta_1 TookOD_{it} + \beta_2 ODSurveys_Last2yrs_{it} \\ & + \beta_3 TookAny_{it} + \gamma \cdot AnySurveys_{it} \\ & + \beta_4 AnyOD_{i,t-1} + \beta_5 Snowball_{i,t-1} + \delta_i + \lambda_t + \varepsilon_{it} \end{aligned}$$

We estimate the model with OLS, and cluster standard errors on panelist.³⁵ The dependent variable is an indicator equal to one if panelist i incurs at least one overdraft fee in month t . This

³³ Lusardi and Tufano's estimates in their Table 4 suggest that only 13% of the U.S. population rate themselves as having below-average overall financial knowledge.

³⁴ The model also includes a dummy equal to one if the month is the panelist's last in the sample. We typically observe transactions for only part of the month in such instances, meaning that overdrafts are less common.

specification models attention as operating on the extensive margin (“does the panelist overdraw?”) rather than the intensive margin (“how many overdrafts, or how much in fees paid?”). We also examine the intensive margin below; to preview those results, we find little evidence of a relationship between overdraft survey-taking and the intensive margin.

The model estimates both short- and long-run relationships between overdrafting and having taken overdraft-related surveys. The contemporaneous variable $TookOD_{it}$ equals one if the panelist took an overdraft survey in the current month and measures the “immediate” effect. The variable $ODSurveys_Last2yrs_{it}$ measures the total number of overdraft-related surveys taken in the last two years. This variable increases with survey-taking and falls over time as overdraft-related surveys recede into the past for a panelist. It specifies what one can think of as a “stock” of attention that builds with survey-taking and decays over time. Below we present results from models with shorter- and longer-horizon specifications of the stock effect and more flexible functional forms on the relationship between the stock variable and overdrafting. The survey stock variable increments by one in the month after taking a survey, in order to cleanly distinguish the stock effects from the immediate effects.

We control for selection into surveys generally by including a contemporaneous indicator $TookAny_{it}$ equal to one if the panelist took *any* survey, including an overdraft-related survey, in the current month. Thus, the coefficient on $TookAny_{it}$ measures the relationship between selection into surveys and overdrafting (along with any causal effect of survey-taking generally), and the coefficient on $TookOD_{it}$ measures the incremental effect of taking an overdraft-related survey relative to taking any survey. $AnySurveys_{it}$ is a set of indicators for the total number of surveys taken. These control flexibly for longer-run selection in surveys: secular trends or dynamics in the dependent variable that are associated with taking surveys or with taking multiple surveys.

The other controls are fixed month/year effects and fixed panelist effects.³⁵ Some models also include a lagged dependent variable $AnyOD_{i,t-1}$ measuring overdrafting in the previous month, and a variable $Snowball_{i,t-1}$ equal to one if the panelist incurred five or more overdraft

³⁵ Within survey months there is variation in survey-taking; in one month, we have variation in survey-taking and survey content, because both an overdraft-related and a generic survey were offered.

³⁶ We have also estimated specifications including leads of overdraft-survey and any-survey taking; these variables tend to be insignificant and do not change the results.

fees in the previous month.³⁷ Both variables are intended to capture the autocorrelation in overdrafting summarized in Table 1, although their inclusion does not change the results.³⁸

Because we rely on within-panelist variation in survey-taking to identify the model, our econometric analysis excludes panelists with no variation in the dependent variable. We exclude the 48% of panelists who never overdraw their accounts, and the 1% of panelists who overdraw every month.

Our identifying assumption is that conditional on the right-hand-side variables there are no differential unobserved secular dynamics in the dependent variable across those who take overdraft-related surveys and those who do not take overdraft-related surveys. Under that assumption our estimate of the causal stock effect of taking overdraft-related surveys is the $ODSurveys_Last2yrs_{it}$ coefficient, and the immediate causal effect is the $TookOD_{it}$ coefficient.³⁹ The identifying assumption seems reasonable given that survey topics are not pre-announced and that we control for selection into survey-taking through the “any survey” variables.

B. Main Results

Table 4 presents results from our primary specifications. We build from simpler to richer specifications reading across columns. The first column includes only the overdraft survey variables. The next adds the “any survey” variables, including $AnySurveys_{it}$ linearly. Tests of the linearity restriction for both the overdraft and any survey stock variables do not reject linearity on the $ODSurveys_Last2yrs_{it}$ variable, but do reject linearity on $AnySurveys_{it}$ (column 3). The “any survey” stock variables are jointly significant and generally indicate that selection into survey-taking is positively correlated with incurring overdraft fees.⁴⁰ Column 4 adds controls for

³⁷ We have also estimated the model with lagged values of the survey variables – one lag in some cases and several in others (up to a model omitting the parametric stock effects and including 24 lags of $Took_OD$ and $Took_Any$). What we learn from those models is very similar qualitatively to what we learn from the more parsimonious models.

³⁸ We have also estimated specifications that control (flexibly) for the number of months a panelist has appeared in the panel. These variables are not significant and tend to reduce the precision of the estimates on the variables of interest. The “months-in-sample” variables are included (and significant) in some of the specifications we discuss later.

³⁹ Mechanically, the immediate effect is the sum of the coefficients on the two immediate survey variables, subtracting any bias from selection into surveys; i.e., the net effect is: $\beta_1 + \beta_3 - \beta_3 = \beta_1$.

⁴⁰ The relationship between generic survey-taking and overdrafting is non-linear: there is little relationship until roughly the fifth survey, then the likelihood of overdrafting rises. For 7-12 generic surveys taken, the likelihood of overdrafting averages 5% higher than baseline. Results beyond that (13+ surveys) are noisy.

other survey content (gift card, contactless card and auto loan questions). The coefficients on these variables are not significant and can be interpreted as a form of placebo test. Column 5 adds measures of lagged overdrafting—the $AnyOD_{i,t-1}$ and $Snowball_{i,t-1}$ variables. Overall, the results are robust to variations in specification. The specification in column 5 is our preferred one for the rest of the analyses.

We find significant immediate effects of taking an overdraft-related survey. In our preferred specification (column 5), the point estimate on the $TookOD_{it}$ coefficient is -0.037 (p-value 0.06) on a sample mean of 0.30, for the estimated reduction in the likelihood of overdrafting of 12% from its mean level. The point estimate on the immediate “any survey” coefficient is positive, with a p-value of 0.09. The lagged overdrafting coefficients are what one would expect given the very strong unconditional positive serial correlation documented in Table 1.

We also find statistically strong and economically meaningful stock effects of taking multiple overdraft-related surveys. The point estimate in column 5 suggests that each overdraft survey taken in the last two years reduces the probability of overdrafting by 0.017. For a panelist taking three surveys, the cumulative effect is -0.051 on a mean of 0.30, which is a 17% decline.

Table A5 provides more detail on the specification and dynamics of the stock effect. The first column is identical to that in Table 4, column 5 (we show only the stock effect variable). The second and third columns include linear specifications of the stock effect with faster decay (“OD surveys taken, last year”) and no decay (“OD surveys taken, ever”). The coefficients on these alternatively specified stock variables are smaller and not statistically significant. The fourth column includes all three variables; the two-year stock effect variable is significant, while the others are not.

Table A5 columns 5-7 replace the linear stock variable with a set of indicators for number of surveys taken and also vary the stock horizon. Column 5 is of greatest interest, as it shows the unrestricted coefficients from the two-year stock effect specification. The linear specification seems to be a good approximation for the survey-by-survey stock effects, and one cannot reject the linear functional form restriction in any model.

Together the results suggest that taking surveys with overdraft-related questions meaningfully reduces the likelihood of overdrafting in the short term and also leads to persistent but not permanent cumulative effects of having taken many surveys. Attention builds with survey-taking and decays over time.

C. Effects on the Intensive Margin of Overdrafting?

Table A6 specifies the dependent variable in several different ways: as an indicator equal to one if the panelist incurred 5+ overdraft fees, as the number of overdraft fees incurred, and as the level of overdraft fees incurred in dollars. In brief, the results show little support for a relationship between survey-taking and the intensive margin of overdrafting. It may be that attention to overdrafting is discrete or simply that our confidence intervals are wide (they do not rule out large reductions commensurate with what we find on the extensive margin).

D. Sample Splits by Education, Income, and Financial Literacy

Table 5 presents results from our preferred specification (Table 4, column 5), split across three panelist characteristics measured in the registration survey: education, financial literacy and income. Because of the noise that arises when splitting the sample too finely, we break the sample into two groups for each of the three baseline variables of interest here. The top row of the table reports sample means for the dependent variable for each group. Sample sizes vary across groups in this table, both because of the category break-points and because of missing data for some panelists (the financial literacy question, in particular, was asked only in a subset of registration surveys).

The results in Table 5 show that the point estimates of treatment effect(s) are uniformly larger among those with less education (relative to more education), low/medium self-assessed financial literacy (relative to high), and lower income (relative to higher income). The lower groups do overdraw accounts more often on average, though not dramatically so, meaning that the *proportional* effects implied by the point estimates are also larger for these groups. For example, in the low-education subsample, the point estimates suggest that a panelist who has taken two surveys has the probability of overdrafting lowered by 0.044pp on a base of 0.290, a 15% reduction.

Evidence on whether the treatment effects differ significantly across the groups is weak. T-tests reject equality only for the stock effect in the education split. Generally, the coefficient estimates are not precise enough for us to make sharper statements about how demographics are correlated with the effects of survey-taking.

We can say with some conviction, however, that the treatment effects are substantial in the low-education, low-literacy and low-income subsamples. This is informative in itself, in that consumer advocates and some policymakers tend to focus on bank fees paid by those groups (see, e.g., <http://www.overdrawnmovie.net/>).

E. More Treatment Effects: Other Useful Variation in Survey Content

Survey content varies in two other useful ways. As we mention above, one survey focuses entirely on overdrafts and potentially provides some information on prices and outside options. One might expect the “overdraft-focused” survey to have different effects than surveys that are merely “overdraft-mentioning.” Other useful variation in survey content comes from questions that are not precisely about overdrafts but which seem plausibly related in an associative way. We identify three categories of such questions (see footnote 20 for examples of question wording). One is spending control, since reducing spending is one way people can engineer overdraft reductions (see the next subsection). Another is monitoring of available checking account balances, given that closer monitoring of balances may be key to overdraft avoidance (see the next subsection). The third category is other bank fees, since overdraft fees represent the majority of bank fees paid by our panelists, raising the possibility that mentioning “bank fee” may trigger the thought of “overdraft fee” in the panelist’s mind.

Table 6 presents results from models that allow for differences between overdraft-mentioning surveys and the overdraft-focused survey and that allow balance-/fee-/spending-mentioning surveys to affect overdrafting too. Column (1) adds the balance/fee/spending survey variables to our main specification (from Table 4, column 5). The immediate effect is not significant, but the stock effect is significant and comparable in point terms to the effect of overdraft-related surveys (compare to Table 4, column 5 and/or to the overdraft stock variables in this table). Column (2) relaxes the restriction that overdraft-mentioning and overdraft-focused surveys have identical effects. Their immediate effects are quite similar in point terms, but the stock effect of the overdraft-focused survey is much larger and significantly different from the stock effect of the overdraft-mentioning surveys. Column (3) refines the specification by restricting the “overdraft-mentioning” and balance/fee/spending surveys to have identical *stock* effects. Column (4) maintains that restriction and also restricts the overdraft-mentioning and overdraft-focused surveys to have identical *immediate* effects, since these restrictions cannot be rejected.

The upshot of this table is that variation in survey content seems to matter in an intuitive and associative way. The survey focused on overdrafts has a greater effect, at least in stock terms. Surveys about topics related to overdrafting also have meaningful stock effects. This result suggests that shocks to attention work in an associative way, cognitively speaking.

F. Mechanisms: How People Engineer Overdraft Reductions

If taking overdraft-related surveys induces less frequent overdrafting, how are panelists engineering those reductions? We address this question in Tables 7 and 8 by analyzing other measures of checking account usage as dependent variables, using our main specification for RHS variables except for changing the lagged dependent variable in accordance with the new LHS variables used here. We also control for tenure in the sample here by including a variable measuring “months in sample” and interacting that variable with the time period during which the panelist entered.⁴¹

The first dependent variable in Table 7 is the total number of monthly checking account spending transactions (column 1). The immediate effect on this variable of taking an overdraft-related survey is negative (implying a 2% reduction) but insignificant. The stock effect is negative and implies a 4% reduction in the transaction count for each overdraft-mentioning survey taken within the last two years. Column 2 shows that we do not find a significant reduction in checking account spending in dollar terms (we use the log of spending to deal with skewness), although the confidence intervals do include substantial changes.

Columns 3 and 4 show that the stock effect on spending transactions operates both when balances are quite close to the overdraft threshold (balance <\$100) and when balances are higher (>=\$100): panelists reduce transactions in both states.⁴² (Note from column 7 that crossing the low-balance threshold is a common occurrence: 83% of panelist-months dip below a \$100 balance at some point.) The point estimate on low-balance transactions is larger in proportional terms, implying a 5% reduction (0.376/7.0) for each overdraft-mentioning survey taken, vs. a 3% reduction (1.428/41.5) for higher-balance transactions.

Balances can be a choice variable too, and so we examine whether panelists take actions that are consistent with a strategy of increasing balances (e.g., maintaining a buffer stock) to avoid

⁴¹ These variables are significant in the models of Tables 7-8 but not in the overdraft models of Tables 4-6.

⁴² We get similar results using other cutoffs such as \$50 and \$20.

overdrafts. Column 5 shows no effects on the log of the dollar value of account credits (deposits and transfers into the account). Columns 6 and 7 show that the likelihood of reaching a low-balance threshold ($< \$100$) does not change following overdraft-related survey-taking.

In all, Table 7 suggests that panelists engineer overdraft reductions by reducing spending transactions (i.e., by managing outflows) but not by increasing inflows.

Table 8 looks at the composition of spending transactions. Columns 1 and 3 show significant and negative stock effects on both debit card and autodebit (ACH) transactions. These results seem intuitive given that balances-at-clearing are relatively hard to monitor and manage for both types of transactions (although the same could be said of checks, and we do not find a significant result there). These results also suggest that consumers may use a mix (within and/or across people) of high-frequency and low-frequency vigilance to avoid overdrafts: reducing debit card use may well involve a series of decisions to not pull out the card, whereas an automatic debit can be cancelled (and the bill switched to manual pay) with a single phone call or web log-in. We do not find significant effects on checks, cash withdrawals, or paydown of credit card balances (for credit cards we use dollar value rather than transaction count, since credit cards require at least one payment per month), although the point estimate on the stock variable is negative in each case.

In all, Table 8 suggests that people respond to overdraft surveys by managing spending transactions more intensively on multiple margins. Whether this heterogeneity in response to attention shocks occurs within-person (i.e., a given person implements multiple strategies) or across-person (different people respond differently) is not something we can pin down.

V. Conclusion

Overdraft fees became an increasingly important source of consumer outlays and bank profits over the 2000s. Our results suggest that limited consumer attention plays an important role in explaining overdrafts.

We find that consumer overdrafting drops sharply after taking surveys that mention overdrafts or topics that consumers might associate with overdrafts: spending control, balance-monitoring, or other bank fees. Plausibly unrelated survey content (on auto loans, gift cards, and contactless cards) does not affect overdrafts. The overdraft surveys have an immediate (same-month) effect and a stock (long-run) effect: attention is built via survey-taking and decays slowly

over time. The stock effect is stronger from the one overdraft-focused survey than from the other surveys that merely mention overdrafts. These effects exist even among the less well-educated and financially literate, an important finding if one wants to target those groups with an attention-based treatment. We find strong effects on the extensive margin (paying any overdraft fee in a given month) only, not on the intensive margin (fee amount paid, conditional on paying a fee). Consumers implement overdraft fee reductions by reducing spending transactions on debit cards and autodebits (ACH), suggesting a combination of high-frequency and low-frequency vigilance on transactions where the balance at the time of transaction clearing is difficult to monitor/forecast. We do not find evidence that consumers respond to the shocks by increasing balances in their checking account.

Our results suggest an important role in household finance for limited consumer attention that is dynamic, associative, and malleable. The surveys we examine do not provide any direct and panelist-specific information about account terms, meaning that the effects we observe are, in principle, distinct from those one would observe following classic disclosure of account terms.

Our results raise some provocative new questions for policy. Might more subtle approaches to drawing attention or conveying information complement, or even substitute for, more standard and direct approaches like information disclosure? Mandated disclosure typically focuses on the one-time, up-front provision of information, and our results suggest that this approach may have only a temporary effect on consumer behavior.⁴³ Ongoing messaging might be more useful in helping consumers attend to their household finances. If so, who should and will provide such messaging? These questions merely skim the surface of the possibilities for further research on limited attention, household finance, and consumer financial protection.

⁴³ Stango and Zinman (2011) highlight the enforcement challenges involved in mandated disclosure in consumer lending. Barr, Mullainathan, and Shafir (2008) make a similar point regarding mandated default options; the lender “moves last” and has incentives to get the consumer to undo the default (e.g., “just initial here” to opt in to standard overdraft service).

REFERENCES

- Agarwal, S., J. Driscoll, X. Gabaix, and D. Laibson. 2009. "The Age of Reason: Financial Decisions over the Lifecycle with Implications for Regulation." *Brookings Papers on Economic Activity* Fall: 51-117.
- Agarwal, Sumit, John Driscoll, Xavier Gabaix, and David Laibson. 2008. Learning in the Credit Card Market. Working Paper, February.
- Alba, Joseph, J. Wesley Hutchinson, and John G. Lynch. 1991. Memory and Decision Making. In *Handbook of Consumer Behavior*, 1-49. Edgewood Cliffs, NJ: Prentice Hall.
- American Bankers Association. 2007. *80 Percent of consumers have not paid overdraft fees in past year, says ABA survey*. Press Release. August 29.
- Barber, Brad, and Terrence Odean. 2008. "All that Glitters: The Effect of Attention on the Buying Behavior of Individual and Institutional Investors." *Review of Financial Studies* 21 (2): 785-818.
- Barr, Michael, Sendhil Mullainathan, and Eldar Shafir. 2008. *Behaviorally Informed Financial Services Regulation*. New American Foundation White Paper.
- Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman. 2010. "What's Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment." *Quarterly Journal of Economics* 125 (1) (February): 263-305.
- Burrhouse, Susan, Patricia Cashman, Robert Cordeiro, Timothy Critchfield, Yan Lee, Victoria Pawelski, and Katherine Samolyk. 2008. *FDIC Study of Bank Overdraft Programs*. November.
- Center for Responsible Lending. 2009. *Overdraft Fees and Opting In: A survey of consumer preferences*. March.
- . 2010. *National Bank Regulator Enabled Overdraft Abuses*. CRL Policy Brief. February.
- Chartrand, Tanya, Joel Huber, Baba Shiv, and Robin Tanner. 2008. "Nonconscious goals and consumer choice." *Journal of Consumer Research* 35 (August): 189-201.
- DellaVigna, S. 2009. "Psychology and economics: Evidence from the field." *Journal of Economic Literature* 47 (2) (June): 315-372.
- Dholakia, Utpal. 2010. A critical review of question-behavior effect research. In *Review of Marketing Research*, 7:147-199. Armonk, NY: M.E. Sharpe.
- Eliasz, Kfir, and Ran Spiegler. 2011. "Consideration Sets and Competitive Marketing." *Review of Economic Studies* 78 (1): 235-262.
- Gabaix, Xavier, and David Laibson. 2006. "Shrouded attributes, consumer myopia, and information suppression in competitive markets." *Quarterly Journal of Economics* 121 (2): 505-540.
- General Accounting Office. 2008. *Federal banking regulators could better ensure that consumers have required disclosure documents prior to opening checking or savings accounts*. January.
- Glaser, Markus, Martin Weber, and Thomas Langer. 2010. "Overconfidence of Professionals and Lay People: Individual Differences within and between Tasks?" *Working Paper*.
- 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.

- Grubb, Michael. 2011. Penalty pricing: Optimal price-posting regulation with inattentive consumers. Working Paper.
- Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman. 2011. Getting to the Top of Mind: How Reminders Increase Saving. Working Paper, January.
- Kolko, Jed. 2010. "A new measure of US residential broadband availability." *Telecommunications Policy* 34 (3) (April): 132-143.
- Koszegi, Botond, and Adam Szeidl. 2011. "A Model of Focusing in Economic Choice." *Working Paper*.
- Lacetera, Nicola, Devin Pope, and Justin Sydnor. 2010. "Heuristic Thinking and Limited Attention in the Car Market." *Working Paper*.
- LeBoeuf, Robin, Eldar Shafir, and Julia Belyavsky Bayuk. 2010. "The conflicting choices of alternating selves." *Organizational Behavior and Human Decision Processes* 111: 48-61.
- Lee, Young Han, and Ulrike Malmendier. 2010. "The Bidder's Curse." *American Economic Review*.
- Lusardi, Annamaria, and Peter Tufano. 2009. Debt Literacy, Financial Experience, and Overindebtedness. Working Paper, March.
- Parrish, Leslie. 2008. *Consumers Want Informed Choice on Overdraft Fees and Banking Options*. Center for Responsible Lending. April 16.
- . 2009. *Overdraft Explosion: Bank fees for overdrafts increase 35% in two years*. Center for Responsible Lending. October 6.
- 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.
- Sela, Aner, and Baba Shiv. 2009. "Unraveling priming: When does the same prime activate a goal versus a trait?" *Journal of Consumer Research* 36 (October): 418-433.
- Stango, Victor, and Jonathan Zinman. 2009. "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.
- . 2011. "Fuzzy math, disclosure regulation, and credit market outcomes: Evidence from Truth-in-Lending reform." *Review of Financial Studies* 24 (2): 506-534.
- Stavins, Joanna. 1999. "Checking Accounts: What Do Banks Offer and What Do Consumers Value?" *New England Economic Review* (April): 3-13.
- Zinman, Jonathan. 2009. "Where is the missing credit card debt? Clues and implications." *Review of Income and Wealth* 55 (2): 249-265.
- Zwane, Alix, Jonathan Zinman, Eric Van Dusen, William Pariente, Clair Null, Edward Miguel, Michael Kremer, et al. 2011. "Being surveyed can change later behavior and related parameter estimates." *Proceedings of the National Academy of Science* 10 (1073): 1-6.