Identifying "Tipping Points" in Consumer Liabilities Using High Frequency Data*

Carlos Garriga
Federal Reserve Bank of St. Louis

Lowell R. Ricketts
Federal Reserve Bank of St. Louis

Don E. Schlagenhauf
Federal Reserve Bank of St. Louis

October 19, 2016

Abstract

The concept of a "tipping point" captures the idea of a threshold where debt is no longer sustainable and becomes a drag on household financial well-being. In the most severe cases, a "tipping point" represents the point at which a household experiences a debt default. This paper constructs different measures based on the dynamics of the monthly debt payment to after-tax income ratio. The preliminary examination using data from the Credit Risk Insight Servicing McDash (CRISM) dataset suggests that some of these measures have some predictive content when compared to alternative risk measures such as a FICO score.

*The views expressed are those of the authors and not necessarily of the Federal Reserve Bank of St. Louis or the Federal Reserve System. This is a preliminary draft, content may change in upcoming revisions.

1 Introduction

In a period where the Federal Reserve bank is debating normalizing monetary policy by raising interest rates, the balance sheet of the average American household is vulnerable in light of recent behavior in credit markets. While the typical household in the U.S. is moving away from the housing market, represented by a declining homeownership rate since 2006, and diving into the credit card market. In the second quarter of 2016, U.S. consumers increased their credit card balances by record high setting of $34.4 billions. This has continued the growing trend of total outstanding credit card balances that started in 2011 that is likely to achieve the $1 trillion mark observed in 2008 by the end of 2016.

In light of this trend, the obvious question is whether this level of borrowing could expose the U.S. economy to another major deleveraging period if growth slows and the labor market deteriorates. This could push households past the tipping point towards delinquencies, defaults, and adjustments of the household balance sheet. Identifying a priori the level of these tipping points where an individual's debt is no longer sustainable, given their current income levels and cost of borrowing, could inform monitoring programs to avoid a debt default outcome. By anticipating a household debt crisis, one can mitigate or eliminate some of the consequences that helped cause the Great Recession.

This paper examines whether the concept of a "tipping point" can be quantified and implemented as a predictor of future default in the credit market. While this notion is often used at the aggregate level, our interest is on how an economist defines this concept at an individual's level. At the micro level whether an individual reaches this point will depend on the size of their credit balances, the cost of borrowing (i.e. the interest rate and the amortization structure of the loan), the disposable income, and liquid net worth from other assets or financial investments. All these different ingredients for an individual are related through the household's budget constraint. Using this constraint, it is possible to derive a monthly debt payment-to-after-tax income ratio that measures this turning point. Our analysis indicates that a "tipping point" depends on an individual's entire balance sheet. When an individual cannot make payment on debt claims, the "tipping point" has been reached.

Armed with this measure, we turn to the data to see if our measure could be implemented and whether it could be a leading indicator for when an individual's debt level is becoming unstable. Since the focus of the paper is on an individual's debt "tipping point," we argue that data at the individual level is appropriate. More importantly, the monitoring goal requires a panel data set at a quarterly or monthly frequency. As a result, we employ a data set that is a combination of Equifax credit bureau data merged with mortgage-level McDash servicing data. The merged data set allows an individual's total leverage

to be monitored at the monthly frequency. Borrower credit information is included for the life of the mortgage loan and starts with six months prior to origination and ends six months following termination. A termination is defined as either paying off the obligation or defaulting on the obligation. The credit history of each borrower is reported. An interesting feature of this data set is that estimated income is provided monthly for each individual. This data allows us to create an individual's monthly debt to income measure, which is suggested by our budget constraint analysis. The preliminary examination of this measure based on the dynamics of the monthly debt payment to income ratio suggests that this metric has some predictive content when compared to alternative metrics based on the dynamics of the FICO score.

The use of the Equifax data to study an individual's behavior with respect to debt is becoming more common. This data is used to empirically examine the behavior of consumer credit as Brown, M., J. Wen and B. Zafar (2016), but also across different debt categories like mortgage in the case of Adelino, Gerardi, and Willen (2013), Brown, Stein, and Zafar (2015), DiMaggio, Kermani, and Ramcharan (2015), Yu and Gerardi () and Lambie-Hanson and Willen(2013), or student debt as in Bricker, Brown, Hannan, and Pence (2015). There is also a growing literature using this empirical micro data to construct structural models of borrowing behavior and the determination of interest rates. Some close references include Low, Meghir, and Pistaferri (2010), Chatterjee and Gordon (2012), Benjamin and Mateos-Planas (2014), Athreya, Tam, and Young (2015), Mitman (2015), and Athreya, Sánchez, Tam, and Young (2016).

The CRISM data contain consumer-level observations on monthly debt payment as well as estimates of personal income. Hence, sufficient data are available to create our suggested "tipping point" measure for individuals over time. More importantly, data are potentially available monthly since 2005. The frequency and time series component is important to evaluating our metric as an advanced signal of severe financial distress, which we define as a bankruptcy and/or a foreclosure. We find that our metric does provide an advanced signal and we assess the quality of that signal using tradional forecasting evaluation methods.

The paper is organized into four major sections. In the first section, we take an economics approach to identify how a "tipping point" may be measured. The second section discusses the CRISM data used as well as restrictions that were imposed. In the third section, we examine the data by income ranking and age for non-defaulters and defaulters. Our point of comparison is December 2007 and December 2014. The fourth section describes how the monthly payment to income ratio behaves prior to and immediately after a financial crisis. This analysis suggests our measure can be a leading indicator. Our initial conclusion is based on "in-sample" analysis, as it is conditional on a known bankruptcy and or foreclosure,

therefore an out-of-sample prediction test is required for robustness. We conclude with psuedo out of sample prediction in a larger sample drawn without the conditional restriction of a financial default. Here, further evidence is presented that is favorable to our measure.

2 What is a "tipping point"?

The increase in household debt has generated ample discussion of the implications of this pattern for household financial stability and the possible adverse consequences for an economy. This discussion requires some metric so that household leverage can be quantified and compared. One measure is the value of household debt to household income. Another possible measure is to look at a household's monthly debt payments and compare this to monthly income. Other measures certainly exist. In this section, we will offer a measure based in the field of economics. Starting from a household's budget constraint, we explore whether a measure of a debt "tipping point" can be determined.

2.1 Simple Measure

We will start this exercise by considering a simple example. We assume that an individual does not have any assets, but has access to credit markets in the form of borrowing. In any month, this individual uses income to buy consumption goods such as food which we denote consumption as C_t as well as make payments on the amount of debt, D_t. This payment is denoted as $r_t D_t$. In order to make such payments, has to earn income, Y_t, perhaps through the labor market, and pay taxes, T_t. The after-tax income for this person is $(Y_t - T_t)$. In addition to this source of revenue, an individual can borrow additional funds. Of course, in the future period, interest payments of this debt will increase. The additional source of funds from borrowing is denoted by ΔD_{t+1}. The budget constraint for this individual is simply

$$C_t + rD_t = (Y_t - T_t) + \Delta D_{t+1}$$

This equation tells us that consumption and debt obligations can be met as long as after-tax income and new debt funding can be obtained, $\Delta D_{t+1} > 0$.

Suppose this individual cannot take on additional debt, $\Delta D_{t+1} = 0$. With this assumption, the individual's budget constraint reduces to:

$$C_t + r_t D_t \ = (Y_t - T_t)$$

If consumption spending is subtracted from both sides of the equation, and then both sides

of the equation are divided by after tax income, one finds

$$\frac{r_t D_t}{(Y_t - T_t)} = 1 - \frac{C_t}{(Y_t - T_t)}$$

What does this equation suggest about "tipping points"? The left hand side measures monthly debt payments relative to *after-tax household income.* This equation tells us that the metric to judge a household's financial stability is the ratio of total monthly (or quarterly) debt payments relative to after-tax monthly (or quarterly) income. If this household is to avoid violating it's budget constraint, the left hand side of the equation must be less than one minus the consumption to disposable income ratio.

2.2 General Measure

While this measure may be appropriate for individuals without assets, we should expand our example. Households have wealth or assets to draw on and this could influence the "tipping point." We will introduce the value of wealth by W_t. These are the asset choices made in the prior period and expressed in last periods prices. This wealth can be very liquid such as cash. Alternatively, wealth could be illiquid in the form of a house or an IRA. In order to convert these assets into cash, one may have to take a 'haircut' on this wealth. We will represent this haircut by the symbol γ. In period t, decisions must be made on the amount of wealth to have in the next period. These assets are purchased at period t prices. The amount of assets purchased in period t will be denoted as W_{t+1}. For this environment, an individual's budget constraint can be written as:

$$C_t + W_{t+1} + r_t D_t = (Y_t - T_t) + (1 - \gamma)W_t + \Delta D_{t+1}$$

This equation can be rewritten as:

$$\frac{r_t D_t}{(Y_t - T_t)} = 1 - \frac{C_t}{(Y_t - T_t)} + \frac{\Delta D_{t+1}}{(Y_t - T_t)} + \frac{(1 - \gamma)W_t - W_{t+1}}{(Y_t - T_t)}$$

This simple modification yields some interesting insights into a generalized measure of the "tipping point". The above equation indicates that the monthly debt payment to after-tax income measure continues to depend on the consumption to after-tax income ratio. However, this measure is also dependent on other factors. One is whether an individual has the ability to take on additional debt. Of course, this decision has implications for next period's debt payment to after-tax income ratio. A second important factor is changes in the wealth or asset position. In other words, the availability of wealth to finance debt obligations is important in the determination of a "tipping point."

The message from this analysis is that a measure of a household's "tipping point" can be couched in terms of a contemporaneous budget constraint using the notion of debt payment to after-tax income. However, the analysis also points out that the debt payment to after-tax measure depends on a individuals ability to access additional resources through their debt position or through the ability to adjust wealth positions. Simply stated, the state of the entire households liquid financial resources are critical. This idea of wealth liquidity is closely connected to the hand-to-mouth rich households discussed by Kaplan and Violante (2014).

3 The Data Set Employed

If the goal is to develop a measure of a "tipping point" at the individual level, we argue that micro-level data must be employed. We rely on a merge of Equifax credit bureau information with mortgage-level McDash servicing data titled Credit Risk Insight Servicing McDash (CRISM). CRISM data are updated monthly, with coverage beginning in June 2005. While the McDash side of the merge is focused on mortgage related debt, consumer debt in other areas is included at the individual level. While not exhaustive, tradeline data for auto, credit card, closed-end second accounts, first mortgages, consumer finance, HELOC, student, retail, and other types of debt are included. A main caveat of the data is that it is biased sample in the sense that all individuals own housing related debt.

The CRISM data do not report asset holdings for individuals, or actual reported income data. Our "tipping point" measure requires individual income data. Researchers have often dealt with this limitation by assigning income to an individual based on average income for a zip code or census code. While not a direct measurement of individual income CRISM does contain modeled personal income. Equifax has developed a new product, known as the *Personal Income Model*, that creates income estimates for individuals based on a national sample of employer provided current employment and income information. The actual procedures employed to generate these estimates are proprietary and not detailed.

The modeled income variable has limitations that are readily apparent. Actual individual income would reflect monthly changes in income from various earnings shocks. Herkenhoff (2012, 2015) has shown that shocks to a household, especially shocks to income, are a critical factor in a debt default outcome. While it is dynamic over time, the income measure we employ does not reflect such shocks. We did examine the properties of the "modeled" income measure we employ. In order to get some sense on how modeled income varies over time, we examined average modeled income in a variety of zip codes. In Figure 1, we present income between 2005 and 2015 for a non-random sample of locations in midwestern United States. The locations vary by wealth, and size. Clayton, Missouri is a wealthy suburb of St.

Louis. As can be seen, income for this zip code illustrates income variability. In contrast, Fergusion, Missouri is more ethnically and economically diverse area than Clayton. While Union, Missouri is in the western edge of the St. Louis MO-IL MSA. For these areas, we find less less variability in the mean income measure and less variability. Flint, Michigan is a relatively large community, with ethnic diversity and lower average income. The average of income for this zip code is relatively flat. The standard deviation of income for this location does indicate some variability. However, the Great Recession doesn't show up as we would expect. The lack of income variance is a potential issue with data. Despite this limitation and the lack of high frequency income variability (idiosyncratic shocks), we use the modeled income variable.

Figure 1 An Examination of Modeled Income

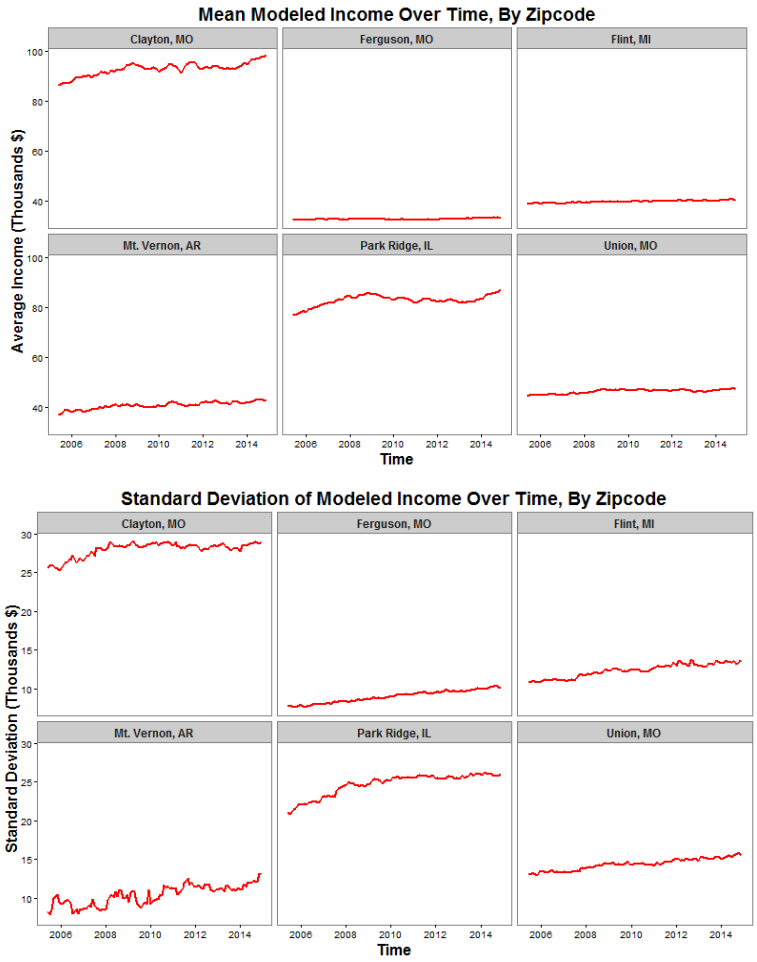

The findings discussed in this paper, with the exception of the final section, are based a three percent random sample of the CRISM dataset. Within that sample we restrict it further to individuals who had at least 30 months worth of observations in which their loans had a match confidence score of at least 0.8 (based on a scale of 0.0 to 9.0), were indicated

as present in the Federal Reserve Bank of New York/ Equifax Consumer Credit Panel at any point, were primary borrowers, had a non-missing value for modeled personal income, and had a non missing value for total debt payments. During the available time series for each individual they were also required to have any kind of bankruptcy with or without a single incidence of foreclosure OR a single incidence of foreclosure without a bankruptcy. Finally, there were instances where within the continuous time series for each individual where either the modeled personal income or total monthly debt payment were temporarily coded as missing. To ensure a continuous time series, the observations for those individuals were dropped from the sample. Altogether our default sample has 8,963,862 monthly observations distributed unevenly across 99,535 consumers. The start date of an individual's bankruptcy is provided in the data, with adjustments made where needed.[1] A foreclosure event was identified by the date when a individual's number of foreclosures changed from zero to one for the first time. If an individual had a bankruptcy or foreclosure on their record prior to 2006, then this individual was deleted from the sample as our focus is on the first default.

4 Summary Statistics in December 2007 and 2014.

In order to get a sense of the data, we examine the panel data set in December 2007 and then again in December 2014. December 2007 preceded the start of the Great Recession. The data should indicate large debt holdings and early signs of the debt and foreclosure crisis. December 2014 is the most recent data and reflects debt levels after an extensive period of deleveraging.[2] We want to examine how consumer liabilities have changed over these two periods among different age groups and income quartiles. Lastly, we want to see how two measures of individual default risk - the FICO score and the ratio of monthly debt payment to income - behave for nondefaulters and defaulters.

We begin our analysis by examining the extensive margin in December 2007. As can be seen, over eighty percent of individuals in the sample hold mortgage debt. The large ownership rate reflects the merge with McDash mortgage servicing data. Among other types of debt, we find ownership shares were higher among individuals which experienced a default.

Our analysis begins by describing the characteristics of the groups in terms of default experience. In order to address this question, we calculated a variety of group averages. If individuals are compared in terms of personal income, individual who did not default had

[1] We checked to make sure that the bankruptcy or foreclosure date is consistent with the actual monthly default date. If an inconsistency was identified, an adjustment was made.

[2] December 2014 was chosen as the cutoff given that there is a lag of about one year in the match that Equifax creates. This is an intentional lag so as to ensure that Equifax has all records available before attempting to match with McDash data. This is meant to void false matches. Because of this, it is not possible to use CRISM to study recently originated loans.

higher income in 2007. Despite this income difference, the monthly debt payment of defaulters was approximately $500 higher. As a result, the monthly payment to income ratio was higher for individuals who defaulted. The fraction of monthly income devoted to debt obligations was approximately 60 percent for nondefaulters, the corresponding fraction for defaulters was approximately 90 percent. The FICO scores suggest defaulters were significantly more risky, which means a lower credit score. This fact is supported by the differences in serious delinquency.

In 2014, the average income of those individuals who did not default was somewhat higher. In contrast, the average income level of those individuals who ended up defaulting was $9,000 lower compared to 2007. We find the monthly debt amount falls for both defaulters and nondefaulters. Especially interesting is the large decline in the average monthly payment for defaulter. The monthly payment to income ratio falls for defaulters in 2014. Defaulters continue to be riskier as suggested by the FICO score.

Table 1: Debt Holder Basic Characteristics

	December 2007		**December 2014**	
Characteristic	**No Default**	**Default**	**No Default**	**Default**
Personal Income	$55,000	$46,000	$60,000	$37,000
Monthly Debt Payment	2,858	3,318	2,758	1,310
Payment/Income Ratio	0.589	0.865	0.552	0.425
No. Seriously Delinquent	3,104	13,369	4,097	8,722
FICO Score	703	628	710	617
Sample Size	14,962	43,040	12,632	24,743

Source: Equifax Credit Risk Insight Servicing McDash (CRISM)

When the data set is reexamined for December 2014, one finding stands out. The fraction of holders of debt by debt category declined appreciably. In 2007, 85 percent of individuals who defaulted had mortgage debt. By 2014, the percent of individuals who defaulted had declined to 37 percent. In 2014, the ownership rate of HELOC debt was high for both individuals who did not default. By 2014, very few individuals who defaulted held HELOC debt. In 2007, individuals who ended up defaulting had higher percentages of auto, credit card, installment, and retail debt than nondefaulters. By 2014, ownership of almost every type of debt–save for student debt–declined.

Table 2: Average Debt Holding for Defaulters and NonDefaulters

	Sample Holding Percentage			
Debt Product	**December 2007**		**December 2014**	
(average)	**No Default**	**Default**	**No Default**	**Default**
Mortgage(First)	88.9	85.3	88.6	37.2
Heloc	18.2	20.4	12.2	3.0
Auto	49.5	58.6	49.1	42.3
Credit Card	82.0	82.3	79.7	52.2
Student	14.0	17.0	17.9	19.1
Installment	12.8	17.7	6.7	3.7
Retail	29.5	38.2	33.7	28.0

Source: Equifax Credit Risk Insight Servicing McDash (CRISM)

While the extensive margin, or the ownership of various debt types is important, the intensive margin, or the amount of debt held, is of more interest. In Table 3, the average debt holding by type is presented in 2007 and 2014. In 2007, individuals who ended up defaulting held more debt on average than those who did not. This holds for both overall debt and the various types with the exception of student debt. For instance, average total debt for defaulters is \$300,754 in 2007, while average total debt for nondefaulters is \$252,072 in the same year. This greater intensive margin for defaulters holds across various types of debt, with the exception of student debt. For example, defaulters held on average approximately \$5,000 more Heloc debt, \$3,500 more auto debt, and \$4,000 more credit debt. This relationship reverses in 2014, in large part because many individuals in the default sample drop out or deleverage due to the default events they experience. Student debt levels remain higher for defaulters than nodefaulters, and the average is higher than in 2007. Given that

student debt is difficult to shed in bankruptcy this may explain those persistent balances.

Table 3: Average Debt Holding

	December 2007		December 2014	
Debt Type	**No Default**	**Default**	**No Default**	**Default**
Total Debt	252,072	300,754	239,033	89,057
Mortgage (First)	210,257	240,546	201,212	65,388
Heloc	11,854	16,688	8,660	1,915
Auto	9,162	12,610	9,938	8,585
Credit Card	7,175	11,452	6,764	2,783
Student Loan	4,174	3,660	6,467	6,903
Retail	383	684	524	431
Sample Size	14,962	43,040	12,632	24,743

Source: Equifax Credit Risk Insight Servicing McDash (CRISM)

We now turn to redistributional effects and look to concentrations of debt by income quartiles and age cohorts. Table 4 offers the average total debt balance by income groups. Debt rises monotonically with income levels for both defaulters and nondefaulters. Interestingly, around 64 percent of the default sample has income between $30,000 to $65,000 while 82 percent of the nondefault sample has income between $15,000 and $45,000. Thus, those individuals that ran into repayment difficulty actually had higher incomes. For each income range, except for the lowest, individuals which defaulted had higher average total debt than their nondefaulting peers. Average balances across groups in 2014 reveals the impact of the extended deleveraging period. Average debt levels for the three lowest income groups declined to the where they are lower than individuals who did not default. For the other income cohorts, defaulters had higher average debt levels. For instance, in the $65,001 to $100,000 income group, the average individual who defaulted held 24 percent higher level of debt.

Table 4: Average Debt, By Income Range

		December 2007		
Income Group	**Sample (%)**	**Default ($)**	**Sample (%)**	**No Default ($)**
Less than $15,000	1.0	74,600	1.0	94,165
$15,001- $30,000	10.0	99,657	30.0	94,932
$30,001- $45,000	31.0	187,857	52.0	148,919
$45,001- $65,000	33.0	374,765	13.0	222,335
$65,001- $100,000	20.0	750,422	4.0	410,549
Greater than $100,001	5.0	1,200,356	1.0	764,764
Total Sample	14,962		24,668	

		December 2014		
Income Group	**Sample (%)**	**Default ($)**	**Sample (%)**	**No Default ($)**
Less than $15,000	1.0	27,125	1.0	57,693
$15,001- $30,000	31.0	25,465	7.0	70,369
$30,001- $45,000	51.0	65,389	28.0	119,375
$45,„001- $65,000	13.0	228,651	34.0	192,849
$65,001- $100,000	4.0	431,377	25.0	347,323
Greater than 100,001	1.0	735,340	9.0	640,019
Total Sample	24,743		12,632	

Source: Equifax Credit Risk Insight Servicing McDash (CRISM)

Similar to previous distributions, Table 5 shows that in 2007 defaulters held more debt across all age groups. Over 40 percent of debt was held by the 36-50 year old age group. This is consistent with the view that housing debt for first time home owners played an important role in the Great Recession. In 2014, average total debt declined drastically over all age groups for defaulters. The average total debt continued to be highest in the 36-50 age cohort.

Table 5: Average Total Debt, By Age Group

	December 2007			
Age Group	**Sample (%)**	**Default ($)**	**Sample (%)**	**No Default ($)**
30-35	25.0	267,175	25.0	235,590
36-50	47.0	315,680	41.0	273,468
51-65	24.0	315,624	26.0	256,681
66+	4.0	259,803	8.0	178,530
Sample Size	42,791		14,874	
	December 2014			
Age Group	**Sample (%)**	**Default ($)**	**Sample (%)**	**No Default ($)**
30-35	8.0	81,458	16.0	206,258
36-50	45.0	93,479	40.0	265,262
51-65	37.0	88,640	32.0	237,938
66+	8.0	78,034	12.0	198,946
Sample Size	24,723		12,591	

Source: Equifax Credit Risk Insight Servicing

Table 6 shows average risk scores, both FICO and monthly debt payment to income ratio, for income groups. For each group, the FICO score was lower in 2007 compared to 2014, in general. Given a FICO score captures risk of loan delinquency, then there is some evidence that risky loans were given across all income groups in 2007. If the spread in income group credit scores are examined, it is difficult to argue a particular income group was riskier. In 2014, the FICO scores of nondefaulters increased slightly in each income group. For the defaulters, the corresponding scores improved in the two lowest income groups, but declined for the other income groups. It's not clear from these findings that lower income groups accounted for greater lending risk prior to the recession.

The other measure of delinquency risk is the average monthly debt payment to income ratio. In 2007, this ratio was higher for defaulters over every income group. In fact, the highest payment to income ratios occur in the highest income groups. By 2007, the debt payment to income ratio declined in every income group for both defaulters and nondefaulters. At the two lowest income groups, the ratio was lower for defaulters than nondefaulters. After the 30,001-45,000 income group, the payment/income ratio was increasing and greater for defaulters. This evidence suggests the payment/income ratio may have limitations as a

"tipping point" measure. This may be the case if the sensitivity to income shocks is a key

driver.

Table 6: Average Risk Scores By Income Group

		December 2007		
Income Group	**Sample (%)**	**Default**	**Sample (%)**	**No Default**
$15,000-30,000	0.18		0.10	
FICO		596		558
Payment/Income		0.24		0.52
30,001-45,000	0.41		0.31	
FICO		678		622
Payment/Income		0.36		0.48
45,001-65,000	0.28		0.33	
FICO		727		661
Payment/Income		0.62		0.49
65,001-100,000	0.12		0.21	
FICO		742		678
Payment/Income		0.83		0.57
100,001-above	0.02		0.06	
FICO		759		716
Payment/Income		0.77		0.70
Sample Size	42991		14,892	
		December 2014		
$15,000-30,000	0.30		0.07	
FICO		592		598
Payment/Income		0.24		0.52
30,001-45,000	0.52		0.26	
FICO		632		684
Payment/Income		0.36		0.48
45,001-65,000	0.13		0.34	
FICO		655		734
Payment/Income		0.62		0.49
65,001-100,000	0.04		0.25	
FICO		665		750
Payment/Income		0.83		0.57
100,001-above	0.01		0.08	
FICO		698		768
Payment/income	0	0.77		70.0
Sample Size	24,477		12,481	

Source: Equifax Credit Risk Insight Servicing McDash (CRISM)

5 The Monthly Debt Payment/Income Ratio as a Measure of a "Tipping Point."

The economic analysis of an individual's budget constraint suggests that the monthly debt payment to monthly (after tax) income is a potential measure of a "tipping point." That discussion also indicates that this measure should be conditioned on the ability to raise additional debt, as well as the availability of liquid assets. Unfortunately, measuring the diverse holdings of assets is beyong the scope of our data. As a result, we will focus on the monthly debt payment to the income ratio that is not conditional on other information. A financial crisis for an individual can take the form of a bankruptcy, a foreclosure, or both. Our data set allows us to identify when these events have occured

5.1 The Case of Bankruptcy

We calculate the monthly debt payment to income ratio for twelve months prior to and after a bankruptcy event for each individual. The data for each individual is then reordered around the bankruptcy event. We then want to examine how this ratio behaves, on average, across individuals. The expectation is that an individual's monthly debt payment to income should increase as the bankruptcy event approaches. We also examine the behavior of the FICO score. One of the strengths of our data is the availability of multiple credit scores, including the FICO score. As the FICO score reflects more information than is contained in the monthly debt interest payment to income ratio, this is a rigorous standard of comparison.

We begin by examining the avergage debt payment to income ratio and the FICO score averaged over the sample of individuals who have a bankruptcy default. As can be seen in Figure 2, the average debt payment to income value twelve months prior to bankruptcy accounts for over 30 percent of monthly income. This ratio rises steadily until a month before bankruptcy when it reaches approximately 40 percent of income. Immediately after bankruptcy, the falls to under 20 percent of income and only gradually increases over the next twelve months.

The analysis of the FICO score is also informative. the average FICO score twelve months before bankruptcy is about 620, indicating a relatively risky individual. The prior hyptohesis is that the FICO score should decline as a bankruptcy event nears. As can be seen, this is exactly what happens. In fact, the decline in the FICO score is much sharper than that oberved in the payment to income ratio. At bankruptcy, the average FICO score is about 560. After the bankruptcy event, the data indicates that the avergae FICO improves, and

actually exceeds the average score twelve months prior to default.

Figure 2: Monthly Payment to Income Ratio (Bankruptcy)

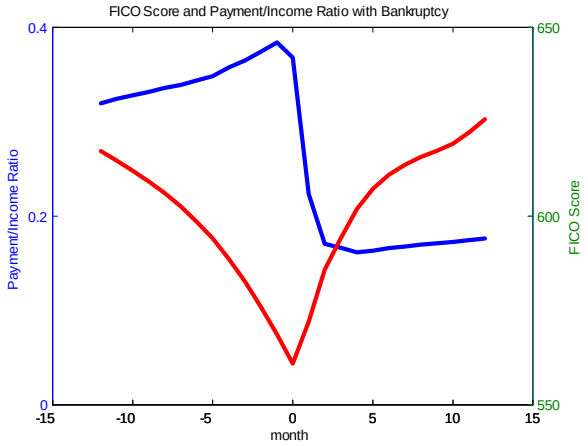

The study of averages can mask heteogeneous effects leading to bankruptcy. We do not have data on various events which would contribute an income shock such as job loss or adverse health shock. As a result, we take the payment to income ratio at the bankruptcy event and reorder the data from lowest ratio to highest ratio. The payment to income ratio is then segmented into quartiles with the first quartile composes of the lowest payment to income ratios. The FICO scores are ordered in a similar manner and also divided into

quartiles. We present these results in Figure 3.

Figure 3: Bankruptcy from a Quartile Perspective

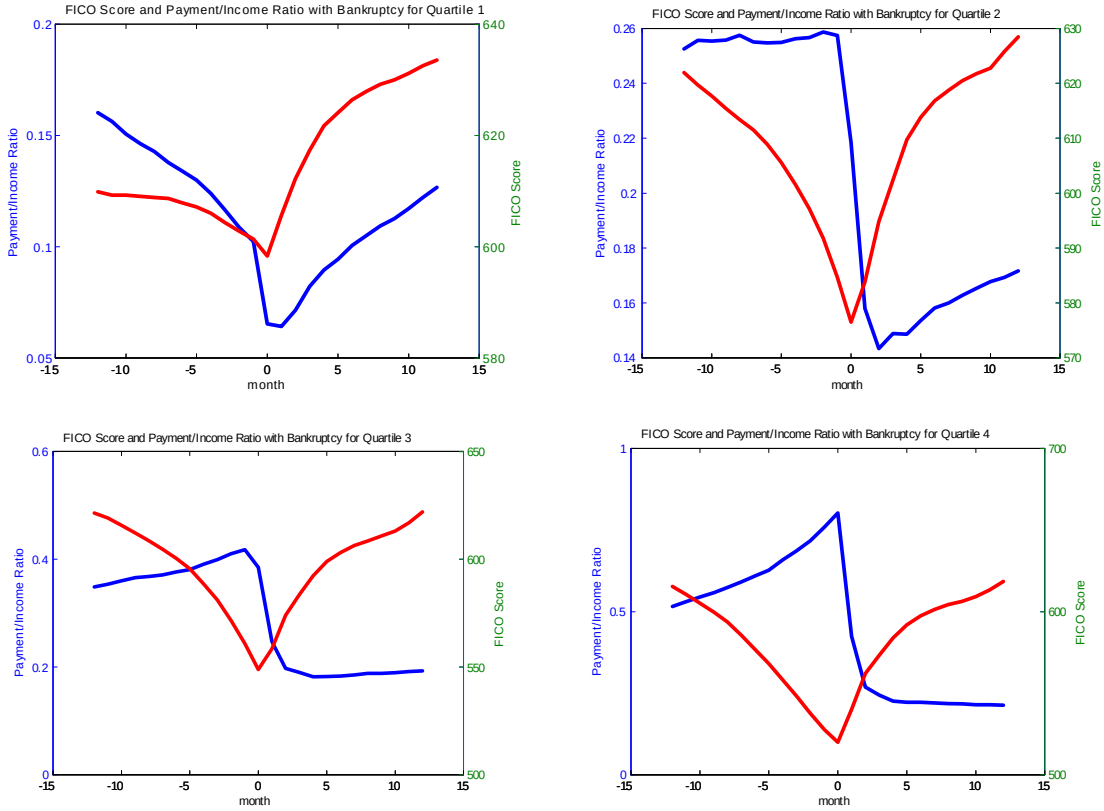

The first quadrant shows individuals with the lowest monthly payment to income ratios. Contrary to our expectations the ratio actually declines for these individuals. Upon furthe examination, we dicovered a large number of individuals reduce their monthly payment or stopped making monthly payments altogether. This is an important point as it indicates that discontinuing payment could also indicate an impending default. The average FICO score moves in a similar fashion as the general case..

As other quartiles are examined, a rising monthly payment to income ratio once again precedes default. In the second quartile, the monthly payment to income ratio increases slightly as the default date approaches. In last two quadrants, the upward trend approaching bankruptcy is more obvious. In the last quadrant, the average monthly debt payment accounts for half of income. At a month before default that ratio has increased to over 75 percent.

5.2 The Case of Foreclosure

We conduct a similar analysis for the foreclosure outcome. The payment to income ratio should be greater than the ratio observed in the bankruptcy case due to the relative value of the underlying housing asset. Figure 4 shows that twelve months prior to foreclosure, the average value of the ratio is slightly greater than 40 percent of income. The ratio increases steadily until three months prior to default. In the fourth period prior to default, the metric indicates that 46 percent of income is allocated to debt repayment. The last three months prior to foreclosure are characterized by a rapid increase in the ratio. Between the four months prior to the default event and the actual date, the monthly payment to income ratio increases by 15 percentage points - a substantial increase.

Figure 4: Monthly Payment to Income: Foreclosure

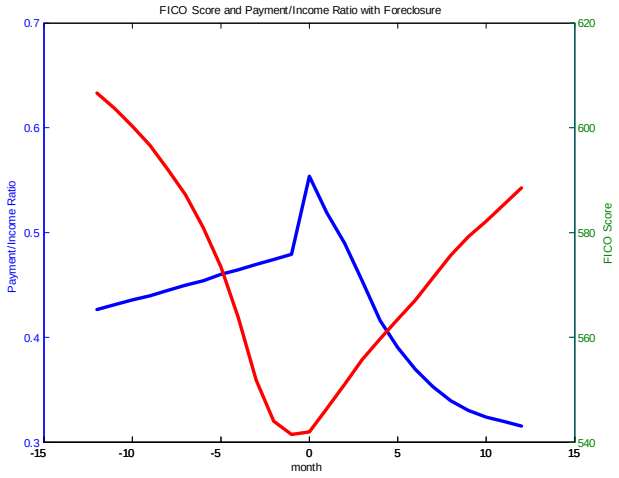

The average FICO score in the foreclosure case is close to 610 for those individuals which end up defaulting. Their FICO score contintually declines until default occurs. For example, two months prior to default the average score declines from 610 to just over 540. The FICO score, which is continually updated, reflects the individual's deteriorating debt payment history. After foreclosure, the FICO score does improve over the next twelve months.

In Figure 5, the monthly debt to income ratio and the FICO scores are presented prior to and after foreclosure. For individuals who are in the lowest quartile of payment to income ratios: twelve months prior the ratio is less than 30 percent and the average FICO score is 590. An increase in the monthly payment to income is not observed. In fact, this ratio is relatively flat with a step decline in the seventh month prior to default. The FICO score gradually declines until mortgage default occurs and then steadily improves.

In the other quadrants, the payment to income ratio increases in value until foreclosure. After the default event, the debt payment ratio increases. The rate of increase in the debt

to payment ratio changed a few months prior to default. This pattern appears in all three cohorts. A closer examination of the data does not give a clear answer as to whether thiis is due to an effort to increase rate of repayment or a decrease in income. The FICO pattern in these three quadrants contines to be approximated by a 'v' pattern.

Figure 5: Foreclosure from a Quartile Perspective

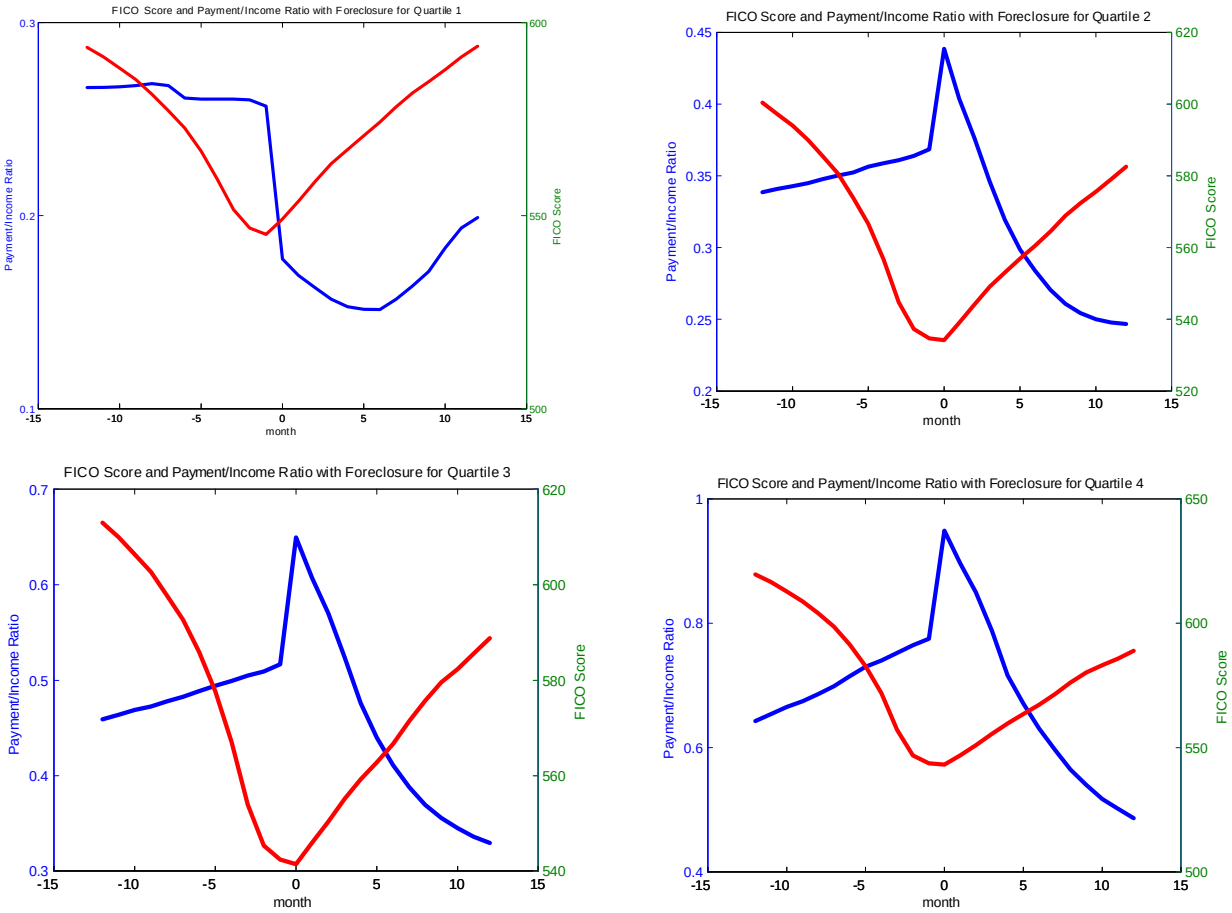

5.3 The Case of Both Bankruptcy and Foreclosure

A number of individuals in our sample both filed for bankruptcy and experienced a foreclosure. In studying these individuals, we did not differentiate which default event came first. In addition, we did not limit these default decisions to a given period in time. Simply, both types of default had to occur. We date the timing of the default to correspond to whichever type of default occurs first.

As in the prior two cases, we begin with an examination of the average, and then examine the quartile distributions. The average situation is presented in Figure 6. A number of facts stand out. First, the average monthly payment to income ratios are higher than in the prior two cases. The average value of this ratio twelve months prior is approximately 65 percent of

income. At default, debt payments account for over 75 percent of income. Second, the debt payment ratio increases as default occurs. As in the foreclosure case, this ratio surges just prior to default. Lastly, the FICO score trend continues to resemble a 'v-pattern.' Twelve months prior to either default, the average individual had a credit score of 615. At default, this score had declined to a value slightly greater than 560. By twelve months after the default, the average score increased by about 30 points.

Figure 6: Monthly Payment to Income: The Joint Case

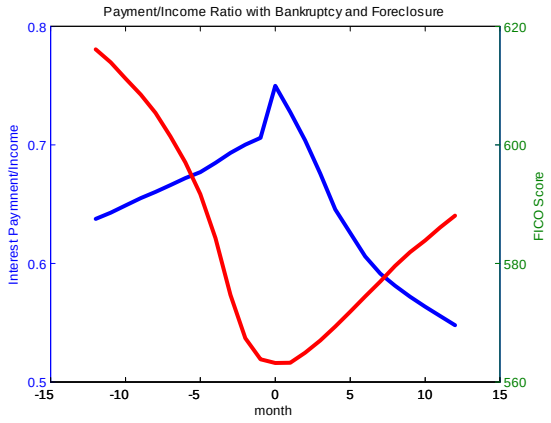

Figure 7 presents the average behavior of the two metrics from a quartile perspective. In the quartile with the lowest payment to income ratios, both the average monthly payment to income ratio and the average FICO can be characterized by 'v-patterns.' As explained previously, this pattern for the payment to income ratio is accounted for by a large number of indiviuals who make small or zero payments. For the other quartiles, the patterns in the

payment to income ratio and the FICO score are consistent with the previous two cases.

Figure 7: Joint Bankruptcy and Foreclosure from a Quartile Perspective

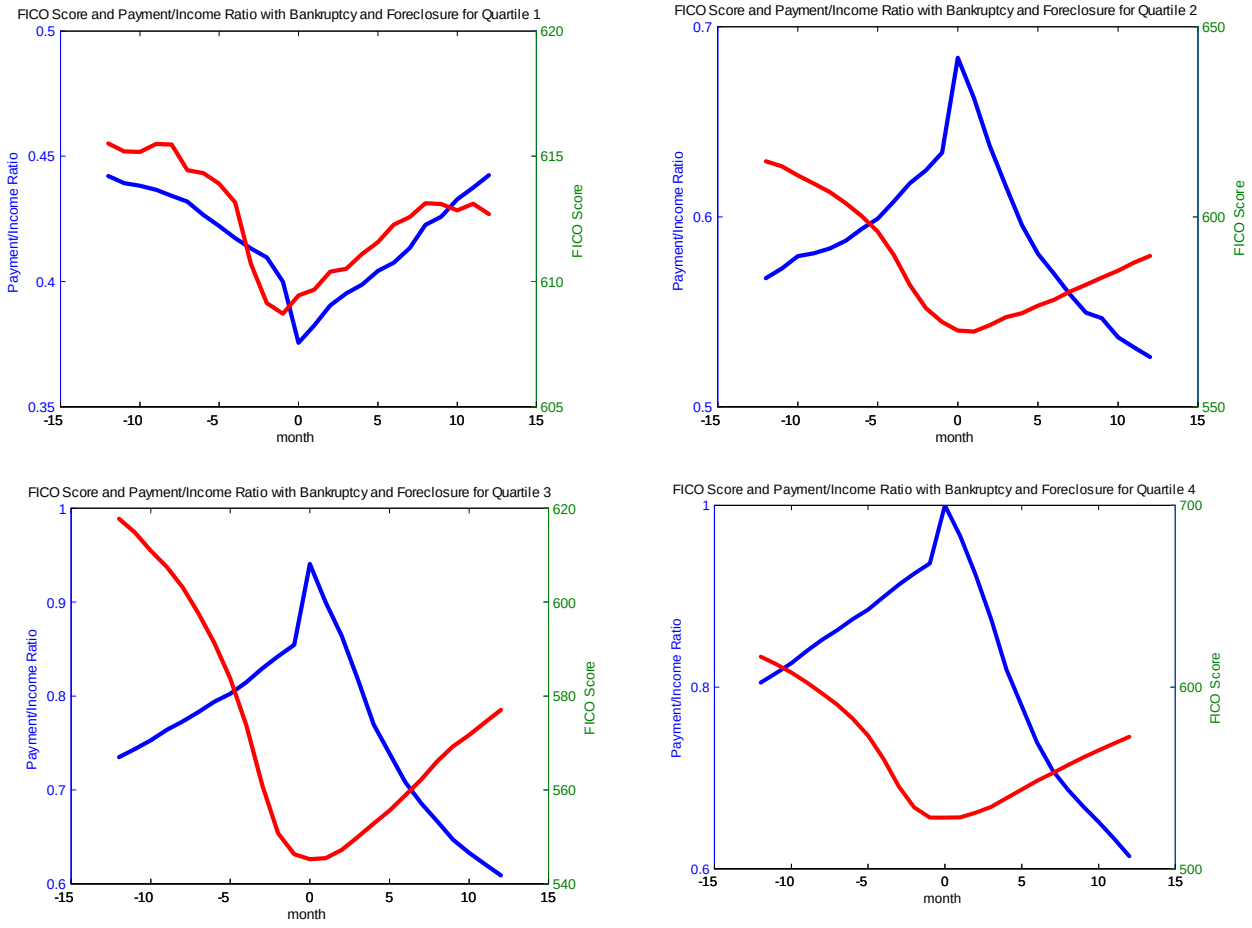

6 Are "Tipping Points" Forecastable?

The prior section examined patterns both before and after a default event. For the monthly payment to income ratio the general upward trend suggests that this indicator could be useful as an early indicator of a household default crisis. Given that there is a break in this pattern around three months prior to the default suggests that the rate of change could be more informative. Similarly, the fact that the FICO score steadily declined prior to an impending default further suggests a potential indicator. In this section, we evaluate whether the clues from the prior analysis could result in an simple indicator for financial default. In addition, the prior section suggests that some ditributional factors may be at play in the behavior of a default indicator.

In this section, we examine the question of whether an indicator should be based on average data. Based on this examination, we conduct a forecasting evaluation of the monthly

payment to income ratio. We conduct this evaluation in terms of a new data sample that include a random sample of both nondefaultors and defaultors. Essentially, we are trying to conduct the evaluation in an out of sample test.

6.0.1 Some Distributional Considerations

In this section, we investigate whether development of the forecasting rules, such as the monthly percentage change in the payment ratio, should be based on the data set that focused on overall averages or the average of some distribution of the data. Figures 9 through 11 are the basis for this decision.

In Figure 9, we focus on the case where an individual only files for bankruptcy. When we examine the data averaged over individuals, we find that the payment to income ratio steadily increased until default occured. Alternatively, the FICO score declined prior to a banktuptcy. Both of these tends stand out again here. However, both trends are relatively 'flat', which casts doubt on finding a powerful indicator. As a result we examine two restricted samples: individuals with the highest 25 percent of payment to income ratios, and individuals with the highest 5 percent of ratios.

As can be seen, the payment to income ratio increases as the sample becomes more selective and has a steeper increase in slope. Similarly, the average FICO score shifts down and declines faster in the restricted cases. Thus, stronger indicators may come from a more selective sample.

In Figure 10 and Figure 11 we construct the same set of curves for the case of only foreclosure and the case of both foreclosure and bankruptcy.

Figure 9

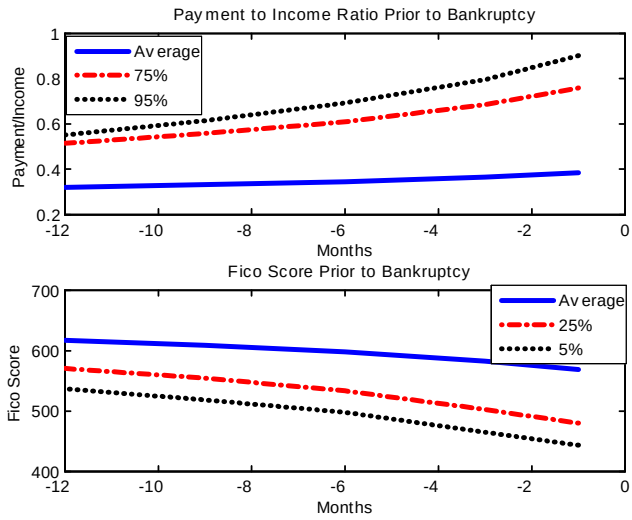

The trends in these figures is consistent with the previous examples. In the next section we perform and evaluate a simple forecasting exercise. Given the results in this section we will assess the monthly rate of change of the monthly payment to income ratio at 3, 6, and 9 months prior to event as an indicator of a default for each of our three cases. We will compare the performance of the predictive forecast with that of the FICO score. In addition, we expand our sample to include individuals who did not have a default. This will allow us to have an out-of-sample evaluation of the monthly payment to income ratio as an indicator of a impending financial crisis.

Figure 10

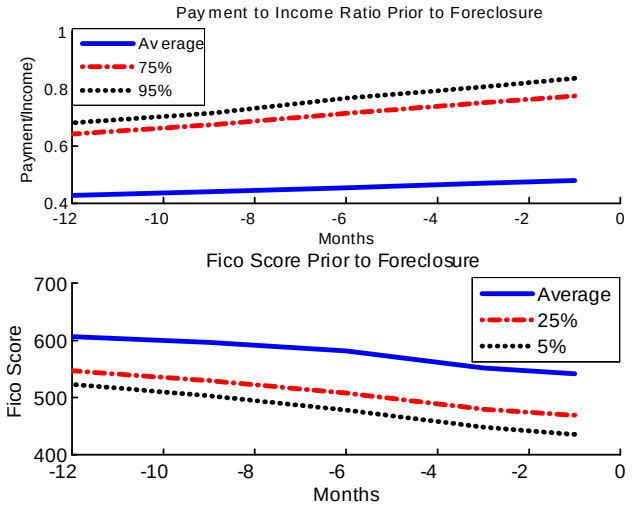

Figure 11

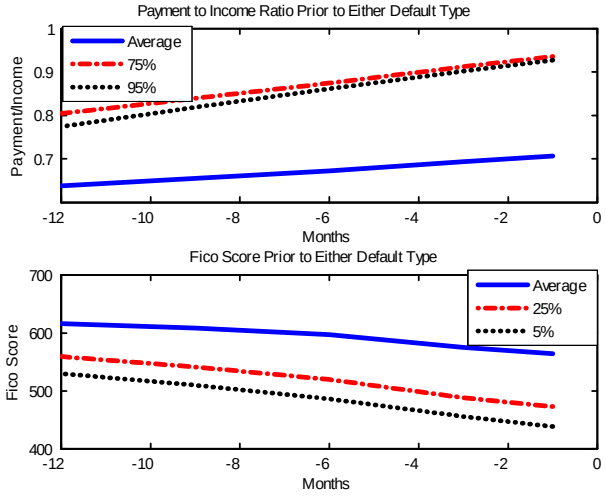

6.0.2 An Evaluation of Forecast Potential

7 Conclusions

In this paper we define a new notion of a household "tipping point" using principles of economics and a novel high-frequency consumer-level dataset. We attempt to create a metric which identifies when a household reaches the point where their debt holdings become a drag on their activity rather than an engine for growth. We begin by posing a question: how does economics help in the identification of such a tipping-point? The answer is a tipping point occurs when an individual no longer an satisfy their individual budget constraint. This is a point where an individual must turn to bankruptcy or a foreclosure to satisfy their budget constraint. The ratio of the total debt payment to after-tax income has proven to be an informative measure. However, we show that this metric should be conditioned on such things as available wealth, the liquidity of the assets, ability to raise additional debt, and the structure of this debt. The depth of our data allows us to assess how this ratio behaves based on age groups and income quartiles. We also assessed the ability of the debt payment to income ratio to predict bankruptcy and foreclosure events.

8 References

Adelino, M., K. Gerardi, and P. Willen (2013), "Identifying the Effect of Securitization and Foreclosure and Modification Rates Using Early Payment Defaults," Journal of Real Estate and Financial Economics, 1-27.

Athreya, K., X. S. Tam and E. R. Young (2012), "A Quantitative Theory of Information and Unsecured Credit," American Economic Journal: Macroeconomics 4, 153–83.

Athreya, K., J. M. Sanchez, X. S. Tam and E. R. Young (2015), "Labor Market Upheaval, Default Regulations, and Consumer Debt," Review of Economic Dynamics 18 (January 2015), 32–52.

Benjamin, D. and X. Mateos-Planas, "Formal versus Informal Default Consumer Credit," Working Paper, October 2014.

Bricker, J. M. Brown, S. Hannan, and K. Pence (2015), "How Much Student Debt is Out There?," FEDS Notes, 2015-08-07, Board Of Governors of the Federal Reserve.

M. Brown, S. Stein, and B. Zafar (2015), "The Impact of Housing Markets on Consumer Debt: Credit Report Evidence from 1999 to 2012," Journal of Money, Credit, and Banking, 47(S1): 175-213.

Brown, M., J. Wen and B. Zafar (2016), "Financial Education and the Debt Behavior of the Young," The Review of Financial Studies, forthcoming.

Chatterjee, S. and G. Gordon (2012), "Dealing with consumer default: Bankruptcy vs garnishment," Journal of Monetary Economics 59 (December 2012), 1–16.

DiMaggio,M., A. Kermani, R. Ramcharan, and E Yu, (2016) "Financial Volatility and Household Consumption," Working Paper.

Gerardi, K., L. Lambie-Hanson and P. Willen,(2015), "Do Borrower Rights Improve Borrower Outcomes? Evidence from the Foreclosure Process," Journal of Urban Economics, 73(1), 1-17.

Herkenhoff, K. F. and L. E. Ohanian, "Foreclosure Delay and U.S. Unemployment," Working Paper, March 2012.

Kaplan, G. and G. L. Violante (2014), "A Model of the Consumption Response to Fiscal Stimulus Payments," Econometrica 82, 1199–1239.

Low, H., C. Meghir and L. Pistaferri (2010), "Wage Risk and Employment Risk over Life-Cycle," American Economic Review 100, 1432–67.

Mitman, K. (2015), "Macroeconomic Effects of Bankruptcy and Foreclosure Policies," Working Paper