

w o r k i n g
p a p e r

14 01

**Consumer Debt Dynamics: Follow the
Increases**

**John Carter Braxton and
Edward S. Knotek II**



FEDERAL RESERVE BANK OF CLEVELAND

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment on research in progress. They may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications. The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or of the Board of Governors of the Federal Reserve System.

Working papers are available at:

www.clevelandfed.org/research.

Consumer Debt Dynamics: Follow the Increases*

John Carter Braxton and Edward S. Knotek II

Consumer debt played a central role in creating the U.S. housing bubble, the ensuing housing downturn, and the Great Recession, and it has been blamed as a factor in the weak subsequent recovery as well. This paper uses micro-level data to decompose consumer debt dynamics by separating the actions of consumer debt increasers and decreasers, and then further decomposing movements into percentage and size margins among the increasers and decreasers. We view such a decomposition as informative for macroeconomic models featuring a central role for consumer debt. Using this framework, we show that variations in borrowing activity among the increasers explain four times as much of the total variation in consumer debt as variations among the decreasers who are shedding debt, whether through paydowns or defaults. We also provide micro-level evidence of a sharp decline in the percentage of increasers during the financial crisis that is qualitatively consistent with a binding zero lower bound on nominal interest rates, and evidence of a cycle in the average size of debt changes among the increasers that is related to rising collateral values pre-crisis coupled with additional financial frictions after the crisis.

Keywords: consumer debt, deleveraging, zero lower bound, increasers, decreasers

JEL codes: E32, E21, D14.

*This paper was previously titled “Consumer Debt Dynamics: How Many versus How Much?”

Suggest citation: John Carter Braxton and Edward S. Knotek II, 2014. “Consumer Debt Dynamics: Follow the Increases,” Federal Reserve Bank of Cleveland, working paper no. 14-01.

John Carter Braxton is at the Federal Reserve Bank of Kansas City. Edward S. Knotek II (corresponding author) is at the Federal Reserve Bank of Cleveland (216.579.2180; edward.knotek@clev.frb.org). The views expressed in this paper are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Cleveland, the Federal Reserve Bank of Kansas City, or the Federal Reserve System.

I. Introduction

Consumer debt played a central role in creating the U.S. housing bubble, the ensuing housing downturn, and the Great Recession, and it has been blamed as a factor in the weak subsequent recovery as well. Recent works by Mian and Sufi (2009, 2010, 2011a), Mian et al. (2013), Midrigan and Philippon (2011), and Dynan (2012) find a significant connection between overleveraged households and economic weakness. Debt also plays a role in papers that focus on the zero lower bound on nominal interest rates: in a world in which the zero lower bound becomes a binding constraint, less-leveraged households may not step up their borrowing to make up for deleveraging on the part of other agents, as in Hall (2011) and Eggertsson and Krugman (2012); the results in Guerrieri and Lorenzoni (2011) are similar.

Recent papers also attempt to make some headway in understanding household debt dynamics. Bhutta (2012) identifies both weak demand and elevated numbers of defaults as key drivers of declines in household mortgage debt since the Great Recession. Haughwout et al. (2012) also identify weak borrowing and paying down existing liabilities as important sources of consumer debt reduction. Mian and Sufi (2011b) find that house price appreciation during the housing bubble led to an increase in borrowing, especially among households with relatively high credit card utilization rates and low credit scores.

This paper seeks to more formally quantify the forces driving consumer debt dynamics using a nationally representative source of micro-level data on consumers' outstanding debt balances that spans 1999Q1 through 2013Q4. Our framework separates “increasers”—consumers who increase their debt balances—from the “decreasers” whose debt balances decline. We then further decompose the contributions coming from each group into the percentage of consumers increasing or decreasing their debt and the average size of the changes

in debt they made. In essence, such a division bears similarities to extensive and intensive margins, respectively.

We show that the actions of the increasers have been highly correlated with changes in consumer debt, even for periods dominated by consumer deleveraging. Variance decompositions show that variations in borrowing activity among the increasers have historically explained at least four times as much of the total variation in consumer debt as variations among the decreasers who are shedding debt, whether through paydowns or defaults. Among the increasers, we find that movements in our size measure explain about twice the variance as movements in the percentage measure, but even the latter are significantly more important than the cumulative actions of the decreasers. Applying these results to the recent past suggests that consumer deleveraging following the Great Recession can be interpreted as a story about consumer “non-leveraging” despite the presence of historically low interest rates.

Our empirical findings are relevant for a range of macroeconomic models that incorporate consumer debt dynamics. Beyond the variance decompositions showing the importance of the increasers, we provide qualitative and quantitative evidence of time variation in the percentage of increasers. The time series evidence on the percentage of increasers is inconsistent with the calibrations from several recent models that feature only two representative agents, an impatient borrower and a patient lender: the percentage of increasers was less than 50 percent during the housing boom and fell further during the financial crisis, consistent with a wider range of heterogeneity. Nevertheless, we document a sharp decline in the percentage of increasers during the financial crisis and low levels during the subsequent slow recovery that provides micro-founded evidence for the intuition of a binding zero lower bound on nominal

interest rates, during which time less leveraged agents fail to step up borrowing to offset deleveraging by others.

One concern is whether tight lending standards may have been the primary reason for such a failure rather than the zero lower bound on nominal rates. To address this issue, we rely on the fact that low credit scores are often interpreted as being correlated with borrowing constraints, as in Gross and Souleles (2002). The data show that consumers who were perceived as being among the safest credits—and therefore should have been less subject to tight credit conditions—failed to ramp up borrowing in the wake of the financial crisis, when short-term nominal rates were constrained by the zero lower bound, consistent with interest rates remaining “too high” to incentivize this least-borrowing-constrained group from leveraging up. But more generally, the deleveraging process has been led by historically weak borrowing activity from the consumers in the middle of the risk distribution who had previously been responsible for the majority of the leveraging during the housing bubble phase.

Finally, we document rapid growth in the average size of debt changes among the increasers during the housing bubble and a steep decline in the size measure during the housing bust. This pattern is again consistent with a binding zero lower bound on nominal rates, in which even lower rates could have incentivized those agents who were borrowing to take on additional leverage while keeping their debt servicing burden unchanged. But fully explaining the complete cycle requires some combination of a collateral channel and additional financial frictions: while house prices and size dynamics are tightly correlated before the crisis, consistent with Mian and Sufi (2011b), movements in collateral values alone cannot account for the post-crisis behavior of the size measure.

The structure of this paper is as follows. Section II describes the data. Section III presents our framework for decomposing consumer debt dynamics. Section IV presents regression and variance decomposition results. Section V shows the robustness of our results to alternative decompositions, and Section VI extends the results to consider a measure of household liquidity constraints. Section VII discusses how our results relate to the recent literature, and Section VIII concludes.

II. Consumer Debt Data

Studying consumer debt has historically been limited by a lack of high-quality, high-frequency, micro-level panel data. We use the Federal Reserve Bank of New York's Consumer Credit Panel (FRBNY-CCP) to help fill this void. The FRBNY-CCP is a nationally representative, ongoing panel dataset that focuses on consumers' debts using a five percent random sample of all individual credit records maintained by Equifax, one of the major credit reporting agencies. This quarterly dataset begins in 1999 and contains information on all aspects of a consumer's outstanding debts—including outstanding balances, past-due balances, and delinquency status—for all forms of debt. We randomly subsample five percent of consumers in the FRBNY-CCP for ease of analysis, allowing us to follow approximately 515,000 consumers on average in every quarter. Along with detailed information on an individual's outstanding debts, the FRBNY-CCP contains information on each individual's birth year, state, county and zip code of residence, and an Equifax risk score. Equifax risk scores are meant to predict the likelihood of a severe delinquency within the next 24 months; scores range from 280 to 850, with lower scores representing greater perceived credit risk—that is, a higher likelihood of a severe delinquency.

The FRBNY-CCP does not contain personally identifying information such as date of birth, Social Security number, or street address of residence. For additional information on the FRBNY-CCP, see Lee and Van der Klaauw (2010).

Our measure of consumer debt includes borrowing in five categories: first mortgage; home equity; auto; credit card; and other.¹ Due to difficulties in matching borrowers with student loans over time, we exclude student loans from the analysis. In calculating a consumer's total outstanding debt, jointly held accounts are given a weight of one half.

The combination of breadth and depth makes the FRBNY-CCP ideal for studying consumer debt dynamics. The representative nature of the FRBNY-CCP allows for computing aggregate debt statistics at the national level. At the individual level, the FRBNY-CCP employs unique identifiers so that individuals can be tracked over multiple quarters and their borrowing histories reconstructed. Observing these borrowing histories can help determine which consumers play an important role in driving consumer debt dynamics.

Unfortunately, one drawback to the FRBNY-CCP as a source of consumer debt records is that the records can contain noise. For example, refinancing can be recorded as a payoff in one quarter and a new balance in a later quarter, or accounts can be transferred or sold and the closing of the old account and opening of a new account may not perfectly coincide; see Lee and van der Klaauw (2010). In addition, the number of consumers varies over time due to entry, exit, or coding errors. It is not always possible to ascertain which explanation holds: some consumers temporarily disappear from the sample only to reappear later, often with a different level of debt.

¹ First mortgage borrowing is defined as a closed-end mortgage with priority over all other mortgages and liens on a property. Home equity borrowing includes both home equity loans and home equity lines of credit. Auto borrowing covers loans provided by banking institutions and automobile dealers for the purpose of purchasing an automobile. Credit card borrowing includes revolving accounts maintained by banks, bankcard companies, national credit card companies, credit unions, and savings & loan associations. Other borrowing is comprised of consumer finance and retail loans, with the latter including clothing, grocery, department stores, home furnishing, and gas loans.

On average, 0.9% of consumers disappear from the sample from one quarter to the next; of these, 26.9% reenter at some point in the future through the end of the sample, and 47.2% of those reenter with a different level of debt than they owed when they left.

Our analysis focuses on changes in debt over year-ago levels. Using four-quarter changes helps to mitigate some of the noise in the high frequency quarterly data. In addition, four-quarter changes abstract from predictable variations in debt over the calendar year from seasonal effects, which are typically omitted or removed in most macroeconomic analyses.

III. Decomposing Consumer Debt Dynamics

The level of total consumer debt D_t is the sum of all individual consumers' debt holdings $d_{i,t}$ in that quarter,

$$(1) \quad D_t = \sum_{i=1}^{N_t} d_{i,t},$$

where the total number of consumers, N_t , varies with time. Given the representative nature of our panel, multiplying the sampled level of total debt and the sampled number of consumers by a constant multiplicative scaling factor $\eta=400$ produces the corresponding aggregate totals. In most of what follows, however, these scaling factors drop out.

Our approach focuses on nominal consumer debt. One rationale for this focus is the lack of a single clear method to deflate consumer debt. Deflating by a measure of consumer prices compares growth in debt to movements in the prices of a broader basket of goods and services; alternatively, deflating by nominal disposable personal income can be useful for illustrating the servicing capability of the debt burden. A second rationale for focusing on nominal debt is that the available data for our analysis are not so long that deflating nominal debt becomes absolutely

necessary. But finally and perhaps most importantly, the focus on nominal debt allows for a closer examination of the individual consumer's choice to increase or decrease debt, notwithstanding movements in aggregate variables; hence, individual decisions are preserved. For example, under the approach of deflating nominal debt by the price level, a small increase in an individual's debt level that is less than the growth in consumer prices would appear as a decrease in debt, which is at odds with the individual's choice problem.

We consider two measures of movements in consumer debt. The first measure is the change in debt per person over its year-ago level:

$$(2) \quad \Delta_t = \frac{D_t}{N_t} - \frac{D_{t-4}}{N_{t-4}}.$$

The second measure is the growth of consumer debt over its year-ago level:

$$(3) \quad g_t = \frac{D_t - D_{t-4}}{D_{t-4}}.$$

While the statistics Δ_t and g_t are straightforward, aggregating from individuals poses challenges. As noted above, a small number of individual consumers enter, exit, and “disappear” from the panel only to “reappear” later, which means that treating missing consumers as having zero debt is sometimes but not always correct. In what follows, we only follow individuals who were present at both the beginning of each four-quarter window ($t-4$) and the end of the four-quarter window (t), ignoring their status in between. Over a given four-quarter window, there are $N_{(t,t-4)}$ such consumers. The total level of consumer debt among those consumers at the end of the window is $D_{(t,t-4),t} = \sum_{i=1}^{N_{(t,t-4)}} d_{i,t}$, and debt at the beginning of the window is

$D_{(t,t-4),t-4} = \sum_{i=1}^{N_{(t,t-4)}} d_{i,t-4}$. We can then rewrite the change in debt per person as

$$(4) \quad \Delta_t = \frac{D_{(t,t-4),t} - D_{(t,t-4),t-4}}{N_{(t,t-4)}}$$

and the growth of consumer debt as

$$(5) \quad g_t = \frac{D_{(t,t-4),t} - D_{(t,t-4),t-4}}{D_{(t,t-4),t-4}}.$$

Equation (4) implies that the change in debt per person is equivalent to the average change in debt, because the number of persons does not change during the window.

Figure 1 plots the two measures. Both measures capture the rise and subsequent fall in consumer debt. In percentage terms, debt growth averaged about 10 percent on a year-over-year basis from the beginning of the sample in 2000 through the end of 2007. Consequently, changes in nominal debt were rising during this time as the total stock of debt rose. Debt growth fell off sharply in 2008 and turned negative in 2009 based on either measure. The ongoing declines in consumer debt in the latter part of the sample represent consumer deleveraging.

Because individuals need to be present at both the beginning and the end of each window, the level of total debt derived from this window construct will necessarily be less than or equal to debt derived from the entire sample; i.e., $D_{(t,t-4),t} \leq D_t$.² In this way, we focus on individual actions within a window and not on population growth that might affect debt dynamics.³

We decompose movements in consumer debt by separating consumers based on what happened to their total outstanding debt level relative to four quarters prior: some consumers increased their outstanding debt since that time, some decreased their debt, and the remainder

² Additionally, the subscript construct is necessary to denote the window under consideration: it is not necessarily the case that $D_{(t,t-4),t} = D_{(t+4,t),t}$.

³ From a practical standpoint, the change in debt per person and the growth of consumer debt are only slightly impacted by the window construct. The largest absolute difference is 1.2 percentage points for the growth of consumer debt. For the change in debt per person, the largest absolute difference is \$1,184.

saw their outstanding debt loads remain constant. Next, we examine two key margins: the percentage of consumers who change their level of debt and the size of those debt changes.

Increasers, Decreasers, and No-changers

At each point in time, consumers can be divided into three groups.

The first group of consumers is the $N_{I(t,t-4)}$ “increasers” whose individual debt levels increased between time $t-4$ and t , comprising the percentage $p_{I,t} = N_{I(t,t-4)} / N_{(t,t-4)}$ of the total number of consumers. The set of increasers at time t owes total debt $D_{I(t,t-4),t} = \sum_{i \in \{N_{I(t,t-4)}\}} d_{i,t}$.

Among the increasers, the average size of debt changes between time $t-4$ and t is

$$s_{I,t} = (D_{I(t,t-4),t} - D_{I(t,t-4),t-4}) / N_{I(t,t-4)}.$$

The second group of consumers is the $N_{D(t,t-4)}$ “decreasers,” comprising percentage $p_{D,t} = N_{D(t,t-4)} / N_{(t,t-4)}$ of the population. This group’s total debt load is $D_{D(t,t-4),t}$ at time t . The average size of debt changes between time $t-4$ and t among the decreasers is

$$s_{D,t} = (D_{D(t,t-4),t} - D_{D(t,t-4),t-4}) / N_{D(t,t-4)}.$$

Finally, the third group of consumers is the “no-changers.” The $N_{NC(t,t-4)}$ no-changers have total debt in period t of $D_{NC(t,t-4),t}$, they account for $p_{NC,t} = 1 - p_{I,t} - p_{D,t}$ of the population, and the average size of debt changes within the group is zero.

Total aggregate debt at time t is the sum of the debt levels of the mutually exclusive groups, $D_{(t,t-4),t} = D_{I(t,t-4),t} + D_{D(t,t-4),t} + D_{NC(t,t-4),t}$. Taking into account this relationship and the definitions above, the change in total consumer debt per person in equation (4) can be rewritten as

$$(6) \quad \Delta_t = p_{I,t} s_{I,t} + p_{D,t} s_{D,t}.$$

The first term on the right hand side of equation (6), $p_{I,t} s_{I,t}$, is the contribution coming from debt changes among the increasers and can be thought of as the “inflow” of consumer debt. The second term, $p_{D,t} s_{D,t}$, is the contribution coming from debt changes among the decreasers and can be thought of as the “outflow.”⁴

The growth of consumer debt is also a function of the percentage and size measures.

Letting $\delta_t = (N_{(t,t-4)} / D_{(t,t-4),t-4})$ be the inverted level of debt per person at the start of the window, the growth of consumer debt can be written as

$$(7) \quad g_t = \delta_t (p_{I,t} s_{I,t} + p_{D,t} s_{D,t}).$$

Detailed Decompositions

We can provide additional detail to the decompositions by differentiating consumers based on whether they were holding positive debt balances at the beginning or the end of the window under consideration.⁵ We denote with superscripts the beginning and ending debt levels among groups of consumers. The number of increasers is the sum of the number of consumers whose individual debt levels went from zero at time $t-4$ to a positive level at time t ($N_{I(t,t-4)}^{0 \rightarrow +}$) and the number of consumers whose debts went from a positive level to a greater positive level

⁴ Bhutta (2012) uses inflows and outflows in a similar fashion to describe changes in mortgage debt among individuals, without examining the percentage and size measures. In addition, Bhutta (2012) examines six non-overlapping two-year windows to consider dollar inflows and outflows.

⁵ Bhutta (2012) classifies consumers whose mortgage debt went from zero at the start of non-overlapping 2-year windows to a positive level at the end of the window as “entrants,” and consumers whose debt went from a positive level to zero as “exiters.” This terminology can create some confusion, however, because these consumers need not be true entrants or exiters: the former may have had positive debt at some point in the past but zero debt at the start of the window being examined, and the latter may hold future debt.

over the same time span ($N_{I(t,t-4)}^{+\rightarrow++}$) such that $N_{I(t,t-4)} = N_{I(t,t-4)}^{0\rightarrow+} + N_{I(t,t-4)}^{+\rightarrow++}$. Similarly, the number of decreasees is the sum of the number of consumers who had positive debt in time $t-4$ and zero debt in period t , and the number of consumers whose debts went from a positive level to a smaller positive level over the same time span: $N_{D(t,t-4)} = N_{D(t,t-4)}^{+\rightarrow 0} + N_{D(t,t-4)}^{++\rightarrow+}$.

The change in consumer debt per person can then be written as the sum of the contributions coming from the mutually exclusive groups of increasees (zero to positive, and positive to greater positive) and decreasees (positive to zero, and positive to smaller positive). In turn, each group's contribution can be decomposed into percentage and size measures:

$$(8) \quad \Delta_t = p_{I,t}^{0\rightarrow+} s_{I,t}^{0\rightarrow+} + p_{I,t}^{+\rightarrow++} s_{I,t}^{+\rightarrow++} + p_{D,t}^{+\rightarrow 0} s_{D,t}^{+\rightarrow 0} + p_{D,t}^{++\rightarrow+} s_{D,t}^{++\rightarrow+}.$$

The percentage of consumers who had unchanged zero debt at both time $t-4$ and t , $p_{NC,t}^{0\rightarrow 0}$, and the percentage of consumers who had positive but unchanged debt at both time $t-4$ and t , $p_{NC,t}^{+\rightarrow+}$, do not enter equation (8) because their respective size measures are zero by definition. Similarly, the growth of consumer debt can be written in terms of percentage and size measures as

$$(9) \quad g_t = \delta_t \left(p_{I,t}^{0\rightarrow+} s_{I,t}^{0\rightarrow+} + p_{I,t}^{+\rightarrow++} s_{I,t}^{+\rightarrow++} + p_{D,t}^{+\rightarrow 0} s_{D,t}^{+\rightarrow 0} + p_{D,t}^{++\rightarrow+} s_{D,t}^{++\rightarrow+} \right).$$

IV. Who Drives Consumer Debt?

We study consumer debt dynamics through a variety of exercises. We begin by considering correlations and simple graphical and statistical relationships between the aggregate variables of interest (the change in debt per person and the growth in consumer debt) and our constructed percentage and size measures. While these analyses provide suggestive evidence on consumer debt dynamics, variance decompositions allow us to quantify the key drivers of consumer debt.

In these exercises, we show that variations in borrowing activity among the increasers have typically explained at least four times as much of the total variation in consumer debt as variations among the decreasers who are shedding debt, whether through paydowns or defaults.

Correlations

As a first pass, we consider correlations between our constructed measures and aggregate debt movements. Table 1 looks at correlates of the change in debt per person (Δ_t) and the growth of consumer debt (g_t), based on the relationships derived in equations (6) and (7). In the simplest versions, the contributions coming from the increasers tend to be more highly correlated than the contributions coming from the decreasers with both the change in debt per person (0.94 versus 0.71, respectively) and the growth of consumer debt (0.84 versus 0.81, respectively). After splitting the contribution terms into their percentage and size components, the percentage measures in most cases exhibit higher correlations than the size measures, but the strength of the correlations does not uniformly favor either the increasers or decreasers.

We next look at the correlations coming from the more detailed breakdowns in equations (8) and (9). Table 2 shows that correlations between the two groups of increasers' contributions (zero to positive, and positive to greater positive) and consumer debt movements are generally high, all above 0.8. The strongest correlations come from the percentage of increasers going from a positive level of debt to a greater positive debt level ($p_{I,t}^{+\rightarrow++}$), at 0.89 for the change in consumer debt per person and 0.95 for the growth of consumer debt. Correlations are somewhat smaller for the contributions coming from the two groups of decreasers. Drilling down, the percentage of consumers who are decreasers and go from positive debt to zero debt ($p_{D,t}^{+\rightarrow 0}$)—

which would include individuals who potentially paid off all debts or defaulted on all debts—is essentially uncorrelated with aggregate debt movements. In addition, the correlations are weak for the average size of debt changes among decreasers who reduce but do not eliminate their debts ($s_{D,t}^{++\rightarrow+}$); this category would be capturing consumers who are amortizing debts on a regular basis as well as those who default on a portion of their debts.

We interpret these correlations as offering some support for the notion that the behavior of the increasers is a key driver of consumer debt dynamics.

Regression Results

We build upon these correlations and consider simple regressions to explain changes in debt per person and growth of consumer debt. Each regression takes the form

$$(10) \quad y_t = \alpha + \beta x_t + e_t,$$

where the dependent variable y_t is either the change in debt per person (Δ_t) or the growth of consumer debt (g_t), the explanatory variable x_t is one of the percentage or size measures from the earlier derivations, and e_t is an error term. We consider each explanatory variable individually to assess its ability to explain movements in consumer debt.

Table 3 presents the adjusted R-squared statistics from each regression as a summary of fit; we omit the coefficient estimates, as we do not find them to be informative. Among the percentage measures, adjusted R-squared is largest for the percentage of increasers going from a positive level of debt to a greater positive level of debt ($p_{I,t}^{+\rightarrow++}$). Explanatory power is rather low for the percentage of increasers who held zero debt at the beginning of the window under consideration ($p_{I,t}^{0\rightarrow+}$) and negligible for the percentage of decreasers going from positive to zero

debt ($p_{D,t}^{+\rightarrow 0}$). Thus, the fractions of consumers “entering” and “exiting” debt markets seem to have minimal descriptive power in explaining aggregate debt movements. The size regressions generally explain less of the variation in the dependent variables. While the fit for the percentage of decreasees going from positive to zero debt was very poor, the largest adjusted R-squared among the average size measures comes from those decreasees who eliminated all their debt ($s_{D,t}^{+\rightarrow 0}$), going from a positive debt load to zero either by defaulting or by completely paying down their previous debt balance. These results suggest that increasees—and in particular those increasees whose debt went from a positive to a greater positive level—play a larger role in driving debt dynamics than the decreasees.

We use the regression results in creating Figure 2, which plots our constructed percentage and size measures. To maximize compatibility, we use the estimated regression coefficients from equation (10) to scale the right axis for each measure, and we plot the change in total consumer debt or the percentage growth of total consumer debt on the left axis.

Visually, the percentages of increasees and decreasees tend to track aggregate consumer debt growth well, measured using either growth rates or average changes. If one period stands out, it is the very tight connection between the percent change in debt and the percentage of increasees from the end of 2007 onward. Looking at key trends, the percentage of increasees was flat-to-declining during the housing bubble, plunged during the Great Recession, and remained at low levels in the subsequent recovery. Meanwhile, the housing bubble run-up was more visible in the average size of debt changes among the increasees: this size measure steadily climbed during the bubble phase, and then declined sharply during the recession. The average size of debt changes among the decreasees provides only limited visual evidence of a sudden increase in

defaults and debt paydowns following the bubble. We discuss the significance of these findings and their implications for recent attempts to model consumer debt dynamics in Section VII.

Variance Decompositions

As a more formal exercise, we employ variance decompositions to explain movements in consumer debt. We focus on explaining the variance of the change in debt per person, $\text{var}(\Delta_t)$. We find roughly similar results when explaining the variance of the growth rate of total consumer debt using equation (7).⁶ We consider four different decompositions.

The first decomposition uses the variance of the contributions from the increasers, $\text{var}(p_{I,t} s_{I,t})$, the variance of the contributions from the decreasers, $\text{var}(p_{D,t} s_{D,t})$, and the covariance term, $\text{cov}(p_{I,t} s_{I,t}, p_{D,t} s_{D,t})$, to explain $\text{var}(\Delta_t)$. In this three-term decomposition, column (1) in Table 4 shows that 61% of the variance is explained directly by the increasers, while only 14% is explained directly by the decreasers; the remainder is attributed to the covariance term, which incorporates a range of factors. On this basis, the variations in borrowing activity among the increasers explain four times as much of the total variation in consumer debt changes as variations among the decreasers. Even if the entire covariance term could be attributed to decreasers, the increasers would nonetheless continue to dominate consumer debt movements.

⁶ These results are available upon request. However, the extreme differences in the magnitudes of the δ_t term and the other terms in equation (7) tend to cause the variance decompositions to be less precise than the one presented in the text.

Column (2) uses a first-order Taylor series expansion of equation (6) to account for the variance of changes in consumer debt.⁷ As we add more terms to the decomposition, the number and importance of the covariance terms grow. Nevertheless, we continue to obtain similar results: the direct variance terms coming from the increasers now explain about half of the total variance, while the direct variance terms coming from the decreasers now explain about 10%, for a five-to-one ratio.

Column (3) splits increasers and decreasers into two groups, depending on whether the consumers began or ended the four-quarter window with zero debt, respectively. In this decomposition, debt changes among consumers who began or ended with zero debt explain essentially none of the variance in changes in consumer debt. Instead, variation in debt changes among the increasers who go from a positive debt level to a greater positive level explains the vast majority, in roughly a six-to-one ratio.

Finally, column (4) presents results from a variance decomposition using a first-order expansion of equation (8). With a large number of terms in the original equation, there are a very large number of covariance terms, and these terms collectively explain a considerable share of the total variance. Nevertheless, even with this fine level of granularity, the increasers whose debt went from a positive to a greater positive level continue to directly explain about six times more of the variation in total consumer debt changes than the decreasers.

Among the increasers, the variance share explained by the average size measures is about twice that of the percentage measures, and both are far larger than the corresponding measures from decreasers. These facts accord with the graphical evidence from Figure 2, which showed

⁷ Letting \bar{x} be the steady state value of x and defining $\hat{x}_t = x_t - \bar{x}$, equation (6) implies

$\hat{\Delta}_t = \bar{p}_I \hat{s}_{I,t} + \bar{s}_I \hat{p}_{I,t} + \bar{p}_D \hat{s}_{D,t} + \bar{s}_D \hat{p}_{D,t} + O_t$, where O_t denotes higher order terms. Squaring both sides yields the variance decomposition. Klenow and Kryvtsov (2008) perform a similar variance decomposition for the frequency and size of price changes.

that the size measures played a large role in the bubble phase run-up in debt in the mid-2000s, while the percentage measures were flat-to-declining. Both the average size of debt changes among increasers and the percentage of increasers fell off sharply during the recession and have remained at low levels in the recovery.⁸ We discuss interpretations of this evidence below.

The variance decompositions offer clear evidence of the importance of the increasers in explaining movements in consumer debt over the last decade and a half.

V. Robustness to Alternative Decompositions

For the sake of robustness, we consider three variants to the results presented above. We first consider an alternative derivation of the growth rate of consumer debt that does not rely on the percentage and size measures. We next change the time period under consideration to focus on consumer debt dynamics after the peak of the housing bubble in 2006; that is, we focus on the period of consumer deleveraging. Finally, we exclude housing-related debt to assess the degree to which our results are directly related to the housing bubble. All three cases provide further evidence of the key role of the increasers in driving consumer debt dynamics.

Robustness: Decomposing the Growth Rate of Consumer Debt

Alternative decompositions of the growth rate of consumer debt are possible beyond equation (7). For example, let $\theta_{I,t} = D_{I(t,t-4),t-4} / D_{(t,t-4),t-4}$ denote the share of total debt owed by

⁸ While the variance decompositions strongly favor the role of the average size of debt changes among increasers in driving consumer debt dynamics over the percentage of increasers, Knotek and Braxton (2012) present a counterfactual exercise in which the low percentage of increasers has a greater impact on consumer deleveraging than the low average size of debt changes among the increasers.

increasers at the beginning of the window, and denote the growth rate of increasers' debt over the same window as $\gamma_{I,t} = (D_{I(t,t-4),t} - D_{I(t,t-4),t-4}) / D_{I(t,t-4),t-4}$. Combined with similar metrics for the share of total debt owed by decreasees ($\theta_{D,t}$) and the growth rate of decreasees' debt ($\gamma_{D,t}$), we can rewrite equation (5) as

$$(11) \quad g_t = \theta_{I,t} \gamma_{I,t} + \theta_{D,t} \gamma_{D,t}.$$

As before, the term $\theta_{I,t} \gamma_{I,t}$ represents the contribution coming from the increasers, and $\theta_{D,t} \gamma_{D,t}$ is the contribution from the decreasees.

It is also possible to do detailed breakdowns analogous to equation (8). Note that the percentage change in debt among consumers whose debt levels go from a positive level to zero is -100% , and let $\tau_{I,t}^{0 \rightarrow +} = D_{I(t,t-4),t}^{0 \rightarrow +} / D_{(t,t-4),t-4}$ be the total *new* debt acquired by consumers who previously had no debt at time $t-4$, expressed as a share of the previous total level of debt. Then the growth of total consumer debt can be written in terms of shares and growth rates as

$$(12) \quad g_t = \tau_{I,t}^{0 \rightarrow +} + \theta_{I,t}^{+ \rightarrow ++} \gamma_{I,t}^{+ \rightarrow ++} - \theta_{D,t}^{+ \rightarrow 0} + \theta_{D,t}^{++ \rightarrow +} \gamma_{D,t}^{++ \rightarrow +}.$$

Table 5 presents results from the four alternative variance decompositions. Columns (1) and (3) show that the increasers alone explain more than 100% of the variance in the growth of consumer debt, $\text{var}(g_t)$, and more than ten times the variance explained by the decreasees. Obviously, the covariance terms in these cases are highly negative, which complicates the interpretation. Columns (2) and (4) are variance decompositions derived from first-order expansions of equations (11) and (12), respectively, and they tell essentially similar stories: variation among the increasers dictates the vast majority of consumer debt dynamics.

We do not find the share measures $\theta_{I,t}$ and $\theta_{D,t}$ to be as intuitive as the percentage measures $p_{I,t}$ and $p_{D,t}$. The percentage measures have a familiar form in that they reveal the

fraction of the population willing and able to take on additional debt over a given time period. Nevertheless, this alternative approach provides a sensitivity check and largely confirms our initial results that the increasers dominate consumer debt dynamics.

Robustness: Crisis and Recovery Period

The above results consider the entire sample of the FRBNY-CCP, which spans 1999Q1 to 2013Q4. One possibility is that these decompositions primarily capture the run-up in debt during the housing bubble, when consumers were rapidly increasing outstanding mortgage debt. As a second robustness check, we decompose the variance of the change in debt per person, $\text{var}(\Delta_t)$, for the period after the housing bubble's peak: 2007Q1 to 2013Q4. This time span is dominated by consumer deleveraging during the financial crisis and subsequent recovery.⁹

Columns (1) and (3) in Table 6 show that variations in debt changes among increasers explain more than half of the variance of the change in debt per person, and variation among the increasers explains more than ten times the variation explained by the decreasers. Detailed decompositions in columns (2) and (4) produce broadly similar results and ratios. The variance terms in the decompositions using the shorter sample tend to explain a smaller share of the total variance than when considering the entire sample. Nevertheless, our main result holds: increasers continued to drive consumer debt dynamics after 2007.¹⁰

⁹ While not presented here, we find roughly similar results when performing variance decompositions for the growth rate in consumer debt using equations (11) and (12).

¹⁰ Starting the variance decomposition considerably later avoids the initial downturn in debt growth and reduces the total variance to be explained— $\text{var}(\Delta_t)$ —in some cases by more than 90%. Nevertheless, the contributions from the increasers continue to explain a greater share of the total variance than the contributions from the decreasers in every subsample that we examined.

Robustness: Non-Housing Debt

The primary analysis in this paper considers a broad definition of debt, including first mortgage debt, home equity borrowing, auto debt, credit card debt, and other debt but excludes student loans. Together, first mortgage debt and home equity borrowing account for over 80% of total debt on average during our sample. In an effort to remove direct effects of the housing bubble and bust on consumer debt dynamics, we consider a variance decomposition that excludes housing-related debt in the form of first mortgage debt and home equity borrowing.¹¹

Columns (5) and (7) in Table 6 show that the contributions from the increasers explain the majority of the total variance in changes in non-housing debt per person and more than three times the contributions from the decreasers. Columns (6) and (8) tell a roughly similar story. In one break from the earlier results, fluctuations in the percentage of increasers play a greater role than fluctuations in the average size of borrowing among the increasers in explaining non-housing debt movements. Putting these results together with those presented earlier in the paper, increasers are better able to explain aggregate movements in debt over time across types of debt.

VI. Extension: Consumer Debt Dynamics and Consumer Risk

Our framework for distinguishing between increasers and decreasers and examining the percentage and size margins within each group allows us to examine other questions, such as how consumers' perceived riskiness has affected consumer debt dynamics. Low credit scores are often interpreted as being correlated with borrowing constraints, as in Gross and Souleles

¹¹ Such a separation is necessarily imprecise: some credit card and other debt may be acquired to furnish a house, and home equity borrowing can support non-housing consumption.

(2002), potentially reflecting borrowers' underlying behavioral characteristics related to self-control, and a number of macroeconomic models featuring consumer debt employ heterogeneous discount rates to derive a segment of the population that owes debt in equilibrium (e.g., Hall 2011, Eggertsson and Krugman 2012, Justiniano et al. 2013).¹² One popular narrative to explain the housing bubble and bust is that the gradual easing of underwriting standards fueled the housing bubble, and the sudden tightening of standards proved the impetus for the housing downturn and financial crisis (e.g., Mian and Sufi 2009, Favara and Imbs 2012; a similar mechanism is at play in Geanakoplos 2010). We take an agnostic view of the causes of the boom and bust to simply let the data speak.

To do so, we group individual consumers into quartiles based on their Equifax risk scores at the beginning of each four-quarter window and then follow the changes in their total debt balances to classify them as increasers or decreasers and to compute the percentage and size margins. Individuals in the first (fourth) risk quartile have the lowest (highest) risk scores and are perceived as being the riskiest (safest) credits based on the Equifax scoring system. By construction, the mutually exclusive quartiles imply that the change in debt per person (Δ_t) is

$$(13) \quad \Delta_t = \sum_{q=1}^4 \Delta_{q,t},$$

where the contribution coming from each quartile q is $\Delta_{q,t} = (D_{q,(t,t-4),t} - D_{q,(t,t-4),t-4}) / N_{(t,t-4)}$.

Figure 3 shows the contributions by risk quartile over the sample period. Consumers in the first risk quartile contributed virtually nothing to aggregate consumer debt movements prior to 2008. However, their debt contributions declined sharply during the financial crisis and into

¹² Mian and Sufi (2011b) use credit scores as one measure of borrowing constraints to explain borrowing against increases in home equity during the run-up phase of the housing bubble. Other measures of borrowing constraints also exist, such as credit card utilization rates (Gross and Souleles 2002, Mian and Sufi 2011b), liquid wealth to income ratios (Zeldes 1989), and debt service ratios (Johnson and Li 2010).

the recovery, reflecting an elevated level of defaults: delinquencies impair credit scores before defaults occur, thus individuals who default are typically in the first quartile. Similarly, individuals in the fourth risk quartile contributed little to aggregate consumer debt growth during the pre-crisis period, and they modestly reduced their outstanding balances during the recovery. Consumers in the middle of the risk score distribution made the largest contributions to the rapid increase of consumer debt during the housing bubble. Following the crisis, these consumers experienced the largest change in behavior, adding very little net debt on average since 2009. Putting these facts together, the deleveraging process appears to be a combination of defaults by individuals in the first quartile; paydowns by individuals in the fourth quartile; and historically low net borrowing activity by individuals in the second and third quartiles.

Each risk score quartile's contribution depends on the percentage and size measures,

$$(14) \quad \Delta_{q,t} = \omega_q \left(p_{q,I,t} s_{q,I,t} + p_{q,D,t} s_{q,D,t} \right),$$

where $\omega_q = 0.25$ is the share of individuals in quartile q and $p_{q,j,t}$ denotes the percentage of individuals within that quartile who either increase ($j=I$) or decrease ($j=D$) their debt. The “inflows” from the increasers in quartile q are the product $\omega_q p_{q,I,t} s_{q,I,t}$; the “outflows” from the decreasers are $\omega_q p_{q,D,t} s_{q,D,t}$.

Figure 4 shows the percentages of increasers and decreasers along with the average size of their debt changes within each risk quartile. While considerable heterogeneity dominates the figure, there are common trends; e.g., both the percentage of increasers and the average size of debt changes among the increasers declined across all risk scores during the recession. As shown in Figure 3, individuals in the second and third quartiles have historically made large positive contributions to consumer debt; Figure 4(a) shows that they have had the highest percentage of increasers, and Figure 4(c) shows they have often increased their borrowing by the

largest amounts when doing so. The percentage of decreasers in the first quartile and the average size of debt declines among decreasers in the first quartile—i.e., among the highest risk individuals—increased sharply during the recession, consistent with an elevated level of defaults.

Table 7 reports variance decompositions by risk quartile. The four variance terms from each quartile's contribution— $\text{var}(\Delta_{q,t})$ —in column (1) explain a minority of the total variance; as indicated above graphically, there are strong covariance contributions that we do not report. The first three risk quartiles explain relatively more of the total variance than the fourth quartile. Looking at increasers and decreasers in column (2), decreasers in the first risk quartile have a larger influence on changes in consumer debt than increasers, but the results are reversed in the second and third risk quartiles. Column (3) shows results from the full variance decomposition. With a very large number of terms in the decomposition, the share of the variance explained by any single variable is small. Nevertheless, the results show that default activity as captured by the average size of debt changes among decreasers in the first quartile explains less of the total variance than the size variables among increasers in the second and third quartiles.

VII. Discussion

Our results highlight the importance of properly capturing the actions of the increasers for matching aggregate consumer debt dynamics. Our quantitative approach shows that even periods of consumer deleveraging are mostly dominated by weak borrowing activity rather than high levels of defaults. Beyond this finding, our empirics are broadly relevant for the new and growing literature on the relationship between consumer debt and the macroeconomy.

First, we use high-frequency, micro-level evidence to document the percentage of consumers who increase their debts, along with qualitative and quantitative evidence of time variation in the percentage of increasers, and we show that capturing this time variation is important in accounting for consumer debt dynamics. A number of recent papers on consumer debt have motivated a household borrowing channel by separating households into “patient” and “impatient” (or “unconstrained” and “constrained”) types based on differing discount factors, with identical households in each type to produce two representative agents. The relative shares of the two agents have then been calibrated using moments from the Survey of Consumer Finances (e.g., Hall 2011, Justiniano et al. 2013) to derive a time-invariant segment of the model population that is borrowing constrained, with the share of constrained consumers calibrated to be greater than 50%.¹³ A period of rising debt in these models would coincide with additional borrowing by all the constrained consumers. By contrast, Figure 2 shows that the percentage of increasers in our sample remained below 50% even during the housing boom and fell below 30% immediately following the recession.¹⁴ Models with only two representative agents cannot match this time series evidence. Thus, matching the quantitative behavior of the percentage of increasers requires additional heterogeneity in either discount rates or some other mechanism.¹⁵

Notwithstanding these quantitative shortcomings, we view our results on the sharp decline in the percentage of increasers around the financial crisis as providing micro-founded empirical evidence of the types of behaviors induced by a binding zero lower bound on nominal interest rates. As set out in papers such as Hall (2011) and Eggertsson and Krugman (2012), the

¹³ Eggertsson and Krugman (2012) employ a similar theoretical construct.

¹⁴ Furthermore, these numbers are an upper bound: while our sample is representative of every consumer with a credit report, it will not include consumers who have never applied for credit and thus are not in the Equifax database. Adding these consumers into our analysis as “no-changers” would cause a downward level shift in the percentage of increasers.

¹⁵ As one alternative, Guerrieri and Lorenzoni (2011) propose a model of uninsurable idiosyncratic risk that allows for a wider range of heterogeneity in household productivity and bond holdings.

presence of the zero lower bound prevents nominal interest rates from falling far enough to incentivize less-leveraged agents to increase borrowing to offset deleveraging by others, which can take the form of a decrease in the number of net borrowers. Notably, we document in Figure 3 that even the fourth risk score quartile—which contains the least risky consumers with the highest scores—was deleveraging in the aggregate following the financial crisis, and Figure 4(a) shows that the fraction of increasers in this quartile was essentially unchanged before and after the crisis. A larger decline in nominal rates should have induced additional borrowing by this group because their position in the risk score distribution would have made them less subject to the tight credit conditions that impeded other potential borrowers after credit conditions tightened post-crisis. Hence, their lack of leveraging can be viewed as evidence of interest rates remaining “too high.”

Finally, we document a strong cycle in the size of debt changes among the increasers, featuring rapid growth in our size measure during the housing bubble and a steep decline in the size measure during the housing bust (see Figure 2(d)). The decompositions find that variations in the average size of debt changes among the increasers are about twice as important as variations in the percentage of increasers in explaining the total variance in consumer debt.¹⁶ The decline in the average size of changes among the increasers around the financial crisis can be partly interpreted through the prism of the binding zero lower bound post-crisis: with even lower rates, more agents could have been incentivized to increase borrowing as described above, and those who did could have borrowed to a greater extent by taking on additional leverage while keeping their debt servicing burden unchanged.

¹⁶ Knotek and Braxton (2012) present a counterfactual in which the average size of debt changes among the increasers played a key role in inflating debt-to-income ratios during the housing boom.

But fully explaining the cycle requires an important role for a collateral channel. Using housing supply elasticities as an instrument, Mian and Sufi (2011b) present evidence of a causal link from house price appreciation between 2002 and 2006 to increased borrowing during that period, without ruling out potential feedback from increased collateral values to house prices (see Mian and Sufi 2009). Because we separate increasers from decreasers, the average size of debt changes among the increasers that we document is potentially a purer reflection of house price dynamics during both the boom and bust given the large role played by housing debt in our measure of total debt excluding student loans, and we indeed find a contemporaneous correlation of about 0.9 between national home prices and the average size of debt changes among increasers over the entire time span.¹⁷ Justiniano et al. (2013) propose a DSGE model in which house prices drive a collateral channel; e.g., as house prices increase, collateral values appreciate, facilitating more borrowing. Over the entire cycle, the micro data are broadly consistent with this collateral channel among the increasers.

However, the very strong correlation between the series breaks down around the financial crisis. We estimate a model for the average size of debt changes among the increasers $s_{i,t}$:

$$(15) \quad s_{i,t} = -472.41 + 11.29^* (\text{house prices}_t) + 0.64^* s_{i,t-1} + e_t, \quad \bar{R}^2 = 0.94, \quad F = 265.19$$

(347.57)
(3.33)
(0.10)

using the mean non-seasonally adjusted sales prices for all existing home sales (house prices_t, expressed in thousands of dollars) from the National Association of Realtors and data from 2000Q1 to 2008Q3 (with HAC standard errors in parentheses). Dynamically simulating this model from 2008Q4 forward by feeding in the actual path of home sales price realizations, we

¹⁷ As a measure of transaction prices, we use the mean non-seasonally adjusted sales price for all existing home sales from the National Association of Realtors; the correlation between this measure and our size variable is 0.85 between 2000Q1 and 2013Q4. However, our basic results are unchanged using alternative measures of home prices: the correlation is 0.92 with the seasonally adjusted quarterly S&P/Case-Shiller U.S. national home price index between 2000Q1 and 2013Q3, and the correlation is 0.88 with the non-seasonally adjusted CoreLogic national house price index between 2000Q1 and 2013Q4.

find that the average size of debt changes among the increasers (the blue line in Figure 5) fell by more than would have been predicted by the decline in home prices and remained low even when home prices rebounded (the dashed line in Figure 5).¹⁸ Thus, movements in collateral prices alone cannot explain the post-crisis period, which may help to explain why the preferred model simulations in Justiniano et al. (2013) fail to exhibit a binding zero lower bound on short-term nominal interest rates as in the data. The post-crisis period provides additional scope for frictions such as increases in collateral requirements and lower leverage ratios, as in the leverage cycle proposed by Geanakoplos (2010).

VIII. Conclusion

Recent work has emphasized the role of consumer debt in both the housing bubble and bust and stressed the broad implications of consumer debt for macroeconomic performance. We present a new framework for decomposing consumer debt dynamics using data at the micro level by separating the actions of increasers and decreasers, and we then further decompose movements into percentage and size margins among the increasers and decreasers.

We find this framework useful for informing macroeconomic models that feature a role for consumer debt. First, we show that the actions of the increasers dominate consumer debt dynamics: the variations in borrowing activity among the increasers explain four times as much of the total variation in consumer debt as variations among the decreasers. This result holds even for the recent period of consumer deleveraging, which rather than being dominated by consumer defaults might better be termed a period of “non-leveraging.” Second, we provide qualitative and quantitative evidence on time variation in the percentage of consumers taking on additional

¹⁸ We observe similar patterns using other house price measures as well.

debt. We view the qualitative evidence of a sharp decline in the percentage of consumers taking on additional debt during the financial crisis and the ongoing low levels thereafter as consistent with a binding zero lower bound on nominal interest rates that has prevented less leveraged consumers—including those in the perceived safest end of the distribution of risk scores—from stepping up borrowing to offset deleveraging on the part of others. The ability of macroeconomic models featuring a consumer debt channel to match these features is an open question; for example, we show that the quantitative evidence on time variation in the percentage of increasers is at odds with a number of such models.

Finally, we document a key cycle in the average size of debt changes among the increasers. Our variance decompositions reveal that while variations in the percentage of increasers are important in explaining consumer debt dynamics, movements in the average size of debt changes among the increasers explain about twice as much of the total variation in consumer debt changes. Home prices primarily drove one part of this cycle, the run-up prior to the crisis. But modeling the crisis and post-crisis period requires additional financial frictions that limited consumer leveraging and caused the strong relationship between house prices and the average size of debt changes among the increasers to break down. The extent to which these frictions interacted with—or were responsible for—the decline in the percentage of increasers at this same time is a key topic for future research.

IX. References

Bhutta, Neil (2012) “Mortgage Debt and Household Deleveraging: Accounting for the Decline in Mortgage Debt Using Consumer Credit Record Data” Finance and Economics Discussion Series 2012-14, Divisions of Research & Statistics and Monetary Affairs, Federal Reserve Board.

Dynan, Karen (2012) “Is a Household Debt Overhang Holding Back Consumption?” *Brookings Papers on Economic Activity* Spring: 299–344.

Eggertsson, Gauti B. and Paul Krugman (2012) “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach” *Quarterly Journal of Economics* 127(3): 1469–1513.

Favara, Giovanni and Jean Imbs (2012) “Credit Supply and the Price of Housing” unpublished manuscript.

Geanakoplos, John (2010) “The Leverage Cycle” In Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds. *NBER Macroeconomics Annual 2009* Volume 24:1–65.

Gross, David B and Nicholas S. Souleles (2002) “Do Liquidity Constraints and Interest Rates matter for Consumer Behavior? Evidence from Credit Card Data” *Quarterly Journal of Economics* 117(1): 149–185.

Guerrieri, Veronica and Guido Lorenzoni (2011) “Credit Crises, Precautionary Savings, and the Liquidity Trap” NBER Working Paper 17583.

Hall, Robert E. (2011) “The Long Slump” *American Economic Review* 101(2): 431–69.

Haughwout, Andrew, Donghoon Lee, Joelle Scally, and Wilbert van der Klaauw (2012) “Has Household Deleveraging Continued?” Federal Reserve Bank of New York Liberty Street Economics Blog, <http://libertystreeteconomics.newyorkfed.org/2012/08/just-released-household-debt-and-credit.html>.

Johnson, Kathleen and Geng Li (2010) “The Debt-Payment-to-Income Ratio as an Indicator of Borrowing Constraints: Evidence from Two Household Surveys” *Journal of Money, Credit and Banking* 42(7): 1373–90.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti (2013) “Household Leveraging and Deleveraging” unpublished manuscript.

Klenow, Peter J. and Oleksiy Kryvtsov (2008) “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?” *Quarterly Journal of Economics* 123(3): 863–904.

Knotek, Edward S., II, and John Carter Braxton (2012) “What Drives Consumer Debt Dynamics?” Federal Reserve Bank of Kansas City *Economic Review* fourth quarter: 31–54

Lee, Donghoon and Wilbert van der Klaauw (2010) “An Introduction to the FRBNY Consumer Credit Panel” Federal Reserve Bank of New York Staff Report no. 479.

Mian, Atif and Amir Sufi (2009) “The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis” *Quarterly Journal of Economics* 124(4): 1449–96.

Mian, Atif and Amir Sufi (2010) “Household Leverage and the Recession of 2007-09” *IMF Economic Review* 58(1): 74–117.

Mian, Atif and Amir Sufi (2011a) “Consumers and the Economy, Part II: Household Debt and the Weak U.S. Recovery” Federal Reserve Bank of San Francisco *Economic Letter* 2011-02.

Mian, Atif and Amir Sufi (2011b) “House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis” *American Economic Review* 101(5): 2132–56.

Mian, Atif, Kamalesh Rao, and Amir Sufi (2013) “Household Balance Sheets, Consumption, and the Economic Slump” *Quarterly Journal of Economics* 128(4): 1684–1726.

Midrigan, Virgiliu and Thomas Philippon (2011) “Household Leverage and the Recession” NBER Working Paper #16965.

Zeldes, Stephen (1989) “Consumption and Liquidity Constraints: An Empirical Investigation” *Journal of Political Economy* 97(2): 305–46.

Table 1: Correlations with Aggregate Debt Movements

	Change in debt per person (Δ_t)	Growth of consumer debt (g_t)
Contributions from increasers	0.94	0.84
Percentage of consumers who are increasers ($p_{I,t}$)		0.94
Average size of changes among increasers ($s_{I,t}$)	0.72	0.54
Contributions from no-changers
Percentage of consumers who are no-changers ($p_{NC,t}$)		-0.86
Contributions from decreaseers	0.71	0.81
Percentage of consumers who are decreaseers ($p_{D,t}$)		-0.93
Average size of changes among decreaseers ($s_{D,t}$)	0.52	0.64

Notes: By construction, the debt change contributions from the no-changers are zero and are uncorrelated with movements in aggregate debt. All statistics are computed using four-quarter windows, as described in the text. The sample runs from 2000Q1 through 2013Q4.

Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

Table 2: Correlations with Aggregate Debt Movements using Detailed Breakdowns

	Change in debt per person (Δ_t)	Growth of consumer debt (g_t)	
Contributions from increasers, zero to positive debt	0.96		0.91
Percentage, zero to positive ($p_{I,t}^{0 \rightarrow +}$)		0.44	0.59
Average size, zero to positive ($s_{I,t}^{0 \rightarrow +}$)		0.59	0.42
Contributions from increasers, positive to greater positive debt	0.94		0.84
Percentage, positive to greater positive ($p_{I,t}^{+ \rightarrow ++}$)		0.89	0.95
Average size, positive to greater positive ($s_{I,t}^{+ \rightarrow ++}$)		0.72	0.55
Contributions from no-changers
Percentage, zero to zero ($p_{NC,t}^{0 \rightarrow 0}$)		-0.77	-0.86
Percentage, positive to same positive ($p_{NC,t}^{+ \rightarrow +}$)		-0.72	-0.72
Contributions from decreaseers, positive to zero debt	0.88		0.91
Percentage, positive to zero ($p_{D,t}^{+ \rightarrow 0}$)		0.12	0.21
Average size, positive to zero ($s_{D,t}^{+ \rightarrow 0}$)		0.86	0.90
Contributions from decreaseers, positive to smaller positive debt	0.66		0.77
Percentage, positive to smaller positive ($p_{D,t}^{+ \rightarrow +}$)		-0.87	-0.94
Average size, positive to smaller positive ($s_{D,t}^{+ \rightarrow +}$)		0.38	0.51

Notes: By construction, the debt change contributions from the no-changers are zero and are uncorrelated with movements in aggregate debt. All statistics are computed using four-quarter windows, as described in the text. The sample runs from 2000Q1 through 2013Q4.

Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

Table 3: Adjusted R-Squared Comparisons

	Dependent variable:	
	Change in debt per person (Δ_t)	Growth of consumer debt (g_t)
Percentage of consumers who are increasers ($p_{I,t}$)	0.740	0.877
Percentage, zero to positive ($p_{I,t}^{0 \rightarrow +}$)	0.181	0.334
Percentage, positive to greater positive ($p_{I,t}^{+ \rightarrow ++}$)	0.787	0.904
Percentage of consumers who are no-changers ($p_{NC,t}$)	0.593	0.732
Percentage, zero to zero ($p_{NC,t}^{0 \rightarrow 0}$)	0.585	0.743
Percentage, positive to same positive ($p_{NC,t}^{+ \rightarrow +}$)	0.504	0.510
Percentage of consumers who are decreasers ($p_{D,t}$)	0.757	0.869
Percentage, positive to zero ($p_{D,t}^{+ \rightarrow 0}$)	-0.005	0.029
Percentage, positive to smaller positive ($p_{D,t}^{++ \rightarrow +}$)	0.756	0.884
Average size of changes among increasers ($s_{I,t}$)	0.505	0.283
Average size, zero to positive ($s_{I,t}^{0 \rightarrow +}$)	0.333	0.163
Average size, positive to greater positive ($s_{I,t}^{+ \rightarrow ++}$)	0.510	0.287
Average size of changes among decreasers ($s_{D,t}$)	0.254	0.398
Average size, positive to zero ($s_{D,t}^{+ \rightarrow 0}$)	0.729	0.814
Average size, positive to smaller positive ($s_{D,t}^{++ \rightarrow +}$)	0.129	0.250

Notes: Regressions take the form $y_t = \alpha + \beta x_t + e_t$, with dependent variable y given by either the change in debt per person or the growth of consumer debt, and the explanatory variable x given by the variable in the first column. The sample runs from 2000Q1 through 2013Q4.

Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

Table 4: Variance Decompositions

Explained by variance of:	Variance share of change in debt per person (Δ_t)			
	(1)	(2)	(3)	(4)
Contributions from increasers	0.61			
Percentage of consumers who are increasers ($p_{I,t}$)		0.16		
Average size of changes among increasers ($s_{I,t}$)		0.32		
Contributions from increasers, zero to positive			0.00	
Percentage, zero to positive ($p_{I,t}^{0 \rightarrow +}$)				0.00
Average size, zero to positive ($s_{I,t}^{0 \rightarrow +}$)				0.00
Contributions from increasers, positive to greater positive			0.56	
Percentage, positive to greater positive ($p_{I,t}^{+ \rightarrow ++}$)				0.15
Average size, positive to greater positive ($s_{I,t}^{+ \rightarrow ++}$)				0.27
Contributions from decreaseers	0.14			
Percentage of consumers who are decreaseers ($p_{D,t}$)		0.02		
Average size of changes among decreaseers ($s_{D,t}$)		0.08		
Contributions from decreaseers, positive to zero			0.00	
Percentage, positive to zero ($p_{D,t}^{+ \rightarrow 0}$)				0.00
Average size, positive to zero ($s_{D,t}^{+ \rightarrow 0}$)				0.00
Contributions from decreaseers, positive to smaller positive			0.10	
Percentage, positive to smaller positive ($p_{D,t}^{+ \rightarrow +}$)				0.02
Average size, positive to smaller positive ($s_{D,t}^{+ \rightarrow +}$)				0.05

Notes: The remainder of the total variance within a given column is explained by covariance terms that are not reported. The sample runs from 2000Q1 through 2013Q4.

Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

Table 5: Alternative Variance Decompositions Using Debt Growth

Explained by variance of:	Variance share of growth of consumer debt (g_t)			
	(1)	(2)	(3)	(4)
Contributions from increasers	1.26			
Share of debt owed by increasers ($\theta_{I,t}$)		0.30		
Debt growth rate among increasers ($\gamma_{I,t}$)		0.50		
Contributions from increasers, zero to positive ($\tau_{I,t}^{0 \rightarrow +}$)			0.00	0.00
Contributions from increasers, positive to greater positive			1.12	
Share of debt owed, positive to greater positive ($\theta_{I,t}^{+ \rightarrow ++}$)				0.26
Debt growth rate, positive to greater positive ($\gamma_{I,t}^{+ \rightarrow ++}$)				0.44
Contributions from decreaseers	0.08			
Share of debt owed by decreaseers ($\theta_{D,t}$)		0.05		
Debt growth rate among decreaseers ($\gamma_{D,t}$)		0.15		
Contributions from decreaseers, positive to zero ($\theta_{D,t}^{+ \rightarrow 0}$)			0.00	0.00
Contributions from decreaseers, positive to smaller positive			0.07	
Share of debt owed, positive to smaller positive ($\theta_{D,t}^{++ \rightarrow +}$)				0.04
Debt growth rate, positive to smaller positive ($\gamma_{D,t}^{++ \rightarrow +}$)				0.14

Notes: The remainder of the total variance within a given column is explained by covariance terms that are not reported. The sample runs from 2000Q1 through 2013Q4.

Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

Table 6: Alternative Variance Decompositions: Post-Housing Bubble and Non-Housing Debt

Variance share explained by variance of:	Post-Housing Bubble				Non-Housing Debt			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Contributions from increasers	0.63				0.83			
Percentage of consumers who are increasers ($p_{I,t}$)		0.08				0.29		
Average size of changes among increasers ($s_{I,t}$)		0.22				0.24		
Contributions from increasers, zero to positive			0.00				0.00	
Percentage, zero to positive ($p_{I,t}^{0 \rightarrow +}$)				0.00				0.00
Average size, zero to positive ($s_{I,t}^{0 \rightarrow +}$)				0.00				0.00
Contributions from increasers, positive to greater positive			0.57				0.73	
Percentage, positive to greater positive ($p_{I,t}^{+ \rightarrow ++}$)				0.08				0.23
Average size, positive to greater positive ($s_{I,t}^{+ \rightarrow ++}$)				0.19				0.20
Contributions from decreaseers	0.06				0.24			
Percentage of consumers who are decreaseers ($p_{D,t}$)		0.01				0.04		
Average size of changes among decreaseers ($s_{D,t}$)		0.03				0.17		
Contributions from decreaseers, positive to zero			0.00				0.01	
Percentage, positive to zero ($p_{D,t}^{+ \rightarrow 0}$)				0.00				0.00
Average size, positive to zero ($s_{D,t}^{+ \rightarrow 0}$)				0.00				0.00
Contributions from decreaseers, positive to smaller positive			0.04				0.21	
Percentage, positive to smaller positive ($p_{D,t}^{+ \rightarrow +}$)				0.01				0.04
Average size, positive to smaller positive ($s_{D,t}^{+ \rightarrow +}$)				0.02				0.17

Notes: Numbers reflect the variance shares of the change in debt per person (Δ). The remainder of the total variance within a given column is explained by covariance terms that are not reported. The post-housing bubble sample runs from 2007Q1 through 2013Q4. The non-housing debt sample runs from 2000Q1 through 2013Q4 and excludes mortgage and home equity debt.

Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

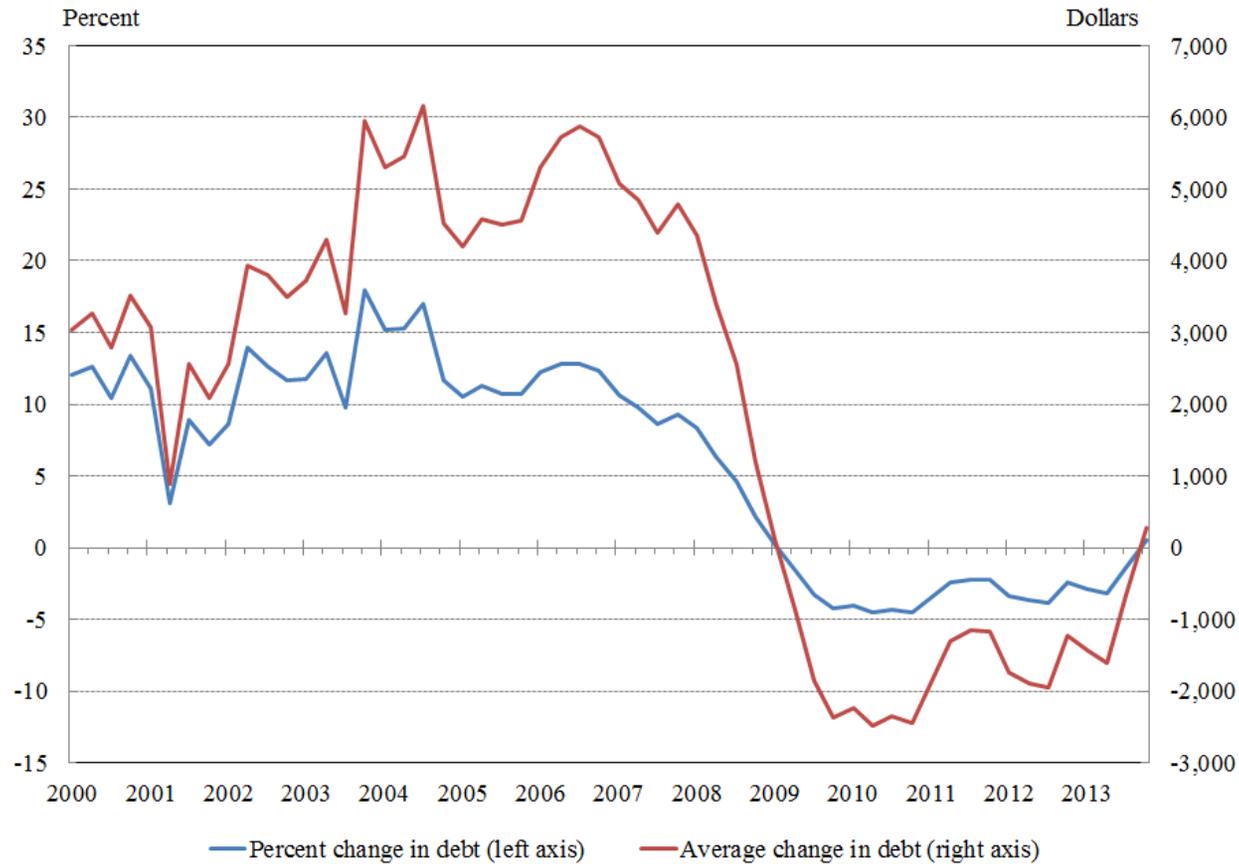
Table 7: Variance Decompositions Based on Risk Quartiles

Explained by variance of:	Variance share of change in debt per person (Δ_t)		
	(1)	(2)	(3)
1st quartile contributions	0.09		
Increases		0.01	
Percentage of increasers			0.01
Decreasers		0.04	
Average size among decreaseers			0.03
2nd quartile contributions	0.09		
Increases		0.09	
Percentage of increasers			0.02
Average size among increasers			0.05
Decreasers		0.00	
3rd quartile contributions	0.07		
Increases		0.07	
Percentage of increasers			0.01
Average size among increasers			0.04
Decreasers		0.01	
4th quartile contributions	0.04		
Increases		0.02	
Average size among increasers			0.01
Decreasers		0.01	
Average size among increasers			0.01

Notes: The remainder of the total variance within a given column is explained by covariance terms that are not reported. Variance terms explaining less than 1% of the total variance in the decomposition in column (3) were omitted from the table. The sample runs from 2000Q1 through 2013Q4.

Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

Figure 1: Growth and Decline in Consumer Debt



Notes: The percent change in debt and the average change in debt are reported over a four-quarter window. Both measures only include consumers who were in the panel at the beginning and end of each window.

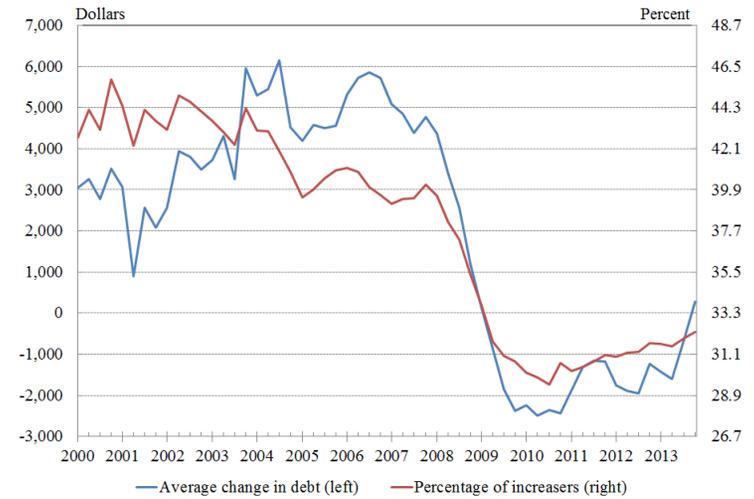
Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

Figure 2: Percentage and Size Measures and Aggregate Consumer Debt Growth

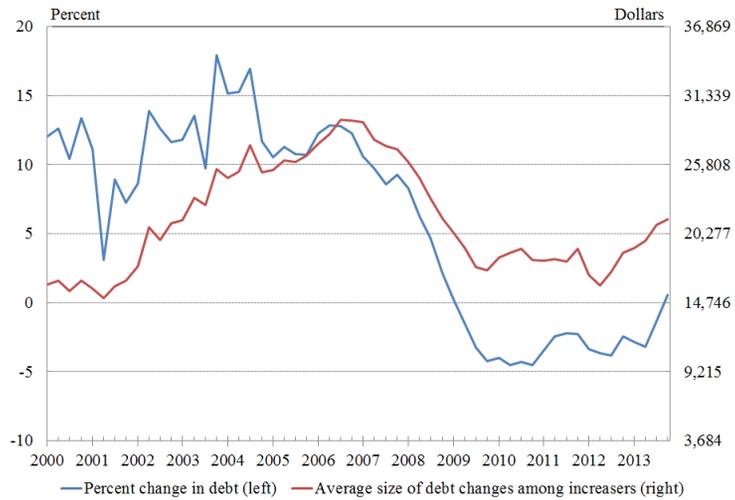
(a) Percent change in debt and percentage of increasers



(b) Average change in debt and percentage of increasers



(c) Percent change in debt and average size among increasers



(d) Average change in debt and average size among increasers

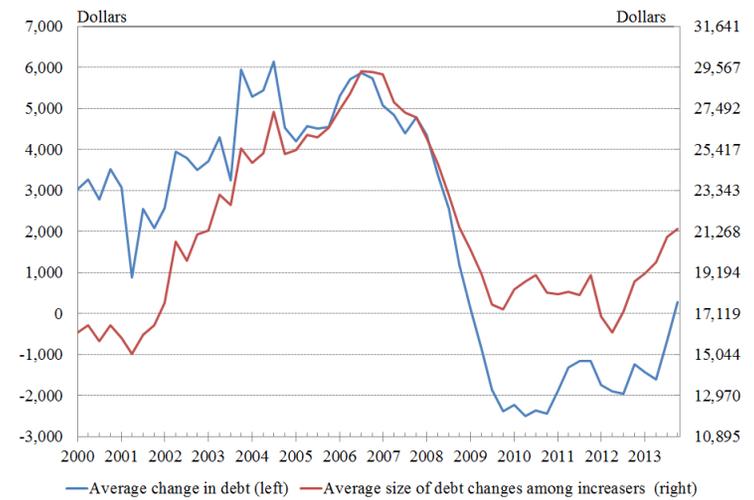
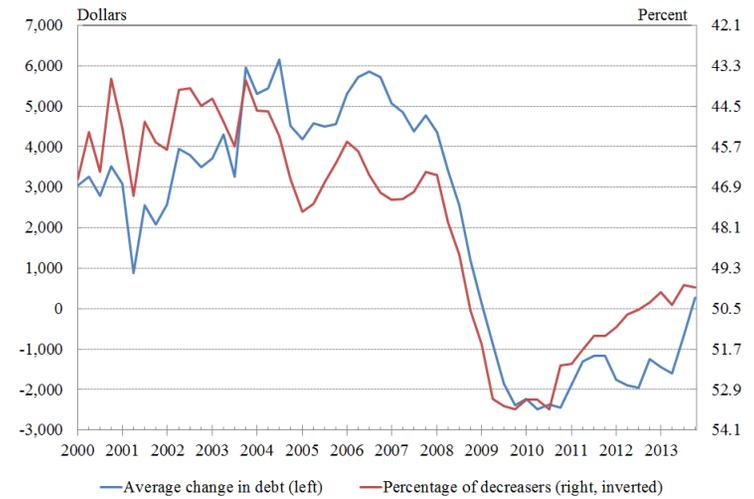


Figure 2 (continued): Percentage and Size Measures and Aggregate Consumer Debt Growth

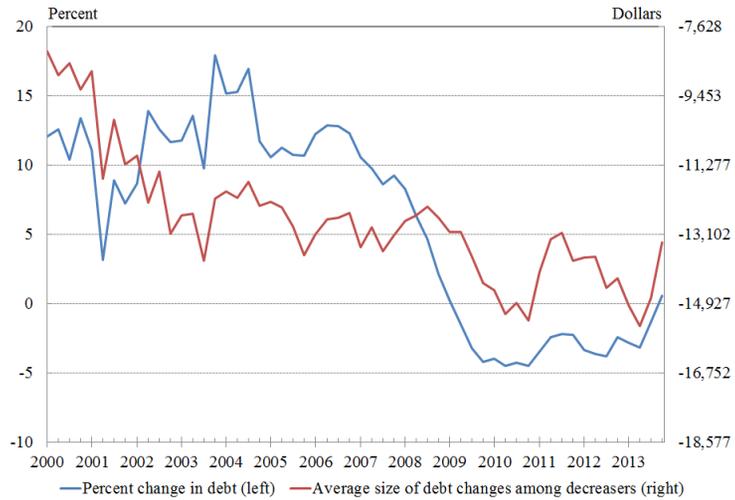
(e) Percent change in debt and percentage of debtors



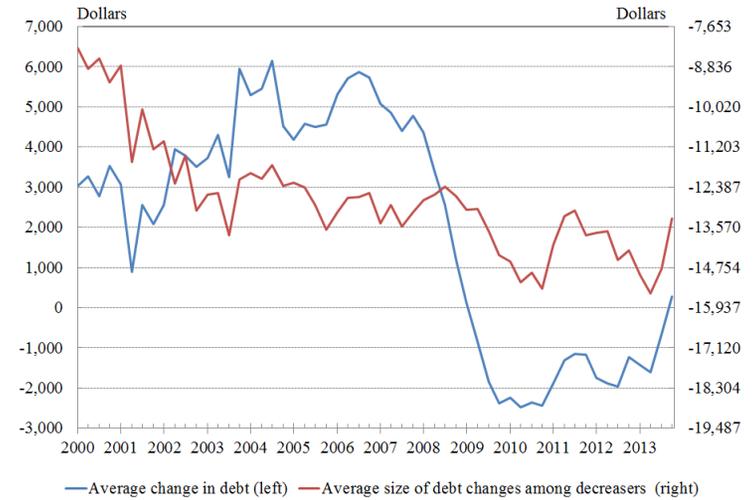
(f) Average change in debt and percentage of debtors



(g) Percent change in debt and average size among debtors



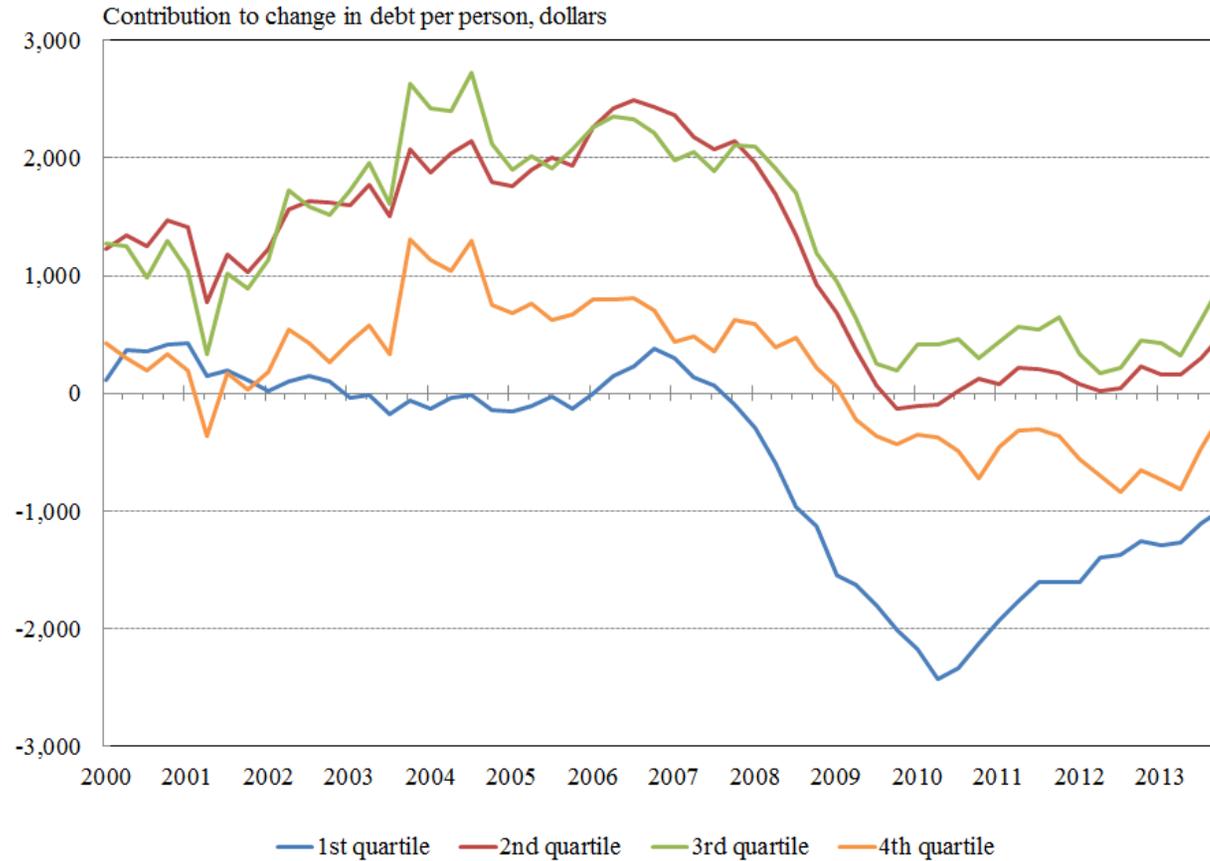
(h) Average change in debt and average size among debtors



Notes to Figure 2: All measures are reported over a four-quarter window and only include consumers who were in the panel at the beginning and end of each window.

Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

Figure 3: Contributions to Changes in Consumer Debt, By Risk Score Quartile

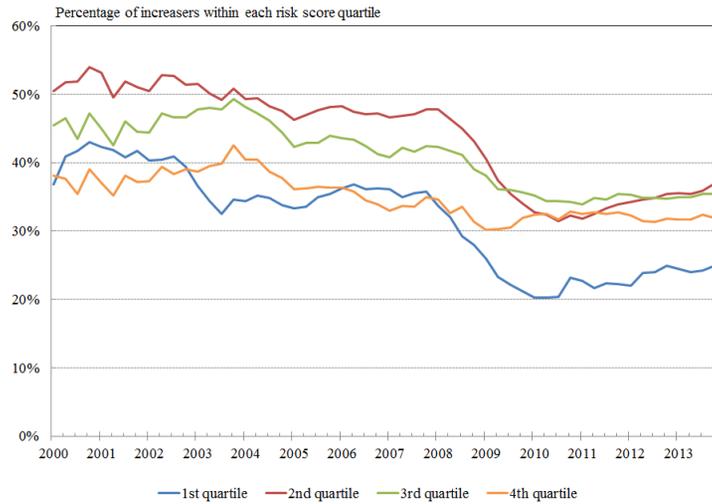


Notes: Summing across the quartiles' contributions produces the change in debt per person at a given point in time. Consumers with the lowest risk scores are in the first quartile; they are perceived as the riskiest credits by Equifax. Consumers with the highest risk scores are in the fourth quartile and are perceived as the safest credits by Equifax. The changes in debt per person are reported over a four-quarter window and only include consumers who were in the panel at the beginning and end of each window. Consumers are sorted by risk score quartile at the beginning of each window.

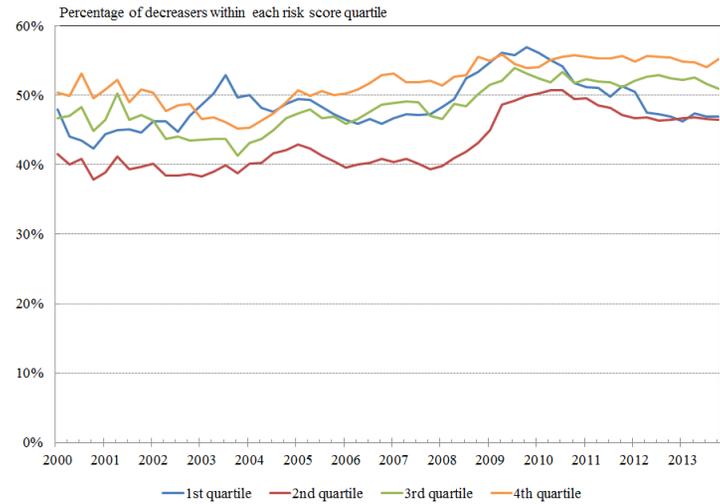
Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

Figure 4: Increasers and Decreasers, By Risk Score Quartile

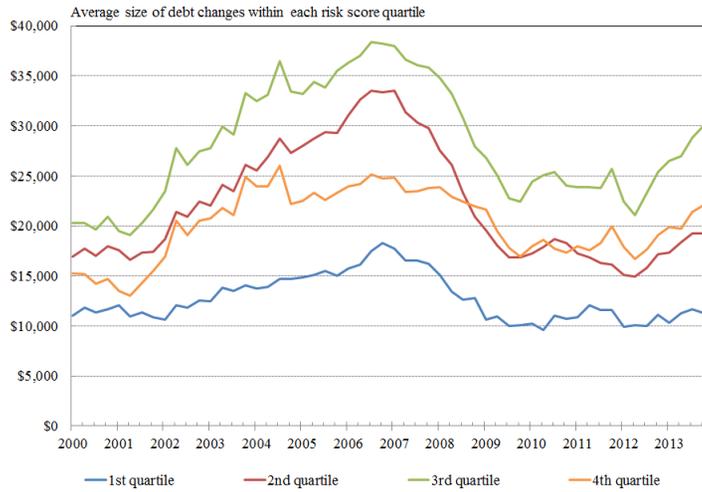
(a) Percentage of increasers



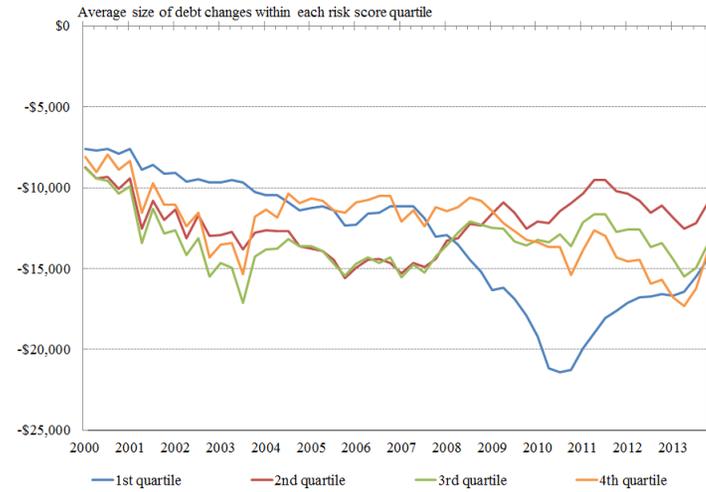
(b) Percentage of decreasees



(c) Average size among increasers

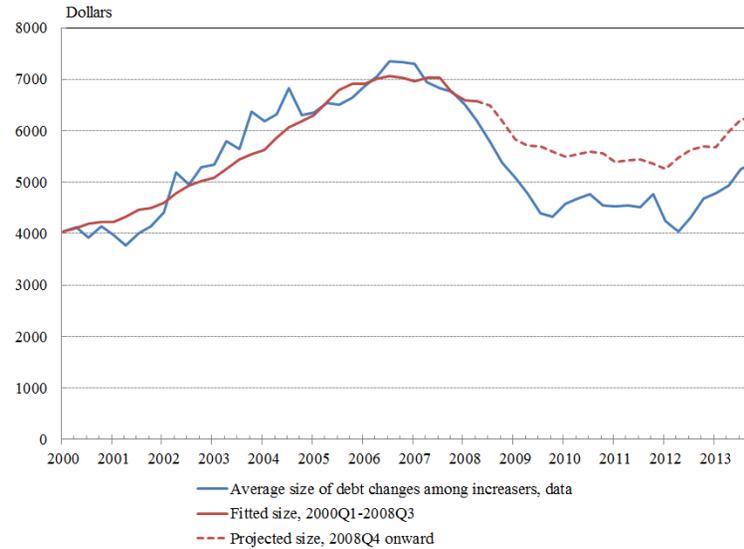


(d) Average size among decreasees



Notes to Figure 4: Consumers with the lowest risk scores are in the first quartile; they are perceived as the riskiest credits by Equifax. Consumers with the highest risk scores are in the fourth quartile and are perceived as the safest credits by Equifax. All measures are reported over a four-quarter window and only include consumers who were in the panel at the beginning and end of each window. Consumers are sorted by risk score quartile at the beginning of each window.
Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax and authors' calculations.

Figure 5: The Average Size of Debt Changes among Increases and House Price Dynamics



Notes: The average size of debt changes among increasers is computed based on consumer debt data. The fitted size line is based on estimating the empirical model in equation (15) over the period 2000Q1 through 2008Q3. The projected size line comes from dynamically simulating the same fitted model using the actual path of mean existing home sales prices over the period 2008Q4-2013Q4.

Sources: Federal Reserve Bank of New York Consumer Credit Panel/Equifax, National Association of Realtors, and authors' calculations.