Stockholm School of Economics
Department of Economics
BE551 Degree Project in Economics
Spring 2023

Navigating the Waves: A Study on the Effect of Income Volatility on Household Debt Accumulation

A difference in difference analysis of the differential borrowing response of high- and low volatility households to the enactment of the Gramm-Leach-Bliley act

Loke Sjödin (25055) and Lukas Segerberg (25065)

Abstract

This study investigates the influence of income volatility on household borrowing behavior, contributing to the broader discussion on the relationship between inequality and the risk of financial crises. While prior studies have generally focused on permanent income inequality, we address the overlooked aspect of transitory income inequality. We use data from the Panel Study of Income Dynamics to calculate the inherent individual earnings volatility of different industries, which is used to represent the income volatility faced by households. By adopting a difference-in-differences methodology, we compare the debt accumulation of households with high and low income volatility during the years preceding the 2008 crisis, against the backdrop of the Gramm-Leach-Bliley Act. The results are inconclusive. We attribute this to possible methodological limitations. Nevertheless, this study contributes to the literature by shedding light on an underexplored dimension of income inequality and its potential impact on credit growth. We call for further research to refine the methodology and to provide more definitive insights into the link between income volatility and household borrowing behavior.

Keywords: Inequality, Income volatility, Gramm-Leach-Bliley Act, Financial crises, Household Debt
JEL: D31, E50, G18.

Supervisor:	Fabio Blasutto
Date submitted:	14 May 2023
Date examined:	24 May 2023
Discussants:	Leonard Xander & Klara Holmer
Examiner:	Johanna Wallenius

Acknowledgements

We would like to express our gratitude to our supervisor Fabio Blasutto. We would also like to extend our thanks to our fellow students and friends at Stockholm School of Economics who provided valuable insight and support.

But, ah, think what you do when you run in debt;
you give to another power over your liberty.

Benjamin Franklin

1. Introduction

The idea that the financial system itself can generate economic instability through endogenous credit booms is not entirely novel. Economists have pointed to the impact of credit on financial fragility since the early 1970s, as evidenced by the work of Minsky (1977) and Kindleberger (1978). The global financial crisis of 2008 highlighted the potentially devastating consequences of unbridled credit expansion. However, in the years leading up to the crisis, credit was not the only economic measure that reached record levels in the US. In 2006, income inequality, as measured by the Gini index, reached a record high in the same year (World Bank, 2023). This is striking, and has raised the question among economists of what role income inequality played in the credit expansion that led up to the crisis.

In the aftermath of the global financial crisis, there was a noticeable deceleration in the growth of U.S. household debt. However, it is now on the rise again, having reached a record 16.9 trillion USD in 2022, (Federal Reserve Bank of New York, 2022). At the same time, there has been a precipitous increase in income inequality, as evidenced by the GINI index exceeding its 2006 level in 2014 and yet again in 2019 (World Bank, 2023). While the relationship between credit and financial crisis has been well-established in the literature, dating back to Fischer (1933), and proven more recently by Schularick and Taylor (2012), the relationship between income inequality and credit growth has been more difficult to prove empirically.

Income inequality can be divided into a dichotomy of permanent inequality and transitory inequality, where permanent inequality is the component most commonly referred to by the term income inequality, and transitory inequality refers to the component driven by idiosyncratic income shocks. However, when income inequality is measured as the variance of income, for instance, using cross-sectional data at a fixed point in time, both permanent and transient inequality are captured. Both empirical and theoretical research on the topic of how income inequality affects credit growth has mostly focused on the effects of permanent income inequality (see for instance Kumhof et al., 2015), leaving a gap in the literature. This is a research gap that we intend to fill.

In this paper, we aim to shed light on the link between inequality and the risk of financial crisis by examining only a specific part of income inequality, namely income volatility - which can be interpreted as the transitory component of inequality - and how it affects the borrowing behavior of households. The increase in income inequality over the last decades, as we explain in our literature review, seems to have been in part driven by an increase in individual income volatility (see for instance Gottschalk et al., 1994; Moffitt & Zhang, 2018). Hence, there is reason to examine the effects of this part of income inequality on credit, and therefore, unlike most of the previous literature on this topic, we choose to focus on individual income volatility to analyze the mechanism of the theoretical relationship between income inequality and credit. We hypothesize that households that face higher income volatility have a greater need for consumption-smoothing credit.

Our research question is thus how income volatility affects household borrowing behavior. We attempt to answer this question by investigating how households, whose primary income earners were employed in industries characterized by high income volatility, adjusted their borrowing practices in the period leading up to the global financial crisis. The unique aspect of our study is that it coincides with a policy shift that we argue led to an increase in credit availability. Therefore, we posit that the intersection of these conditions provides a valuable lens to understand the impact of income volatility on household debt accumulation.

We use data from the Panel Study of Income Dynamics on earnings of household heads and the industry they worked in to construct a measure of the intrinsic income volatility of different industries[1]. After imposing sample restrictions to ensure that observed earnings volatility is as idiosyncratic as possible, using a measurement period of 10 years, we calculate the average individual earnings volatility of each of the twelve industries and compare the industry averages to the total sample average earnings volatility, and index industries as either high or low earnings volatility industries depending on if they exceed or subceed the total sample average.

[1] The industries available from the PSID dataset are Agriculture, Forestry, and Fisheries; Construction; Manufacturing; Transportation, Communication, and Public Utilities; Wholesale and Retail Trade; Finance, Insurance, and Real Estate; Business and Repair Services; Personal Services; Professional and Related Services; and Public Administration.

The purpose of using industry-level income earnings volatility to capture the income volatility faced by individual households is to minimize the risk of our measure of income volatility capturing idiosyncratic fluctuations, rather than intrinsic volatility. By focusing on industry-specific income volatility, we ensure that the groups characterized by high and low volatility are structurally differentiated, as they are established based on the inherent attributes of the industries in which individuals are employed, rather than each individual's realized income volatility, which may be driven by self-imposed life choices.

As a consequence of utilizing this approach to delineate the high and low income volatility households, we exclusively capture one channel of income volatility that households face, specifically the earnings volatility of the household head. Additionally, we are unable to account for volatility-reducing income sources, such as unemployment insurance or social security benefits, due to the absence of data pertaining to these variables. These limitations partly detach our measure of volatility from the income foundation that households rely on to make borrowing decisions. However, we will show that an alternative specification, where households are categorized based on their realized individual volatility rather than the industry group they belong to, leads to an arbitrary categorization.

We then estimate the difference in increased borrowing between households that faced low and high income volatility during the years leading up to the global financial crisis. We employ a difference-in-differences model with a treatment group consisting of households whose head worked in an industry characterized by high earnings volatility, and a control group of households whose household head worked in an industry characterized by low earnings volatility. The treatment is the policy intervention of the Gramm-Leach-Bliley Act (GLBA). The GLBA, implemented in 2000, replaced the Glass-Steagall act, and facilitated the consolidation of investment and commercial banks, paving the way for the creation of financial conglomerates. (Grant, 2010) It expanded the range of credit products offered to consumers by enabling banks to engage in proprietary trading and invest in riskier financial products (Grant, 2010; McDonald, 2017; White, 2010). We argue that this led to increased competition in the credit industry, enhancing credit accessibility. Furthermore, we hypothesize that households' reactions to the GLBA varied based on their income volatility, due to the need to cushion their consumption against idiosyncratic income shocks, as well as potentially differing levels of prior credit access. The treatment effect will thus according to our hypothesis be much stronger for the treatment group.

Our main model specification revealed a greater increase in borrowing in response to increased credit availability for the individuals that worked in high volatility industries than for the control group; however, this increase was statistically insignificant, preventing us from rejecting the null hypothesis that the household head's employment in a high volatility industry does not impact household borrowing behavior. This finding, while statistically insignificant, would suggest a potential positive relationship between income volatility and borrowing behavior, aligning with theories proposed by Krueger and Perri (2006). If our estimates had been statistically significant, it would have indicated that individuals in high volatility industries responded to increased credit availability by borrowing more than their counterparts in low volatility industries. However, a robustness check to the methodology of indexing industries as having high or low inherent earnings volatility gave a (statistically insignificant) estimated effect of the opposite sign. We interpret our failure to establish a relationship between income volatility and borrowing behavior not as evidence of an absence of such a relationship, but rather as indicative of a flawed methodology. Among the issues in our methodology are unobservable factors such as attitude towards risk, educational background, and the presence of spillover effects, as well as potential endogeneity of treatment.

The structure of this paper is as follows: Section 2 provides background on the topic and an overview of the current literature pertaining to the interplay between inequality, credit, and financial crisis. Section 3 delineates the existing research gap within the empirical literature that this study seeks to address, and elucidates how our work contributes to the broader literature. In section 4 we describe our data. Section 5 is dedicated to the methodology that we have adopted for our study. The empirical results derived from our difference-in-differences models are presented in Section 6. In Section 7, we interpret the economic and statistical significance of our findings. Section 8 engages in a critical discussion of the limitations of our methodology, particularly focusing on why we were unable to conclusively establish a causal link between earnings volatility and credit behavior. Lastly, Section 9 provides a summary and conclusion, encapsulating the main findings and implications of our study.

2. Background and Literature Review

2.1. Credit Affects the Risk of Financial Crisis

Since the 1970s, economists such as Minsky (1977) and Kindleberger (1978), have highlighted the role of credit in financial vulnerability. Many more recent studies have found empirical evidence that credit is a robust predictor of financial crises. One particularly noteworthy study is Schularick and Taylor (2012), which provides compelling evidence of the link between credit growth and financial crises. The paper analyzed data from 14 advanced economies between 1870 and 2008 and tested, among other things, one element of the credit view argument, namely that financial crises can be seen as “credit booms gone wrong”. The study found that rapid credit growth was a robust predictor of financial crises. According to the authors, credit booms are often associated with financial instability because they lead to a buildup of debt, which can trigger a crisis when borrowers are unable to repay their loans - which is in fact what occurred in the global financial crisis of 2008. Credit aggregates indeed contain valuable information about the likelihood of future financial crises.

2.2. The Nexus of Inequality, Credit, and Financial Crisis

While the relationship between rapidly rising credit and financial crisis is well-established in the literature, the role that income inequality plays in this context is still the subject of some debate and requires further examination. As shown below, there is some evidence to suggest that income inequality can exacerbate financial instability, but the causal relationship between inequality and credit remains an open question and requires further research.

There are several theoretical explanations for why inequality might drive household credit growth. A negative relative income shock to households could motivate increased borrowing as a means of consumption smoothing. However, Friedman's (1957) permanent income hypothesis suggests that increased leverage is only a rational response if the income shock is perceived to be temporary. If the income shock is perceived as permanent, it would be irrational for households to respond to negative income shocks by increasing their leverage. However, there are reasons to not expect consumers to abide strictly by the permanent income theory. First, there may be errors in expectations - households may incorrectly determine the nature of the shock. Furthermore, there may be a behavioral explanation for irrational behavior in relation to the permanent income hypothesis. One explanation is that growing income inequality may stimulate a “keeping-up-with-the-joneses” effect, also referred to as the relative income hypothesis, a theory that contradicts the permanent income hypothesis. Both terms describe the phenomenon where an individual's consumption behavior is influenced by the consumption behavior of others around them, especially those with higher income or social status. If individuals are unable to “keep up with the Joneses” with their income, they might turn to borrowing. Dijk et al. (2010) have found support for this notion using a model that predicts the ways that savings rates respond to changes in both one's own and others’ permanent income. They poignantly formulate the core idea of the relative income hypothesis: “The pain of enduring lower relative living standards today can be experienced directly. In contrast, the pain of enduring lower relative standards in the future can only be imagined.”

As for empirical evidence, Bertrand and Morse (2013) provide evidence of "trickle-down consumption”: Using CEX data, they found that middle-income households would have saved 2.6-3.2% more by the mid-2000s if top incomes had grown at the same rate as median income, suggesting that rising income and consumption at the top of the income distribution since the early 1980s have led lower-income households to consume a larger share of their income. While household debt was not the topic of their study, it lends further credibility to the relative income hypothesis.

2.3. The Savings Glut of the Rich

One frequently posited causality chain suggests that increased (permanent) income inequality leads to increased borrowing by those experiencing a decline in income share, while simultaneously promoting lending by those with rising income shares, essentially resulting in the less affluent borrowing from the more affluent (See for instance Kumhof et al., 2015; Mian et al., 2019; Rajan, 2011). This surge in credit is then argued to subsequently elevate the risk of a financial crisis. Kumhof et al. (2015) developed a theoretical model with an internally consistent mechanism that captures this phenomenon, wherein top earners, following a series of positive income shocks, consume a relatively smaller portion of their wealth and instead channel it into loans to bottom earners. Their model implies that this mechanism enables bottom earners to sustain their consumption levels, and that the consequent expansion of the bottom earners' debt-to-income ratio induces financial instability, ultimately increasing the likelihood of a financial crisis. Mian et al. (2019) provide empirical evidence for this theory by examining who ultimately holds U.S. household debt as a financial asset. They were able to show that the primary source of financing for U.S. households debt is U.S. households, and when allocating financial debt along the wealth distribution, they found that net household debt positions since the 1980s suggest that rich Americans increasingly financed the borrowing of non-rich Americans.

2.4. Evidence to the Contrary

In this context, we would be remiss not to mention factors other than income inequality that may be driving demand for credit. When Bordo and Meissner (2012) studied the effects of income inequality on credit growth, using data from 14 advanced economies over a period of 80 years, they found that after controlling for interest rates and GDP growth, the effect of rising top income shares on credit was no longer robust. They find little evidence connecting credit booms and financial crises to increasing inequality. Instead, according to Bordo and Meissner, the two key determinants of credit booms are economic expansions and low interest rates, which they argue is more consistent with the broader literature on credit cycles.

While Bordo and Meissner's study is informative, it should not be taken as the final word on the subject. Several other studies in the literature suggest that increasing income inequality may put downward pressure on interest rates, as the high saving rates of the rich leads to a downward pressure on aggregate demand in the economy, as Mian et al. (2019) note. See for example, Fitoussi and Saraceno (2010), Stiglitz (2016), and Rachel and Summers (2019). If income inequality is driving the savings glut of the rich, it may be playing a role in reducing interest rates, which would lead Bordo and Meissner to obtain a downward biased estimate of the relationship between income inequality and credit.

2.5. Volatility as a Source of Inequality

An alternative explanation for why Bordo and Meissner were unable to identify a correlation between inequality and credit may lie in their focus on an inappropriate form of inequality. As explained in the introduction, when income inequality is measured as the dispersion of income by employing cross-sectional data at a fixed point in time, both the permanent and transitory elements of inequality are captured. The transient aspect of inequality can be characterized as individual income volatility. Even if all individuals had identical lifetime earnings, some variance in the distribution of incomes would exist at any given moment due to labor market mobility. As individuals transition between jobs, varying durations of unemployment may result in reduced income for a particular year. Other factors contributing to individual income volatility include performance-based remuneration, irregular work schedules, seasonal employment, policy shifts, re-education, demographic changes, and more.

2.6. Trends in Income Volatility

In a seminal study on the topic of income inequality and volatility, Gottschalk et al. (1994) used data from the Panel Study of Income Dynamics (PSID) to show that the literature had “missed a critical aspect of the widening wage distribution”. There was a growing body of evidence that showed that the variance of income was widening, which was, and still is, often interpreted as growing income inequality in terms of permanent earnings. Gottschalk and Moffitt, thanks to the PSID, could follow the same individuals over a long period of time, and were thus able to split income variance into permanent and transitory components. They defined permanent variance as the variance of an individual’s mean wage over a 9-year time period, and the transitory variance as each individual’s variance around his own mean during the same period. By splitting income inequality into permanent and transitory variance, they were able to show that there had been a significant increase in the variance of each individuals’ earnings between 1970-78 and 1979-87, and crucially, that this observed increase in the variance of short-term changes in earnings accounted for as much as one-third of the increase in the variance of earnings from the 1970s to the 1980s. Possible drivers of the change that took place in that time include more individualistic wage setting as a consequence of deregulation as well as the demise of pattern bargaining, increased job-changing rates, and skill-biased technological change.

More studies have followed up the topic using survey data to suggest that labor earnings volatility has continued to increase in the U.S. after the study period of Gottschalk et al. (1994). Gottschalk and Moffitt (2009) found that volatility increased in the 1980s and remained at the same level until 2004, while Elmendorf et al. (2012) found that volatility continued to increase until the late 2000s. Moffitt and Zhang (2018) found, using PSID data and a rolling nine-year window for male head earnings, that gross volatility has had a three-phase trend, increasing from the 1970s to the 1980s, stabilizing in the 1990s through the early 2000s, and again increasing strongly during the Great Recession.

However, other studies, like Kopczuk and Saez (2010) have found a smaller increase in transitory earnings in the 1970s, and that the increase reverted in the late 1980s and 1990s. Kopczuk and Saez admit that the difference between their findings and that of Gottschalk and Moffitt could be a result of different sampling techniques. For instance, Gottschalk and Moffitt focused exclusively on white males. Different sampling techniques and methods indeed yield conflicting answers regarding the overall trend in individual income variance over the last decades. Studies using U.S. administrative data have typically found that volatility has remained constant since the mid-1980s, as found by Dahl et al. (2011). Bollinger et al. (2011) also found conflicting evidence on the topic. Using both survey and administrative tax data for the years 1995-2015, they found that earnings volatility was stable over the period for continuously working men in both survey and tax data. Their results suggest that there is evidence of stable transitory volatility, but rising permanent volatility over the past two decades.

In summary, the trend in individual earnings volatility has sparked scholarly debate, with a substantial body of research indicating a notable rise from the 1970s onwards. Although some studies have reported more tempered increases or fluctuations in this trend, the preponderance of evidence suggests that individual earnings volatility has increased over the last decades. However, the argument by Gottschalk and Moffitt attributing one-third of the income variance increase from the 1970s to the 1980s to income volatility does not necessarily apply to later periods. This is because studies, even those documenting an upward trend in earnings volatility, imply a potential slowdown in this trend from the 1990s to the early 2000s, even as the total cross-sectional variance of incomes continued to rise, as illustrated by Bollinger et al. (2011).

2.7. Income Volatility and Credit Behaviour

Krueger and Perri (2006) relied on this literature, documenting an increase in individual income volatility, to focus on this aspect of inequality to explain the relationship between income inequality and credit behavior. They utilize data from the Consumer Expenditure Survey to demonstrate that the rise in income inequality in the United States since the 1980s has not been proportional to observed increases in consumption inequality. To investigate the relationship between income inequality and consumption, the authors develop a "Debt Constraints Model" (DCM), wherein agents engage in risk-sharing contracts and have the option to default on their obligations, with the consequences of forfeiting assets and future risk-sharing opportunities. By increasing the income volatility parameter in the DCM, the researchers demonstrate that agents shield their consumption from idiosyncratic income shocks by engaging in risk-sharing through the use of non-collateralized, state-contingent credit, dependent on the ability of financial markets to adapt to increasingly volatile individual income. Krueger and Perri contend that if the escalation in income inequality is partly attributable to increased income volatility, households' demand for credit as a means to buffer their consumption against idiosyncratic income shocks should increase. Furthermore, their argument that income volatility should increase the demand for consumption-insulating credit supports our conjecture that individuals who face high income volatility should be disproportionately affected by a policy change that decreases borrowing constraints. The authors conclude by advocating for further empirical research using micro-level data to examine the underlying mechanisms connecting income volatility and credit behavior, which aligns precisely with the objective of our paper.

2.7.1. Other Links between Income Volatility and Credit

The impact of increased individual income volatility on borrowing behavior remains a complex issue, with theoretical and empirical evidence suggesting potential influences in both directions. Traditional economic behavior models postulate that income volatility leads to a reduction in borrowing and an increase in precautionary saving. This concept aligns with the rationale explaining why households experiencing substantial income fluctuations might resort to credit usage, as individuals with irregular income streams essentially have two options for maintaining consumption stability: they could either accrue savings during periods of prosperity when their income surpasses the average, drawing upon these reserves during leaner times when income falls short, or alternatively, they could resort to borrowing during periods of lower income.

Adding complexity to the situation, households confronted with high income volatility may exhibit heightened risk aversion, curtailing their propensity to incur debt due to the looming threat of potential default. Concurrently, income volatility can adversely impact a borrower's perceived creditworthiness, thereby making lenders more reticent to extend credit. Carroll and Samwick (1998) provide empirical support for the precautionary savings mechanism. Using PSID data, they regress households' wealth on a measure of uncertainty. They found evidence supporting the notion that households tend to engage in precautionary saving as a response to income uncertainty.

On a related note, but examining a different facet of financial behavior, Chang et al. (2022) used administrative panel data to examine how earnings volatility affects portfolio choices. They found that households reduce the share of risky financial assets when income risk increases, reducing their overall risk exposure. Both studies highlight the tendency for individuals with volatile incomes to eschew risk, albeit via different mechanisms. These findings could be used to support an argument that households that face uncertain income are more risk-averse, negatively impacting their willingness to leverage themselves.

Once more, empirical findings in the literature seem to diverge from key economic hypotheses. Specifically, the Permanent Income Hypothesis, as conceived by Friedman (1957), does not encompass the notion of precautionary savings. According to this theory, individuals save to sustain consumption in scenarios such as retirement; however, an increase in volatility does not translate into increased savings, as that the underlying utility function does not reflect prudence. Friedman's seminal contribution to economic theory suggests that households allocate a constant fraction of their permanent income for consumption, independent of wealth status. Therefore, according to this hypothesis, savings rates are expected to remain unaffected by variations in household income and maintain consistency over time.

2.7.2 Psychological Effects of Income Volatility

Beyond the need for households with high income volatility to insulate their consumption from idiosyncratic income shocks, there may be psychological factors that connect income volatility and borrowing behavior. A study by DeVoe et al. (2020) found a connection between income volatility and financial impatience. Their findings suggest that individuals experiencing more income volatility, including a greater frequency of either income spikes or dips, are more financially impatient than those who experience income volatility of a lesser magnitude. Economic and psychological research has previously found correlations between financial impatience and credit card debt, which may suggest that individuals who face more volatile income should utilize credit card debt more extensively. For instance, Meier and

Sprenger (2010) found that individuals with present-biased preferences are more likely to have credit card debt, and moreover, even after accounting for disposable income, other socio-demographic characteristics, and credit constraints, these individuals tend to have significantly greater amounts of credit card debt.

Furthermore, volatile income might lead to increased financial insecurity, creating a decision environment that discourages long-term planning of personal finances. In a study by Peetz et al. (2021) the authors studied the causal effect of income volatility on financial impatience. They found that after only a 30-min task of simulated volatile income, participants were less likely to save by postponing an immediate payout for a higher, later payout than participants who completed similar tasks with stable payouts. Their results also support the findings by Fisher (2010) and Pew Charitable Trusts, (2017); individuals who experience income volatility tend to have shorter saving horizons, are less likely to report saving, and have lower motivation to save for retirement. Income volatility has also been associated with higher risk of mortgage delinquency and missing bill payments (Farrell & Greig, 2019; Diaz-Serrano, 2005). In a more general perspective, this volatility tends to be linked with financial strains and lower self-reported financial well-being (DeVoe et al., 2020). While there is evidence to suggest that individuals experiencing income volatility also have more negative financial outcomes, studies linking volatility in income and financial outcomes may be providing evidence of a correlational relationship, leaving out the possibility that it is the behavior of individuals who make worse financial decisions that leads to higher income volatility.

2.8. The Gramm-Leach-Bliley Act

The Glass-Steagall act, also recognised as the Banking Act of 1933, was a federal law in the U.S which aimed to separate investment banking activities from commercial banking activities. It was enacted in response to the Great Depression to prevent speculative activities involving risk-taking which had contributed to the collapse of various banks. It remained in effect for over six decades until it was partially repealed by the Gramm-Leach-Bliley act (GLBA) which was enforced in 2000. The latter allowed commercial banks to once again engage in investment banking activities such as securities underwriting and dealing (Grant, 2010).

Many critics refer to the GLBA as an important factor leading up to the financial crisis in 2008 (Grant, 2010; McDonald, 2017; White, 2010). It allowed for the consolidation of investment and commercial banks and removed barriers to the creation of financial conglomerates. We believe this allowed more banks to engage in proprietary trading and invest in riskier financial products such as mortgage-backed securities and derivatives. Hence, banks and other financial institutions could offer a wider range of credit products to consumers, in terms of personal loans, mortgages and credit cards. We further believe this increased competition in the credit industry which conclusively resulted in greater access for credit consumers.

3. Our Contribution to Existing Literature

There is already evidence of a causal effect of credit on financial stability; however, the link between income inequality and credit remains elusive. As observed in the literature on this topic, studies employing macroeconomic variables to examine cross-sectional correlations, such as Bordo and Meissner (2012), have been unable to find convincing evidence supporting the notion that inequality is a driver of financial crisis. We argue that both inequality and financial instability are influenced by numerous factors, making it challenging to understand the causality chain on an aggregate level in macro cross-sectional studies. By using more precise measures that are more closely aligned with economic theory, we aim to elucidate the causality chain between income inequality and financial instability by examining only a specific part of the mechanism underlying this potential link.

Given that numerous economists have posited increased income inequality as a possible explanation for the rapid credit growth in recent decades, particularly in the lead-up to the global financial crisis, we believe that the mechanism of this link merits further empirical investigation. If income inequality does drive increased borrowing, we aim to identify which aspect of inequality is responsible. If Krueger and Perri's (2006) explanation - that increased credit can be partially attributed to a response to greater income risk (volatility) - is accurate, the increase in credit can be interpreted as improved risk-sharing among groups, and is arguably less problematic than other explanations documented in our literature review.

Our contribution is thus to clarify how inequality is linked to financial instability and to thereby enhance the understanding of the mechanisms underpinning the findings of this literature. We achieve this in two ways: first, by examining one specific type of inequality - volatility - which is expected to impact consumer credit; second, by not focusing on how income volatility affects overall financial stability, which is influenced by numerous factors. Instead, we investigate credit, which is known to affect financial instability, but is more closely related to consumer behavior. Furthermore, we investigate not only the relationship between earnings volatility and total debt, but also the relationship between earnings volatility and credit card debt, as the psychological literature suggests a specific relationship between volatility and this kind of borrowing, If a correlation between income inequality and financial instability exists, we hope to determine the reasons behind this correlation and its origins. A robust theoretical framework underlies the link between volatility and credit; therefore, if we can find evidence supporting the notion that consumers with higher income volatility have a higher tendency to increase their leverage, we can assert that this channel - linking the need for consumption smoothing of households and credit - is important for explaining the macro correlation between inequality and financial instability.

4. Data Description

In regard to the gathering of data, the PSID will be used. The PSID, or the Panel Study of Income Dynamics, is a U.S. household panel survey that began in 1968, making it the longest running longitudinal household survey in the world. The PSID has been instrumental in providing researchers with insights into microeconomic behavior, and in relation to our topic, it has frequently been used to study trends in income volatility, as it allows for the possibility of following the same individuals over time. The PSID transitioned from annual to bi-annual surveys in 1997. We thus only have biannual data after 1997. For consistency, since our period of study overlaps 1997, we exclusively use biannual data throughout the entire paper.

4.1. Variable Description

The PSID code and the interview questions for all our variables are given in section A.4 in the Appendix. The crucial PSID variables for our analysis are the industry variable, earnings of household head, and three debt variables. The industry variable is an ordinal variable collected each year with the range 0-12, where each value corresponds to the industry that the household head worked in during that survey year, and 0 represents unemployment. The industries are based on the 3-digit industry code from the 1970 Census of Population. The twelve industries defined are: Agriculture, Forestry, and Fisheries; Mining; Construction; Manufacturing; Transportation, Communications, and Other Public Utilities; Wholesale and Retail Trade; Finance, Insurance, and Real Estate; Business and Repair Services; Personal Services; Entertainment and Recreation Services; Professional and Related Services; and Public Administration.

The earnings variable we use is the wages or salaries of the household head before taxes or transfers, henceforth referred to as earnings for brevity. It does not include bonuses, overtime, tips, or commission. The data on earnings is lagging one year, meaning survey data on earnings in e.g. 1991 relates to income earned in 1990, etc.. The household head has historically been used to refer to the husband in a heterosexual married couple and to a single adult of either sex. The PSID replaced the term 'household head' with the term 'Reference Person' in 2017, but this change was not retroactive, meaning that for the purpose of our paper, in which we use data from before 2017, we continue to use the term head.

We further construct our own debt variable, henceforth referred to as Total Debt, as the sum of the three debt variables we have available to us in our PSID dataset; First Mortgage, Second Mortgage, and Other Debt. The mortgage variables represent the principal left on an interview subject's first and second mortgage, and the variable Other Debt captures any other debt that the individuals in the household may have, including credit card charges, medical or legal bills, and loans from relatives, but importantly, not vehicle loans. This is a weakness of the sample, given that vehicle loans are one of the biggest categories of non-mortgage debt according to the New York Fed quarterly report on household debt and credit (2022).

4.2. Limitations of our PSID sample

Our sample encompasses the years 1990 to 2017; however, due to data limitations, the industry variable is only available until 2001, and data on debt is accessible from 1999 onwards. This has implications for our methodology, as we explain in the Research Design section.

Participation in the PSID is voluntary, meaning there is sample attrition. Fitzgerald et al. (1998) found evidence to suggest that attrition rates are positively correlated with past income volatility, which might result in PSID having families who are more stable than the population at large over time. This means we have a case of selective attrition, which might bias our estimation of the relationship between income volatility and debt. It would be ideal to address this issue by for example employing inverse probability weighting, but we lack both the time and the data required for doing so. We therefore accept this as a limitation on our study, and instead drop individuals who do not have full participation.

In the analysis of income volatility across different industries, it is important to consider the issue of imputation in the PSID. A fraction of earnings values are imputed in the PSID for various reasons, such as "do not know" responses, refusals to answer, implausible values indicating response error, and more (Moffitt & Zhang, 2018). The imputation procedures for income in the PSID have evolved over time, using different methods and variables depending on the type of income being imputed (Duffy, 2011). Unfortunately, our PSID dataset does not give us the ability to identify which observations are imputed, and therefore, we cannot directly examine the potential impact of imputation on our analysis. We rely on the findings of Moffitt and Zhang (2018) to support our decision to ignore this aspect in our estimation of income volatility across industries. They examined the sensitivity of their analysis of income volatility to imputed income observations and found that there was little difference in trends when including or excluding imputed observations, suggesting that imputed observations are either ignorable or the imputation process adequately corrects for any nonignorability.

The PSID is biannual, meaning we lose half of the yearly earnings realizations. Furthermore, even yearly intra-year income volatility can mask month-to-month variation. For example, a worker in a seasonal industry may earn the majority of their yearly income in a select few months of the year, and have low earnings in the rest of the year. Any PSID measure of earnings volatility will thus not capture the full income variation that some individuals experience.

The reliability of the earnings reported in the PSID surveys has at times been questioned. However, Bound et al. (1994) examined the reliability of PSID earnings data by comparing payroll records to PSID-worded survey responses from the same individuals and found a reliability ratio of 0.75, which is relatively high.

5. Research Design

We want to split the sample into individuals who face high and low earnings volatility, respectively, to study differences in their behavior. An intuitive approach would be to split the sample using the realized earnings volatility of each individual, and categorize them as high or low volatility individuals depending on whether they belong to the upper or lower half of the entire sample's distribution of individual earnings volatility. However, such an approach would likely capture earnings volatility that is less inherent, and more susceptible to endogenous, self-imposed factors. It could be, for example, that individuals in the upper half of the volatility distribution in any given period were simply extremely lucky or unlucky in that specific period, or that they for whatever reason chose to change their behavior in some way during the measured period, such as going on parental leave, taking time off from work to study or travel, etc., which would impact their earnings volatility. The intrinsic earnings volatility that individuals face would be difficult to disentangle from endogenous, self-imposed volatility.

We argue that industries inherently exhibit diverse income volatility profiles due to their unique structural characteristics and external market conditions. For instance, certain industries may subject their workers to more individual earnings volatility due to the prevalence of project-oriented and seasonal employment, or performance based pay, resulting in a non-uniform income pattern. Additionally, some industries may be more sensitive to macroeconomic variations, rendering workers susceptible to income shocks due to changes in the broader economic environment. Conversely, other industries may offer more stability in terms of earnings. They may to a greater extent offer salaried contracts and be subject to stricter regulation, ensuring a more steady income stream regardless of economic climate. Therefore, it is expected that individual earnings volatility would be inherently higher in some industries than others.

Using industry-level income volatility to index individuals as high or low volatility ensures that the high and low volatility groups are structurally distinct from each other, as they are defined based on the inherent characteristics of the industries in which individuals work. This approach allows us to better capture the type of volatility that affects entire occupational categories, such as changes brought on by technological advancements or foreign competition, rather than volatility caused by specific changes in personal circumstances.

Using industry-level income volatility thus helps to disentangle intrinsic earnings volatility from self-imposed volatility, providing a more robust predictor of borrowing behavior. We visualize the benefit of using industry-level income volatility in section 5.1.1. in figure 3.

As a consequence of using this method to create the high and low volatility groups, we must use the income of the household head, who is the individual for which we have data on the industry they worked in, to create the measure of earnings volatility that households face, rather than for example the total household income. This is a drawback because a household is likely to make borrowing and saving decisions based not only on the income of the household head, but the total household income. Furthermore, many households rely on public payments or transfer payments, such as unemployment insurance or social security. Such transfers can reduce volatility of disposable income, as they often are interdependent with labor earnings. Our measure of volatility does not capture these payments, which further detaches it from the income that households use as a basis for financial decision making. These issues are likely to obfuscate the potential link between earnings volatility and borrowing behavior.

On the other hand, utilizing a different measure of earnings, such as household income that includes earnings from other household members could make it more difficult to ensure that volatility is idiosyncratic. A significant portion of the volatility in women's earnings is likely attributable to their propensity to take parental leave more frequently and leave employment for child-rearing and domestic responsibilities. This could potentially influence our results, as for example households with more children may exhibit greater volatility in household income. Notably, this form of volatility would not be idiosyncratic, as having children in most cases is a deliberate choice that households plan for.

Due to the unavailability of industry data after 2001 and debt data prior to 1999, we cannot examine the transitions between industries after 2001 or the growth of debt before 1999. These limitations necessitate a bifurcated methodology, relying on two distinct samples for our investigation. In the first part, we employ a sample of household heads from 1991 to 2001 to determine the inherent volatility of various industries and categorize them as high or low volatility, thus creating our treatment and control group. The second part of our analysis involves a separate sample of households from 1999 to 2007, which is used to estimate the impact of employment in the previously identified high or low volatility industries on credit behavior. In this second part, based on the findings of the first part, we construct our treatment and control group based on whether individuals work in high or low volatility industries, and the econometric specification of the difference in difference model.

Our approach requires us to make certain assumptions. First, we assume that the relative relationships between the earnings volatility of different industries remain consistent between the two sample periods. Second, regarding the second part of our study, we assume that the industry an individual was employed in during 2001 persists throughout the remainder of the study period. This second assumption restricts the number of post-treatment periods we can include in the regression, as the likelihood that individuals remain in the same industry decreases with each year that we move away from 2001. The choice of focusing on the period from 1999 to 2007 is thus in part a consequence of this sample limitation, however, it is also a logical choice to study debt growth leading up to the global financial crisis, which serves as a prime example of a debt-driven financial crisis.

5.1. Part One - Determining High and Low Volatility Industries

We want to use earnings from as many years as possible to construct the measure of average industry-level earnings volatility, to ensure that the averages accurately represent the inherent income volatility of the different industries. Since we only have data on the industry variable between 1991 and 2001, we use these years, and require that they have full participation to get a consistent measure of biannual intra-year income volatility. We then remove individuals who were students or retired during this period, following Dahl et al. (2012). This serves to avoid capturing income variance that can be associated with planned transitions. Since we do not have data on whether an individual is a student or retired, we assume that household heads who are aged lower than 25 or above 60 fall within these categories. Following many other studies of individual income volatility, we also drop the top and bottom 1% earners over the period to limit the effect of outliers, following among others Carroll and Samwick (1998). We then construct a new time-invariant variable that represents the industry that an individual spent the majority, if not all of the years, working in during the period. We drop individuals who were unemployed for the majority of the period, as the income volatility of individuals who were unemployed for the majority of the period cannot be representative of any one industry. Table A.1 in the appendix reports how our restrictions impacted the sample.

While many studies on income volatility drop individuals that have zero income in any year, following Dynan et al. (2012), we do not. As they argue, an analysis of income volatility should include even drastic events such as job loss, because they likely have a large impact on the welfare of households, and arguably represent the most important cases of volatility, seen from the household's perspective. If individuals are more likely to lose their jobs and remain jobless for an extended period in certain industries, our measure of the income volatility of those industries should capture that fact. However, this choice forces us to manipulate our dataset, as our measure of volatility requires that we take the logarithm of income, which is not defined for zero values. Furthermore, as mentioned in section 4.2., the PSID measure of household head earnings does not capture other sources of income, such as income from social security or unemployment insurance and/or benefits. Thus, we argue that the high incidence of 0 yearly earnings in the PSID is likely not representative of the real outcomes for individuals. We therefore bottom code household head earnings that are near or equal to zero to a more realistic lower bound. Specifically, we follow the method adopted by Jensen and Shore (2015), using earnings from a half-time job (1000 hours per year) at the real equivalent of the 1991 federal minimum wage ($4.25 per hour). This also serves the purpose of preventing extreme observations of our measure of volatility caused by movements around $0.

Next, we construct a measure of the individual earnings volatility of each household head over the period. We use a standard variance formula and take the logarithm of earnings. Taking the logarithm is motivated by a desire to treat income shocks of the same proportion equally. For instance, an income shock from $30 000 to $15 000 will likely have a greater welfare impact on a household than a shock from $80 000 to $65 000. Let y be the logarithm of earnings; y_{it} be the value of y for individual i in year t; and $\overline{y}$ be the mean of y_{it} over period T_i. The volatility of earnings (v) is computed as follows:

$$\sigma_v^2 = \frac{1}{T_i - 1} \sum_{t=1}^{T_i} (y_{it} - \overline{y})^2$$

Many other studies use measures of volatility that are symmetric with respect to the measure of earnings, however we have no need for symmetry, as our focus is not to study the trends or various aspects of earnings volatility, but rather to use it to be able to categorize industries as having high or low inherent earnings volatility. Furthermore, most other studies have used more sophisticated methods of measuring income volatility, for instance using error component models that decompose income into persistent and transitory components, where volatility is the short-term fluctuations around the trend, as pioneered by Benus and Morgan (1972). For the sake of time and complexity, we do not make this distinction, and instead measure what could be termed 'gross' volatility. In the sections that follow, we refer to the measure σ_v^2 as earnings volatility, or simply volatility.

We discount the earnings of the household heads to 1990 dollars, using the CPI from the U.S. Bureau of Labor Statistics, which helps to ensure that the wage growth that can be explained by inflation does not enter our measure of volatility. Remember that while the period we use to study volatility of industries is 1991-2001, the data on income is lagging one year, meaning 1991 income is actually the income earned in 1990.

Table 1 reports summary statistics of our gross volatility measure over the period 1990-2000, and figure 1 visualizes the results in a histogram for the entire sample.

Table 1

Summary Statistics for 10-year Variance of Log Earnings

Statistic	Value
Mean	0.297
Median	0.072
SD	0.484
Min	0.000
Max	3.953
N	1994.000

Note: Data from the PSID, 1991 - 2001

Figure 1

Histogram of 10-year Earnings Volatility of Household Heads

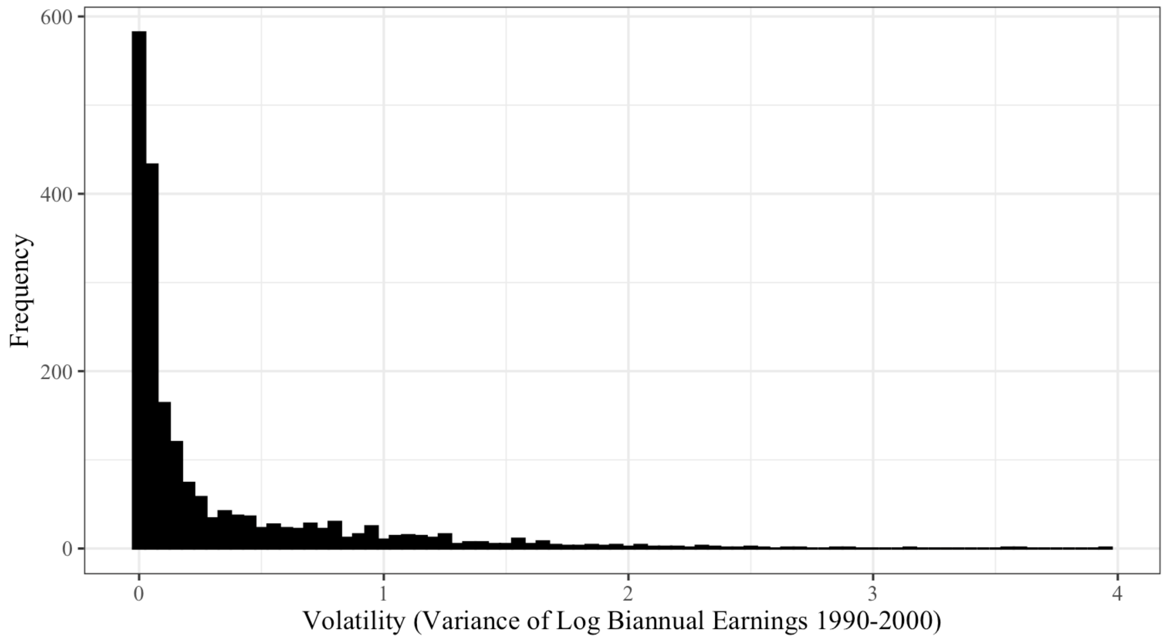

Note: Data from the PSID, 1991-2001

Next, we categorize the industries that individuals work in as high or low volatility industries, depending on whether the mean earnings volatility of individuals who worked in a given industry exceeds the total sample mean volatility.

Two of the twelve industry categories, namely the Mining industry and the Entertainment & Recreation industry, have very small samples (18 and 13 individuals, respectively). As these groups are too small for us to confidently estimate the average level of income volatility in them, we exclude them from our sample.

An issue with the approach of using industry-level earnings volatility to index individuals as facing high or low earnings volatility is that not all individuals stay in the same industry over the entire period. Table 2 shows the frequency of individuals changing industries during the period 1990-2000.

Table 2

Number of Individuals Who Change Industry During the Period 1990-2000

Number of Industry Changes	Number of Individuals
0	987
1	365
2	348
3	187
4	81
5	26

Note: Data from the PSID, 1991 - 2001

We could handle this issue by dropping individuals that work in more than one industry over the period 1990-2000. However, this strategy would lead to a substantial reduction in sample size, from 1,994 to 987 individuals. This diminished sample size could potentially undermine the robustness of our results, as several industry groups would become too small for a confident estimation of their inherent earnings volatility. In addition, the group of individuals who remain in a single industry over a decade might not accurately reflect the characteristics of the broader population. These individuals might possess unique attributes, distorting the representativeness of our sample. Given these considerations, we opt to classify an individual's industry based on the industry in which they spent the majority of the measured period. This approach allows us to maintain a more substantial and representative sample.

Table 3 reports the mean and median earnings volatility of each industry, as well as the number of individuals who spent the majority of the period working in them. The histogram in figure 2 visualizes the categorization of industries that follows from the results given in table 3.

Table 3

Table of Mean Volatility of Each Industry

Industry	Mean volatility	Median volatility	Number of individuals
Agriculture Forestry, Fisheries	0.6157727	0.4324182	57
Construction	0.4239578	0.1921785	198
Manufacturing	0.2315074	0.0514800	492
Transportation, Communications, Public Utilities	0.2159056	0.0481822	209
Wholesale and retail trade	0.2763633	0.0842558	270
Finance, Insurance, Real Estate	0.4549212	0.1200509	110
Business and Repair Services	0.4537766	0.1642409	83
Personal Services	0.4457261	0.1818677	31
Professional and Related Services	0.3044943	0.0588891	343
Public Administration	0.1696322	0.0343376	201
Total Sample	**0.2972651**	**0.0723381**	**1994**

Note: Data from PSID, 1991 - 2001

Figure 2

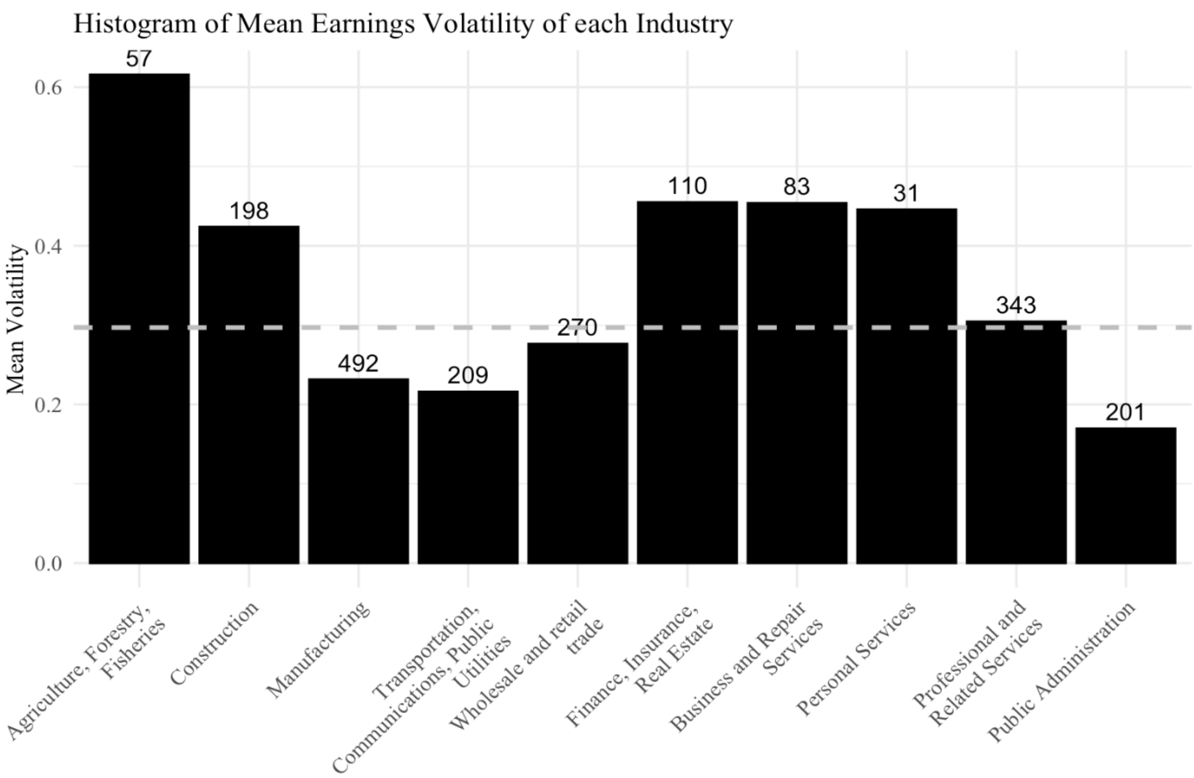

Note: Data from the PSID, 1991 - 2001

The dashed line in figure 2 represents the total sample mean earnings volatility, which is 0.297. We categorize Industries above this line as high volatility industries, and industries below this line as low volatility industries. We thus categorize Agriculture, forestry, and fisheries; Construction; Finance, insurance, and real estate; Business and repair services;

personal services; and Professional and Related Services as high volatility industries. Henceforth, we refer to individuals who work in these industries as the high volatility group, and others as the low volatility group.

For completeness, we include a sensitivity analysis in section A.5 in the Appendix , where we present the results of an alternative approach that excludes individuals who change industries. Importantly, our classification of industries into high and low volatility categories remains robust, demonstrating no significant sensitivity to the exclusion of individuals who transition between industries.

As becomes evident in figure 2, two substantial industries are proximate to the cutoff: Wholesale and Retail Trade, positioned slightly below, and Professional and Related Services, positioned slightly above. The proximity of these two industries, which collectively represent 30% of the entire sample, to the cutoff may result in an arbitrary classification as high or low volatility. Moreover, the categorization of all other industries remains robust when altering the central tendency measure and cutoff to the median instead of the arithmetic mean; however, these two industries switch positions. Wholesale and Retail Trade has a median volatility surpassing the total sample median volatility, while Professional and Related Services has a median below the total sample median volatility. This robustness check is important, because as can be inferred from both figure 2 and table 3, the mean volatility of both the total sample and each industry is significantly higher than the respective median volatility because there is a significant positive skew. No apparent justification exists for favoring one methodology over the other; therefore, to ensure the robustness of our approach of utilizing the arithmetic mean, we offer an alternative specification of the models in the second part, wherein we exclude these two industries, henceforth referred to as the robustness check sample.

Moving forward with the main sample, we provide summary statistics of the high and low volatility groups as defined above in table 4 and visualize the different realized earnings volatilities of individuals in the high and low volatility groups, respectively, with the box plots in figure 3.

Table 4

Summary Statistics of 10-year Biannual Earnings Volatility for High and Low Volatility Groups

Statistic	Low Volatility Group	High Volatility Group
Max	3.595	3.953
Mean	0.228	0.395
Median	0.055	0.123
Min	0.001	0.000
N	1172.000	822.000
SD	0.398	0.571

Note: Data from the PSID, 1991 - 2001

Figure 3

Box Plots of the Earnings Volatility of Household Heads in Low and High Volatility Industries

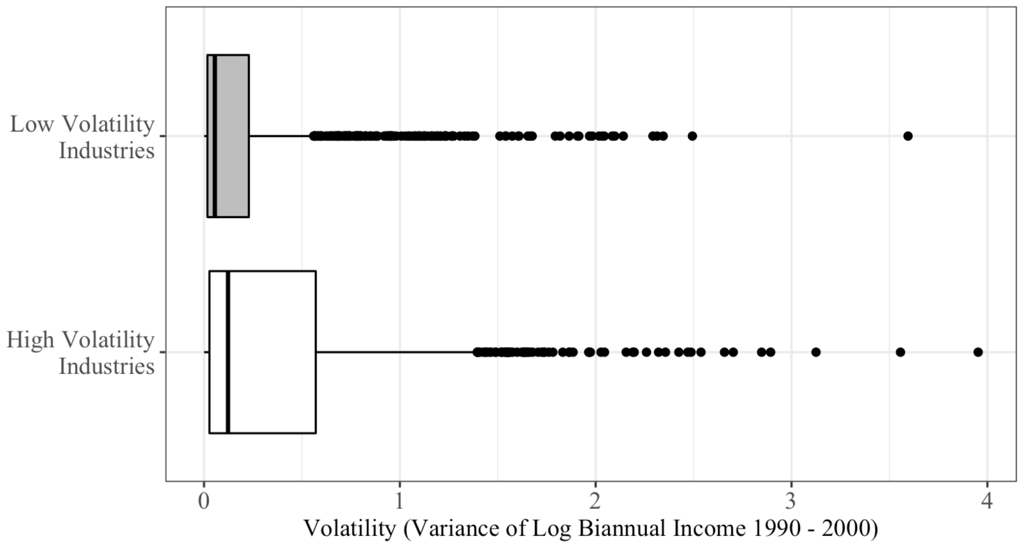

Note: Data from the PSID, 1991 - 2001

To visualize the argument that employing realized individual earnings volatility directly as an index for high or low volatility individuals might result in arbitrary classification, we present Figure 4. In the first plot of Figure 4, individuals are indexed as high or low volatility if their realized individual volatility between 1991 and 1995 is above or below the sample median individual volatility during the same period, cutting the sample in half. Subsequently, the rolling three-year mean variance of log earnings for the two groups is plotted separately. In the second plot of Figure 4, we apply the industry-level volatility method as previously described, but with a shorter time frame to facilitate comparison between the two methods.

The rationale for utilizing a reduced period for these plots is to illustrate how volatility evolves over time for both groups.

Figure 4

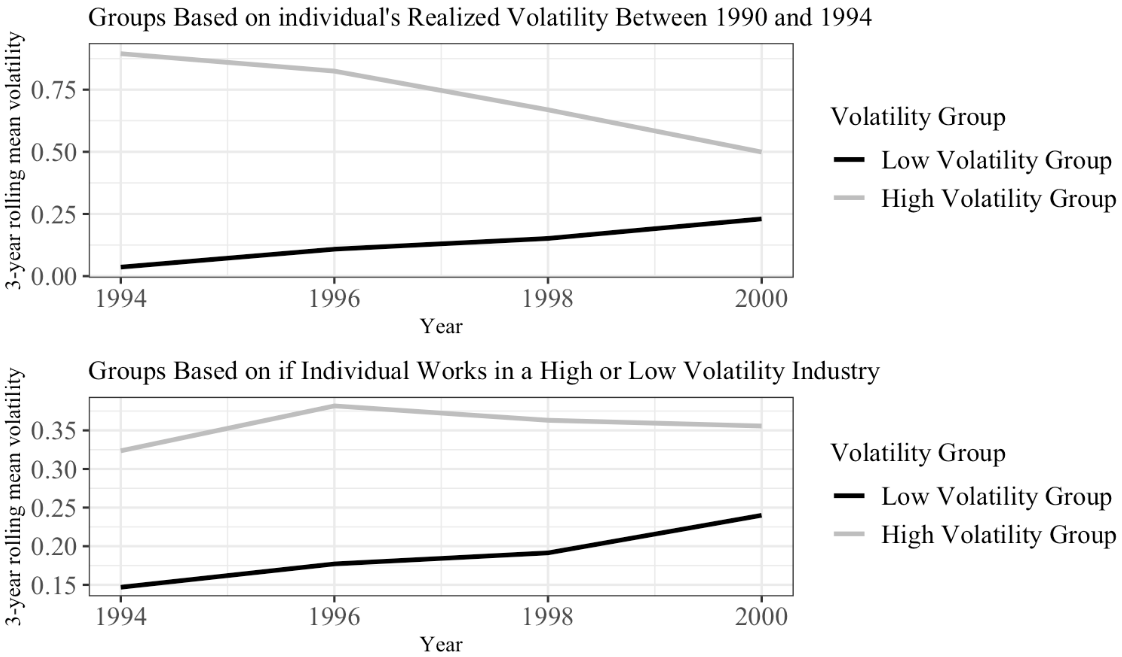

Note: Data from the PSID, 1991 - 2001

The plots in Figure 4 illustrate that directly using realized individual earnings volatility for classification results in arbitrary distinctions between individuals facing high or low volatility, as the difference in earnings volatility between the groups generated in this manner lacks structural consistency, as demonstrated by the strong convergence between the groups over time in the first plot in figure 4. By employing the industry-level method, the disparity in income volatility between the groups is evidently less sensitive to time, as they do not exhibit such a pronounced convergence, although the overall difference between the groups is smaller. This outcome is anticipated because numerous individuals in high volatility industries may not experience volatile income and vice versa, as illustrated in the box plots in figure 3. This figure further motivates our decision to employ industry-level volatility as a metric for distinguishing high or low volatility individuals.

5.2. Part Two - Estimating the Causal Effect of Working in a High Volatility Industry

We now turn toward the estimation of the effect of working in a high volatility industry. The motivation for employing a difference-in-differences (DiD) model is primarily based on such a model's ability to estimate a causal effect. Another key strength of the DiD methodology is its ability to control for time-invariant unobserved heterogeneity, which may be confounding the link between earnings volatility and borrowing. As our treatment and control groups are defined based on their employment in high and low volatility industries, respectively, they are likely to differ systematically in ways that we cannot observe in our data. By comparing the changes over time between these two groups, the DiD model helps to mitigate the influence of such characteristics.

Additionally, the DiD method allows us to exploit the natural experiment provided by the enactment of GLBA. The implementation of the GLBA can be interpreted as a shock of exogenous nature to credit availability, enabling us to compare borrowing behavior before and after the enactment. This temporal variation adds a layer of control for possible confounding factors.

Importantly, our hypothesis posits that households experiencing higher income volatility are more susceptible to income shocks, thereby increasing their demand for credit as a mechanism to buffer their consumption against the unpredictability of income, an argument bolstered by Krueger and Perri (2006). Consequently, we anticipate that these households exhibit a heightened latent demand for credit to facilitate consumption-smoothing. This underpins our hypothesis that reactions to the GLBA would differ according to the degree of income volatility, with households facing higher earnings volatility being more responsive due to their need to insulate consumption against idiosyncratic income shocks and potentially varying levels of previous credit access. This results in a stronger treatment effect for the high volatility group relative to those with more stable incomes. The DiD model is thus well suited to test this hypothesis, as it focuses on estimating differential changes over time between groups.

We estimate the level of income volatility faced by households using the earnings volatility of the primary income earners in these households, which is in turn gauged based on the industries in which they are employed. We categorize households as high volatility - forming our treatment group - if the industry of employment for the household head exhibits high inherent earnings volatility. The borrowing behavior of these households is then compared to that of our control group, which consists of households where the primary income earner is employed in industries characterized by low inherent earnings volatility.

As we only have observations of our debt variable from 1999 and later periods, we are confined to this period. We create a new subsample of our PSID dataset, placing the same sample restrictions as in part 1. This includes the requirement for full participation from 1999 to 2007, exclusion of students and retirees, and the removal of top and bottom 1% earners. As in part 1, we bottom code household head income to the equivalent of working 1000 hours with the federal minimum wage of $5.15, which was the federal minimum wage throughout the whole period 1999-2007, and discount nominal dollar variables to their real value in the first year, in this case 1999. This only has relevance because we control for income in our DiD regressions.

As we only have data on the industry variable up to 2001, we have to assume that individuals stay in the industry they worked in in 2001. We drop individuals who were unemployed in 2001, as well as the industry groups that were too small to estimate their inherent volatility, i.e. the individuals who worked in the Mining industry or the Entertainment industry in 2001. we index the individuals that worked in one of the industries that were identified as high volatility in part 1 in 2001 as high volatility by constructing a dummy variable that is equal to 1 if an individual worked in one of the industries Agriculture, Forestry, and Fisheries; Construction; Finance, insurance, and real estate; Business and repair services; personal services; or Professional and Related Services.

The groups in high and low volatility industries should ideally be as similar as possible in all other aspects other than the inherent volatility they face. We therefore provide summary statistics on income of the high volatility and low volatility groups in table 5, as we expect that from the variables we have available to us, income is the most important determinant of borrowing behavior. Summary earnings and income statistics of each individual industry can be found in section A.6 of the Appendix.

Table 5

Summary Statistics of Household Head Earnings and Total Household Income

Volatility Group	count	Mean Household Income	Median Household Income	Mean Head Earnings	Median Head Earnings
Low	1558	64117.76	54402.61	40341.52	34000.00
High	1307	66508.38	52388.33	39550.27	30724.21

Note: Data from the PSID, 1991 - 2001

In terms of earnings and household income, our volatility groups appear to exhibit relatively comparable characteristics. However, one notable distinction is observed in the high volatility group, where there is a greater positive skewness in the distribution of household head earnings, as evident from the fact that the median value deviates further from the mean value in this group.

5.3. Econometric Specification

Our econometric specification can be written as follows:

$$Y_{it} = \alpha + \beta_1 \times D_i + \beta_2 \times Post_t + \gamma \times D_i \times Post_t + \delta X_{it} + \varepsilon_{it}$$

where Y_{it} represents the total debt (or other debt, in the alternative specification) of individual i at time t; D_i is the time-invariant dummy variable for individuals working in high volatility industries in 2001 (1 if high volatility, 0 otherwise); $Post_t$ is a dummy variable equal to 1 for years after 1999, when the policy intervention was enacted, and 0 otherwise; $D_i \times Post_t$ is the interaction term between the high volatility industry dummy variable and the post 1999 control dummy. X_{it} is a vector of control variables, including log of yearly income, number of children, etc., for individual i at time t. ε_{it} is the error term. γ is the DID estimator, which captures the causal effect of working in a high volatility industry on borrowing behavior after the treatment in 2000 and thus the coefficient of interest in our study.

The vector of control variables include the logarithm of yearly earnings of the household head, the number of children in the household, and sex of the household head. These are all variables that we expect may correlate with borrowing behavior, the most important of which being income. The number of children in the household is also relevant as larger families likely tend to live in larger houses or apartments, thus likely requiring larger mortgages. There are more variables for which we have no data that we ideally would like to control for, however we defer this issue to the discussion section of the paper.

The choice of appropriate standard errors is crucial to obtain valid inference. We use cluster-robust standard errors clustered on the individual level. This helps dealing with autocorrelation of individuals, as well as heteroscedasticity across groups. Autocorrelation could otherwise violate the assumption that error terms are uncorrelated across observations. If we did not use clustered standard errors, the standard errors of the estimated coefficients would likely be biased, making hypothesis tests invalid. Robust clustered standard errors further allow for dependence within clusters, i.e. within individuals, while maintaining the assumption of independence between clusters.

As a standard robustness check, we test for multicollinearity of our control variables by estimating the Variance Inflation Factor (VIF) of each covariate in our regressions. Higher VIF values indicate higher multicollinearity between the independent variables. A common rule of thumb is that a VIF value above 10 indicates a potential multicollinearity problem. The VIF of every covariate is provided in section A.3 of the Appendix. The results give us no reason to suspect problematic multicollinearity between our independent variables.

6. Results

In this section, we present the results from our regressions. The results are divided into two primary segments: the main sample and the robustness check sample. For the main sample, two key variables are examined - 'total debt' and 'other debt' - which are depicted in separate tables. Each table comprises four models, with each model progressively adding a covariate to the regression, thereby allowing us to trace the incremental impact of each variable on the estimates.

In the second section of the results section, we present findings from the robustness checks sample. This section replicates the regressions from the main sample, but omits the two industries whose high or low volatility index was not robust to a change in the measure of central tendency used for the categorization of industries.

In addition, to further enhance our understanding, we provide graphical illustrations of how the average 'total debt' and 'other debt' evolved during the period 1999-2007 for individuals employed in high and low volatility industries. These plots, presented separately for both the main and robustness check samples, offer a visual representation of the trends and differences in borrowing behavior.

6.1. Main Sample Results

Table 6

Total Debt Regression Results - Full Sample

	Dependent variable: Total debt			
	(1)	(2)	(3)	(4)
Post	26,285.830***	26,184.900***	27,645.920***	27,665.220***
	(1,661.233)	(1,643.190)	(1,653.415)	(1,653.888)
Logarithm of Income		8,910.225***	8,655.387***	8,619.087***
		(1,467.193)	(1,452.604)	(1,459.672)
Number of children in household			7,450.216***	7,448.384***
			(1,300.478)	(1,300.874)
Sex of household head				-3,512.943
				(8,088.488)
DiD Estimate	1,702.018	1,213.605	1,145.060	1,141.755
	(2,650.295)	(2,628.180)	(2,596.088)	(2,596.690)
Observations	13,594	13,594	13,594	13,594
R^2	0.039	0.045	0.051	0.051

Note: Data from the PSID, 1999 - 2007 $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Table 7

Other Debt Regression Results - Full Sample

	Dependent variable: Other Debt			
	(1)	(2)	(3)	(4)
Post	2,417.950***	2,420.705***	2,309.109***	2,309.508***
	(379.336)	(379.523)	(389.931)	(391.775)
Log Income		-481.040	-463.177	-463.928
		(398.217)	(398.388)	(399.225)
Number of Children in household			-552.302*	-552.348*
			(315.164)	(315.065)
Sex of household head				-71.870
				(2,045.820)
DiD Estimate	-75.666	-47.978	-39.005	-39.100
	(585.216)	(586.160)	(586.528)	(586.808)
Observations	14,088	14,088	14,088	14,088
R^2	0.004	0.004	0.005	0.005

Note: Data from the PSID, 1999 - 2007 *p<0.1; **p<0.05; ***p<0.01

Figure 5

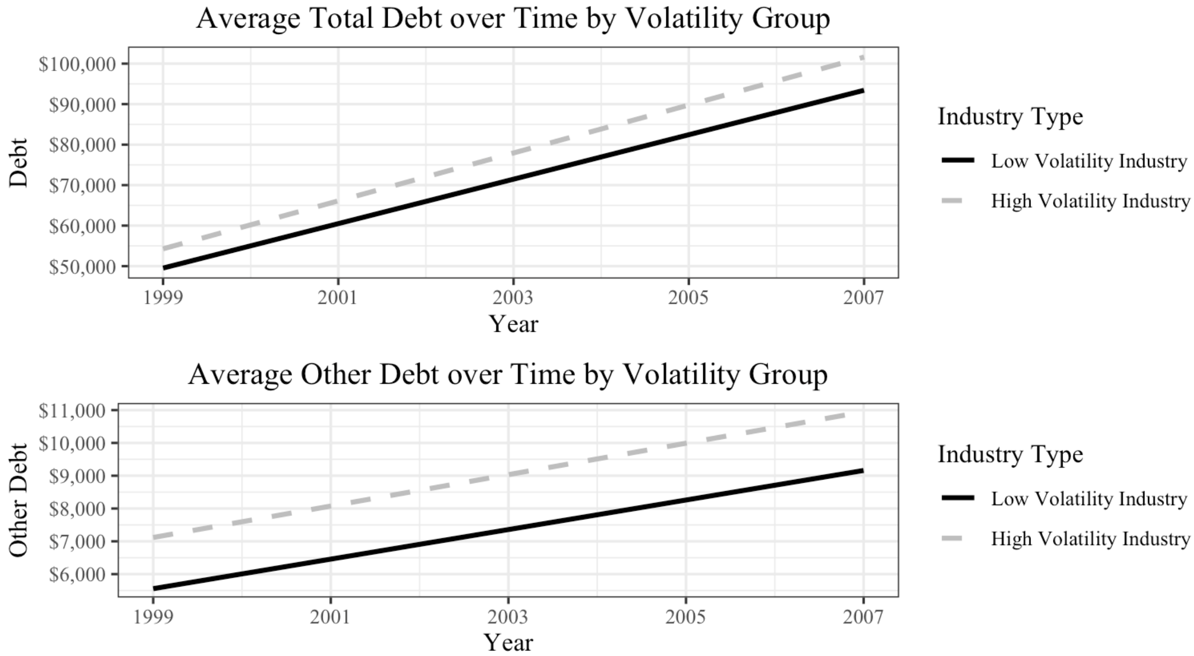

Note: Data from PSID, 1999 - 2007

Figure 5 presents the debt plots from the main sample, N = 2685.

6.2. Robustness Check Sample Results

In this section, we present the results from the same model specifications after dropping the industries that were not robust to a change in the methodology of the measure of central tendency of volatility.

Table 8

Total Debt Regression Results - Alternative Sample

	Dependent variable: Total Debt			
	(1)	(2)	(3)	(4)
Post	28,453.020***	28,394.220***	30,283.920***	30,296.830***
	(1,991.104)	(1,973.187)	(2,020.197)	(2,020.680)
Log Income		9,086.063***	8,850.738***	8,818.987***
		(1,812.855)	(1,793.633)	(1,807.934)
Number of children in household			7,961.000***	7,958.502***
			(1,728.313)	(1,728.823)
Sex of household head				-2,551.069
				(10,077.750)
DiD Estimate	-550.657	-1,066.074	-1,768.707	-1,767.689
	(3,410.175)	(3,387.255)	(3,352.348)	(3,352.478)
Observations	9,042	9,042	9,042	9,042
R^2	0.041	0.047	0.054	0.054

Note: Data from the PSID, 1999 - 2007 — *p<0.1; **p<0.05; ***p<0.01

Table 9

Other Debt Regression Results - Robustness Check Sample

	Dependent variable: Other Debt			
	(1)	(2)	(3)	(4)
Post	2,340.895***	2,340.958***	2,209.281***	2,195.096***
	(453.207)	(453.254)	(473.452)	(473.741)
Log Income		-28.005	-14.605	19.555
		(489.569)	(491.207)	(494.373)
Number of children in household			-529.204	-526.051
			(407.884)	(407.974)
Sex of household head				2,753.107
				(2,046.942)
DiD Estimate	530.166	531.721	585.437	585.596
	(714.702)	(718.941)	(721.918)	(721.798)
Observations	9,380	9,380	9,380	9,380
R^2	0.004	0.004	0.005	0.005

Note: Data from the PSID, 1999 - 2007 — *p<0.1; **p<0.05; ***p<0.01

Figure 6

Note: Data from the PSID, 1999 - 2007

Figure 6 presents the debt plots from the robustness check sample, N = 1909.

7. Interpretation of Results

7.1 Main Sample DiD on Total Debt - Table 6

The results of our baseline specification using the full sample are provided in table 6. The coefficient of the interaction term between post GLBA enactment and working in a high volatility industry is our estimated causal effect of working in a high volatility industry on total debt during the period. The estimated coefficient, termed DiD estimate, is $1,702 in the regression without any of the covariates, but drops to around $1,200 when the covariates are introduced. As expected, the covariates income and number of children in the household both seem to be relevant for explaining borrowing behavior, with high economic and statistical significance. However, the sex of the household head does not seem to be an important determinant of the household's borrowing behavior.

None of these estimated treatment effects (DiD estimates) are particularly economically significant, but more importantly, they are all completely statistically insignificant, with a standard error more than twice as high as the coefficient when covariates are introduced. We are therefore unable to reject the null hypothesis that the causal effect of the household head of a household working in a high volatility industry on the borrowing behavior of households is different from zero.

While the statistical power of our estimate is too low to lay any causal claim, it is still worthwhile to consider the fact that we found a positive estimate in all specifications of the regression. If it had been statistically significant, our results from the main specification regression using the full sample with total debt as dependent variable, would suggest that Individuals who work in industries with higher inherent earnings volatility responded to an increase in credit availability by increasing their borrowing by more than their counterparts in industries with low inherent earnings volatility. Such a finding would support the work of Krueger and Perri (2006), who used increased individual earnings volatility to explain the relationship between increased inequality on credit growth.

7.2 Main Sample DiD on Other Debt - Table 7

Table 7 reports our results from when we run the same regression using Other Debt, which captures non-housing related debt, as the dependent variable instead of total debt, for the purpose of investigating the potential psychological relationship between earnings volatility and credit card debt. The results are weak in this specification as well. The estimate of the causal effect is even not particularly economically significant in this case either, at between - $75.7 and - $39, depending on how many covariates we include. Interestingly, it has the opposite sign from what the literature would suggest. However, the standard errors are an order of magnitude larger than the estimates, and thus we fail to reject the null hypothesis that the causal effect is different from zero.

7.3 Interpretation of Debt Plots From Main Sample - Figure 5

In figure 5 we display the average Debt and Other Debt, respectively, over time for individuals in high and low volatility industries. The plots do not reveal any surprising differences in trends between the groups, given the results of the regressions. However, we see that the intercept of the mean debt level, both for total debt and other debt, is higher for the low volatility industries. In any case, we are interested in the difference in trends and not the levels in particular.

7.4 Robustness Check Sample DiD on Total Debt - Table 8

The results of our baseline specification using the robustness check sample when we have dropped two industries are provided in table 8. The DiD estimate in this case is negative regardless of which covariates are introduced. With all covariates included, it is - $1,768, which should be interpreted as the amount of debt that households' whose head worked in a high volatility industry accumulated over the period in excess of the control group. It is concerning that this estimate is of the opposite sign from the main sample estimate. However, the standard errors are very large, at around twice to three times the size of the coefficient when covariates are introduced, which means that the results are again completely statistically insignificant.

As we found was the case for the main sample results, the covariates income and number of children in the household still both seem to be relevant for explaining borrowing behavior, with high economic and statistical significance. Again, the sex of the household head does not seem to be an important determinant of households' debt accumulation.

7.5 Robustness Check Sample DiD on Other Debt - Table 9

Table 9 reports the results from when we run the same regression using Other Debt using the robustness check sample. The signs of the DiD estimates are again opposite to the results from the main sample regressions on other debt, with positive coefficients instead of negative coefficients. They are much more economically significant using this specification, at between $530 and $586 in all four models. However, yet again the results are not statistically significant, although in this case the standard error is not as large as in the other model specifications. Unlike the results from the main sample regression on other debt, these estimates are in line with what the literature suggests regarding the relationship between income volatility and credit card borrowing.

7.6 Interpretation of Debt Plots From Alternative Sample - Figure 6

The plot of total debt over time in figure 6 again does not reveal anything surprising given the results of the difference-in-difference estimates using the alternative sample. However, the plot of other debt over time corroborates the positive estimated treatment effect. We see that average other debt increased at a higher rate for the households whose heads worked in a high volatility industry, thus providing some superficial support for the notion that individuals or households that face higher income volatility may be more inclined to take on credit card debt.

8. Discussion

The main sample regression results demonstrate a modest positive causal effect of employment in high volatility industries on total debt accumulation over the period under study. However, the robustness check regressions - where we removed industries that were not consistent when altering the central tendency measure used to categorize industries into high and low volatility groups - displayed a minor negative causal effect. This inconsistency alone would hinder us from concluding that our investigation uncovered empirical evidence supporting a causal connection between earnings volatility and credit behavior. However, outcomes of both the main and robustness check regressions were statistically insignificant, further amplifying the ambiguity of our findings. In light of these results, we cannot assert a definitive causal relationship due to the lack of statistical significance.

There are two plausible interpretations for these findings. The first is that there in fact is no relationship between earnings volatility and credit behavior. The second is that our methodological or sampling approach may have been flawed, thus inhibiting our ability to establish a relationship. Given our initial conviction that a relationship should exist, and our acute awareness of the issues in our methodology, we will use the discussion section to focus on potential reasons why our methodology may have been unsuccessful in generating statistically significant and robust evidence.

8.1. Threats to the Validity of the DiD Design

The critical assumption of a difference-in-differences study is the parallel trends assumption, which refers to the assumption that in absence of the treatment, the treatment and control group would have followed the same trend. This assumption is necessary for unbiased difference-in-differences estimation. In our case, this assumption posits that households whose heads work in high volatility industries would have increased their leverage at the same rate as the households whose heads work in low volatility industries after the enactment of the GLBA. This assumption is by definition untestable as we cannot observe the counterfactual conditional expectation. A widely used method to provide support for the parallel trends assumption is to simply display the raw data prior to the treatment period and show that the treatment and control group follow similar trends. As we have explained, we do not have data on the debt levels of our sample before 1999, which is the only pre-treatment period. We are thus unable to make this analysis. This is a serious weakness of our difference-in-differences approach. However, even if we were able to display the pre-treatment trends, there is no way to be sure that counterfactual trends in borrowing behavior would have been the same for the groups post-treatment without further assumptions about the predictive power of pre-treatment trends.

Another critical threat to the parallel trends assumption is the exogeneity of treatment. If treatment is endogenous, i.e. it is non-randomly assigned and influenced by unobservable factors correlated with borrowing behavior, the parallel trends assumption is violated. In our context, this refers to potential self-selection into industries that have higher or lower income volatility based on factors that are also correlated with borrowing behavior. Sources of such self-selection bias could for instance be that individuals with higher education levels or different risk preferences are more likely to choose to work in industries with high income volatility, and that these unobserved factors are correlated to borrowing behavior. That would be a source of self-selection issues, leading to endogenous treatment and biased estimates. However, we argue that the choice of which industry to work in should not be strongly correlated to risk preferences, as the degree of industries' inherent income volatility is not directly observable to individuals when making career choices. Moreover, this is a choice that is based on many other factors, such as interests, skills, educational background, and geographic location. Therefore, we believe that the endogeneity concerns in our study are mitigated to a certain extent.

To further address potential endogeneity concerns, we should have controlled for more observable characteristics that might be correlated with both industry choice and borrowing behavior, such as education and potentially even risk preferences, although the latter is inherently more difficult to obtain reliable data on. For instance, the high_volatility dummy may be endogenous, as individuals may choose to work in high or low volatility industries based on preferences or abilities. This could bias the estimates of the treatment effect and we thus acknowledge that unobserved factors may affect our results.

Another threat to the validity of the DiD design is the presence of spillover effects. Individuals in high volatility industries clearly face more earnings volatility, on average, than individuals in low volatility industries. However, as can be seen in figure 3 and in the summary statistics of income volatility for our high and low volatility groups in table 4, and as can be deduced intuitively, there are many individuals in the high-volatility industries that experienced low earnings volatility, and vice versa. The presence of spillover effects contaminates our treatment and control groups, which can lead to biased estimates of the treatment effect, making causal inference less reliable. The implication for future research is that it is important to address the issue of spillover effects by using alternative identification strategies, restricting the sample to a more clearly defined treatment and control group, or including additional control variables to account for potential sources of spillover. One concrete way to address this would be to use a matching or propensity score method to control for observable differences between the two groups, an approach that we unfortunately did not have time to pursue.

As the industries that were eliminated in our robustness check sample had mean and median volatility closest to the total sample mean and median, they are likely the ones affected most strongly by spillover effects. The effect of spillover effects is thus likely more mitigated in our robustness-check sample, which showed a (statistically insignificant) negative DiD estimated coefficient of the effect of income volatility on total debt accumulation. This provides an argument for the case that this is the specification we should rely on for causal inference rather than the main sample estimate, which was positive. However, as mentioned, the inconsistency of the main and robustness check samples, as well as the statistical insignificance of the DiD estimates, prevent any such claims.

8.2. Issues with our Measure of Income Volatility

A weakness of our study is the fact that PSID is biannual, meaning we lose half of the yearly earnings realizations. Furthermore, even yearly intra-year income volatility can mask month-to-month variation. For example, a worker in a seasonal industry may earn the majority of their yearly income in a select few months of the year, and have low earnings in the rest of the year. Any PSID measure of earnings volatility will thus not capture the full income variation that some individuals experience. However, this limitation should not be detrimental to our efforts to find significant evidence, as it is most probably larger income shocks that would create a need for consumption insulating credit rather than smaller monthly fluctuations.

Another weakness of our volatility measure is the fact that gross volatility only captures deviations around the mean earnings over the period, thus potentially biasing it toward capturing volatility that occurred in the beginning and end of the measured period, if incomes were growing. Taking the gross volatility of different industries could thus capture shocks that are not transitory, but permanent. Therefore, if some industries experienced higher wage growth over the period, these increases will enter our measure of volatility, potentially exaggerating the inherent earnings volatility of these industries. As correct indexation of industries as characterized by high or low earnings volatility is crucial to our methodology, this could be contributing to our inability to find significant results.

Another aspect of weakness stems from the fact that we are using the mean volatility as the cutoff. It's possible that using a different cutoff or defining the groups based on a different measure of income volatility could yield different results. Therefore, it could be worth exploring alternative ways to define the groups and testing the sensitivity of the results to these definitions since the definitions could be interpreted as rather arbitrary.

8.3. Issues Related to the PSID Sample

Another aspect that is important to address is the issue of sample attrition, which as we mentioned in section 4.2. seems to be correlated with past earnings volatility. This issue may introduce selection bias into the study, as those with higher income volatility, potentially those working in high volatility industries, may be underrepresented in the sample. This selectivity may bias our estimates of the DiD estimate, as the individuals most affected by the treatment may be systematically missing from the sample. Our sample restriction, which dropped individuals lacking data for any year in the study period, may have exacerbated this issue. To mitigate this issue in future research, it would be beneficial to employ strategies aimed at minimizing data loss, such as imputation methods for missing data, or adjustments to the sample restrictions to include individuals with partial data. This would ensure a more representative sample, enhancing the robustness of the findings.

Furthermore, when we run the regression in part 2 we have to remove individuals who were unemployed in 2001 because we don't know from which industry they were unemployed. If we had data on the industry they worked in after 2001, we may have found that many of these individuals went back to work in a given industry, to which they “belonged” even in 2001 despite being unemployed. As we motivated previously, being unemployed is an important outcome and should be counted towards the volatility of industries. This could be contributing to the insignificance of our results.

9. Conclusion

We were unable to identify statistically significant evidence indicating that households with heads employed in high volatility industries exhibit different borrowing behavior when credit access increases. This absence of evidence, however, should not be misconstrued as proof to the contrary. Our methodology contains several limitations that collectively diminish the likelihood of obtaining statistically significant results. The most critical of these limitations include reliance on the household head's attributes alone for determining the volatility index, despite households increasingly depending on multiple earners, as well as our rudimentary measure of volatility. It is plausible, and indeed probable, that our conclusion regarding which industries are characterized by high or low inherent income volatility would differ if a more sophisticated approach to estimating individual earnings volatility were employed. However, even if that were the case, the industry-level volatility may not be individual-specific enough for an analysis of this kind, where spill-over effects make it more difficult to find statistically significant and unbiased estimates.

There are compelling theoretical justifications for the impact of earnings volatility on household borrowing decisions. While our study has not been able to establish a relationship, this should not deter further investigations into this subject area. We encourage researchers to persist in examining the relationship between volatility and borrowing behavior and hope that our research can offer guidance on potential pitfalls to avoid.

10. Reference list

Benus, J., & Morgan, J. (1972). Income instability as a dimension on welfare *Proceedings of the American Statistical Association, 102–6*. http://www.asasrms.org/Proceedings/y1972/Income%20Instability%20As%20A%20Dimension%20Of%20Welfare.pdf

Bertrand, M., & Morse, A. (2016). Trickle-Down Consumption, *The Review of Economics and Statistics, MIT Press, 98(5): 863-879.* https://www.nber.org/papers/w18883

Bordo, M. D., & Meissner, C. M. (2012, March 1). Does inequality lead to a financial crisis?, *Journal of International Money and Finance, Elsevier,* 31(8): 2147-2161. https://www.nber.org/papers/w17896

Borio, C., & Lowe, P. (2002, July 2). Asset prices, financial and monetary stability: Exploring the Nexus, *No 114, BIS Working Papers, Bank for International Settlements.* https://www.bis.org/publ/work114.html

Carroll, C., & Samwick, A. (1998). How Important Is Precautionary Saving? The Review of Economics and Statistics, 80(3): 410-419 https://www.jstor.org/stable/2646749

Chang, Y., Hong, H. J., Karabarbounis, M., Wang, Y., & Zhang, Tao. A. (2022) Income volatility and portfolio choices, Review of Economic Dynamics, 44(1): 65-90, https://ideas.repec.org/a/red/issued/20-409.html

Dahl, M., DeLeire, T., & Schwabish, J. (2011). Estimates of Year-to-Year Volatility in Earnings and in Household Incomes from Administrative, Survey, and Matched Data. *Journal of Human Resources. University of Wisconsin Press, 46(4): 750-774.*https://econpapers.repec.org/article/uwpjhriss/v_3a46_3ay_3a2011_3aiv_3a1_3ap_3a750-774.htm

Diana, F., Greig, F., & Hamoudi, A. (2019). The Evolution of the Online Platform Economy: Evidence from Five Years of Banking Data. *AEA Papers and Proceedings*, 109: 362-66. https://www.aeaweb.org/articles?id=10.1257/pandp.20191040

Diaz-Serrano, L. (2015). Labor income uncertainty, skewness and homeownership: A panel data study for Germany and Spain*, Journal of Urban Economics*, 58(1): 156-176 https://econpapers.repec.org/article/eeejuecon/v_3a58_3ay_3a2005_3ai_3a1_3ap_3a156-176.htm

Duffy, D. (2011). 2007 PSID Income and Wage Imputation Methodology. *Institute for Social Research University of Michigan. Technical Series Paper #11-03 https://psidonline.isr.umich.edu/publications/Papers/tsp/2011-03_Duffy.pdf*

Dynan, K., Mian, A., & Pence, K. M. (2012). Is a Household Debt Overhang Holding Back Consumption?. *Brookings Papers on Economic Activity*, 299–362. http://www.jstor.org/stable/23287219

Federal Reserve Bank of Minneapolis. (n.d.). *Consumer Price Index, 1913-*. https://www.minneapolisfed.org/about-us/monetary-policy/inflation-calculator/consumer-price-index-1913-

Federal Reserve Bank of New York. (2022). *Household Debt and Credit Report.* https://www.newyorkfed.org/microeconomics/hhdc

Fisher, I. (1933). *The debt-deflation theory of great depressions*. Econometrica, 1(4): 337-357. https://doi.org/10.2307/1907327

Fisher, P. J. (2010). Income uncertainty and household saving in the United States. *Family Consum. Sci. Res. J.* 39, 57–74. doi: 10.1111/j.1552-3934.2010.02045.x

Fitoussi, J., & Saraceno, F. (2010). Inequality and Macroeconomic Performance, No 2010-13, *Documents de Travail de l'OFCE, Observatoire Francais des Conjonctures Economiques (OFCE).* https://econpapers.repec.org/paper/fcedoctra/1013.htm

Friedman, M. (1957). *Theory of the Consumption Function.* Princeton University Press. http://www.jstor.org/stable/j.ctv39x7zh

Gottschalk, P., & Moffitt, R. (2009). The Rising Instability of U.S. Earnings. *Journal of Economic perspectives. 23 (4): 3-24.* https://www.aeaweb.org/articles?id=10.1257/jep.23.4.3

Gottschalk, P., Moffitt, R., Katz, L. F., & Dickens, W. T. (1994). The Growth of Earnings Instability in the U.S. Labor Market. *Brookings Papers on Economic Activity*, *1994*(2): 217–272. https://doi.org/10.2307/2534657

Grant, K. J. (2010). What the Financial Services Industry Puts Together Let No Person Put Asunder: How the Gramm- Leach-Bliley Act Contributed to the 2008 - 2009 American Capital Markets Crisis, *Journal Publications Florida A&M University College of Law. 73 Alb. L. Rev. 371* https://commons.law.famu.edu/cgi/viewcontent.cgi?article=1088&context=faculty-research&httpsredir=1&referer=

Hyman P., & Minsky, Ph. D. (1977). A theory of systemic fragility. *Bard Digital Commons.* https://digitalcommons.bard.edu/hm_archive/231/

Karen, D., Elmendorf, D., & Sichel, D. (2012). The Evolution of Household Income Volatility. *The B.E. Journal of Economic Analysis & Policy, 12 (2): 1-42.* https://econpapers.repec.org/article/bpjbejeap/v_3a12_3ay_3a2012_3ai_3a2_3ap_3a1-42_3an_3a3.htm

Kindleberger, C. P. (1978). *Manias, Panics, and Crashes: A History of Financial Crises.* New York: Basic Books.

Krueger, D., & Perri, F. (2006). Does Income Inequality Lead to Consumption Inequality? *The Review of Economic Studies, 73(1): 163–193.* https://www.jstor.org/stable/3700621

Kumhof, M., Rancière, R., & Winant, P. (2015). Inequality, Leverage, and Crises. *The American Economic Review*, *105*(3): 1217–1245. http://www.jstor.org/stable/43495415

Kopczuk, W., Saez, E., & Song, J. (2010) Earnings Inequality and Mobility in the United States: Evidence From Social Security Data Since 1937. *The Quarterly Journal of Economics, 125(1): 91–128.* *https://eml.berkeley.edu/~saez/kopczuk-saez-songQJE10mobility.pdf*

Levine, A. S., Frank, R. H., & Dijk, O. (2014), Expenditure Cascades, *Review of Behavioral Economics, 1(1–2): 55-73.* http://dx.doi.org/10.1561/105.00000003

McDonald, O. (2019). *Holding bankers to account.* Manchester University Press.

Meier, S., & Sprenger, C. (2010). Present-Biased Preferences and Credit Card Borrowing. *American Economic Journal: Applied Economics*, 2(1): 193-210 https://www.aeaweb.org/articles?id=10.1257/app.2.1.193

Mian, A., Straub, L., & Sufi, A. (2019). The saving glut of the rich and the rise in household debt. *scholar.harvard.edu.* https://www.rba.gov.au/publications/workshops/research/2019/pdf/rba-workshop-2019-sufi.pdf

Moffitt, R. A., & Zhang, S. (2018, March 12). Income volatility and the PSID: Past research and new results. *AEA Papers and Proceedings*, 108: 277-80 https://www.nber.org/papers/w24390

Morris, P., Hill, H., Gennetian, L., Rodrigues, C., & Wolf, S. (2015). Income Volatility In U.S. Households With Children: Another Growing Disparity Between The Rich And The Poor? *University of Wisconsin* https://www.irp.wisc.edu/resource/income-volatility-in-u-s-households-with-children-another-growing-disparity-between-the-rich-and-the-poor/

Panel Study of Income Dynamics, public use dataset. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2023).

Peetz, J., Robson. J., & and Xuereb, S. (2021). The Role of Income Volatility and Perceived Locus of Control in Financial Planning Decisions. *Front. Psychol.* 12:638043. https://www.frontiersin.org/articles/10.3389/fpsyg.2021.638043/full

Rachel, L., & Summers, L. (2019). On Falling Neutral Real Rates, Fiscal Policy, and the Risk of Secular Stagnation, *Brookings Papers on Economic Activity*, Spring, 1-76. https://www.brookings.edu/wp-content/uploads/2019/03/On-Falling-Neutral-Real-Rates-Fiscal-Policy-and-the-Risk-of-Secular-Stagnation.pdf

Rajan, R. G. (2011). *Fault lines: How hidden fractures still threaten the world economy*. Princeton, New Jersey.

Schularick, Moritz, & Alan M. Taylor. (2012). Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008. *American Economic Review*, 102 (2): 1029-61. https://www.aeaweb.org/articles?id=10.1257%2Faer.102.2.1029

Shin, D., & Solon, G. (2011). Trends in men's earnings volatility: What does the Panel Study of Income Dynamics show? *Journal of Public Economics, 95, (7-8): 973-982.* https://econpapers.repec.org/article/eeepubeco/v_3a95_3ay_3a2011_3ai_3a7-8_3ap_3a973-982.htm

Stiglitz, E. J. (2016). How to Restore Equitable and Sustainable Economic Growth in the United States. *American Economic Review, 106 (5): 43-47.* https://www.aeaweb.org/articles?id=10.1257/aer.p20161006

The World Bank. (n.d.). *Gini Index - United States.* https://data.worldbank.org/indicator/SI.POV.GINI?end=2019&locations=US&start=1974

West, C., Whillans, A. V., & DeVoe, S. (2020). Income Volatility Increases Financial Impatience. *Cambridge, MA: Harvard Business School* https://www.hbs.edu/ris/Publication%20Files/21-053_54843a1e-7599-4981-a829-bac3234830cc.pdf

White, J, L. (2010). Credit rating agencies and the financial crisis: Less regulation of CRAs is a better response, *Journal of international banking law. 30(3)* https://www.stern.nyu.edu/sites/default/files/assets/documents/con_039549.pdf

Ziliak, J., Hardy, B., & Bollinger, C. (2011). Earnings volatility in America: Evidence from matched CPS. *Labour Economics,* 18(6): 742-754. https://econpapers.repec.org/article/eeelabeco/v_3a18_3ay_3a2011_3ai_3a6_3ap_3a742-754.htm

Appendix

A.1. Sample restrictions imposed in part 1

Table A.1

Sample Restrictions for Part 1

Restrictions	Households.Eliminated	Households.Remaining
Full sample	0	8775
Full participation	5389	3386
No household head younger than 25 in 1991	183	3203
No household head older than 60 in 2001	919	2284
Remove top and bottom 1% earners	90	2194
Remove households whose head was unemployed for majority of period 1991-2001	169	2025
Remove households whose head worked in mining or entertainment	31	1994

Note: Data from the PSID, 1991 - 2001

A.2. Sample restrictions imposed in part 2

Table A.2

Sample Restrictions for Part 2

Restrictions	Households.Eliminated	Households.Remaining
Full sample	0	6799
Full participation	1961	4838
No household head younger than 25 in 1999	310	4528
No household head older than 60 in 2007	1175	3353
Remove top and bottom 1% earners	189	3164
Remove households whose head was unemployed in 2001	253	2911
Remove households whose head worked in mining or entertainment in 2001	46	2865
Remove households whose head worked in Wholesale and Retail Trade or Professional and Related Services in 2001	956	1909

Note: Data from the PSID, 1999 - 2007

A.3. Variable Inflation Factor Test

Table A.3

Variance Inflation Factor (VIF) Test Results

Variable	VIF
post	1.86
sex	1.09
Number of children in the household	1.01
Log of household head earnings	1.07
DiD estimator	5.83

Note: Data from the PSID, 1999 - 2007.
This table presents the VIF test results for the independent variables in the main model specification of our model, i.e. with the full sample N = 2865, including the Wholesale and Retail Trade and Professional and Related Services industries.

A.4. PSID Variable description

Table A.4

Variable	Description/Interview Question	PSID Label
First Mortgage	Do you have a mortgage or loan on this property? About how much is the remaining principal on this mortgage?	A24 REM PRINCIPAL MOR 1
Second Mortgage	About how much is the remaining principal on this mortgage?	A24 REM PRINCIPAL MOR 2
Other Debt	Aside from the debts that we have already talked about, like any mortgage on your main home or vehicle loans–do you (or anyone in your family) currently have any other debts such as for credit card charges, student loans, medical or legal bills, or on loans from relatives? If you added up all of these debts (for all of your family), about how much would they amount to right now?	W39 VALUE ALL DEBTS
Earnings of household head	How much did you (head) earn altogether from wages or salaries in [year], that is, before anything was deducted for taxes or other things?	G13 WAGES/SALARY OF HEAD
Total Household income	This variable is the sum of the variables below: taxable income of head and wife, transfer income of head and wife, taxable income of other family unit members (OFUMs), transfer income of OFUMs, and Social Security income.	TOTAL FAMILY INCOME
Industry	What is your/head's main occupation? What kind of business or industry is that in?	B10 MAIN INDUSTRY: 3 DIGIT (HD-E)
Relationship to head	10 if head	RELATION TO HEAD [year]
Age of household head	This variable represents the actual age of the head of the FU. The range of values is usually from 18 through 98, although in rare cases a person under 18 might become head.	AGE OF HEAD
Number of children in household	This variable represents the actual number of persons currently in the FU who are neither Head nor Wife/'Wife' from newborns through those 17 years of age, whether or not they are actually children of the Head or Wife/'Wife.	# CHILDREN IN FU
Sex of household head	Sex of [year] head	SEX OF HEAD

Note: Data from PSID variable search

A.5. Table of mean volatility of industries after dropping industry changers

Table A.5.

Table of Mean Volatility of Each Industry After Dropping Individuals Who Change Industry

Industry	Mean volatility	Median volatility	Number of individuals
Agriculture Forestry, Fisheries	0.8036715	0.6137808	33
Construction	0.3806429	0.1646996	94
Manufacturing	0.1750840	0.0326326	223
Transportation, Communications, Public Utilities	0.1794014	0.0258698	103
Wholesale and retail trade	0.2558452	0.0669453	105
Finance, Insurance, Real Estate	0.3331874	0.0835890	60
Business and Repair Services	0.5413398	0.3745900	20
Personal Services	0.4156473	0.1081748	17
Professional and Related Services	0.3058651	0.0439133	216
Public Administration	0.1315902	0.0212858	116
Total Sample	**0.2694050**	**0.0467150**	**987**

Note: Data from PSID, 1991 - 2001

Figure A.5.1

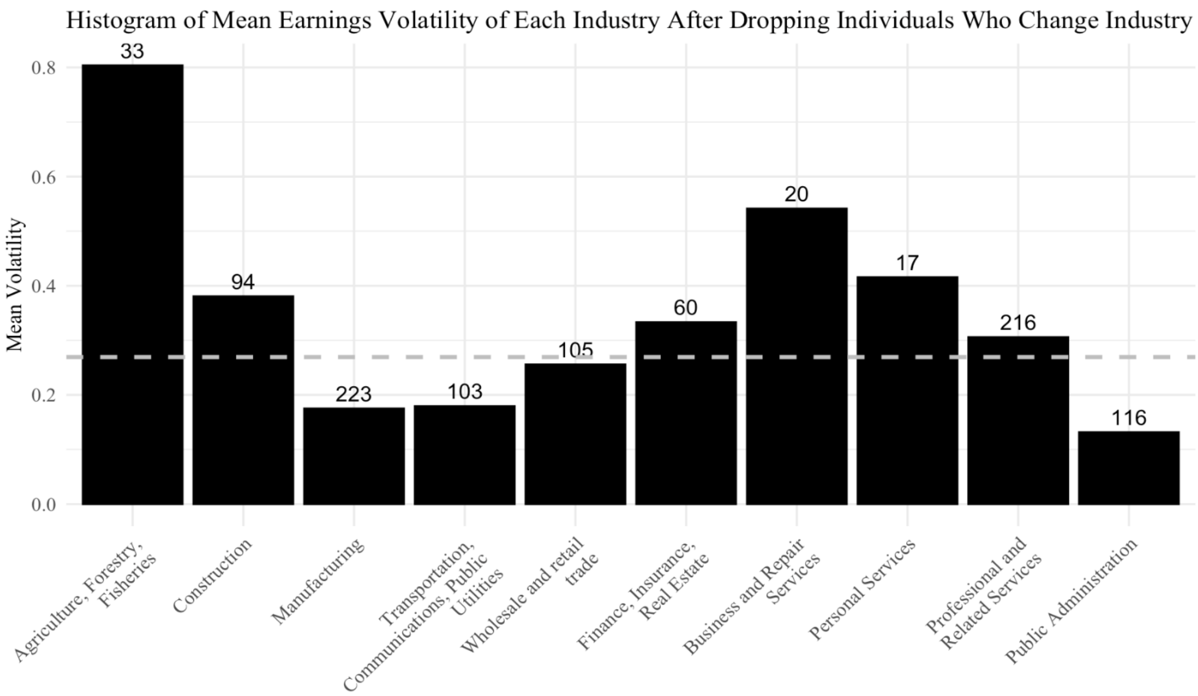

Note: Data from PSID, 1991 - 2001

A.6. Mean and Median Head Earnings and Household Income of Each Industry

Table A.6

Summary Statistics of Household Head Earnings and Total Household Income per Industry

Industry	count	Mean Household Income	Median Household Income	Mean Head Earnings	Median Head Earnings
Agriculture Forestry, Fisheries	92	61326	44062	28788	21357
Construction	238	67241	55826	37947	31064
Manufacturing	583	62364	53602	40937	33689
Transportation, Communications, Public Utilities	281	70837	61612	45126	38054
Wholesale and retail trade	426	57523	47459	33270	27623
Finance, Insurance, Real Estate	145	80956	63886	49613	36242
Business and Repair Services	231	73702	59831	46129	33958
Personal Services	71	41276	26513	22066	17668
Professional and Related Services	530	63371	49586	38860	32000
Public Administration	268	71371	61422	45271	41901

Note: Data from the PSID, 1991 - 2001