Crowdsourcing Financial Information to Change Spending Behavior

Francesco D'Acunto* Boston College

Alberto G. Rossi† Georgetown University

Michael Weber University of Chicago & NBER‡

This version: December 2019

Abstract

We test the effects of providing individuals with crowdsourced spending information about their peers (individuals with similar characteristics) through a FinTech app. Users change their spending in the direction of their peers asymmetrically—those who spend more than their peers cut their spending by 9%, whereas others only increase their spending by 1%. We confirm these baseline facts in a regression discontinuity design that compares similar users who are assigned to different peer groups due to the sharp definitions of peer-group thresholds. Users whose peer groups are more precisely tailored to their characteristics, and hence more informative, more liquid users, lower-income users, and users whose pre-signup spending was farther away from that of peers react more. Discretionary spending including cash withdrawals, which are commonly used for incidental expenses and anonymous transactions, drive the effects. We argue that Bayesian updating, peer pressure, or the fact that bad news loom larger than (equally-sized) good news can barely explain all these facts at once.

JEL: D12, D14, D91, E22, G41.

Keywords: FinTech, Learning, Beliefs and Expectations, Peer Pressure, Financial Decision-Making, Saving, Consumer Finance.

This research was conducted with restricted access to data from *Status* Money. The views expressed here are those of the authors and do not necessarily reflect the views of *Status* Money. We are indebted to Majd Maksad and Korash Hernandez (the founders of *Status* Money Inc.) for providing us with invaluable feedback and comments at each step of this project in addition to allowing us to use the company's data. All errors are our own. Weber gratefully acknowledges financial support from the University of Chicago Booth School of Business and the Fama Research Fund at the University of Chicago Booth School of Business.

*Carroll School of Management, Boston College, Chestnut Hill, MA, USA. e-Mail: dacuntof@bc.edu

†McDonough School of Business, Georgetown University, Washington, DC, USA. e-Mail: agr60@georgetown.edu

‡Booth School of Business, University of Chicago, Chicago, IL, USA and NBER. e-Mail: michael.weber@chicagobooth.edu.

1 Introduction

Low savings limit the wealth accumulation of households, who often reach the time of retirement holding inadequate financial resources to maintain their pre-retirement lifestyle (e.g., see Banks, Blundell, and Tanner, 1998; Bernheim, Skinner, and Weinberg, 2001; Lusardi and Mitchell, 2007; Olafsson and Pagel, 2018b;). Channels that contribute to this phenomenon, whether neoclassical or not, include liquidity constraints (Zeldes, 1989; Jappelli and Pagano, 1994), hyperbolic discounting (Laibson, 1997), limited attention (Madrian and Shea, 2001; Carroll, Choi, Laibson, Madrian, and Metrick, 2009), expectations-based reference-dependent preferences (Pagel, 2017), and the lack of financial literacy (Van Rooij, Lusardi, and Alessie, 2012; Chalmers and Reuter, 2012; Lusardi and Mitchell, 2014; Lusardi and Mitchell, 2017).

These potential explanations suggest that households might have little information about the income, spending, and savings rates that would guarantee the appropriate wealth accumulation before retirement. A rule of thumb inexpert households might exploit is to obtain information about saving norms while observing the behavior of peers, whether or not peers' spending and saving levels are optimal for the household (D'Acunto, Malmendier, Ospina, and Weber, 2018). But the overall spending and savings of peers are mostly unobserved, and hence households can barely learn about the prevailing savings rates of those with similar incomes and demographic characteristics (Lieber and Skimmyhorn, 2018). Households might even infer wrong saving rates by observing only the most visible and conspicuous component of peers' spending (Han, Hirshleifer, and Walden, 2019; Charles, Hurst, and Roussanov, 2009).

If this information friction was material, disclosing the spending of peers with similar income and other demographic characteristics might change individuals' beliefs about the appropriate spending and savings rates. This update would happen irrespective of whether peers' savings rates are optimal, as long as agents believe the signal they receive is credible and valuable (Gargano and Rossi, 2018; Gargano, Rossi, and Wermers, 2017) and as long as agents are attentive to their personal finances (Olafsson and Pagel, 2017). Moreover, this information might affect individuals' beliefs and choice both directly, through learning about others' spending, and indirectly, through peer pressure, that is, the concern of lagging behind with respect to peers. For the case of households' and investors' financial decisions, existing research is split on whether peer effects are material (e.g., see Duflo and Saez, 2003; Bursztyn, Ederer, Ferman, and Yuchtman, 2014; Chalmers, Johnson, and Reuter, 2014; Ouimet and Tate, 2019; Maturana and Nickerson, 2018; Maturana and Nickerson, forthcoming) or immaterial (Beshears, Choi, Laibson, Madrian, and Milkman, 2015; Lieber and Skimmyhorn, 2018).

In this paper, we study the effects of providing households with crowdsourced information about their peers' spending through a free-to-use FinTech application (app) called *Status*. Upon subscribing to the app, users provide a set of demographic characteristics, which include their annual income, age, homeownership status, location of residence, and location type. *Status* also obtains credit scores via credit reports. Using transaction-level data from a large representative sample of US consumers outside the app, *Status* computes the average monthly spending of consumers with similar characteristics as the users (*peers*). Moreover, users link their credit, debit, and other financial accounts to the app.[1]

Status then produces easy-to-grasp graphics that compare the evolution of the users' monthly spending with the evolution of the peers' spending. Figure 1 is an example of the graphics *Status* users see on their homepage.[2] These graphics give users simple and immediate feedback on whether their spending is higher or lower than peers' spending. Displaying this crowdsourced information in an easy-to-understand setting is a crucial feature of *Status*.

A desirable feature of the data is that *Status* obtains the pre-signup history of transactions of a large fraction of users through their financial institutions, and hence we can compute directly users' pre- and post-signup average monthly spending. In the baseline analysis, we find that the users who see that they overspend relative to peers at the time of signup reduce their monthly spending, whereas those who underspend increase their spending. The reaction, though, is severely asymmetric across the positive and negative domains. On average, users who overspend relative to peers reduce their seasonally-adjusted spending by $237 per month

[1]As we discuss in detail below, we only consider users with actively linked accounts and whose number of linked accounts does not change over time.

[2]Figure A.1 and Figure 3 show other graphics, and section 2 describes details about the setting.

Figure 1. Graphics Comparing Users' and Peers' Spending on *Status*' Homepage

around the adoption of the app. Instead, users who underspend increase their seasonally-adjusted spending by $71. Moreover, the distance from the level of peer spending relates monotonically to the reaction—significant overspenders cut their spending by $563, or 9% of their income. Significant underspenders only increase their spending by $91, or 1% of their income.[3]

The baseline analysis cannot rule out a set of endogeneity concerns. For instance, users might have already decided to cut or increase their spending before signing up (e.g., see D'Acunto, Prabhala, and Rossi, forthcoming). They might have signed up to enjoy other features of the app, such as the income-aggregation function or the possibility of setting dynamic targets for consumption and savings. They might have changed their spending irrespective of the information they received about peers. Although it seems barely plausible that the predefined change in spending lines up exactly with the peer spending users observe after signup, the baseline analysis cannot rule out this concern definitely.

Moreover, the baseline analysis cannot rule out that the baseline effects are, at least in part, mechanical, in the sense that users who spent more than usual in the months before signup

[3]Here, by significant we mean clients who under- or over-spend their peers by approximately 50% in either direction.

might revert to their typical level of spending after signup, irrespective of the information they obtained from the app, and vice versa for underspenders.[4]

To tackle these endogeneity concerns directly, we propose a sharp regression discontinuity identification strategy that exploits the fact that *Status* constructs peer groups based on preset ranges of income. Two users with similar incomes, but one with income slightly below a threshold and the other at the threshold, are matched to different peer groups and hence obtain different information about their peers' spending. Because users do not know this rule, they cannot strategically manipulate their reported income to avoid receiving negative news about their spending relative to peers.

For example, consider one of the yearly income thresholds the app uses—$100K. Suppose user A declares she earns $99K, whereas user B reports $101 K. Although these reported incomes only differ by $2K, users A and B will be matched to different peer groups. User B will face a peer-spending value that is the average of the transactions of otherwise similar consumers earning between $75K and $99K, whereas user A will observe a higher peer-spending value—the average of the transactions of consumers earning between $100K and $149K. We show that observables are not economically or statistically distinguishable across users just below or above the income thresholds. Crucially, we show that the amounts spent before signup are also similar around all thresholds, and hence mechanical mean reversion of spending cannot drive the results in this setting.

The regression-discontinuity strategy confirms our baseline results. Users assigned to a peer group whose spending is lower cut their spending more than users who are almost identical on observables, but are assigned to a peer group with higher spending.

We then move on to consider the heterogeneity of the baseline results based on the informativeness of the signal about peers. Note that *Status* does not say that the average behavior of peers is optimal in any respect. Yet, users might assume that the average behavior of peers includes information about their own optimal spending behavior (Galton, 1907; Wolfers and Zitzewitz, 2004; Da and Huang, forthcoming). If users believe that the signal about peers is

[4]Note that this mechanical interpretation could explain the baseline directions of spending changes, but it could barely explain the asymmetric reaction between overspenders and underspenders.

informative, those who observe a more informative signal should react more than others.

To bring this conjecture to the data, we first note that peer groups require e minimal number of individuals in the representative US cohorts to which the app has access. For this reason, in some cases users will see that all their characteristics are matched to a tight peer group, whereas in other cases, only some of the characteristics are tightly matched. If users believe the information is valuable, the tighter is the peer group, the more the user will find the signal about peers' spending informative. And, indeed, users whose peers groups are tightly tailored to their characteristics react more than other users.

Moreover, users who observe more peers in their daily lives might already have good information about peers' spending, and hence might react less to the information at signup relative to users who observe less peers. To assess this conjecture, we compare urban and rural users. Intuitively, rural users are exposed to less people in their daily lives and hence also observe less individuals with their same characteristics. Urban users, instead, have a dense population around them and likely a higher number of similar individuals. Consistent with this conjecture, we find that rural users react more to the peer information they obtain at signup relative to urban users.

We also consider heterogeneity in reaction based on users' liquidity and income. Renters react substantially more than home owners, users with low credit scores more than users with high credit scores, and low-income users more than high-income users.

We then move on to assess which spending categories users adjust more after they obtain information about peers' spending. Consistent with the presence of frictions in spending, the whole margin of adjustment comes from discretionary spending relative to non-discretionary spending, which households can barely reduce.[5] In particular, cash withdrawals show a dramatic drop after signup for users who spend more than their peers relative. Cash is mainly used for incidental expenses (Bagnall, Bounie, Huynh, Kosse, Schmidt, Schuh, and Stix, 2014) and for transactions consumers want to keep anonymous (Acquisti, Taylor, and Wagman, 2016), which implies that among the discretionary expenses users might reduce the ones that are least

[5]As we discuss below, non-discretionary spending includes groceries, fees, mortgage payments, and tuitions. Discretionary spending includes outside food and drink spending, clothes, entertainment, travels, and cash withdrawals.

likely to benefit their whole household instead of discretionary expenses that might benefit the whole households, like outside food and drinks and entertainment.

In the last part of the paper, we exploit our setting to dig deeper into the economic channels that might explain users' reactions to information about peer spending. An important feature of the app is that upon signup, users are not only exposed to information about peers, but also to information about the national average spending in the US as well as users' own average monthly income computed by the app based on inflows in users' accounts. We exploit this feature in the analysis of channels. For this part, we cannot use the regression-discontinuity design, because the other pieces of information households obtain through the app do not change systematically based on the quasi-random income threshold rule.

The first channel we consider is the *wisdom of the crowds*, under which users should update their beliefs about the optimal spending rate after observing information about their peers. For this channel to be relevant, users need to believe the information that *Status* provides is an informative signal about their optimal spending rate irrespective of whether it is or not. This channel could explain some of the facts we document, and especially the stronger reaction to more informative signals. At the same time, this channel can barely explain the asymmetric reaction of overspenders relative to underspenders. Under the *wisdom of the crowds*, users at the same distance from peers in either direction should react in opposite directions by the same amount.

The second channel we consider is peer pressure—individuals might obtain disutility from behaving worse than their peers. In this case, overspending might be perceived as a negative behavior because it reduces users' financial health with respect to peers. This channel can explain the reaction of overspenders but can barely explain why underspenders, who are not behaving worse than their peers, would react at all.

We then assess the extent to which overreaction to negative news might explain our results in full. Under this channel, users learn from peers' spending, but negative news about the difference between own and peers' spending loom larger than positive news. Although this channel has the potential to explain the asymmetric reaction of overspenders and underspenders, we find that users react more to negative information about peers than to negative information about average US consumers or about overspending with respect to their own income. The peer component is thus crucial to obtain a reaction, and users do not just react to the most negative piece of news they obtain once they sign up.

In terms of economic channels, we conclude that only a combination of the three channels we consider above can explain all our results in full.

We do not claim that that users' reaction to peer information is optimal in any respect, or that the information they obtain captures the actual spending of their peers—the very notion of peer is hard to define precisely and could be based on many different dimensions relative to the ones our app uses. What our paper finds is that once individuals observe the spending information of a large crowdsourced representative cohort that looks similar to them based on observables, and that is labeled as "peer group" in the app, individuals react systematically to this information. Further research should be devoted to study the optimal design of FinTech tools based on crowdsourced information to provide tailored advice for each user. For instance, whether including more categories than the ones we studied (e.g., amount of mortgage debt or student debt outstanding) might have a stronger effect on behavior is an interesting avenue for future research.

Note also that our analysis cannot speak to the external validity of the results. Because *Status* is marketed as a tool that provides comparisons with peers, the population that selects into using this service might be more sensitive than the average US consumers to the differences between their spending and peers' spending. Users might thus react more to the peer information than the average US consumer. We see our setting as an ideal laboratory to study the channels through which learning about unknown peers might affect individual behavior, and as in any other reduced-form analysis, we cannot make claims about the predicted effects of providing information about peers to the broader population.

The persistence of the effects of providing information about peers is also an aspect further research should assess. Within the time frame we observe, which includes about three months around the signup to *Status*, we do not detect any dissipation or reversal of the effects. Whether peer information might have long-lasting effects on savings will require observing longer time series than are currently available.

2 Institutional Setting

In this section, we discuss the characteristics of *Status*, the procedures users face in order to sign up, and the information they observe after signup.

2.1 Purpose of the app and Signup Process

Status is an app designed to help individuals make more informed decisions in the personal finance space. Similar to other apps in the US and abroad, *Status* provides an income aggregation feature, which allows users to visualize their full balance sheet in a simple and dynamic way and to access a set of metrics about their finances, which would be hard to compute on the household say without a holistic view of all their debit, credit, and investment financial accounts. The unique feature of *Status* is that is also shows users how individuals similar to them in terms of observable characteristics, which we label *peers*, manage their finances. In particular, the app shows statistics about the spending, the interest rates peers earn on their savings and pay on their loans, and what credit cards they use, among others. The information about peers is crowdsourced from transaction-level data the app accesses of a large population of individuals representative of the US population. The app thus aims at making easy for users to access an elaboration of big data about the transactions of a large representative population, which represents information that users would not be able to access on their own.

It is important to stress that the peers in our setting are not individuals that interact socially with users, and users do not know these individuals personally. This feature is crucial to our analysis, because if we were showing users information about the spending and savings of their friends and acquaintances, users would largely know this information already through social interactions and hence the extent to which providing such information would be expected to affect users' behavior is unclear. Instead, users know that the peers are a representative population of individuals with certain demographic characteristics, and that peer statistics are based on elaborating big data underlying the transactions of such individuals. This is also the reason why we do not label our tests as estimates of peer effects, which the literature attributes to the extent to which observing the behavior of individuals with whom one socially interacts changes one's behavior (e.g., see Bailey, Cao, Kuchler, and Stroebel, 2018). Rather, we estimate the extent to which obtaining crowdsourced information about the behavior of similar individuals using data and elaboration techniques that individual users are unable to access on their own changes users' behavior.

To sign up to *Status*, users provide their date of birth, their annual income, and their housing type—whether they own or rent the home in which they live. Users are then prompted to insert their address and the last four digits of their social security number. This information allows the app to connect to the credit bureau that returns all of the user's credit-score-related information.[6] Finally, the app asks users to link their checking and savings accounts, their credit-card accounts, as well as taxable and non-taxable accounts.

For each user, the app constructs a peer group based on the user's age, income, location, credit score, and housing type. Peer groups are constructed to be as precise as possible subject to the constraint that each group should have at least 5,000 individuals. The trade-off is that coarse groups may not be too informative, because they might contain individuals to whom users do not relate. On the other hand, spending patterns constructed using too few individuals may be too noisy and provide non-credible information. Note that for the sake of testing whether users react to information about peers' spending, whether such information is accurate or inaccurate is not relevant as long as users think the information they obtain contains an informative signal about their optimal spending rate. As we will show below, users seem to understand the trade-off of peer group characteristics in terms of informativeness, because users whose peer groups are more tightly matched to their characteristics react more to the information about peer spending.

In Figure 2, we provide an example of the screenshot that *Status* users observe about their

[6] We as researchers do not observe any individually-identifiable information about *Status*' users.

own characteristics (Panel (a)) and the characteristics based on which the peer group is defined (Panel (b)). In this fictitious example, the user is 42 years of age, has an annual income of \$140K, lives in New York, has a credit score of 769, and is a renter. The peer group constructed for this user contains individuals whose age ranges between 40 and 49, whose income ranges from \$100K and \$150K, who live in New York City, pay rent, and have a credit score that ranges between 720 and 779.

Main features of the app

Once the user is enrolled, the app automatically retrieves information from the users' savings and investment accounts. The app stores all transactions and investment returns and computes the user's net worth as the difference between assets and liabilities. To give the reader a sense of the information users observe, we describe the content of the home page below.

The main feature of the home page is comparing the user's spending with their peers' spending. Figure 1 in the introduction displays the vivid graphics that compare the users' own daily spending based on daily transactions with the projected average daily spending of the peer group and the national US average. The screenshot is taken as of October 30. On the top, the plot shows the user's total spending, which turns out to be \$17,799, together with the average peer spending, \$8,651, and the national average, \$4,222. The blue line presents the user's cumulative spending over the course of the month until October 30. It also presents a forecast of total spending until the end of the month. On the same graph, the light and dark red lines presents the peers and national average cumulative spending over the month. The app also displays as a grey dotted line the user's average income, \$10,204. As a final piece of information, the app explicitly tells the users how they are doing in terms of spending for the current month.

Note the users' spending is based on their own daily transactions. Recall that peers' and US national average information are computed using a proprietary algorithm that aggregates spending information for a large sample of US consumers whose transactions *Status* observes. These transactions are aggregated at the monthly level and projected linearly for each day of the month, and not computed on a daily basis. This difference in the frequency at which users' and peers' data are elaborated is not relevant to the scope of our inquiry unless the difference in frequency and timing of the pieces of information makes certain users believe the information *Status* gives them is not credible. But in this case, we would observe no reaction of users to peers' information irrespective of their distance from peers' spending. If anything, this feature of the app might reduce the average reaction on the side of users.

The bottom of the home page reports more comprehensive statistics regarding users' debts, assets, net worth, and credit score (see Figure 3). Our fictitious user has a debt of $37,393, which is compared to peers' debt of $13,429 and a national average debt of $50,297. On top of this information, the app tells the user that the interest rate he/she is charged is competitive with the national average. The user has $40,839 in assets, compared to $45,759 for the peers and $119,934 for the national average. The interest rates earned on the accounts are competitive for two of the three accounts, but not for the third. The third quadrant reports the information for net worth, which is simply computed as the difference between assets and liabilities, and the fourth quadrant reports information for the user's credit score. He/she has a credit score of 769, the peers average is 754, and the national average is 630.

3 Data

Status collects and displays large amounts of information from and to their users. Some of this information is calculated on demand based on users' requests and is not stored on their servers. For this study, we accessed a subset of the information that *Status* collects.

First, we observe a set of demographic characteristics as of the time users sign up to *Status*. We do not observe these characteristics over time. In addition to users' unique identifiers, we observe the date on which the user joined *Status*. As far as demographics are concerned, we observe several self-reported dimensions including users' age, credit score, gross income, whether the user owns or rents the house in which he/she lives, as well as the zip code and city in which the user lives.

The second set of characteristics we observe refers to users' peers. For each user, we observe the characteristics of the peer group that *Status* computes as an average of the individual characteristics of individuals with demographic characteristics similar to the user's. Note *Status* does not use the characteristics of other app users to construct peer groups, but instead uses external proprietary data on a representative set of US consumers. This procedure rules out that any selection in the types of consumers that sign up might be captured in the average demographics of peer groups. The peer demographics we observe include the average credit score of the peer group, the average debt, the average value of assets, the average net worth, and the average income. We also observe the range of credit scores and income for the peers. Moreover, peer groups are constructed separately across geographic locations and rural versus urban areas. For each peer group, we also observe the number of individuals that enter the group.

Third, we observe a set of variables that capture the usage of *Status* accounts by account holders. These variables include the number of account logins by users during the first, second, and third month after signing up, the number of external financial accounts users had linked as of August 21, 2018, the number of external financial accounts users had linked at the time of sign-up, as well the asset balance, the debt balance, the savings balance, and the balance from linked investment accounts, all measured as of August 21, 2018.

Finally, we observe data on users' and peer groups' spending amounts and spending categories. Specifically, we observe users' total spending and peers' average spending over the first, second, and third month before users signed up for *Status* as well as the total spending over the first, second, and third month after users signed up for *Status*. Spending is broken down into categories based on classifying the vendor related to each transaction. The transactions that cannot be classified are labeled as other expenses. Observed categories include checking-account withdrawals, auto and gas, education, entertainment, fees, gifts and charity, groceries, health and medical, home improvement, housing, loans, restaurants, shopping, travel, utilities and bills, and other expenses. We use these categories to classify each spending amount into discretionary or non-discretionary spending in each month. Discretionary spending includes checking-account withdrawals, entertainment, restaurants, shopping, travel, and fees. Non-discretionary spending includes groceries, utilities and bills, health and medical, auto and gas, and education. Because we cannot ultimately classify the remaining categories into discretionary and non-discretionary, we exclude them from the analysis of the effects of information about peers on spending broken down by spending categories.

3.1 Sample Selection

To ensure that our working sample includes individuals for which we meaningfully observe inflows and outflows before and after signup, we select the raw sample in a few steps.

First, we only include users that have linked at least one bank account to the app at signup. Second, we only include users whose number of linked account does not change through the 90 days after signup, which is the longest horizon we use in the analysis. This step is important, because users might link one account at signup and start linking other accounts over time. In this case, we might miss the reaction of users in accounts that were not linked at the time of signup.

Third, we only include users with at least one monthly login to the app after signup. This step excludes that because of users' inaction, the number of linked accounts in our working sample drops and/or the app fails to download information from users' accounts, which would appear as a misleading drop in inflows and spending.[7]

Fourth, we only include users who spend at least $100 monthly throughout the period we have them in the working sample. This step ensures that we do not keep individuals who are not actively using the accounts they link for spending purposes.

Finally, we verify that no users are duplicates, that is, individuals who might have opened different accounts at different points in time, for instance because they forgot their login credentials. Repeated accounts would overestimate the statistical significance of our results. To identify duplicates, we search for individual accounts with the same exact balances at the end of each month. We do not detect any duplicates based on this procedure.

[7] We thank Michaela Pagel for suggesting this possibility.

3.2 Summary Statistics

Table 1 reports the basic characteristics of the clients in our sample. For each variable, we report the number of observations, average, and standard deviation. The first three variables are demographic characteristics: *Age*, *Credit Score*, and *Home Ownership*. The average client is 30 years old, with a standard deviation of seven years, indicating *Status* users are rather young. The average credit score is 728, higher than the average US credit score of 687. Thirty-eight percent of users are homeowners, which is below the US average, in line with the fact our sample is, on average, younger than the US average consumer.

The average client earns approximately $90,000 per year, with a large standard deviation of $61,000, suggesting our sample spans individuals with varying levels of income. The majority of the *Status* users have a positive net worth. The average assets are $42,462, whereas the average debt—including credit-card debt—equals $29,971.

Figure 4 reports the distribution of monthly spending by income quartiles. We highlight two main facts that suggest our data align with intuition and are reliable. First, monthly spending increases with income. Across the four income groups, average spending equals $2,200, $3,409, $4,470, and $6,974. Second, the within-group standard deviation of spending increases with income. Higher-income individuals have more varied levels of spending than lower-income individuals, which is consistent with low-income individuals facing spending constraints.

4 Signup and Spending: Baseline Results

Our baseline analysis tests whether the two pieces of information users receive at signup—whether they spend more or less than their peers, and how different their spending is with respect to peers' spending—have any effects on subscribers' subsequent spending behavior. We first compute the overall spending for each subscriber for the 60 days before signup and the 60 days after signup, and measure the change in aggregate spending across the two periods. Because spending is cyclical, we deduct the average change in spending across all users from the change in spending of each user. We refer to this quantity as seasonally adjusted spending in some cases and simply as spending in other cases.

Figure 5 is a graphical depiction of the change of users' spending before and after signup as a function of their distance from the value of peer spending users observe at signup. Subfigure (a) reports the results for changes in spending, whereas subfigure (b) reports the results for changes in spending, normalized by income. Each binned scatterplot divides the 17,500 users in 100 groups.

Figure 5 documents three features of the raw data. First, both underspenders and overspenders relative to their peers appear to converge to the value of peer spending in the 60 days after signup, relative to before signup. The second fact is a substantial asymmetric sensitivity of users' change in spending to their distance from peers' spending based on whether the group of users spend more or less than their peers. Third, the distance of each group of users from the amount of peer spending is monotonically related to users' change in spending—the further the group is from the peers' spending level, the higher the change in their spending, irrespective of the sign.

As an aside, note the average subscriber underspends compared to her peers. This detail is likely driven by the fact that peers' spending is computed in one specific month, July 2017. To assess whether the baseline reaction survives once we account for this and other systematic differences between the change in spending of all users and of peers, we estimate the effect in a simple multivariate setting. In Table 2, we regress the change in spending of each user on the amount of peer spending as well as a constant, which constant captures any systematic difference between the change in spending of all subscribers and peers.

Panel A of Table 2 shows that, in the raw data, we find the average subscriber who overspends with respect to her peer group reduces her spending after signing up, by an average of \$474/2=\$237 per month in the first 60 days after signup. Users who underspend compared to their peers instead increase their monthly spending by \$142/2=\$71.

To allow a more appropriate comparison across subscribers with different levels of income, we normalize the change in aggregate spending by the subscribers' income to make sure systematic differences in the propensity to spend across income levels do not drive any results.

The results, reported in Panel B Table 2, suggest overspenders reduce their spending by 3% of their income, whereas underspenders increase their spending by 1% of their income.

5 Causality? Regression Discontinuity Design

The baseline results do not rule out a few concerns about endogeneity in our setting. For instance, the endogenous timing of signup might be correlated with users' intentions to change their spending behavior, irrespective of the information about peers they obtain at signup. In this case, we would assign a causal interpretation to the exposure to peer information that would be unwarranted.

Another potential concern is that the reaction we attribute to information about peers might be mechanical, that is, users who happened to spend unusually more than their standard monthly expense just before signup and mean revert to their typical levels of spending after signup would be more likely to be classified as overspenders relative to peers and would mechanically appear as reducing their spending after signup. Similarly, users who spent unusually little just before signup would be more likely to be classified as overspenders and would appear as increasing their spending after signup.

Below, we will document a few regularities that indirectly question the plausibility of these concerns. For instance, users for which peer information is more informative react more; the distance from peers' spending relates monotonically to users' reaction; and, we estimate an asymmetric reaction above and below peers' spending, including the non-parametric estimate of a kink in users' reaction around peers' spending amounts. Yet, all these features of our results only speak to the concerns above indirectly. Before moving forward, ideally we would rule out these concern directly by comparing the reactions of users who have similar incomes and spend similar amounts before signup, but observe different information about peers at signup.

To tackle this challenge, we exploit a feature of the design of peer groups that allows a sharp regression discontinuity design (RDD) for identification purposes.

The intuition behind the design is that *Status* matches users with peer groups based on whether users fall above or below a set of annual income thresholds that are pre-specified and of which users are not aware when they sign up. The annual income thresholds are \$35K, \$50K, \$75K, \$100K, and \$150K. Because income is a continuous variable, small differences in income capture similar subscribers. And yet, users who report similar incomes but end up falling below or above an unknown threshold will be matched to different peer groups, and will face different information about peers' spending.

For example, consider a user who reports an annual income of \$99K and one who reports an annual income of \$101K. These two users would be quite similar with respect to their annual income, and we can test directly that they are similar along the set of observables we have, which include demographics as well as, crucially, the amount of spending in the months before signup. We can thus test directly that a mechanical mean reversion of users' spending cannot drive any results in this setting.

Whereas users around the thresholds are similar, the peers group to which they are matched, and hence the information about peers' spending, can be substantially different. In our example, the user who reports an income of \$99K receives information about the average spending of peers whose income is between \$76K (just above the immediately lower threshold) and \$99K, whereas the similar user who reports an income of \$101K receives information about the average spending of peers whose income is between \$100K and \$149K (just below the immediately higher threshold).

The user just below the threshold faces a peer group that spends, on average, substantially less than the group of the users just above the threshold. The treatment effect of peers' spending information is thus systematically different for the two users, and especially the user just below the threshold will be likely to appear as an overspender relative to her peers, whereas the user just above the threshold will be likely to appear as an underspender relative to her peers.

For this RDD design, we need to restrict the sample to a group of users that are close enough to the thresholds so that they do not differ on observables. Moreover, we need to have a large enough mass of users both below and above each threshold. Because some users tend to report rounded values of annual income, the mass of users just above the thresholds (which includes the exact threshold value) is higher than the mass of users at uneven values around the threshold.[8] In our identification sample, we thus include users with an income up to $6K below the threshold and users with an income up to $2K above the threshold.

Before proceeding with the analysis, we verify that we cannot reject the null hypothesis that users that end up just below and just above the thresholds differ on observables. The crucial difference we want to rule out is that users in different groups systematically spend more or less than others before signup, which would rule out the possibility of a mechanical effect of the treatment due to mean reversion in spending.

We provide evidence against this difference in Figure 6. For each threshold, we plot the average monthly spending in the 3 months before signup (in $ thousands) for users around the threshold at intervals of $500 of yearly income. Around each average, we report intervals whose size if one third of a standard deviation above and below the average. We can see vividly that we cannot reject the null that users above and below each threshold spent different amounts before signup. The highest threshold ($150K) is the one for which average pre-spending is more varied across income levels, largely because the amount of users with such high yearly income is not high and hence averages are based on a small number of observations. At the same time, even for this threshold we do not detect systematic economic or statistical patterns.

Armed with our RDD and our identification sample, we estimate a set of two-stage-least-square regressions. The first stage consists of the following specification:

$$Peer\ Spending_i = \alpha + \beta\ Dummy\ Above_i + \zeta Spending\ Before_i + \epsilon_i, \tag{1}$$

where *Peer Spending_i* is the peer-spending value user i sees at signup, and *Dummy Above_i* is a dummy variable for whether the user's income at the threshold or just above it. Based on the design, we predict that $\hat{\beta} > 0$—users at the threshold will be matched to a peer group whose spending is higher relative to those of the peer group of users below the threshold.

[8] In robustness tests, we show that the effects are similar if we exclude all rounders from the analysis, because earlier research shows that rounding relates to characteristics like cognitive abilities that might determine the extent of reaction to information (see D'Acunto, Hoang, Paloviita, and Weber, 2018)

In the second stage, we use the instrumented *Peer Spending_i* in equation 1 as the main covariate in the following specification:

$$\Delta Spending_i = \alpha + \gamma\ \widehat{Peer\ Spending}_i + \zeta Spending\ Before_i + \epsilon_i, \tag{2}$$

where $\Delta Spending_i$ is the change in consumption before and after signing up.

Even in the second stage, we predict that $\hat{\gamma} > 0$—the higher the spending of peers, the less likely the user is an overspender and the less likely she is to cut her spending. Conversely, the more likely she is to be an underspender and hence the more likely she is to increase her spending.

Before discussing the estimates of the two-stage least-square specifications, we verify that the baseline endogenous result we documented in the full sample also holds in the RDD sample. The RDD sample is a selected group and if the baseline result did not hold in this group of users the RDD analysis would not be meaningful. In Table 3, the columns to the left report the estimates for the association between peer spending and users' change in spending normalized by income for the three months after signup, relative to the three months before. Indeed, we find that the higher is the peer spending users observe at signup, the higher (i.e., the less negative in case users overspend) is the change in their spending. The baseline result ranges between 3.3 and 6.5 percentage points of the ratio of change in spending to income, based on whether we control for observables or not.

Moreover, the baseline effect holds both within the subsample of users below the income threshold and users above the income threshold. This test verifies that users on the two sides of the threshold behave exactly like the average user in the sample, that is, they increase their spending by more and are less likely to cut their spending whenever they are matched to a peer group with a higher average spending.

We report the estimates for equations (1) and (2) in Table 4, where the right columns control directly for individual-level observables. Irrespective of whether we control for observables, which indeed should not matter if the RDD assumptions hold, we find that (i) the quasi-random assignment rule is not a weak instrument for peers' spending (F-stats above 500) and

(ii) the difference between the change in spending over income for users that see a higher level of peer spending and users that see a lower level of peer spending is positive, consistent with the baseline result in the full sample, and amounts to 8 percentage points.

6 Heterogeneity: Informativeness of Signal, Liquidity, Distance from Peers

We move on to exploit the regression-discontinuity design to assess the potential heterogeneity of the average effect across dimensions that should be relevant if the information about peers users obtain at signup is relevant to their behavior.

6.1 Informativeness of Signal

First and foremost, if users indeed reacted to the information about peers, their reaction should be stronger if the signal they obtained was more informative. To proxy for informativeness, we exploit the fact that, to make peer group averages well behaved and meaningful, *Status* imposes a minimal number of 5,000 underline observations to construct the peer groups. Based on this rule, if a peer group tightly tailored to all the users' characteristics with at least 5,000 underlying observations does not exist, the app compares the user to a broader set of peers by enlarging the range of values around the user's characteristics.

For instance, suppose that two users have the same characteristics under all dimensions except their location, which might be Manhattan for user A and Helena, MT for user B. Whereas the app would easily find 5,000 observations in Manhattan with the same characteristics of user A, it is likely to miss the same number for Helena, MT. For user B, then, the app would enlarge the geographic area of peer comparison to the overall state of Montana, and if not enough observations were still available, to the overall United States. Crucially, users know the characteristics of the peer group, and hence user A knows that she is compared to similar peers living in her same location, whereas user B knows that she is compared to peer that are similar in most respect except for the location in which they live. Intuitively, then, if users attach any information value to the peer signal, user A should believe that the signal she gets is quite informative, whereas user B might think that the signal she gets is not as informative about the actual amount of spending of her peers.

Based on this intuition, we split our regression-discontinuity sample into users whose characteristics are all matched to the closest possible ranges and users for which at least one characteristic is matched to a broader range of peers, and hence for whom the signal should be less informative. We then estimate the two-stage least-square specification in equations (1) and (2) above separately among these two subsamples. The results are reported in the left columns of Table 5. Consistent with our conjecture, users for whom the peer spending signal is more precise drive the baseline spending reaction. In fact, the difference in the change in spending after signup between those above and below the thresholds is 5 times larger for users that observe precise peer groups, and the difference is not statistically different from zero for other users.

A second dimension under which the signal about peer spending might be more or less informative is the extent to which users could have gathered information about their peers before signing up to *Status*. To obtain variation, we use the type of residential location as a proxy for the density of information. Specifically, we conjecture that users who live in rural areas might have a lower number of peers, i.e. people that look similar to them under most characteristics, to observe on a regular basis and hence from whom they could infer spending information before signup. Instead, users who reside in urban location are exposed to more peers in their daily lives and hence to a higher density of information about peers before signup. If the signal is informative, it should thus affect the behavior of rural users more than the behavior of urban users, which is what we document in the right columns of Table 5.[9]

[9]Note that the urban and rural subsamples do not sum up to the full regression-discontinuity sample. This discrepancy arises because *Status* allows for a third location category labeled "suburban". This category includes locations that are at times highly urbanized, such as the suburbs of large metropolises, or quite rural. For this reason, we just exclude this group from the split analysis. When we looked at this group separately, we find an effect of similar size as the effect for the urban subsample.

6.2 Liquidity and Income

We also assess users' reaction based on their liquidity and income. We report the results of this analysis in Table 6.

In the leftmost columns of Table 6, we consider separately renters and homeowners. The idea is that homeowners might have substantial mortgage balances outstanding and hence hold more illiquid wealth relative to renters and might be less able to adjust their spending due to a high share of nondiscretionary consumption in the form of mortgage payments and property taxes (for a debate on the extent to which liquidity relates to hand-to-mouth behavior, see Kaplan, Violante, and Weidner (2014) and Olafsson and Pagel (2018a)). We might thus observe that renters are more likely to react to the provision of information to peers relative to constrained homeowners. Consistent with this intuition, we find that the effect of information on peer spending is economically large and statistically different from zero for renters, whereas it is not different from zero either economically or statistically for homeowners.

In the middle columns, we propose another proxy for the extent to which households might be indebted and hence might face a large share of non-discretionary consumption in their bundles—credit scores. Intuitively, users with higher credit scores are more likely to hold debt than users with lower credit scores, who might be shut off from borrowing. We thus compare users in the bottom two quintiles and the top two quintiles by credit score. We find that the effect is substantial for low-credit-score users but not for others.

Although the distribution of users by credit scores is not the same as the distribution by income, the split we proposed above might in part capture systematic differences among low-income users and high-income users. We therefore also compare users in the lowest two income thresholds of the RDD—$35K and $50K—to users in the highest two thresholds—$100K and $150K. In the rightmost columns, we find that the reaction to information about peer spending is dramatic for low-income users,[10] whereas we detect barely any reaction for high-income users. In fact, income appears to be the demographic characteristic that captures

[10] Note that because signing up to *Status* requires that users have bank accounts, the lowest possible levels of income in our sample are higher those of low-income households in the broader US population.

the largest difference in users' reaction to information about peers' spending.

6.3 Distance from Peers and Change in Spending

So far, we have focused on the direction and size of the reaction of users based on whether they overspend or underspend relative to peers. But another raw-data fact we documented above is that the distance of a user from the average peer spending seems monotonically related to her reaction. We thus also want to assess whether this feature survives in a multivariate setting.

By construction, we cannot assess the heterogeneity by distance using the RDD framework, because we constructed the sample such that users are all close to the peer-group high and low thresholds. Hence, all users in the RDD are the farthest possible from the average of the peer group to which they are matched, and all users on the same of each threshold have the same distance from the peers.

To assess the role of distance from peers in users' reaction, we therefore go back to consider the main sample in the baseline analysis and we estimate the following set of linear specifications by OLS:

$$\Delta Spending_i = \beta_0 + \beta_1 \; Distance \; from \; Peers_i + \boldsymbol{\gamma}' \boldsymbol{x}_i + \epsilon_i, \tag{3}$$

where we standardize the distance to peers so that the β_1 coefficient can be interpreted as the association between a standard-deviation increase in $Distance \; from \; Peers_i$ and the change in spending after users signup to *Status*. The vector of controls $\boldsymbol{x}_i$ contains *asset balance*, *income*, *home ownership*, *credit score*, *age*, *age-squared*, and *debt balance.*

We estimate this specification separately for users above and below the spending of their peer group, if the distance from peers' spending relate to the spending reaction of users, we should observe that $\hat{\beta_1} < 0$, irrespective of whether we consider users that overspend or underspend relative to their peers. This prediction arises because large overspenders would cut their spending more than overspenders closer to the peers, whereas large underspenders (for whom the distance from peers is more negative) would increase their spending by more than underspenders closer to the peers, whose distance is higher as it is less negative.

The results in Table 7 align with this prediction. Indeed, the distance to peers' spending is monotonically related to users' change in spending in both directions. Subscribers far away from the average spending of their peer group are the ones that change their spending by more relative to other users, irrespective of the direction.

The coefficients on the distance from peers' spending confirms the asymmetry in the extent of users' reaction based on whether they overspend or underspend. Whereas a one standard deviation higher distance among overspenders relates to a drop in spending over income by about 12 percentage points, the same one standard deviation higher distance only relates to a lower increase in spending by about 1.2 percentage points among underspenders.

In the RDD heterogeneity analysis, we find that income is one of the main moderators of users' reaction to information about peer spending, that is, low-income users react substantially, whereas high-income users barely react at all. We thus also consider the heterogeneity of the effect of distance from peer spending based on users' income.

In Table 8, we repeat the analysis of equation (3) adding a full set of interactions with dummies capturing the quartile of users' income. We consider quartiles in this case, because in the full sample we do not have income thresholds as in the RDD sample. In Table 8, the coefficient attached to $Distance$ captures the size of the spending reaction by users who belong to the top quartile of the sample based on income and who lie one standard deviation away from other users. Summing up this baseline coefficient to the one estimated separately for each other income group provides the estimated reaction of users in those group. Indeed, the effect of distance on users' reaction is substantially larger for lower income users relative to other users, and it declines monotonically with income, even though we reject the null that this effect is zero for any individual income group.

We also depict this effect graphically in Figure 7, in which, for each income quartile of our main sample, we depict the baseline reaction to spending based on users' position relative to the spending of their peers. The four panels of Figure 7 repeat the plots in Figure 5 for each income group separately. We use the same y-axis scale for all four plots, so that the difference in the change in spending over income can be readily compared across income groups. Consistent with the multivariate results of Table 8, we find that the reaction for both users above and below the peers is stronger for low-income users, and the relationship between the distance from the peers and the size of the reaction also increases as users' income decreases.

6.4 Reaction by Spending Categories

The results computed so far are estimated using clients' total spending. We now exploit the richness of the categorization of transactions into spending categories we observe in the data.

As a first pass, we categorize spending into discretionary and non-discretionary spending, as described in section 3. Intuitively, we would expect that most of the users' reaction in terms of change in spending involves discretionary spending, because users can barely reduce non-discretionary spending and might have no reason to increase it.

We re-estimate the baseline results separately for the two types of consumption. The results, reported in Figure 8, suggest that, as conjectured, the vast majority of spending changes are related to changes in discretionary spending. As shown in subfigure (a), overspending users cut their discretionary spending substantially more than underspenders. Subfigure (b) shows instead that individuals barely react in terms of non-discretionary spending. The regression line is flat both above and below zero, indicating investors do not adjust their non-discretionary consumption.

Although many of the individual spending categories do not display much of a reaction – some categories are noisy – at least two categories display intriguing results. The first is checking-account withdrawals. As shown in subfigure (a) of Figure 9, checking withdrawals respond dramatically to information about peer spending in both directions. This phenomenon might occur for a number of reasons. Cash is mainly used for incidental expenses (Bagnall, Bounie, Huynh, Kosse, Schmidt, Schuh, and Stix, 2014) and for transactions consumers want to keep anonymous (Acquisti, Taylor, and Wagman, 2016). The latter group might include both legal and illegal entertainment expenses. One interpretation of this result might be that individuals limit their spending on vices once they discover they overspend relative to their peers, although the data at hand do not allow us to ultimately pin down how users employed the cash they withdrew before signing up to *Status*.

The second spending category we consider is the amount spent to service loans and credit-card debt, reported in subfigure (b) of Figure 9. Individuals seem more reluctant to take out loans and might reduce their borrowing through credit cards when they find they are overspending relative to their peers.

7 Potential Mechanisms: Bayesian Learning, Peer Pressure, Overreacting to Negative News

In this section, we discuss the economic channels that might help explain the set of facts we have documented so far. As we discussed in the introduction, three non-mutually-exclusive channels might contribute to the results.

The first is a neoclassical channel—Bayesian updating. Users might believe crowdsourced information about peers' spending contains valuable information regarding the optimal spending rate and might update their beliefs accordingly. Even if any individual peer might not be optimizing their spending based on the users' own characteristics, users might think the average spending of a large group of peers provides a valuable signal. We label this channel *wisdom of the crowds* (e.g., see Da and Huang, forthcoming). This channel does not involve any non-standard assumptions about users' preferences or beliefs and could explain both the convergence of users' spending to peers' spending, as well as the monotonic relationship between the distance of users from their peers and the size of the reaction – convergence requires a stronger reaction the further away users' spending is from the spending of their peers.

At the same time, the *wisdom-of-the-crowds* channel can barely explain the asymmetry of the reaction based on whether users overspend or underspend relative to their peers. Under the *wisdom-of-the-crowds* interpretation, the reactions of users should be similar in absolute value and symmetric with respect to the kink – the point of zero distance from peers' consumption – whereas we observe a substantially stronger reaction by users who overspend relative to users who underspend. Thus, the *wisdom-of-the-crowds* channel cannot fully explain all our results.

One could consider a non-Bayesian alternative of this channel, which we label *conformism*. Under conformism, individuals obtain utility from mixing with the crowd and reducing their idiosyncrasies relative to their peers. But in this case, to explain the asymmetric reaction around the kink, we would still need to assume that conforming to peers from a worse starting point looms larger to individuals than conforming to peers from a better starting point.

The second channel we consider is *peer pressure*. By *peer pressure*, we mean that individuals dislike to perform worse than their peers. In the context of spending, if users were told they overspent relative to peers, they might want to amend this behavior and cut their spending, because they obtain utility from perceiving their financial health is not worse than that of their peers. The version of *peer pressure* we propose can help explain the strong reaction by users who overspend relative to peers, but is unlikely to explain the convergence of underspenders to their peers' level of spending. Underspenders perform better than their peers in terms of financial health, and hence if *peer pressure* were the only channel at play, they would not change their behavior after they sign up to *Status*.

We do not consider the predictions of a keeping-up-with-the-Joneses' framework, because in our setting users do not know the individuals that constitute their peer groups and hence they cannot obtain utility from showing higher consumption than them. And, in any case this setting would be unable to explain why overspenders react at all, because overspenders would already be "beyond the Joneses" in terms of spending.

The third channel we consider is *overreaction to negative news*. This channel is a modification of the *wisdom-of-the-crowds* channel that adds a non-Bayesian assumption regarding individuals' reaction to learning from information to account for our results in both the overspending and underspending domains. *Overreaction to negative news* suggests that individuals learn from the information we provide them, as if peers' spending is a valuable signal, but negative news loom larger for them than equally-sized positive news. This channel predicts that both overspenders and underspenders react to obtaining information about their peers, but overspenders react more than underspenders at the same distance from their peers. In principle, this channel could explain all our baseline facts.

Disentangling the three channels above in field data, which include no randomized exposure to different pieces and types of news, is challenging. We propose a set of tests and arguments to assess the potential role of one or more of the channels and their relative magnitude.

First, recall that Bayesian learning can plausibly explain the reaction of underspenders. We could thus conjecture that the size of the reaction we document in the underspending domain is the effect of Bayesian learning. At the same time, the *wisdom-of-the-crowds* channel predicts a symmetric reaction around the kink for overspenders and underspenders. We could thus use the size of the reaction in the underspending domain to obtain a lower bound for the size of the reaction of overspenders due to non-Bayesian channels. This lower bound is the difference between the size of the reaction we document and the size of the reaction of underspenders. Although we do not have a structural model to interpret the magnitudes of the reactions in our paper, Panel B of Table 2 documents the absolute value of the normalized change in users' monthly spending is three times larger for overspenders than for underspenders. Under our conjecture, this result would suggest non-Bayesian channels might explain most of overspenders' reaction.

We propose a set of direct tests aimed at disentangling the two non-Bayesian channels we propose – *peer pressure* and *overreaction to negative news*. These tests use a feature of *Status* we have not exploited so far. As Figure 1 in the introduction shows, *Status* users observe information not only about their own spending and the spending of peers, but also about the (i) average spending of all US households and (ii) their own average monthly income.

Under the *peer-pressure* channel, we should find that overspending users' reaction in terms of reducing their spending should be most sensitive to the distance of their spending from their peer group. The reaction should be less sensitive to the distance between overspending users and the average US household or users' own average monthly income. This prediction stems from the fact that reacting to overspending with respect to one's own income has nothing to do with comparing oneself with peers. Moreover, the information about peers is explicitly labeled as such, and *Status* is marketed as providing crowdsourced and tailored information about one's own peers based on similar demographic characteristics. Users should thus interpret this piece of information as more representative of peers' spending than the information about the average US household.

Under the *overreaction-to-negative-news* channel, instead, users should react most to the worst piece of news they obtain from *Status*, that is, the information that is furthest away from their spending among peers' spending, average US households, and average income.

Across the four panels of Table 9, we regress overspending users' change in spending on the distance of their pre-signup spending from four different points – peers' spending (Panel A), the average US household's spending (Panel B), users' average income (Panel C), and the maximum distance among these three (Panel D). Across columns, we start with the results for the full sample and exclude alternatively the top decile, quintile, or tercile of the sample to ensure none of our results are driven by outliers or extreme reactions. Across the board and for each subsample, the coefficients attached to the distance between users' spending and peers' spending are systematically larger than any of the other coefficients. In particular, the coefficients on the distance from peers are about three times as large as those on the distance from the average US household and more than 50% larger than those on the average users' income and the maximum distance across any three values.

8 Conclusions

We study the effects of providing individuals with crowdsourced information about similar but unknown individuals' spending (*peers*) through a FinTech app. First, users who spend more than the peers reduce their spending, and users who spend less increase their spending, which suggests that users find the signals they receive informative to their behavior. Moreover, users' reaction is severely asymmetric—overspenders cut spending substantially more than underspenders increase it. We argue that this result is not fully consistent with Bayesian updating, but might be driven by peer pressure or the fact that bad news looms larger than (equally-sized) good news.

In terms of heterogeneity of the reaction, the reaction is larger when the signal about peers is more informative, for more liquid users, for lower-income users, and for users that are farther away from their peers in terms of pre-signup spending. Also, discretionary spending drives the reaction in both directions and especially cash withdrawals, commonly used for incidental expenses and transactions for which individuals want to maintain anonymity. Users thus are more likely to cut personal expenses than expenses that might benefit their whole household, such as travel expenses or expenses to consume food and drink at restaurants.

An interesting questions for future research is whether providing information about peer spending has permanent or transitory effects, which would have opposite implications in terms of promoting policies that disseminate information about others' spending behavior. Moreover, our setting does not allow testing whether reacting to peers' spending is optimal or not for users. To the extent that cash withdrawals are mainly capturing incidental expenses or anonymous transactions for illicit goods and services, the results seem to suggest that this socially undesirable type of spending categories react the most to peer information.

References

Acquisti, A., C. Taylor, and L. Wagman, 2016, "The economics of privacy," *Journal of Economic Literature*, 54, 442–92.

Bagnall, J., D. Bounie, K. Huynh, A. Kosse, T. Schmidt, S. Schuh, and H. Stix, 2014, "Consumer cash usage: a cross-country comparison with payment diary survey data," .

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel, 2018, "The economic effects of social networks: Evidence from the housing market," *Journal of Political Economy*, 126, 2224–2276.

Banks, J., R. Blundell, and S. Tanner, 1998, "Is there a retirement-savings puzzle?," *American Economic Review*, pp. 769–788.

Bernheim, B. D., J. Skinner, and S. Weinberg, 2001, "What accounts for the variation in retirement wealth among US households?," *American Economic Review*, 91, 832–857.

Beshears, J., J. J. Choi, D. Laibson, B. C. Madrian, and K. L. Milkman, 2015, "The effect of providing peer information on retirement savings decisions," *The Journal of finance*, 70, 1161–1201.

Bursztyn, L., F. Ederer, B. Ferman, and N. Yuchtman, 2014, "Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions," *Econometrica*, 82, 1273–1301.

Carroll, G. D., J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick, 2009, "Optimal defaults and active decisions," *The quarterly journal of economics*, 124, 1639–1674.

Chalmers, J., W. T. Johnson, and J. Reuter, 2014, "The effect of pension design on employer costs and employee retirement choices: Evidence from Oregon," *Journal of Public Economics*, 116, 17–34.

Chalmers, J., and J. Reuter, 2012, "How do retirees value life annuities? Evidence from public employees," *The Review of Financial Studies*, 25, 2601–2634.

Charles, K. K., E. Hurst, and N. Roussanov, 2009, "Conspicuous consumption and race," *The Quarterly Journal of Economics*, 124, 425–467.

Da, Z., and X. Huang, forthcoming, "Harnessing the wisdom of crowds," *Management Science*.

D'Acunto, F., D. Hoang, M. Paloviita, and M. Weber, 2018, "IQ, Expectations, and Choice," *Working Paper*.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber, 2018, "Shopping, Inflation Expectations, and Household Behavior," *Working Paper*.

D'Acunto, F., N. Prabhala, and A. G. Rossi, forthcoming, "The promises and pitfalls of robo-advising," *The Review of Financial Studies*.

Duflo, E., and E. Saez, 2003, "The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment," *The Quarterly journal of economics*, 118, 815–842.

Galton, F., 1907, "Vox populi (The wisdom of crowds)," *Nature*, 75, 450–451.

Gargano, A., and A. G. Rossi, 2018, "Does it pay to pay attention?," *The Review of Financial Studies*, 31, 4595–4649.

Gargano, A., A. G. Rossi, and R. Wermers, 2017, "The freedom of information act and the race toward information acquisition," *The Review of Financial Studies*, 30, 2179–2228.

Han, B., D. Hirshleifer, and J. Walden, 2019, "Visibility Bias in the Transmission of Consumption Beliefs and Undersaving," Discussion paper, National Bureau of Economic Research.

Hansen, B. E., 1996, "Inference when a nuisance parameter is not identified under the null hypothesis," *Econometrica: Journal of the econometric society*, pp. 413–430.

Hansen, B. E., 2000, "Sample splitting and threshold estimation," *Econometrica*, 68, 575–603.

Hansen, B. E., 2017, "Regression kink with an unknown threshold," *Journal of Business & Economic Statistics*, 35, 228–240.

Jappelli, T., and M. Pagano, 1994, "Saving, growth, and liquidity constraints," *The Quarterly Journal of Economics*, 109, 83–109.

Kaplan, G., G. L. Violante, and J. Weidner, 2014, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014, 77–138.

Laibson, D., 1997, "Golden Eggs and Hyperbolic Discounting*," *The Quarterly Journal of Economics*, 112, 443–478.

Lieber, E., and W. Skimmyhorn, 2018, "Peer effects in financial decision-making," *Journal of Public Economics*, 163, 37–59.

Lusardi, A., and O. S. Mitchell, 2007, "Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth," *Journal of monetary Economics*, 54, 205–224.

Lusardi, A., and O. S. Mitchell, 2014, "The economic importance of financial literacy: Theory and evidence," *Journal of economic literature*, 52, 5–44.

Lusardi, A., and O. S. Mitchell, 2017, "How ordinary consumers make complex economic decisions: financial Literacy and retirement readiness," *Quarterly Journal of Finance*, 7, 1750008.

Madrian, B. C., and D. F. Shea, 2001, "The power of suggestion: Inertia in 401 (k) participation and savings behavior," *The Quarterly journal of economics*, 116, 1149–1187.

Maturana, G., and J. Nickerson, 2018, "Teachers Teaching Teachers: The Role of Workplace Peer Effects in Financial Decisions," .

Maturana, G., and J. Nickerson, forthcoming, "Real Effects of Workers' Financial Distress: Evidence from Teacher Spillovers," *Journal of Financial Economics.*

Olafsson, A., and M. Pagel, 2017, "The ostrich in us: Selective attention to financial accounts, income, spending, and liquidity," Discussion paper, National Bureau of Economic Research.

Olafsson, A., and M. Pagel, 2018a, "The liquid hand-to-mouth: Evidence from personal finance management software," *The Review of Financial Studies*, 31, 4398–4446.

Olafsson, A., and M. Pagel, 2018b, "The retirement-consumption puzzle: New evidence from personal finances," Discussion paper, National Bureau of Economic Research.

Ouimet, P., and G. Tate, 2019, "Learning from coworkers: Peer effects on individual investment decisions," *NBER Working Paper*.

Pagel, M., 2017, "Expectations-Based Reference-Dependent Life-Cycle Consumption," *The Review of Economic Studies*, 84, 885–934.

Van Rooij, M. C., A. Lusardi, and R. J. Alessie, 2012, "Financial literacy, retirement planning and household wealth," *The Economic Journal*, 122, 449–478.

Wolfers, J., and E. Zitzewitz, 2004, "Prediction markets," *Journal of economic perspectives*, 18, 107–126.

Zeldes, S. P., 1989, "Consumption and liquidity constraints: an empirical investigation," *Journal of political economy*, 97, 305–346.

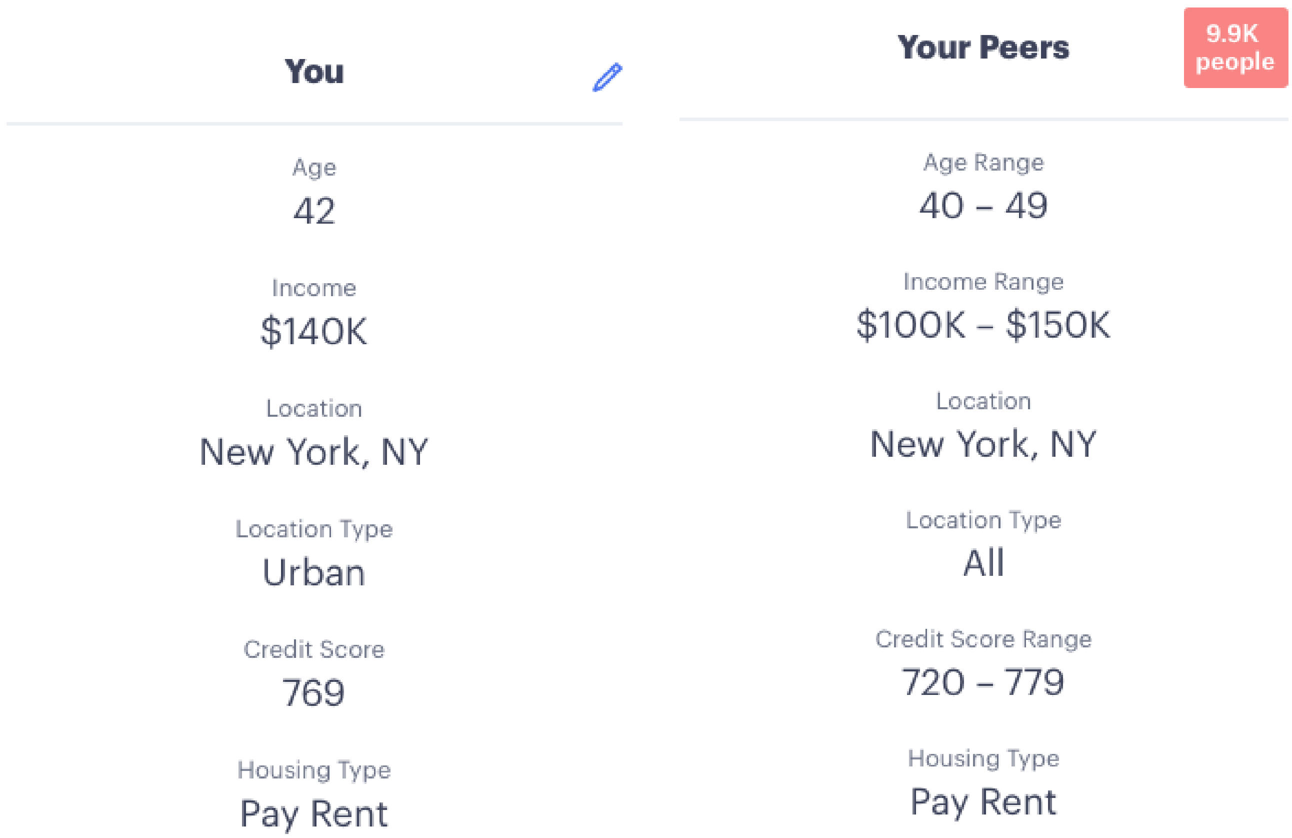

(a) User Profile

(b) Peer Group Information

Figure 2
Peer Group for a Sample Account

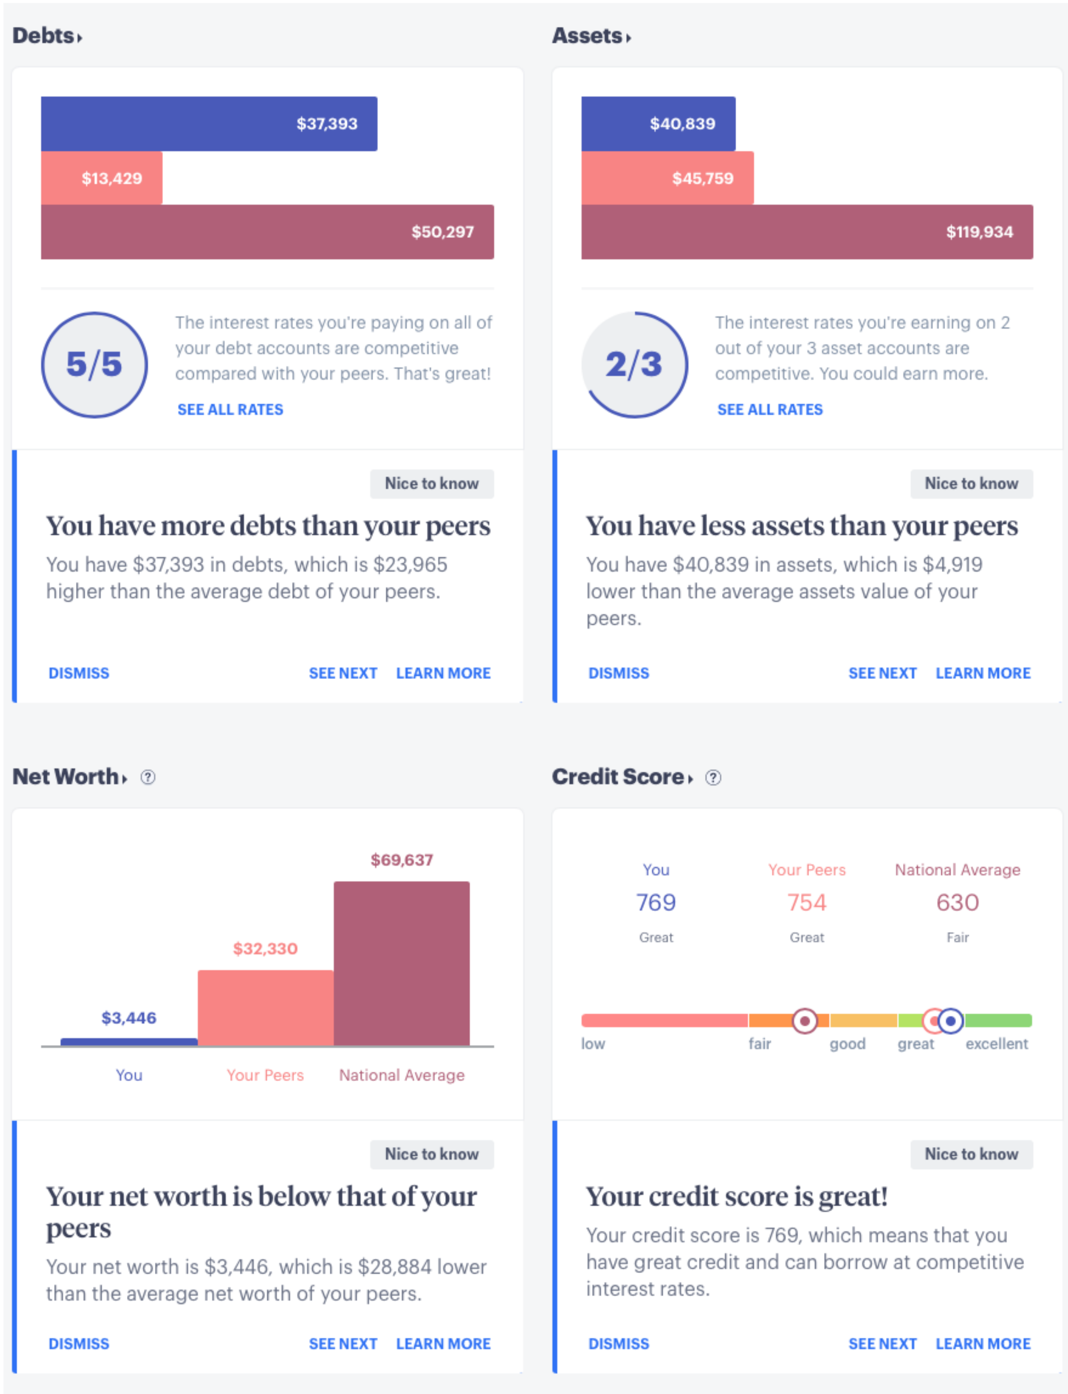

Figure 3
Status Home Page

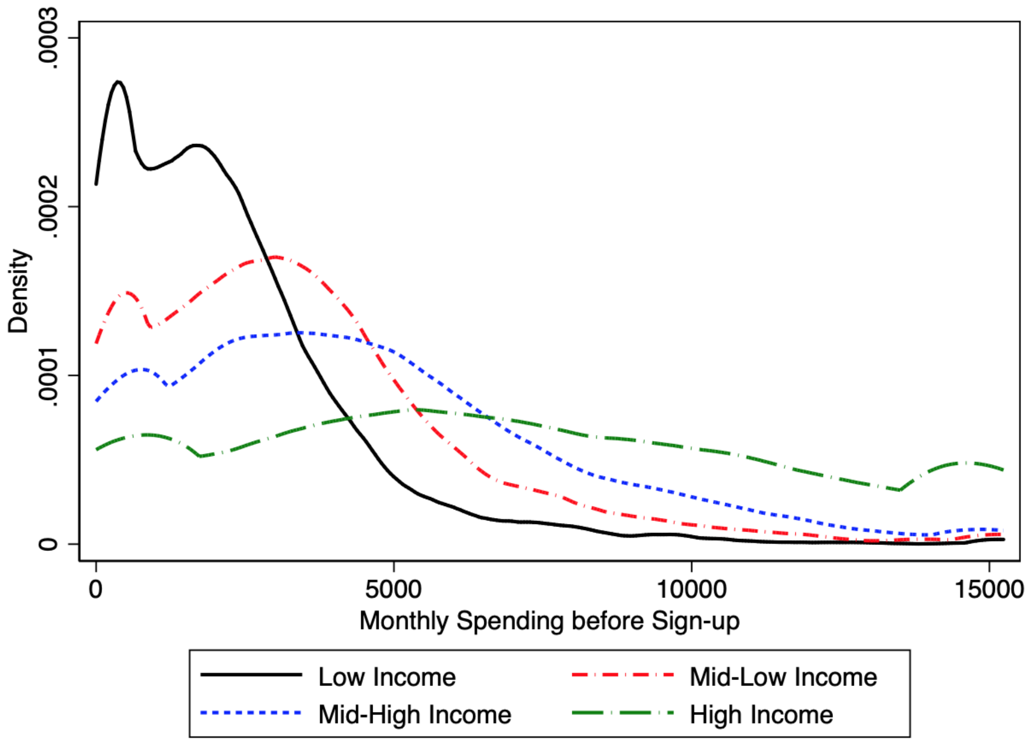

Figure 4
Distribution of Monthly Spending by Income Quartiles

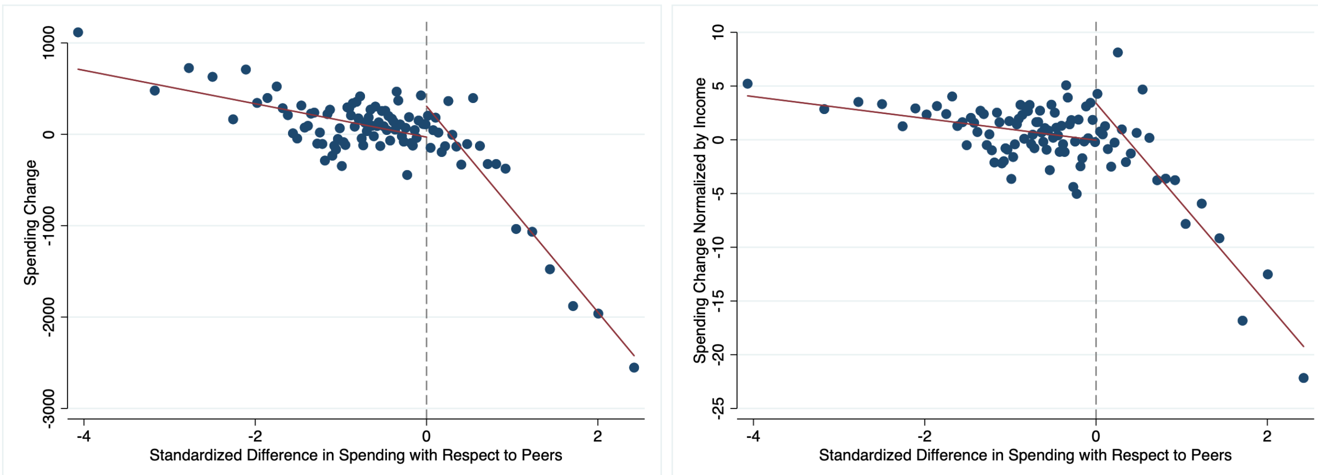

(a) Changes in Spending

(b) Changes in Spending Normalized by Income

Figure 5
Distance from Peers' Spending and Changes in Spending after Signup

This figure shows binned scatterplots of changes in overall consumption after signing up for *Status* and differences in consumption between individuals and their peer group at the time of sign-up. The x-axis measures the difference in consumption with respect to peers, normalized by its standard deviation. The y-axis in subfigure (a) reports results for dollar changes in spending, computed using two months before and after signup. The y-axis in subfigure (b) normalizes the changes in consumption by income. The binned scatterplot divides the 17,500 users in 100 groups. In addition to the scatterplot, we report in red the fitted values of a threshold regression that estimates different linear regression coefficients below and above the zero threshold.

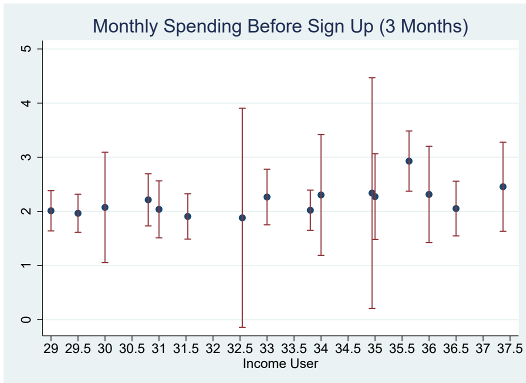

(a) Income Threshold $35,000

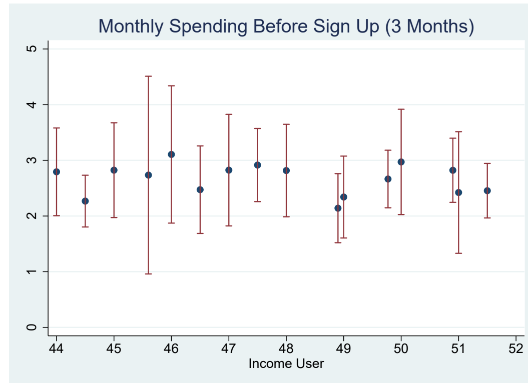

(b) Income Threshold $50,000

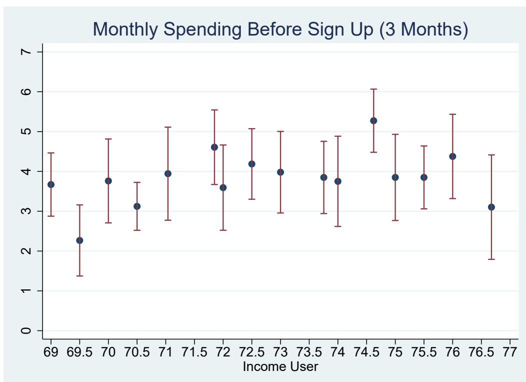

(c) Income Threshold $75,000

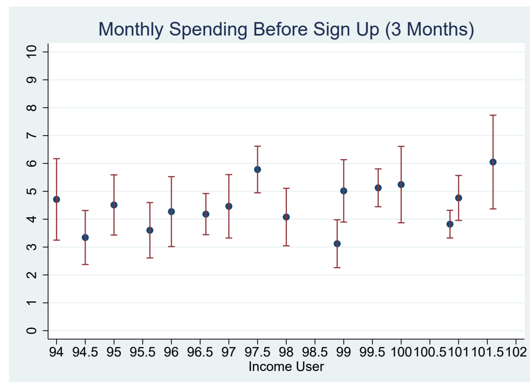

(d) Income Threshold $100,000

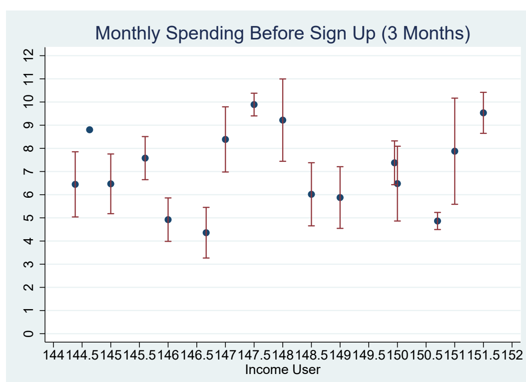

(e) Income Threshold $150,000

Figure 6
Average Monthly Spending in Three Months before Sign-up around Income Thresholds

This figure reports the average monthly spending in the three months before signup to Status by income around the different income thresholds we use in our identification sample. Bar around each point estimates indicate a third of a standard deviation in either direction.

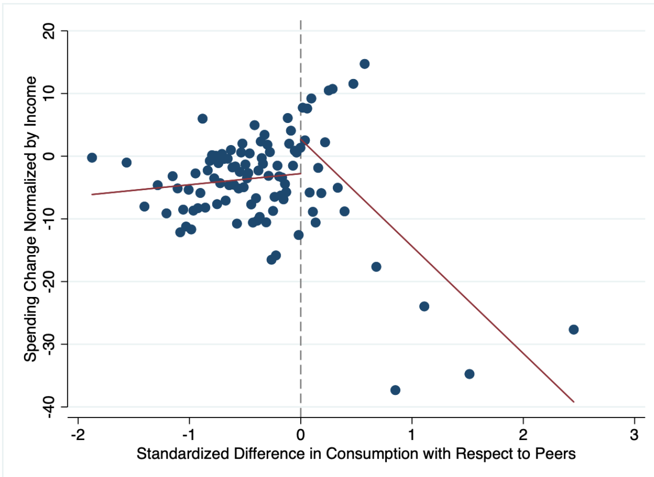

(a) First Income Quartile

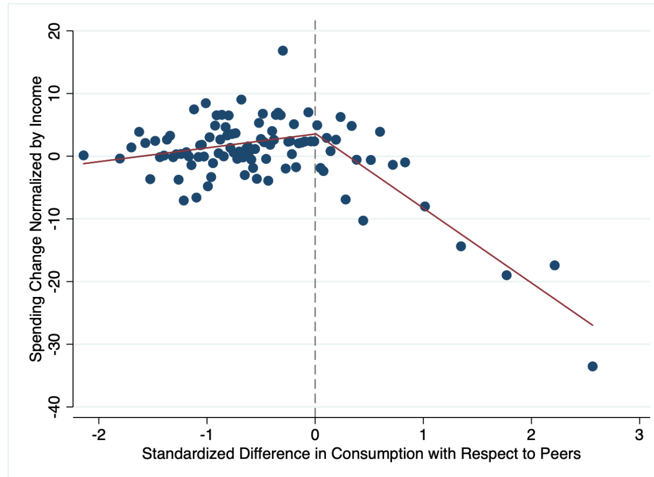

(b) Second Income Quartile

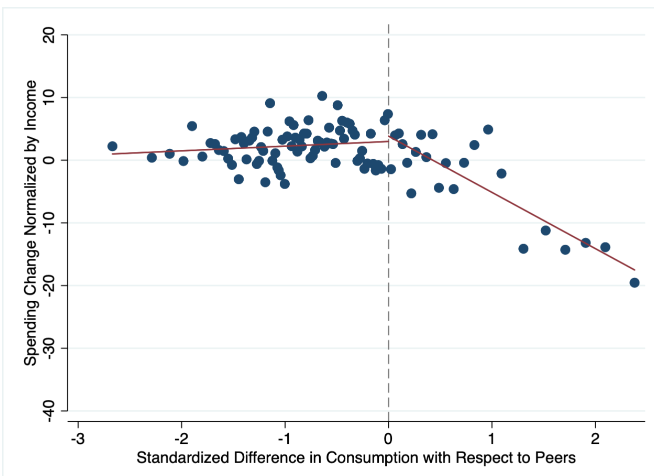

(c) Third Income Quartile

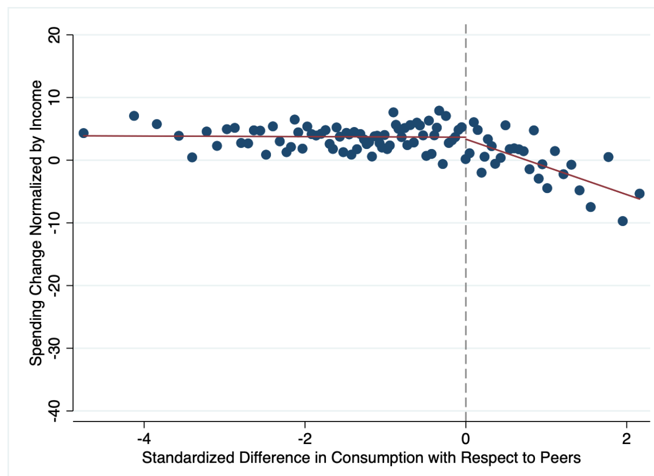

(d) Fourth Income Quartile

Figure 7
Distance from Peers' Spending and Changes in Spending after Signup—by Income Quartiles

This figure shows binned scatterplot of changes in overall consumption after signing up for *Status* and differences in consumption between individuals and their peer group. In all subfigures, the x-axis measures the difference in consumption with respect to peers, normalized by its standard deviation. The y-axis reports results for dollar changes in spending normalized by income, computed using two months before and after sign-up. Each subfigure reports the results for an income quartile and the binned scatterplot divides the users in each income quartile in 100 groups. In addition to the scatterplot, we report in red the fitted values of a threshold regression that estimates different linear regression coefficients below and above the zero threshold.

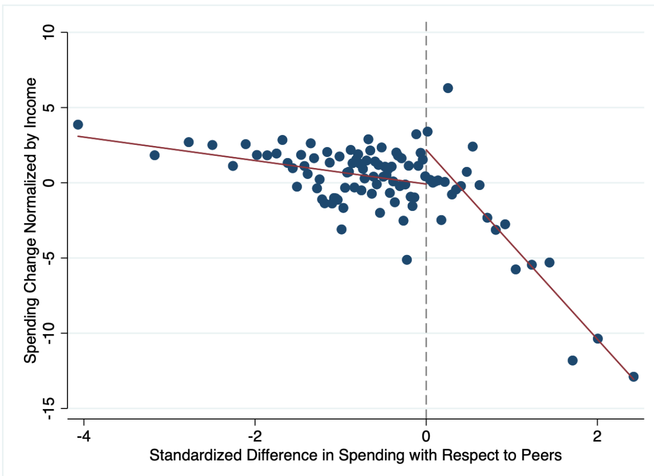

(a) Discretionary Spending

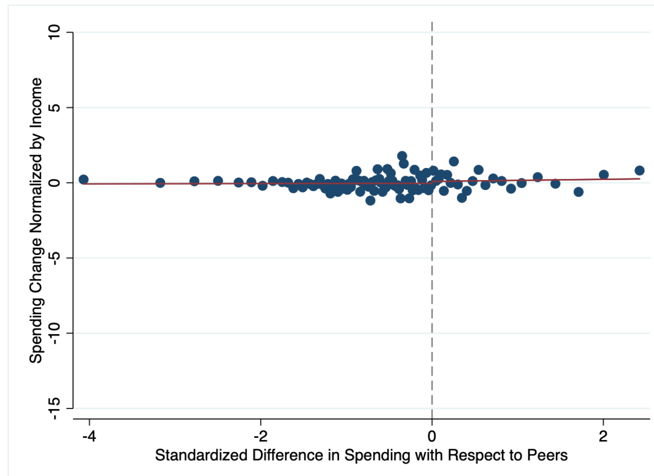

(b) Non-Discretionary Spending

Figure 8
Distance from Peers' Spending and Changes in Spending after Signup—Discretionary and Non-discretionary Spending

This figure shows binned scatterplots of changes in discretionary and non-discretionary consumption after signing up for *Status* and differences in consumption between individuals and their peer group. In all subfigures, the x-axis measures the difference in consumption with respect to peers, normalized by its standard deviation. The y-axis reports results for dollar changes in spending normalized by income, computed using two months before and after signup. Subfigure (a) reports the results for discretionary consumption. Subfigure (b) reports the results for non-discretionary consumption. Each binned scatterplot divides the 17,500 users into 100 groups. In addition to the scatterplot, we report in red the fitted values of a threshold regression that estimates different linear regression coefficients below and above the zero threshold.

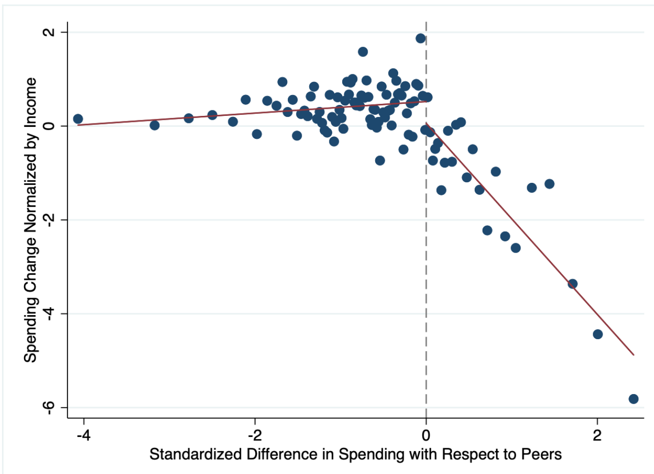

(a) Checking Account Withdrawals

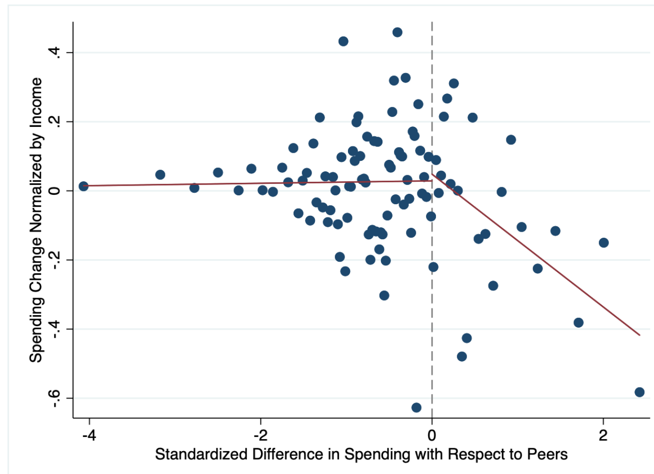

(b) Consumer Loan Fees and Interest (incl. Credit Cards)

Figure 9
Distance from Peers' Spending and Changes in Spending after Signup—Withdrawal from Checking Accounts and Consumer Loan Fees

This figure shows binned scatterplots of changes in checking-account withdrawals and changes in the loans taken out by users after signing up for *Status* and differences in consumption between individuals and their peer group. In all subfigures, the x-axis measures the difference in consumption with respect to peers, normalized by its standard deviation. The y-axis reports results for dollar changes in spending normalized by income, computed using two months before and after sign-up. Subfigure (a) reports the results for checking-account withdrawals. Subfigure (b) reports the results for fees and interest paid on consumer loans the users takes, which includes credit-card debt. Each binned scatterplot divides the 17,500 users in 100 groups. In addition to the scatterplot, we report in red the fitted values of a threshold regression that estimates different linear regression coefficients below and above the zero threshold.

Table 1. Summary Statistics

	Observations	Mean	St. Dev.
Age	17,673	30	7
Credit Score	16,335	728	84
Home Ownership	17,676	0.38	0.49
Annual Income ($)	17,598	90,055	61,796
Assets ($)	15,325	42,462	68,066
Debts ($)	12,332	29,971	64,637
Monthly Spending (Total) ($)	17,676	4,334	4,073
Monthly Spending (Discretionary) ($)	17,676	2,772	2,906
Monthly Spending (Non-Discretionary) ($)	17,676	680	679
Monthly Spending (Other) ($)	17,676	882	1,475

This table reports summary statistics of the main variables used in the paper. For each variable, we report the number of observations, the average, and the standard deviation.

Table 2. Spending Changes after Signing up for *Status*

Panel A. Dollar-Value Changes in Spending

	Below Peers		Above Peers	
	Coeff.	t-stat	Coeff.	t-stat
ΔSpending	142.24***	(5.84)	-474.01***	(-7.81)
Observations	13,596		4,080	

Panel B. Spending Changes Scaled by Income

	Below Peers		Above Peers	
	Coeff.	t-stat	Coeff.	t-stat
ΔSpending	0.924***	(4.27)	-3.079***	(-5.25)
Observations	13,596		4,080	

This table presents results for changes in spending after signing up for *Status*. Panel A reports results for dollar changes in spending, and Panel B scales the changes in spending by income. Within each panel, changes in spending are computed for clients with below-peer spending in columns 1 and 2, and for clients with spending above peers in columns 3 and 4. Spending changes are computed using two months before and after sign-up. To account for cyclicality in monthly spending, we deduct from the change in spending of each client the average change in spending across all the clients who signup in the same month.

Table 3. Change in Spending after Signup in the Regression Discontinuity Sample

	Baseline	Controls	Below Peers	Above Peers
Peer Spending	0.0333***	0.0654***	0.0810***	0.0513***
	(2.71)	(5.15)	(3.49)	(2.98)
Spending before	-0.261***	-0.297***	-0.334***	-0.257***
	(-21.19)	(-21.70)	(-17.89)	(-12.68)
Homeowner		0.105***	0.141***	0.0531
		(3.72)	(3.98)	(1.14)
log of Credit Score		-0.174***	-0.244***	0.0452
		(-2.76)	(-3.50)	(0.34)
log of Age		0.195***	0.206**	0.199*
		(2.98)	(2.47)	(1.88)
log of Asset balance		0.0119**	0.0257***	-0.011
		(2.00)	(3.48)	(-1.09)
log of Debt balance		0.0079	-0.0039	0.0232**
		(1.35)	(-0.52)	(2.50)
Observations	6,745	4,761	1,896	2,865

This table reports results for the sensitivity of the ratio of spending post signup to the ratio of spending pre signup to peer consumption, spending pre sign-up as well as controls. Spending is computed using three months before and after signup. Regression estimates are computed for all users in columns 1 and 2, users with below-peer spending in column 3 and and for users with above-peer spending in column 4.

Table 4. Regression Discontinuity Design

	First Stage	Second Stage	First Stage	Second Stage
Spending before	0.242***	-0.275***	0.294***	-0.301***
	(21.68)	(-17.11)	(20.76)	(-16.07)
Above dummy	0.641***		0.619***	
	(28.02)		(22.97)	
Peer spending		0.0823**		0.0785*
		(2.160)		(1.955)
Homeowner			-0.198***	0.108***
			-6.494)	(3.692)
log of Credit Score			0.0561	-0.175***
			(0.821)	(-2.770)
log of Age			-0.0479	0.195***
			(-0.675)	(2.975)
log of Asset balance			0.0396***	0.0113*
			(6.168)	(1.825)
log of Debt balance			0.0240***	0.00761
			(3.791)	(1.291)
Observations	6,745	6,745	4,761	4,761
IV F-stat		785.2		527.8

This table reports results for a two-stage-least-squares identification strategy that compares users just below and users at each of the income thresholds that *Status* Money uses to define peer groups, that is, \$35K, \$50K, \$65K, \$75K, \$100K, and \$150K. For each threshold, we only keep users who are at the lower threshold of the group as well as users who are in the lower-income group, but are within \$3K of the threshold. Taking the \$100K threshold as an example, we only keep users who declare \$100K in annual income as well as users with an income between \$97K and \$100K. We then estimate the following first-stage specification:

$$Peer\ Spending_i = \alpha + \beta\ Dummy\ Above_i + \epsilon_i,$$

where $Peer\ Spending_i$ is the peer-spending value for user i and $Dummy\ Above_i$ is a dummy variable for whether the income is exactly equal to the threshold value. In the second stage, we use the instrumented $Peer\ Spending_i$ of the first stage as the main covariate in the following specification:

$$\frac{num\ Spending_{i,post}}{Spending_{i,pre}} = \alpha + \beta\ \widehat{Peer\_Spending\_i} + \epsilon_i,$$

where $\frac{num\ Spending_{i,post}}{Spending_{i,pre}}$ is the ratio of post and pre consumption in the three months around signup.

Table 5. Regression Discontinuity Design: Heterogeneity by Informativeness Peer Groups

	Peer Group		Information Density	
	Precise	Imprecise	Rural	Urban
Peer spending	0.153***	0.0306	0.266**	0.0598
	(2.61)	(0.63)	(2.13)	(1.41)
Spending before	-0.277***	-0.274***	-0.326***	-0.232***
	(-10.98)	(-13.09)	(-8.84)	(-11.87)
Controls	X	X	X	X
Observations	2,588	4,157	1,278	3,056

This table reports results for a two-stage-least-squares identification strategy that compares users just below and users at each of the income thresholds that *Status* Money uses to define peer groups, that is, \$35K, \$50K, \$65K, \$75K, \$100K, and \$150K. For each threshold, we only keep users who are at the lower threshold of the group as well as users who are in the lower-income group, but are within \$3K of the threshold. Taking the \$100K threshold as an example, we only keep users who declare \$100K in annual income as well as users with an income between \$97K and \$100K. We then estimate the following first-stage specification:

$$Peer\ Spending_i = \alpha + \beta\ Dummy\ Above_i + \epsilon_i,$$

where $Peer\ Spending_i$ is the peer-spending value for user i and $Dummy\ Above_i$ is a dummy variable for whether the income is exactly equal to the threshold value. In the second stage, we use the instrumented $Peer\ Spending_i$ of the first stage as the main covariate in the following specification:

$$\frac{num\ Spending_{i,post}}{Spending_{i,pre}} = \alpha + \beta\ \widehat{Peer\_Spending\_i} + \epsilon_i,$$

where $\frac{num\ Spending_{i,post}}{Spending_{i,pre}}$ is the ratio of post and pre consumption in the three months around signup and directly report the second stage estimate for different sample splits.

Table 6. Regression Discontinuity Design: Heterogeneity by Liquidity and Income

	Homeownership		Credit Score		Income	
	No	Yes	Low	High	Low	High
Peer spending	0.118**	-0.0108	0.165*	-0.00155	0.309*	0.0158
	(2.48)	(-0.17)	(1.65)	(-0.03)	(1.71)	(0.65)
Spending before	-0.353***	-0.206***	-0.360***	-0.207***	-0.418***	-0.158***
	(-14.45)	(-9.85)	(-7.85)	(-9.16)	(-13.60)	(-12.98)
Controls	X	X	X	X	X	X
Observations	4,303	2,442	2,065	2,052	2,980	1,940

This table reports results for a two-stage-least-squares identification strategy that compares users just below and users at each of the income thresholds that *Status* Money uses to define peer groups, that is, \$35K, \$50K, \$65K, \$75K, \$100K, and \$150K. For each threshold, we only keep users who are at the lower threshold of the group as well as users who are in the lower-income group, but are within \$3K of the threshold. Taking the \$100K threshold as an example, we only keep users who declare \$100K in annual income as well as users with an income between \$97K and \$100K. We then estimate the following first-stage specification:

$$Peer\ Spending_i = \alpha + \beta\ Dummy\ Above_i + \epsilon_i,$$

where $Peer\ Spending_i$ is the peer-spending value for user i and $Dummy\ Above_i$ is a dummy variable for whether the income is exactly equal to the threshold value. In the second stage, we use the instrumented $Peer\ Spending_i$ of the first stage as the main covariate in the following specification:

$$\frac{num\ Spending_{i,post}}{Spending_{i,pre}} = \alpha + \beta\ \widehat{Peer\_Spending\_i} + \epsilon_i,$$

where $\frac{num\ Spending_{i,post}}{Spending_{i,pre}}$ is the ratio of post and pre consumption in the three months around signup and directly report the second stage estimate for different sample splits.

Table 7. Distance from Peers' Spending and Spending Changes after Signup

	All		Below Peers		Above Peers	
	Value	t-stat	Value	t-stat	Value	t-stat
Distance	-4.24***	(-10.85)	-1.20**	(-2.14)	-11.95***	(-5.87)
Asset Balance	0.97***	(3.39)	0.87***	(2.73)	2.38**	(2.42)
Income	-0.40	(-0.40)	-0.77	(-0.63)	3.00	(1.21)
Home Ownership	1.95*	(1.94)	2.63**	(2.15)	-0.20	(-0.07)
Credit Score	0.00	(0.29)	-0.00	(-0.17)	-0.01	(-0.28)
Age	0.04	(0.09)	-0.45	(-1.02)	1.19	(0.95)
Age^2	-0.00	(-0.04)	0.01	(1.13)	-0.01	(-0.88)
Debt Balance	0.42**	(2.46)	0.42**	(2.30)	0.25	(0.46)
Constant	-11.52	(-0.88)	6.79	(0.50)	-70.59*	(-1.82)
Observations	9,597		6,826		2,771	

This table reports results for the sensitivity of spending changes to peer consumption. We estimate the following simple linear regression by ordinary least squares:

$$\Delta\ Spending_i = \beta_0 + \beta_1\ Distance\ from\ Peers_i + \boldsymbol{\gamma}'\boldsymbol{x}_i + \epsilon_i,$$

where $\Delta\ Spending_i$ is the change in consumption of individual i after signing up for *Status*, $Distance\ from\ Peers_i$ is the difference in consumption between individual i and the average spending of his/her peer group at the time of signup, and $\boldsymbol{x}_i$ is a vector of control variables. Spending changes are computed using two months before and after signup and are scaled by income. To account for cyclicality in monthly spending, we deduct the average change in spending across all users who signup in the same month from the change in each user's spending. $Distance\ from\ Peers_i$ is standardized so that the coefficient estimates represent the relation between spending changes and a standard-deviation increase in $Distance\ from\ Peers_i$. The vector of control variables contains the following: $Asset\ Balance$, the user's total assets at the time of sign-up; $Income$, the users' income; $Home\ Ownership$, an indicator variable for whether the user is a homeowner; $Credit\ Score$, the user's credit score at the time of signup; Age and Age^2, the user's age and squared age; and $Debt\ Balance$, the debt balance at the time of signup. Within each panel, regression estimates are computed for all users in columns 1 and 2, for users with below-peer spending in columns 3 and 4, and for users with above-peer spending in columns 5 and 6.

Table 8: Distance from Peers' Spending and Spending Changes after Signup: Heterogeneity by Income

	All		Below Peers		Above Peers	
	Value	t-stat	Value	t-stat	Value	t-stat
Distance	-2.282***	(-7.19)	-0.884*	(-1.89)	-5.196**	(-2.56)
Distance × Income_1	-6.319**	(-2.20)	2.239	(0.72)	-29.591***	(-2.69)
Distance × Income_2	-4.662***	(-4.25)	0.500	(0.32)	-7.257*	(-1.65)
Distance × Income_3	-2.735***	(-3.99)	-0.094	(-0.10)	-5.379	(-1.64)
Constant	-4.742	(-0.65)	6.186	(0.87)	-22.891	(-0.88)
Other Controls	✓		✓		✓	
Observations	12,256		9,247		3,009	

This table reports results for the sensitivity of spending changes to peer consumption. We estimate the following simple linear regression by ordinary least squares:

$$\Delta\ Spending_i = \beta_0 + \beta_1\ Distance_i + \sum_{j=1}^{3} \delta_j\ Distance_i \times Income_{i,j} + \boldsymbol{\gamma}'\boldsymbol{x}_i + \epsilon_i,$$

where $\Delta\ Spending_i$ is the change in spending of individual i after signing up for *Status*, $Distance_i$ is the difference in spending between individual i and the average spending of his/her peer group at the time of signup, and $\boldsymbol{x}_i$ is a vector of control variables. Spending changes are computed using two months before and after signup and are scaled by income. To account for cyclicality in monthly spending, we deduct the average change in spending across all users who signup in the same month from the change of each user's spending. $Distance_i$ is standardized so that the coefficient estimates represent the relation between spending changes and a standard-deviation increase in $Distance_i$. The vector of control variables contains the following: *Asset Balance*, the user's total asset quartile dummy at the time of signup; *Income*, the income quartile dummy; *Credit Score*, the credit score quartile dummy at the time of signup; *Debt Balance*, the debt balance quartile dummy at the time of sign-up. *Age* and Age^2, the user's age and squared age; and *Home Ownership*, an indicator variable for whether the user is a homeowner. We report the estimated coefficient estimates for $Distance_i$ and the interaction between $Distance_i$ and the quartile dummies of the control variables. In all cases, the base case is the fourth quartile. Within each panel, regression estimates are computed for all users in columns 1 and 2, for users with below-peer spending in columns 3 and 4, and for users with above-peer spending in columns 5 and 6.

Table 9. Interpretation: Peer Pressure and Overreaction to Negative News

	Panel A. Distance from Peer Spending							
	Full Sample		Exc. Top Decile		Exc. Top Quintile		Exc. Top Tercile	
	Value	t-stat	Value	t-stat	Value	t-stat	Value	t-stat
Dist. from Peers	-9.342***	(-10.57)	-8.512***	(-12.48)	-9.343***	(-13.83)	-10.693***	(-15.41)
Observations	4,077		3,669		3,261		2,718	

	Panel B. Distance from Average US Spending							
	Full Sample		Exc. Top Decile		Exc. Top Quintile		Exc. Top Tercile	
	Value	t-stat	Value	t-stat	Value	t-stat	Value	t-stat
Dist. from Avg US	-3.652***	(-5.49)	-0.874	(-1.53)	-0.502	(-0.86)	-3.065***	(-4.78)
Observations	3,652		3,287		2,922		2,435	

	Panel C. Distance from Average Monthly Income							
	Full Sample		Exc. Top Decile		Exc. Top Quintile		Exc. Top Tercile	
	Value	t-stat	Value	t-stat	Value	t-stat	Value	t-stat
Dist. from Avg Income	-6.456***	(-12.32)	-6.415***	(-17.10)	-7.141***	(-19.92)	-8.464***	(-23.61)
Observations	8,086		7,282		6,473		5,388	

	Panel D. Maximum Distance from Peer Spending, US Average Spending, and Average Monthly Income							
	Full Sample		Exc. Top Decile		Exc. Top Quintile		Exc. Top Tercile	
	Value	t-stat	Value	t-stat	Value	t-stat	Value	t-stat
Maximum Distance	-4.478***	(-10.22)	-3.362***	(-10.54)	-3.338***	(-10.76)	-5.587***	(-17.22)
Observations	5,151		4,636		4,121		3,434	

This table reports results on the economic channels driving the effects we document in the paper. Across the four panels, we regress overspending users' change in spending on the distance of their pre-signup spending from four different points: peers' spending (Panel A), the average US household's spending (Panel B), users' average income (Panel C), and the maximum distance between these three (Panel D). Within each panel, we report results for the full sample, as well as results that exclude the top decile, quintile, and tercile of observations.

Online Appendix:

Crowdsourcing Financial Information to Change Spending Behavior

Francesco D'Acunto, Alberto G. Rossi, and Michael Weber

Not for Publication

A.1 Estimating the Kink's Location Non-parametrically

Another way to speak to the potential endogeneity issues in our setting it that we estimate non-parametrically whether a kink in users' change in spending after signup exists, and if yes whether the position of this kink is close to the value of peer spending users observe. If

. But if individuals were not basing their reaction only on the value of peers' consumption that *Status* shows them, the actual threshold might fall at a value different from 0. For instance, because *Status* provides information not only about peers' spending, but also about users' own average monthly income as well as average US consumer spending, users might react to a combination of these pieces of information. Although this possibility would still entail an effect of providing users with information on their spending, in this case, we would not be able to conclude users react to information about peers. Moreover, our results so far do not allow testing whether the regression slope coefficients are statistically different below and above the threshold.

We address these concerns in two ways. In this section, we estimate the location of the threshold non-parametrically using two complementary approaches. In the next section, we estimate the effect of the distance of users' spending from points other than peer spending (average monthly income and average US consumer spending) on users' spending change after signup.

To estimate the location of the threshold non-parametrically, the first approach builds on Hansen (1996, 2000). It estimates a threshold model with unknown threshold. To build intuition, consider the case of one regressor. The threshold regression estimates the optimal threshold for a linear model that has different intercept and slope estimates below and above the threshold. Hansen (1996) also proposes a test for whether the coefficient estimates below and above the threshold are statistically different from each other.

For the second approach, we follow Hansen (2017) and estimate a regression kink model with unknown threshold. This model is similar to the one described above, but does not allow for discontinuities. The approach is thus similar to estimating a linear spline model that has a single endogenously determined node. Hansen (2017) also develops the asymptotic theory to make statistical inference about the threshold.

A.1.1 Threshold Regression Results

We estimate the threshold regression model on the full set of 17,673 observations and report the results in Panel A of Table A.1. The first two columns report the linear regression on the full sample. Columns 2 and 3 (4 and 5) repeat the estimates below (above) the endogenously determined threshold.

The threshold is precisely estimated to be 0.235, with a 95% confidence interval of [0.233; 0.237]. The heteroskedasticity-consistent Lagrange multiplier test for a threshold developed by Hansen (1996) rejects the null of no threshold with a p-value of 0.00.

The estimated coefficient equals -1.01 (significant at the 1% level) for the customers below peer consumption. The coefficient is instead -11.09 (significant at the 1% level) for those above peer consumption. Subfigure (a) of Figure A.2 presents a binned scatterplot of the threshold regression estimates.

Kink Regression Results

Panel B of Table A.1 reports the results for the regression with endogenous kink.The threshold is estimated at 0.546, with a 95% confidence interval of [0.34; 0.77], and the null of no-threshold is rejected with a p-value of 0.00. The constant is not statistically different from zero.

The coefficient on the spending difference equals -0.726 below the threshold and is statistically different from 0. The coefficient above the threshold is instead 15 times larger (in absolute value), as the coefficient equals -11.197. This result indicates, once again, that investors who overspend are much more responsive to peer-group-spending information relative to individuals who underspend.

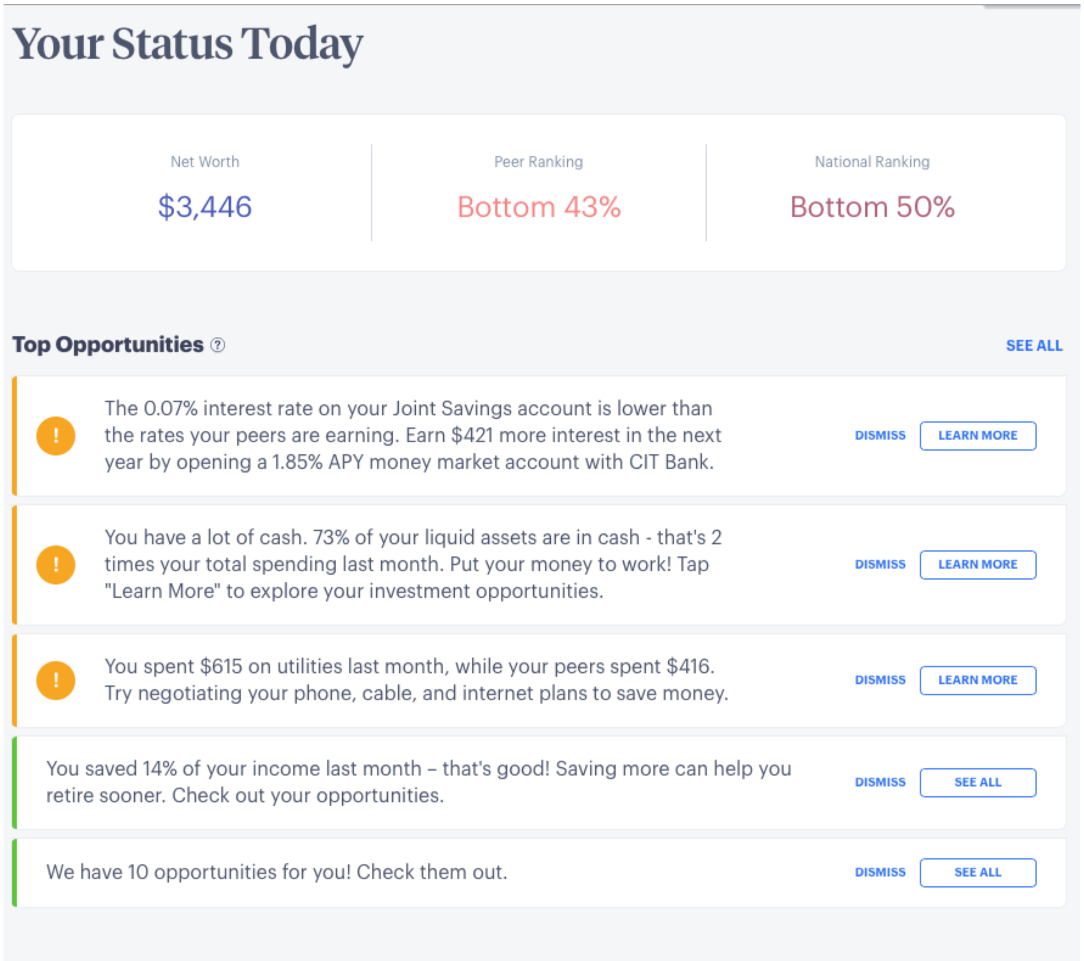

Figure A.1. ***Status*** **Home Page**

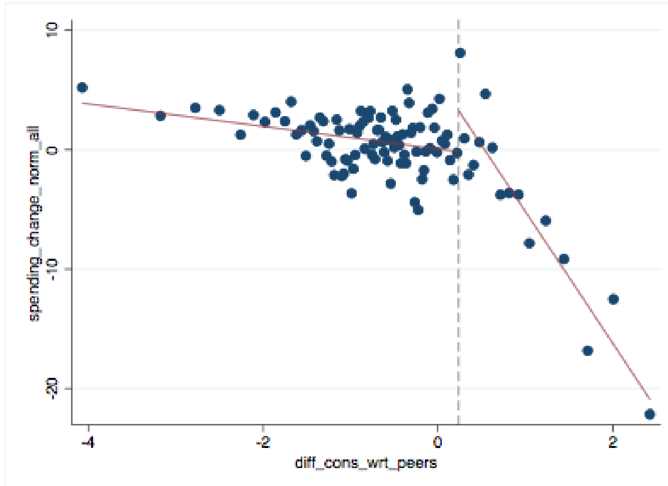

(a) Threshold Regression with Unknown Threshold

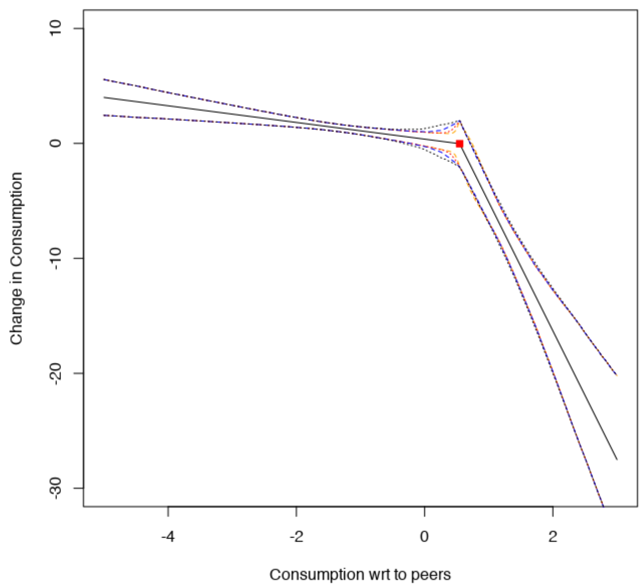

(b) Kink Regression with Unknown Threshold

Figure A.2
Distance from Peers' Spending and Changes in Spending after Signup—Endogenous Threshold Models

This figure reports the fitted values of a threshold regression model, with the optimal threshold estimated using the procedure in Hansen (2000) in subfigure (a). Subfigure (b) reports the fitted values of a kink regression model with the optimal threshold estimated using the procedure in Hansen (2015). In addition to the fitted values, subfigure (b) reports 90% confidence intervals.

Table A.1. Endogenous Threshold Models

Panel A. Threshold Regression Results

	All		Below Threshold		Above Threshold	
	Value	t-stat	Value	t-stat	Value	t-stat
Distance from Peers	-2.52***	(-11.79)	-1.01***	(-3.96)	-11.09***	(-7.81)
Constant	-1.43***	(-5.82)	0.07	(-0.19)	5.94***	(4.27)
Observations	17,673		14,846		2,827	

Threshold Estimate = 0.235; Confidence Interval = [0.233, 0.237]

Hansen (1996) Lagrange Multiplier for threshold: p-value = 0.00

Panel B. Kink Regression Results

	Coeff.	t-stat	Low CI	High CI
Constant	-0.026	-0.05	-0.95	0.89
Below Threshold	-0.726***	-3.02	-1.13	-0.32
Above Threshold	-11.197***	-5.80	-15.15	-7.24

Threshold Estimate = 0.546; Confidence Interval = [0.34, 0.77]

Hansen (2015) Wald test for threshold: p-value = 0.00

This table reports results for endogenous threshold regressions estimating the sensitivity of spending changes to peer consumption. In Panel A, we report the results for the threshold regressions of Hansen (2000). The procedure automatically selects the optimal threshold and estimates unconstrained linear regressions below and above the threshold. In addition to the regression coefficient estimates, we report results for the threshold estimates, the confidence interval for the threshold, and the p-value of the Hansen (1996) Lagrange Multiplier test for the presence of a threshold. Panel B reports the results for the regression kink model with an unknown threshold proposed in Hansen (2017). The procedure automatically selects the optimal threshold and estimates a piecewise linear regression model that is continuous at the threshold. In addition to the parameter estimates, we report results for the threshold estimates, the confidence interval for the threshold, and the p-value of the Hansen (2017) Wald test for the presence of a threshold.