Student Loan Supply, Parental Saving & Portfolio Allocation

Emiel Jerphanion*

October, 2019

Abstract

I show that an expansion of student loan supply affects parents' saving decisions and portfolio allocation. By exploiting policy-induced variation on expected student aid, I find a 2.2 pp increase in the parental saving rate, from 4.9% to 6.1%. The mechanism that drives this result is the positive effect of student aid on students' college enrollment. Consistent with this interpretation, I find a disproportionate increase in college enrollment for children of families affected by the reform. The positive saving response is largest among lower- and middle-income families and for parents with strong saving preferences. A placebo test validates that the effect is absent in families without children. Moreover, I show that affected parents shift the allocation of saving flows towards riskier assets.

Keywords: Student loans, household finance, saving behavior, portfolio choice

JEL classification: D13, D14, E21, G11, I22

*Emiel Jerphanion: Finance Department, Tilburg University. I am very grateful to my advisors Fabio Braggion and Rik Frehen for invaluable comments and support. I also like to thank Pat Akey, Julio Crego, Adeline Delavande, Jens Kvaerner, Wilbert van der Klaauw and Nicola Pavanini as well as all seminar participants at Tilburg University and Young Finance Scholar Conference. I am responsible for all remaining errors or omissions. Please sent correspondence to e.jerphanion@tilburguniversity.edu.

1 Introduction

The swift rise in student aid in recent years made educational debt the largest non-mortgage liability for U.S. households (Brown et al., 2014). From 2000-2001 to 2016-2017 the total (average per student) aid grew from $108.5B to $240.1B ($9,762 to $16,343).[1] This exponential growth has attracted the interest of economists and policymakers, as high levels of educational debt may adversely affect students' future consumption, investment and personal default decisions (Dearden et al., 2008; Rothstein and Rouse, 2011; Fos et al., 2017; Krishnan and Wang, 2019; Goodman et al., 2019; Mueller and Yannelis, 2019; Mezza et al., 2019). While a growing literature studies the relationship between student aid and graduate outcomes, much less is known about potential effects on families' intertemporal choices. This paper is the first to demonstrate the causal effect of student aid supply on parental saving behavior. The parental saving response to the rise in student loans potentially has important implications for the allocation of assets within households, and more broadly, as total household saving corresponds to 48.9% of U.S. national saving, to the distribution of wealth in the economy.[2]

Parental saving is intimately linked to the provision of educational financing as 70% of parents accumulate financial wealth, using both saving instruments and financial markets, to finance college expenses (Fidelity, 2018). The parental saving decision is characterized by a trade-off between consumption smoothing and expected college attendance of their offspring. Economic theory provides two contrasting mechanisms through which student aid levels directly affect parental saving decisions. On the one hand, the supply of student aid *reduces* parental savings since family wealth is a substitute for student loans in alleviating credit constraints of students. On the other hand, the provision of student aid lowers the entry barrier of marginal college entrants. Since the effective costs of obtaining the college

[1] These figures are from the annual reports 'Trends in Student Aid' published by CollegeBoard, based on administrative data from the U.S. department of education and income tax returns. All figures are expressed in 2017 Dollars.

[2] This percentage is calculated using data of the U.S. Bureau of Economic Analysis on gross private and public saving for the second quarter in 2019

premium are reduced, marginal students, who would have previously not attended university, enroll for college. As college attainment becomes a positive NPV investment for marginal students, parents *increase* their savings to cover the remaining unmet financing needs.[3] Similarly, attendance of a superior quality college becomes a positive NPV investment and induces parents to save more. Ultimately, it is an empirical question which of these effects dominates.

Measuring the relationship between student aid and parental saving behavior is particularly challenging for two reasons. First, establishing the direction of causality is difficult since the demand for student aid is a function of parental savings. For example, students from families that accumulated little wealth receive more grants and have access to cheaper (subsidized) student loans. A second challenge is that any observed cross-sectional differences in saving behavior could simply reflect systematic differences in unobserved heterogeneity (Gale and Scholz, 1994; Attanasio and Brugiavini, 2003). In particular, households with a high 'taste' for saving may accumulate more wealth and exhibit a lower demand for student loans. An ideal empirical analysis requires a shock in the supply of student aid that is exogenous to both household demand for student loans and saving behavior. I address these challenges by exploiting the passage of the 1992 higher education amendment act as a plausibly exogenous shock to the supply in student aid in the United States. This reform removed home equity from the set of assets that are 'taxed' by the federal aid formula. Consequently, many students suddenly receive higher levels of federal aid solely because the family balance sheet contains large home equity holdings (Dynarski, 2003b). This setting allows me to directly address the first empirical challenge since the reform induced variation in expected student aid eligibility that is exogenous to the demand for student loans. In the empirical analysis I combine this legal change with detailed household-level panel data on income, wealth accumulation, stock market participation, and demographic information.

[3] Long and Riley (2007) show that these remaining unmet needs are substantial. They estimate an average unmet need of $7,195 of students after accounting for family's expected contribution and the receipt of all aid.

Given the importance of unobserved heterogeneity in saving behavior, it is crucial to account for the household-specific saving effect. I address the second empirical challenge by comparing within household saving changes, therefore the estimated difference in saving behavior can be plausibly attributed to student loan supply (Gormley and Matsa, 2013).

The results consistently show that parents *increase* their savings after an expansion in student aid. This change is economically sizable since a one-standard deviation increase in exposure to student aid yields a 2.2 percentage point increase in the fraction of income saved by the mean household. These finding suggests that the marginal college entrance effect dominates the substitution effect. This interpretation is further corroborated by the finding that college enrollment disproportionately increases in families that are more affected by the reform. I estimate that college enrollment increases by 12 percentage points. This result provides evidence that the expansion of student aid programs succeeded in its primary goal to promote access to post-secondary education for students that would otherwise be unable to attend college.

The results on positive parental saving responses are robust to using an alternative identification strategy that exploits the notion that expected student aid amount sharply increases if siblings are likely to attend college simultaneously. A placebo test validates that the saving response to student aid supply is absent in families without children. Cross sectional tests reveal that the effects are largest among lower- and middle-income families and for college educated parents. Furthermore, I find more substantial saving responses for parents that identify themselves as savers. Documenting this new relationship between student aid supply and families' portfolio decisions improves our understanding of the intergenerational effects of student aid supply, and more broadly, contributes to explaining the observed heterogeneity in households' saving and portfolio decisions.

In sum, this paper is the first to provide systematic evidence on the relationship between the supply of student loans and financial decisions of parents. I make three contributions to the literature. First, the results add to a growing literature that studies spill-over effects

of student aid provision on non-educational outcomes (Bleemer et al., 2017; Goodman et al., 2019; Scott-Clayton and Zafar, 2019; Mezza et al., 2019; Di Maggio et al., 2019). While most of these papers focus on consumption and investment decisions at the graduate level, I demonstrate that student loan supply also adversely affects parental lifetime consumption. A small number of closely related papers show that parental wealth decreases the demand for student loans (Ionescu, 2009; Brown et al., 2011; Mondragon et al., 2017; Hotz et al., 2018; Abbott et al., 2019). I contribute to this literature by demonstrating how student aid supply affects the accumulation of parental wealth. Second, I contribute to the literature on the role of family composition in household portfolio decisions (Kennickell and Starr-Mccluer, 1997; Barnea et al., 2010; Bogan, 2015; Addoum, 2017; Olafsson and Pagel, 2018; Ke, 2018; Hertzberg, 2019). While the empirical papers in this literature focus mostly on financial decision making between spouses, there is some theoretical work that suggests that the arrival of children affects savings and portfolio choices because of future college expenses (Love, 2009; Hubener et al., 2015). While I do not directly test these models, the empirical results in this paper are consistent with the theoretical prediction that parents incorporate expected college expenses of their children in their savings decision. Finally, the mechanism I identify in this paper is directly relevant to the current policy debate regarding the optimal design of federal loan programs to stimulate college enrollment while minimizing the consequences for consumption smoothing (Lochner and Monge-Naranjo, 2011; Hanushek et al., 2014; Lochner et al., 2018; Abbott et al., 2019). I document a new spillover effect of student aid supply on parental wealth accumulation and asset allocation. This finding is relevant for a nascent literature that studies the relationship between household saving, asset allocation and wealth inequality (De Nardi and Fella, 2017; Bach et al., 2018; Fagereng et al., 2018).

The remainder of this paper is organized as follows. Sections 2 and 3 describe the policy environment and data sources respectively. Section 4 discusses the identification strategy. Section 5 present the empirical results and section 6 concludes.

2 Institutional Background

The federal student aid programs are the most important source of college financing in the U.S. There are two primary components that determine a prospective student's eligibility for federal financial aid: a measure of the family's ability to contribute financially and tuition fees. To access federal aid, prospective college students must fill out an application form (FAFSA) that collects detailed information on household income, assets and family composition. These inputs are used by the U.S. Department of Education to estimate the dollar amount a family can pay out of pocket to cover college expenses, called the expected family contribution (EFC). The intuition is that high income families and households with large asset holdings are able to cover more expenses. Important to note is that the EFC formula also includes demographic factors like family size, age of parents and other family members' enrollment in post-secondary education. For instance, the EFC sharply decreases as the number of college-going family members increases (Brown et al., 2011). The eligibility for subsidized loans is determined by the difference between EFC and the cost of attendance. There is an annual limit that caps the amount that students are allowed to borrow.

In 1992, the Higher Education Act (henceforth: HEA) changed the EFC formula. The main feature of HEA was that it removed home equity as a resource from which parents can pay college expenses. As a result of its passage many students who were previously ineligible for federal aid were able to qualify for need-based aid since their EFC decreased dramatically. The goal of this reform was to promote access to post-secondary education for students from lower- and middle-income families. The policymakers argued that rising college costs combined with increasing restrictions on eligibility for federal aid squeezed out middle-income students who could neither qualify for support nor afford to pay outright (Hannah, 1996). This reform became effective starting from the 1993-94 academic year. The resulting increase in federal student aid allocation was immediate and dramatic. Figure 1 shows that, after adjusting for inflation, the total amount of student aid increased by 43% from 1992-93 to 1994-95. Furthermore, this graph shows that this increase is completely driven by a jump

in student loan volumes. Since there is substantial variation between households in the share of home equity in total wealth holdings, the reform introduced heterogeneity in the intensity families were affected by this change in legislation. The main beneficiaries were households with most of their wealth in home equity. This reform constituted a significant increase in student aid eligibility as Dynarski (2003b) estimates that before HEA, each dollar of home equity reduced the federal aid eligibility by three to six cents of families on the margin of receiving aid.

Figure 1: Total Annual Federal Student Loan and Grant Volume

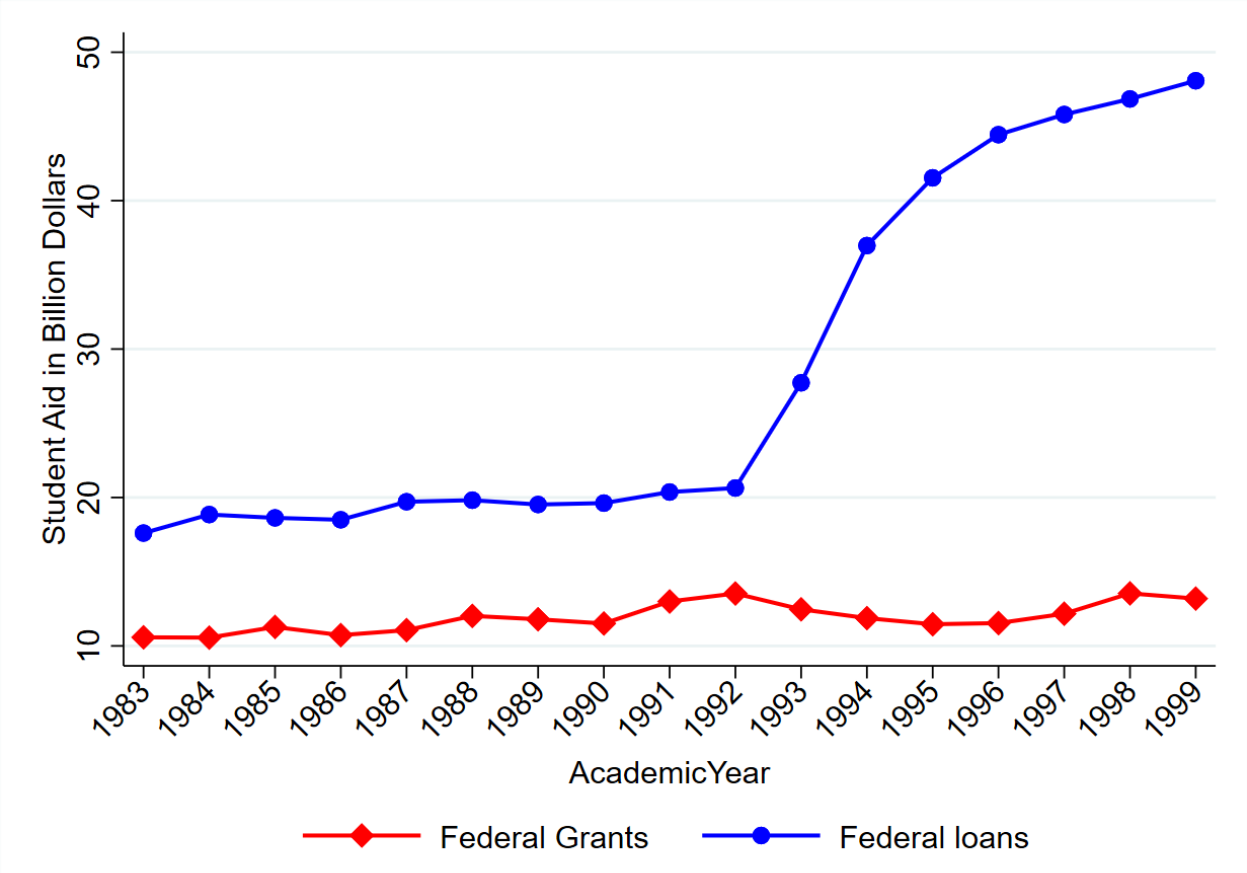

This graph plots the total amount of annual student aid, split by grant aid and student loans. The figure displays the student loan expansions after the introduction of HEA, while the total grant aid remains constant. All figures are from the annual reports 'Trends in Student Aid' published by CollegeBoard and expressed in 2017 Dollars.

3 Data

The primary data come from the Panel Study of Income Dynamics (PSID) for the years 1984 to 1999. The PSID is an annual panel data survey, which contains detailed information on family income, housing, family structure and other demographics. In additional wealth supplements, households are asked questions on their net worth and financial asset holdings for the years 1984, 1989, 1994, and 1999. These wealth supplements allow me to construct measures of household saving and stock market participation. The PSID is particularly well suited for this analysis since it enables me to track individual households over time and exploit within-household variation in saving behavior. This allows me to include household fixed effects that capture unobservable risk preferences, beliefs and other time-invariant characteristics.

The basic sample includes every household that owned a house between 1984 and 1999. I exclude renters from this sample since, by construction, they have no housing equity wealth. Since I am interested in changes in parental saving behavior after the introduction of HEA, an ideal analysis requires households that could have started saving for college expenses before this reform. Therefore, I follow families with a child between 5 and 15 in 1992.[4] I apply several filters to this sample to obtain my final dataset. I start with the definition that households are unique families if the head of household remains the same over the period 1984 to 1999, and drop households that exhibited a change in head of household as is standard in the literature.[5] In several cases households report a very low house value, therefore I follow the suggestion of Gerardi et al. (2010) and eliminate all observations for which the reported house value is below $5,000. I exclude a household-year observation if the household head is retired in that particular year. To eliminate gross outliers from the sample I follow Juster et al. (2006) and trim the top and bottom percentile of each wealth

[4]The exclusion of families with no children at the time of the reform also mitigates a potential concern that the student aid supply might affect household fertility choices.

[5]The exclusion of these households addresses the concern that a change in household head affects the unobserved heterogeneity in saving behavior. In total there are 155 households in which the head changes, however the main results are not sensitive to their exclusion.

component, income and home equity.[6] After deleting observations with missing values for income, wealth or demographic variables the baseline sample consists of 3,111 observation, in 1,207 unique households.

3.1 Household Saving Measures

Households' saving is calculated following the *active saving* approach of Juster et al. (2006). This measure captures the change in total household wealth *minus* capital gains for housing and financial assets, inheritances and gifts received *plus* the value of debt repayment. The active saving approach is particularly well suited to measure changes in saving behavior, because capital gains (passive saving) are not included. For example, household wealth accumulation may reflect revaluation of assets that are independent from an active saving decision. Since my analysis focuses on changes in actual parental saving, I eliminate these capital gains to obtain a more precise measure of the true saving intention of a household (Dynan et al., 2004). Naturally, higher-income households may have the ability to save more, therefore I normalize total household saving by the total family income. More formally, I define a saving rate for household i at time t:

$$SavingRate_{i,t} = \frac{\sum_{j=1}^{J} ActiveSaving_{t-1,t}^{i,j}}{Income_{t-1,t}^{i}} \tag{1}$$

where the sum of accumulated wealth in all assets (j) over the years $t-1$ to t is divided by total income of household i over the same period. I consider a wide range financial and real assets[7], however I exclude home equity as saving vehicle because I use variation in home equity to define treatment exposure. The measurement of active saving of household i in asset j ($ActiveSaving_{t-1,t}^{i,j}$) depends on the presence of potential capital gains in that

[6] All results remain with winsorizing the top and bottom percentile.

[7] The PSID contains information on real estate other than home equity, a farm or private business, vehicles, checking and saving accounts, money market funds, certificates of deposit, government saving bonds, and treasury bills, individual retirement accounts, stocks, mutual funds, investment trusts, bonds, cash value of life insurance, valuable collections and total non-mortgage debt.

particular asset. The exact method is described below, however the main intuition is to exclude potential capital gains by measuring net flows. Since this wealth data was gathered in 5 year intervals, the household saving rate ($SavingRate_{i,t}$) is defined for the periods 1984 to 1989, 1989 to 1994 and 1994 to 1999.

For assets where capital gains do not play a major role according to the PSID classification, I define active saving as the difference between the market value in period t and its value in period $t-1$. More specifically, I compute equation 2 using reported values of households' saving and checking accounts, bond holdings, vehicle values and consumer debt.

$$ActiveSaving^{i,j}_{t-1,t} = V^{i,j}_{t} - V^{i,j}_{t-1} \tag{2}$$

$AactiveSaving^{i,j}_{t-1,t}$ represents the active saving of household i in asset j over time period $t-(t-1)$ where $V^{i,j}_{t}$ is simply the reported value of asset j in time t.

For assets with potentially large capital gains, such as stocks, IRA accounts, other real estate, and investments in a private businesses, active saving of the household over a certain period is defined in equation 3. The value of these assets can change because it is sold or purchased (active saving) or the price of the asset changes (capital gains or passive saving). To isolate the active saving, I compute the net flows of these assets. For instance the amount of equity purchased by the household during a 5 year period *minus* any money they received from selling stocks. In contrast, increases in financial wealth due to stock price appreciation do not count as active saving. Hence active saving for these assets are equivalent to the change in value of holdings *minus* capital gains.

$$ActiveSaving^{i,j}_{t-1,t} = I^{i,j}_{t-1,t} - R^{i,j}_{t-1,t} \tag{3}$$

Again, $ActiveSaving^{i,j}_{t-1,t}$ represents the active saving of household i in asset j over time period $t-$*(t-1)*. $I^{i,j}_{t-1,t}$ is the net amount invested in asset j over period $t-1$ to t, like the amount of stocks purchased in the previous 5 years in the sample above. Conversely, $R^{i,j}_{t-1,t}$

represents the total value that the household sold of asset j. This value captures selling equity, but also the money received from a full or partial sell of the households' interest in a private business or real estate.

To validate the main results I also use an alternative measure of household saving behavior proposed by Cronqvist and Siegel (2015). They measure savings as the change in a household's total non-housing wealth, and scale this amount by the disposable income over the same period. Therefore the saving rate of household i is defined as:

$$SavingRate_{i,t} = \frac{\Delta NetWorth^i_t - \Delta HomeValue^i_t}{Income^i_{t-1,t}} \tag{4}$$

where $NetWorth^i_t$ is the sum of the wealth value in all asset classes at year t *minus* total debt. $HomeValue^i_{t-1,t}$ measures capital gains in housing value over a 5 year period, excluding households that moved between two consecutive periods.

3.2 Descriptive Statistics

Table 1 reveals several key facts about the distribution of saving and income in the sample. The average annual savings rate of households in the sample is 4.9%. This figure is roughly comparable to the average annual U.S. personal saving rate in this period estimated by Parker (1999). Using data from the National Income and Product Account (NIPA) he shows that this period was characterized by a decline in personal savings rate from 6% to below 1%, with a stable period of 5% between 1988 and 1993. The estimated saving rates closely mirror this trend as the average savings rate for the periods 1985-1989, 1990-94 and 1995-99 are 6.5%, 5.3% and 2.2% respectively. Furthermore, Table 1 reveals that 34% of the households in the sample is not accumulating wealth. This is consistent with the notion that many U.S. households are not saving (Lusardi et al., 2001). Note that I am likely to overestimate the

fraction of non-savers in the sample since I ignore accumulation of housing wealth while this is an important saving vehicle for many households. An important feature is that I

Table 1: **Descriptive Statistics**

This table reports descriptive statistics for my main variables of interest. I report mean, median 10th percentile and 90th percentile for all observations at the household-year level. $\text{SavingRate}_{i,t}$ is total annual savings divided by income as defined in section 3.1. 'No Saving (d)' is a dummy variable that equals one if the household has non-positive savings in a given year. Similarly 'Equity Participation (d)' is an indicator variable that equals one if the household holds any stocks in publicly held corporations or mutual funds in a given year, including equity in IRAs. $\frac{HomeEquity^i_{1989}}{NetWorth^i_{1989}}$ is the fraction of home equity wealth of the total wealth in 1989, before the reform. 'Annual Family Income' is defined as the 5 year average income and '5yr Income Volatility' is the volatility of annual income over these 5 year periods. Furthermore, I include the number of children in the household ('Number of Children') and age of the head of the household ('Age (years)'). 'Entrepreneur (d)' and 'College Degree (d)' are dummies that equal one if the head of the household own a business or holds a college degree at a given year respectively. Marital transitions within the household are defined as dummies that equal one if the head of the household got married ('Married (d)') or got divorced ('Divorced (d)') in a given year. Finally, a dummy equals one if the head of the household is of black ethnicity ('Black (d)').

	Mean	*SD*	p_{10}	p_{50}	p_{90}	Obs.
$\frac{HomeEquity^i_{1989}}{NetWorth^i_{1989}}$	0.596	0.249	0.249	0.620	0.935	3,111
$\text{SavingRate}_{i,t}$	0.049	0.151	-0.073	0.028	0.189	3,111
No Saving (d)	0.346	0.476	0	0	1	3,111
Equity Participation (d)	0.464	0.498	0	0	1	3,111
Annual Family Income ($k)	57.48	30.08	25.32	51.36	96.19	3,111
5yr Income Volatility ($k)	12.48	13.54	3.09	8.36	24.54	3,111
Number of Children	1.801	1.106	0	2	3	3,111
Age (years)	43.087	7.006	34	43	52	3,111
Entrepreneur (d)	0.231	0.422	0	0	1	3,111
College Degree (d)	0.370	0.483	0	0	1	3,111
Black Ethnicity (d)	0.193	0.395	0	0	1	3,111
Marital Transitions						
Married (d)	0.021	0.142	0	0	0	3,111
Divorced (d)	0.014	0.117	0	0	0	3,111

Figure 2: Distribution of Home Equity Share

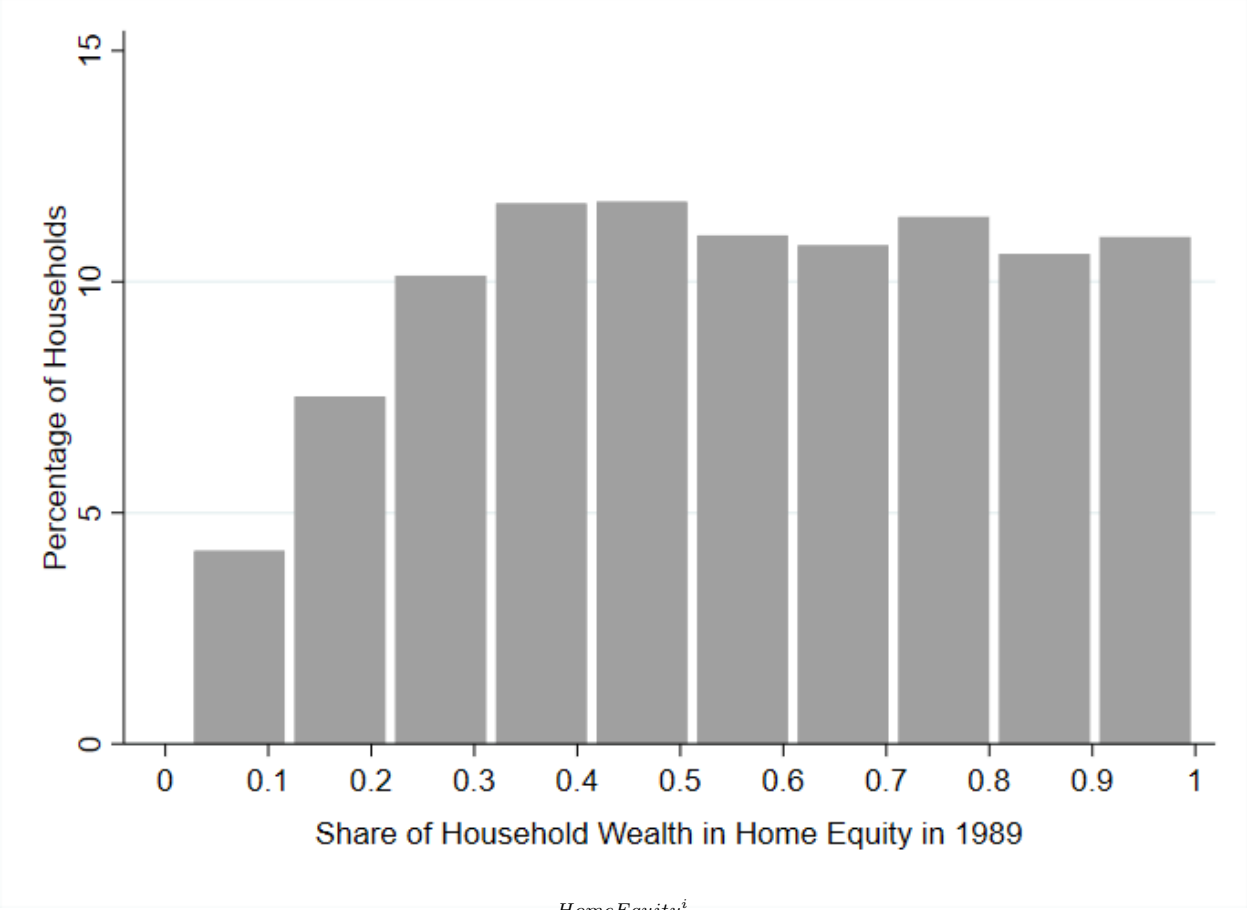

This histogram shows the distribution of $\frac{HomeEquity^i_{1989}}{NetWorth^i_{1989}}$, that captures the exposure to student loan expansion. The figure shows a left skewed distribution that is roughly uniform in the upper 80 percent. The vast majority of households holds a portion of their wealth in home equity, however there is substantial variation in the relative share.

observe substantial heterogeneity in the fraction of home equity by total household wealth in 1989. This variation implies a wide range in household-specific student aid growth that is crucial for my identification. Figure 2 confirms this observation as it reveals that households are roughly evenly distributed among all possible fractions of home equity by total household wealth in 1989.

Furthermore, Table 1 shows that the average household consists of roughly 2 children, a 43 year old head and an average (median) annual family income of $57.48k ($51.36k). The table reports few marital transitions, however this is at the household-year level. At the household level these are substantially larger as 8.8% of the families gets married and 5.6%

experiences a divorce during the sample period.

4 Empirical Strategy

While the legal expansion of student aid constitutes an economy-wide shock, I propose to isolate its effect on family finances by studying differential post reform changes across households. The removal of home equity in the federal aid formula induces variation between households in expected student aid. I construct a household-specific treatment intensity based on the share of household wealth represented by home equity shortly before the reform. This empirical approach is similar to Lucca et al. (2018), who study the effect of federal student aid expansion on tuition fees using variation in treatment intensities among universities. This differences-in-differences specification eliminates the potential concern that a trend in tuition fees increased students' reliance on federal loans over time (Lochner and Monge-Naranjo, 2011). Equation 5 shows the baseline regression I estimate using the PSID panel data.

$$SavingRate_{i,t} = \alpha_i + \alpha_{st} + \beta I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}} + \lambda C_{i,t} + \epsilon_{i,t} \tag{5}$$

The dependent variable is saving rate of household i over a five year period (from *t-1* to t), as defined in section 3.1. $I[HEA]_t$ is an indicator that equals zero if t is 1994 or earlier and one afterwards. Since the timing of the wealth supplements in the PSID only allows me to measure household savings over the periods 1984-1985, 1989-1994 and 1994-1999, I define only the latter period as post-treatment. Even though the treatment indicator is not perfectly aligned with the timing of HEA, this specification biases against finding any parental saving response to the student loan expansion. The coefficient of interest is β, which captures the effect of student aid supply on the saving propensity of the parents. Since families accumulate most of the wealth to cover college expenses before the children go to university, this coefficient can be interpreted as the savings response of families to

expected student loan supply available to their children. Furthermore, $C_{i,t}$ is a vector that contains time-varying household level characteristics that affect savings decisions. These controls include income volatility, household size, age of the head, age squared and three indicator variables that equal one if the head owns a business in the previous year, got married or got divorced. By including household fixed effects (α_i) I effectively study the effect of HEA on within-household saving behavior. Hence, it is unlikely that my results are driven by unobserved heterogeneity in time-invariant characteristics that correlate with saving behavior and demand for debt. For instance, religious households tend to save more, while they borrow less (Guiso et al., 2003). In this specification, I absorb any state-level shocks by including state $\times$ year fixed effects (α_{st}). The mean effect on the population over the sample period is absorbed by the inclusion of α_{st}. All regressions cluster standard errors at the household level, since observations are unlikely to be independent within households.

A crucial assumption for this estimator to be valid is that treatment and control groups follow parallel trends in absence of the reform. Unfortunately PSID only started collecting wealth data in 1984, therefore I am unable to extent the saving rate measure to periods before 1989. In order to examine the similarity between saving trends of the treatment and control groups I rely on a different PSID question that asks respondents whether the household has any savings. Although this is a crude measure of saving behavior, an indicator for positive household savings is used more often (e.g. González and Özcan (2013)). In Figure 3, I plot the average share of households with positive savings from 1971 to 2003. I split the sample by the median level (0.62) of the treatment indicator $\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$. The graph shows that the two groups follow parallel trends prior to the reform.[8] After the introduction of HEA, the propensity to save of affected families exponentially increases relative to the control group.

The balancedness tests in Table 2 show that parents with above median levels of home equity are comparable to the below-median parents. Families in both groups have household heads at roughly the same age and similar college attainment. Furthermore, I find no differ-

[8]Since wealthier households, on average, hold a lower share of their wealth in home equity, the control group consistently has a higher saving propensity.

Figure 3: Percentage of Households with Positive Savings (1994=1)

This graph plots the share of households that report having any savings, such as savings accounts or government bonds. I split the sample by 0.62 (median) of the treatment indicator $\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$. 'High Equity Share' represents the group of households with above median levels of home equity by total wealth and 'Low Equity Share' represents below median households.

ence in family size between the two groups. The similarities in these characteristics reassure that treatment is not correlated with life-cycle factors that also explain saving behavior. There are substantial wealth differences between both groups. This finding is intuitive as wealthier families typically hold wealth in multiple assets, whereas poorer households primarily rely on housing wealth (Fagereng et al., 2019). In unreported results I demonstrate that the findings remain unchanged with the inclusion of wealth-quintile fixed effects.

Table 2: **Balancedness**

This table describes the characteristics of parents with above median levels of home equity ($\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$) and below-median parents. The final column reports t-tests comparing the means of above- and below-median equity shares.

	High Equity Share	Low Equity Share	
	Mean (SD)	Mean (SD)	t(H,L)
Age (years)	43.53	42.52	4.00
	(0.17)	(0.18)	
Number of Children	1.79	1.82	-0.82
	(0.03)	(0.03)	
College Degree (d)	0.38	0.36	0.95
	(0.01)	(0.01)	
Household wealth ($K)	101.75	178.7	-13.84
	(2.38)	(5.48)	

5 Results

5.1 Savings Response

This section studies the effect of expected student loans on parental saving. Table 3 suggests a strong causal effect of student aid expansion on parents' savings. The point estimate of the interaction term is consistently positive and statistically significant at the 1% level. The magnitude of the estimates remain stable if I saturate the model with fixed effects. The point estimate drops somewhat when I include household-fixed effects, which suggests that households with more housing wealth also have higher saving rates. Overall, the effect is economically sizable: a one-standard-deviation increase in expected student aid $\left(\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}\right)$ leads to an increase in annual savings rates by 2.1-2.3 percentage points (=[0.096×0.249]). This increase changes fraction of income saved by an average affected family from 4.9% to 6.1%. To provide an insight on the magnitude of this increase in dollar amounts, I make a simple back-of-the-envelope calculation. Relative to the average annual income in the sample, a 2.3 percentage point increase means that, each year, parents save an additional

Table 3: **Savings response of Families**

This table reports the results for the difference-in-difference regression specification of equation 5, with standard errors clustered at the household level in the parentheses. $I[HEA]_t$ takes the value of one after the introduction of HEA, and $\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$ measures the household's exposure to the reform as the fraction of housing equity wealth of total wealth before the reform. I control for time-varying household characteristics as the number of children, age of the households head, age squared, and dummies that equal one if the household hold owns a business, holds a college degree, got married or divorced. Furthermore, I control for income shocks by including 5 year income volatility. In the first column I also include a dummy hat equals one of the head of the household is black. I also report the number of observations (N). I control for household fixed effects ($HouseholdFE$) and state times year fixed effects ($State \times YearFE$).

	(1) SavingRate$_{i,t}$	(2) SavingRate$_{i,t}$	(3) SavingRate$_{i,t}$	(4) SavingRate$_{i,t}$
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$	0.096***	0.085***	0.088***	
	(0.023)	(0.025)	(0.028)	
$I[HEA]_t \times I\left(\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}} \geq p50\right)$				0.031**
				(0.015)
$\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$	-0.097***			
	(0.013)			
$I[HEA]_t$	-0.088***	-0.070***		
	(0.016)	(0.020)		
Number of Children	0.004*	0.011*	0.018***	0.017***
	(0.002)	(0.006)	(0.006)	(0.006)
Entrepreneur (d)	0.027***	0.046***	0.064***	0.066***
	(0.007)	(0.015)	(0.016)	(0.015)
Age (years)	-0.004	0.006	0.001	0.002
	(0.004)	(0.006)	(0.006)	(0.006)
Age^2	0.000	-0.000	0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Married (d)	0.002	-0.034	-0.035	-0.038
	(0.016)	(0.029)	(0.031)	(0.030)
Divorced (d)	0.013	0.019	0.022	0.019
	(0.021)	(0.037)	(0.043)	(0.043)
5yr Income Volatility	-0.000	-0.000**	-0.000**	-0.000**
	(0.000)	(0.000)	(0.000)	(0.000)
College Degree (d)	0.000	0.034	0.006	0.007
	(0.005)	(0.039)	(0.040)	(0.041)
Black (d)	-0.008			
	(0.006)			
Household FE	No	Yes	Yes	Yes
State × Year FE	No	No	Yes	Yes
N	3111	3003	3000	2980

amount of 1,374$. In column 4, I validate the positive saving response of parents using a dichotomous indicator that equals one for families with an above-median home equity-wealth ratio.

Consistent with the documented decline in household saving rates in the 1990s (Parker, 1999; Lusardi et al., 2001), I find a strong negative effect in the post HEA period. Also notable is the sizable positive effect of entrepreneurship on household saving rates. This result confirms previous findings that entrepreneurial risk increases households' savings (Quadrini, 2000; Gentry and Hubbard, 2004). The coefficients of age are negligible since there is little variation in age is left after including time trends.

5.2 College Enrollment

This section sheds further light on the mechanism behind the positive parental saving response. The human capital investment models in the spirit of Becker (1962) show that a student will only invest in college if the net rate of return is greater than the market rate of return one would earn by investing the forgone earnings and the direct cost of college. Since the provision of student aid reduces the effective costs of college, this investment problem can become a positive NPV investment for marginal students. As many student face non-tuition costs that are not fully covered by the federal student aid they receive (Long and Riley, 2007), they rely on additional financing sources such as parental resources. Hence, parents of marginal students might increase their savings after a credit expansion because the probability of attending college increases for children on the margin of enrollment. A testable implication of this economic mechanism is that one should observe a disproportional increase in college enrollment in affected families. I match the sample data with information on educational attainment of children living in the household from the PSID 'Childbirth and Adoption History' File. This shrinks the sample size since I am only able to match a subsample of parents to the educational information of their children. The PSID contains no questions that directly ask for college enrollment, therefore I follow the standard proce-

dure in the literature and measure enrollment as having completed more than 12 grades of schooling.

I follow the empirical specification of Lovenheim (2011) and estimate the following linear probability model:

$$Enrollment_{i,k,t} = \alpha_t + \alpha_s + \beta I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}} + \lambda C_{i,k,t} + \epsilon_{i,k,t} \tag{6}$$

where $Enrollment_{i,k,t}$ is a dummy variable equal to one if child k in household i enrolls in college in year t. Important to note is that the unit of observation changes from the family to the student since a household can have multiple children. The coefficient of interest is β, which captures the effect of student aid supply on the college enrollment of children. The specification also includes a vector of student and household characteristics ($C_{i,k,t}$). I include gender and ethnicity dummies to control for student characteristics. Furthermore, I control for total family income, the number of children, age of head of household and a dummy that equals one if the head of the family graduated college. Finally, I include college year fixed effects (α_t) and state fixed effects (α_s). In this specification I cannot include household fixed effects since the treatment exposure is measured at the family level.

Table 4 report the results. The coefficients consistently show a increase in college attendance of children in households with more exposure to the student aid expansion. The point estimates suggest a substantial increase in college attendance of 12 percentage points. This finding confirms previous evidence showing that student aid supply increases college enrollment (McPherson and Schapiro, 1991; Van der Klaauw, 2002; Nielsen et al., 2010; Solis, 2017). The disproportional increase in college enrollment provides additional evidence consistent with the economic mechanism that families increase their savings to support marginal college entrants. This mechanism is in line with the argument of Dynarski (2003a) that student aid provision can have a 'threshold effect' by assisting students to cross the hurdle of college entry.

Table 4: **College Enrollment**

This table reports the results of the test whether affected families experience a disproportional increase in college enrollment. More specifically, I estimate equation 6 on the college enrollment decision of children given the household-variation in treatment exposure. $I[HEA]_t$ takes the value of one after the introduction of HEA, and $\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$ measures the household's exposure to the reform as the fraction of housing equity wealth of total wealth before the reform. I control for number of children, income volatility, age of the households head, age squared, and dummies that equal one if the household hold owns a business, got married or divorced. Furthermore, I include college year fixed effects (α_t) and state fixed effects (α_s). All standard errors are clustered at the household level and reported in the parentheses. Finally, I also report the number of observations (N).

	$Enrollment_{i,k,t}$	$Enrollment_{i,k,t}$	$Enrollment_{i,k,t}$
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$	0.103*	0.115*	0.218**
	(0.062)	(0.065)	(0.098)
Controls	Yes	Yes	Yes
College Year FE	No	Yes	No
State FE	No	Yes	No
State × College Year FE	No	No	Yes
N	1,481	1,481	1,459

5.3 Alternative Identification Strategy

To provide additional evidence for the positive saving response I use an alternative identification strategy that relies on the notion that student aid amount sharply increases if there are multiple college-going family members (Brown et al., 2011). To be more specific, I define a treatment indicator that equals one if a household contains at least two siblings with a birth spacing less than 4 years and zero if the household contains no 'overlapping' children. Since most college degrees require 4 years, the overlapping sibling indicator proxies for higher (per student) expected student aid. My alternative specification is then as follows:

$$SavingRate_{i,t} = \alpha_i + \alpha_{st} + \beta I[HEA]_t \times I[SiblingOverlap]_i + \lambda C_{i,t} + \epsilon_{i,t} \tag{7}$$

The main coefficient of interest remains β, the effect of sibling overlap (which proxies for expected financial aid) on saving behavior after student aid expansion. Panel A in Table 5 presents the results. Unless otherwise mentioned, I suppress all control variables for brevity.

Table 5: **Alternative Specifications**

This table reports alternative empirical specifications that validate the results of Table 3. Panel A shows the results from an alternative identification strategy that relies on the notion that student aid amount granted sharply increases if there are multiple college-going family members. More specifically, I estimate a difference-in-difference regression specification with a treatment indicator that equals one if a household contains at least two siblings with a birth spacing less than 4 years and zero if the household contains no 'overlapping' children ($I[SiblingOverlap]_i$). The treatment indicator is interacted with '$I[HEA]_t$', that takes the value of one after the introduction of HEA. Panel B reports the results for the difference-in-difference regression specification of equation 5 with the alternative saving rate measure described in section 3.1. This measure captures the change in non-housing wealth between to period and is scaled by the total income of the family over the same period. In both panels I control for number of children, income volatility, age of the households head, age squared, and dummies that equal one if the household hold owns a business, holds a college degree, got married or divorced. Furthermore, I include household fixed effects ($HouseholdFE$) and state times year fixed effects ($State \times YearFE$). All standard errors are clustered at the household level and reported in the parentheses. Finally, I also report the number of observations (N).

	Panel A: Alternative Identification		
	$\text{SavingRate}_{i,t}$	$\text{SavingRate}_{i,t}$	$\text{SavingRate}_{i,t}$
$I[HEA]_t \times I[SiblingOverlap]_i$	0.040***	0.030*	0.038**
	(0.015)	(0.016)	(0.016)
Controls	Yes	Yes	Yes
Household FE	No	Yes	Yes
State × Year FE	No	No	Yes
N	3261	3122	3119
	Panel B: Alternative Saving Measure		
	$\frac{\Delta NetWorth^i}{Income^i_{t-1,t}}$	$\frac{\Delta NetWorth^i}{Income^i_{t-1,t}}$	$\frac{\Delta NetWorth^i}{Income^i_{t-1,t}}$
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$	0.167***	0.095**	0.167***
	(0.046)	(0.044)	(0.047)
Controls	Yes	Yes	Yes
Household FE	No	Yes	Yes
State × Year FE	No	No	Yes
N	3376	3328	3324

The number of observations slightly grows since I can include families where the home equity-wealth ratio is not properly defined (at least one of the inputs is negative). The results show that the main result, a positive saving response, remains unchanged using this alternative identification, albeit somewhat smaller as in Table 3.

In panel B of Table 5 I find a similar positive saving response using changes in non-housing wealth. Since this alternative measure includes capital gains, it is a more comprehensive definition of saving. However, as discussed in section 3.1 the true intentions of households are likely to be better captured by the active saving rate that exclude these capital gains.

5.4 Heterogeneity and Robustness

5.4.1 Treatment Heterogeneity

I analyze treatment effect heterogeneity to better understand which households drive the increase in parental savings. Consistent with a credit constraints interpretation, the observed effect is largest among lower- and middle-income families. In particular, the first four columns in Table 6 present the differential saving effect by income quartiles. I observe the largest effect in the lowest quartile, and the coefficient becomes statistically insignificant for the top half of the income distribution. This result is in line with previous findings that credit constraints in lower-income families reduce the demand for higher education (Acemoglu and Pischke, 2001; Sun and Yannelis, 2016). Additionally I examine the treatment effect heterogeneity in parental education. A large literature shows parental education increases the probability a child's college attendance (e.g. Black and Sufi (2002)) and influences the amount of educational savings (e.g. Hossler and Vesper (1993)). Column 5 shows the results of the triple difference approach that measures the differential saving response of college educated parents. I find that parents with a college degree save more compared to lower educated parents. This is consistent with the argument that a higher probability of college enrollment and a better ability to estimate future college costs, increases the expected return of the child's college graduation for more educated parents.

Table 6: **Treatment Effect Heterogeneity**

This table reports the results of analyzing treatment effect heterogeneity to better understand which households drive the positive effect of student aid supply on parental savings. I estimate a triple difference regression specification that adds a third difference using time-invariant group indicators. $I[HEA]_t$ takes the value of one after the introduction of HEA, and $\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$ measures the household's exposure to the reform as the fraction of housing equity wealth of total wealth before the reform. To examine heterogeneity in the parental saving response by family income, I create dummies for all quartiles in the income distribution in 1991 ($I[IncomeQuartile]_i$). The triple interaction term measures the difference in saving responses between families in a particular income quartile and families in the rest of the income distribution, for instance lower-income families compared to middle- and top-income families ($I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}} \times I[IncomeQuartile1]_i$). Additionally I test whether college educated parents responded differently to the introduction of HEA. I define a dummy variable that equals one if at least one of the parents is a college graduate. '$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}} \times I[ParentsCollege]_i$' captures the difference in saving responses between college-graduate parents and parents without a college degree. I control for number of children, income volatility, age of the households head, age squared, and dummies that equal one if the household hold owns a business, got married or divorced. Furthermore, I include household fixed effects ($HouseholdFE$) and state times year fixed effects ($State \times YearFE$). All standard errors are clustered at the household level and reported in the parentheses. Finally, I also report the number of observations (N).

	$\text{SavingRate}_{i,t}$	$\text{SavingRate}_{i,t}$	$\text{SavingRate}_{i,t}$	$\text{SavingRate}_{i,t}$	$\text{SavingRate}_{i,t}$
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}} \times I[IncomeQuartile1]_i$	0.312***				
	(0.093)				
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}} \times I[IncomeQuartile2]_i$		0.071***			
		(0.025)			
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}} \times I[IncomeQuartile3]_i$			0.003		
			(0.012)		
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}} \times I[IncomeQuartile4]_i$				-0.008	
				(0.014)	
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}} \times I[ParentsCollege]_i$					0.152***
					(0.051)
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$	-0.013	-0.004	0.005	0.006	-0.021
	(0.015)	(0.016)	(0.016)	(0.016)	(0.019)
Controls	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes
N	3002	3002	3002	3002	3002

Table 7: **Cross Sectional Robustness Checks**

This table reports the results of testing cross sectional implications that naturally follow from the hypothesized relationship. In the first two columns I estimate the difference-in-difference regression specification of equation 5 on a sample of families without children, that should not be affected by HEA. In the final two columns I split the sample by different saving attitudes. Column 3 limits the sample to families where the household head indicated in a 1972 survey that he/she rather 'save for the future' than 'spent money today'. In column 3 I limit the sample to households that prefer to spent. $I[HEA]_t$ takes the value of one after the introduction of HEA, and $\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$ measures the household's exposure to the reform as the fraction of housing equity wealth of total wealth before the reform. I control for number of children, income volatility, age of the households head, age squared, and dummies that equal one if the household hold owns a business, got married or divorced. Furthermore, I include household fixed effects ($HouseholdFE$) and state times year fixed effects ($State \times YearFE$). All standard errors are clustered at the household level and reported in the parentheses. Finally, I also report the number of observations (N).

	$\text{SavingRate}_{i,t}$	$\text{SavingRate}_{i,t}$	$\text{SavingRate}_{i,t}$	$\text{SavingRate}_{i,t}$
	Families with no Children		Saving Preference	Spending Preference
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$	-0.029	0.030	0.148***	-0.055
	(0.027)	(0.080)	(0.058)	(0.045)
Controls	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
State × Year FE	No	Yes	No	No
N	884	877	245	239

5.4.2 Cross sectional Robustness Checks

This section tests important cross sectional implications that naturally follow from the hypothesized relationship. The effect of student aid supply on parental savings should not affect all households equally. If the provision of student aid induces parents to save more because it raises the expected marginal college returns, the relationship should be absent in families with no children. I test this implication by estimating equation 5 on a sample of childless families. The first two columns in Table 7 report the results. I find that families without children that have a high share of wealth in home equity clearly do not adjust their saving behavior because they have no exposure to student aid.

Additionally, one would expect the saving response to be largest for households with more positive saving attitudes. I measure attitudes towards saving using a question in the

1972 PSID survey that asks respondents whether they prefer to *'save for the future'* or *'spent money today'*.[9] The responses are coded as a five-point Likert scale. I follow Knowles and Postlewaite (2005) and focus on households that indicated a clear preference (disregard indifferent respondents). Since this measure is only available for households that were already in the PSID sample in 1972, the total number of observations shrinks to approximately 500. In the final two columns of Table 7 I split this sample by saving attitudes. The results show that the saving response is driven by families with a positive saving attitude.

5.5 Wealth Allocation

The results in previous sections consistently show that parental saving increases after a positive shock to student aid supply. While the ultimate objective is to explore the dynamics of wealth accumulation, recently the attention has shifted more towards the allocation of savings because it introduces heterogeneity in rates of return on household savings (Campbell, 2016; Bach et al., 2018; Fagereng et al., 2019). Therefore, in this section I focus on the natural question whether the positive saving response affects the allocation of wealth between riskless and risky assets. Table 8 reports the average change in stock market participation and household leverage after the introduction of HEA. A number of interesting findings emerge from Table 8. The first column shows that parents save by paying off debts after student loan provision. This finding suggests an intergenerational transfer of household leverage from parents to children.

A second result is that, after an increase in student aid supply, parents increase their holdings in the equity markets. As is common in the finance literature (e.g. Guiso et al. (2008) and Giannetti and Wang (2016)) I measure equity market participation using an indicator variable that equals one if the household holds a any stocks at a given time. This includes both directly held stocks, and indirect equity holdings via investment trusts, mutual funds and retirement accounts. The linear probability model estimates in Column 2 show

[9]The exact question was phrased: 'Would you rather spend your money and enjoy life today, or save more for the future?'

Table 8: **Asset Allocation**

This table reports the average change in asset allocation after the introduction of HEA. In the first column I consider the ratio of household debt over non-housing wealth. The second column examines equity participation using a dummy that equals one if the household is active in the stock market. The final two columns examines the average change in proportion of equity in the household's portfolio of cash, bonds and equity ('Risky Share'). The empirical specification of the difference-in-difference regression is similar as in Table 3. $I[HEA]_t$ takes the value of one after the introduction of HEA, and $\frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$ measures the household's exposure to the reform as the fraction of housing equity wealth of total wealth before the reform. I control for number of children, income volatility, amount of non-housing wealth, age of the households head, age squared, and dummies that equal one if the household hold owns a business, holds a college degree, got married or divorced. Furthermore, I include household fixed effects ($HouseholdFE$) and state times year fixed effects ($State \times YearFE$). All standard errors are clustered at the household level and reported in the parentheses. Finally, I also report the number of observations (N).

	(1) Household Leverage over total Wealth	(2) Equity Participation	(3) Risky Share	(4) Risky Share
$I[HEA]_t \times \frac{HomeEquity_i^{1989}}{NetWorth_i^{1989}}$	-1.409***	0.259***	0.111**	0.272*
	(0.452)	(0.076)	(0.055)	(0.141)
Controls	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
N	2,887	2,903	2,703	793

that equity market participation among affected parents increases. In the final two columns I examine the proportion of the liquid financial portfolio invested in equity of the full sample (column 3) and conditional on equity participation in the previous period (column 4). This ratio is a common measure of financial risk taking in household finance (e.g. Calvet and Sodini (2014)). I find that affected parents tilt their portfolio towards risky assets. These findings suggest that the provision of student aid have an additional impact on household wealth accumulation through household portfolio returns induced by a change in allocation of assets.

6 Conclusion

This paper examines intergenerational effects of the current rise in student loans. The exponential growth in student debt attracted the interest of economists and policymakers, as high levels of educational debt may adversely affect students' future consumption, investment and personal default decisions. While some economists have suggested that the rise in student loans could also have large ramifications for the saving and portfolio choices of students' families (Amromin and Eberly, 2016; Mondragon et al., 2017), this paper is the first to provide systematic evidence on this relationship. I exploit policy-induced variation in expected student aid to estimate its casual effect on parental saving behavior. My results show that parents *increase* their savings after an expansion in student aid. This change is economically sizable since a one-standard deviation increase in exposure to student aid yields a 2.2 percentage point increase in the fraction of income saved by affected families. The mechanism that drives this result is the anticipation by parents of the positive effect of student aid on college enrollment of their children, i.e. the college investment NPV becomes positive for students on the margin of college attendance. Parents increase their savings to cover the remaining unmet financing needs in college expenses after receiving student aid. Consistent with this interpretation, I show that college attendance disproportionately increases for families affected by the reform. Furthermore, I find that student credit expansion shifts parental wealth allocation towards riskier assets.

My findings point to a previously undocumented and non-trivial intergenerational impact of student credit. The parental saving response to the rise in student loans potentially has important implications for the allocation of assets within households, and more broadly, the distribution of wealth in the economy. As this study illuminates, in addition to the effect of educational credit on students' future consumption, policymakers should also consider parents life-time consumption.

References

Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Giovanni L Violante, "Education Policy and Intergenerational Transfers in Equilibrium," *Journal of Political Economy*, 2019. Forthcoming.

Acemoglu, Daron and Jörn-Steffen Pischke, "Changes in the wage structure, family income, and children's education," *European Economic Review*, 2001, *45* (4-6), 890–904.

Addoum, Jawad M, "Household portfolio choice and retirement," *Review of Economics and Statistics*, 2017, *99* (5), 870–883.

Amromin, Gene and Janice Eberly, "Education financing and student lending," *Annual Review of Financial Economics*, 2016, *8*, 289–315.

Attanasio, Orazio P and Agar Brugiavini, "Social security and households' saving," *Quarterly Journal of Economics*, 2003, *118* (3), 1075–1119.

Bach, Laurent, Laurent E Calvet, and Paolo Sodini, "From Saving Comes Having? Disentangling the Impact of Saving on Wealth Inequality," *Swedish House of Finance Research Paper 18-8*, 2018.

Barnea, Amir, Henrik Cronqvist, and Stephan Siegel, "Nature or nurture: What determines investor behavior?," *Journal of Financial Economics*, 2010, *98* (3), 583–604.

Becker, Gary S, "Investment in human capital: A theoretical analysis," *Journal of Political Economy*, 1962, *70* (5, Part 2), 9–49.

Black, Sandra E and Amir Sufi, "Who goes to college? Differential enrollment by race and family background," *NBER Working Paper 9310*, 2002.

Bleemer, Zachary, Meta Brown, Donghoon Lee, and Wilbert van der Klaauw, "Tuition, jobs, or housing: what's keeping millennials at home?," *Federal Reserve Bank of New York Staff Report 700*, 2017.

Bogan, Vicki L, "Household asset allocation, offspring education, and the sandwich generation," *American Economic Review*, 2015, *105* (5), 611–15.

Brown, Meta, Andrew Haughwout, Donghoon Lee, Joelle Scally, and Wilbert Van der Klaauw, "Measuring student debt and its performance," *Federal Reserve Bank of New York Staff Report 668*, 2014.

_ , John Karl Scholz, and Ananth Seshadri, "A new test of borrowing constraints for education," *Review of Economic Studies*, 2011, *79* (2), 511–538.

Calvet, Laurent E and Paolo Sodini, "Twin picks: Disentangling the determinants of risk-taking in household portfolios," *Journal of Finance*, 2014, *69* (2), 867–906.

Campbell, John Y, "Restoring rational choice: The challenge of consumer financial regulation," *American Economic Review*, 2016, *106* (5), 1–30.

Cronqvist, Henrik and Stephan Siegel, "The origins of savings behavior," *Journal of Political Economy*, 2015, *123* (1), 123–169.

De Nardi, Mariacristina and Giulio Fella, "Saving and wealth inequality," *Review of Economic Dynamics*, 2017, *26*, 280–300.

Dearden, Lorraine, Emla Fitzsimons, Alissa Goodman, and Greg Kaplan, "Higher education funding reforms in England: The distributional effects and the shifting balance of costs," *Economic Journal*, 2008, *118* (526), F100–F125.

Di Maggio, Marco, Ankit Kalda, and Vincent Yao, "Second Chance: Life without Student Debt," *NBER Working Paper 25810*, 2019.

Dynan, Karen E, Jonathan Skinner, and Stephen P Zeldes, "Do the rich save more?," *Journal of Political Economy*, 2004, *112* (2), 397–444.

Dynarski, Susan M, "Does aid matter? Measuring the effect of student aid on college attendance and completion," *American Economic Review*, 2003, *93* (1), 279–288.

_ , "Loans, liquidity and schooling decisions," *Harvard University Mimeo*, 2003.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri, "Heterogeneity and Persistence in Returns to Wealth," *Econometrica*, 2018. Forthcoming.

_ **, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik**, "Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains," *Princeton University / BI Norwegian Business School / University of Oslo Working Paper*, 2019.

Fidelity, "College Savings Indicator," Technical Report 2018.

Fos, Vyacheslav, Andres Liberman, and Constantine Yannelis, "Debt and Human Capital: Evidence from Student Loans," *Boston College / NYU Stern / Chicago Booth Working Paper*, 2017.

Gale, William G and John Karl Scholz, "IRAs and household saving," *American Economic Review*, 1994, *84* (5), 1233–1260.

Gentry, William M and R Glenn Hubbard, "Entrepreneurship and Household Saving," *The BE Journal of Economic Analysis & Policy*, 2004, *4* (1), 1–57.

Gerardi, Kristopher S, Harvey S Rosen, and Paul S Willen, "The impact of deregulation and financial innovation on consumers: The case of the mortgage market," *Journal of Finance*, 2010, *65* (1), 333–360.

Giannetti, Mariassunta and Tracy Yue Wang, "Corporate scandals and household stock market participation," *Journal of Finance*, 2016, *71* (6), 2591–2636.

González, Libertad and Berkay Özcan, "The risk of divorce and household saving behavior," *Journal of Human Resources*, 2013, *48* (2), 404–434.

Goodman, Sarena, Adam Isen, and Constantine Yannelis, "A day late and a dollar short: Liquidity and household formation among student borrowers," *Federal Reserve Board Finance and Economics Discussion Series Paper 2018-025*, 2019.

Gormley, Todd A and David A Matsa, "Common errors: How to (and not to) control for unobserved heterogeneity," *Review of Financial Studies*, 2013, *27* (2), 617–661.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, "People's opium? Religion and economic attitudes," *Journal of Monetary Economics*, 2003, *50* (1), 225–282.

_ , _ , and _ , "Trusting the stock market," *Journal of Finance*, 2008, *63* (6), 2557–2600.

Hannah, Susan B, "The Higher Education Act of 1992: Skills, constraints, and the politics of higher education," *Journal of Higher Education*, 1996, *67* (5), 498–527.

Hanushek, Eric A, Charles Ka Yui Leung, and Kuzey Yilmaz, "Borrowing constraints, college aid, and intergenerational mobility," *Journal of Human Capital*, 2014, *8* (1), 1–41.

Hertzberg, Andrew, "Time-consistent individuals, time-inconsistent households," *Journal of Finance*, 2019. Forthcoming.

Hossler, Don and Nick Vesper, "An exploratory study of the factors associated with parental saving for postsecondary education," *Journal of Higher Education*, 1993, *64* (2), 140–165.

Hotz, V Joseph, Emily E Wiemers, Joshua Rasmussen, and Kate Maxwell Koegel, "The role of parental wealth and income in financing children's college attendance and its consequences," *NBER Working Paper 25144*, 2018.

Hubener, Andreas, Raimond Maurer, and Olivia S Mitchell, "How family status and social security claiming options shape optimal life cycle portfolios," *Review of Financial Studies*, 2015, *29* (4), 937–978.

Ionescu, Felicia, "The federal student loan program: Quantitative implications for college enrollment and default rates," *Review of Economic Dynamics*, 2009, *12* (1), 205–231.

Juster, F Thomas, Joseph P Lupton, James P Smith, and Frank Stafford, "The decline in household saving and the wealth effect," *Review of Economics and Statistics*, 2006, *88* (1), 20–27.

Ke, Da, "Who Wears the Pants? Gender Identity Norms and Intra-Household Financial Decision Making," *University of South Carolina Working Paper*, 2018.

Kennickell, Arthur B and Martha Starr-Mccluer, "Household saving and portfolio change: evidence from the 1983–89 SCF panel," *Review of Income and Wealth*, 1997, *43* (4), 381–399.

Knowles, John and Andrew Postlewaite, "Do Children Learn to Save from Their Parents," *Population Aging Research Center Working Paper WPS 05-07, University of Pennsylvania*, 2005.

Krishnan, Karthik and Pinshuo Wang, "The Cost of Financing Education: Can Student Debt Hinder Entrepreneurship?," *Management Science*, 2019, *65* (10), 4522–4554.

Lochner, Lance J and Alexander Monge-Naranjo, "The nature of credit constraints and human capital," *American Economic Review*, 2011, *101* (6), 2487–2529.

_ , Todd Stinebrickner, and Utku Suleymanoglu, "Parental Support, Savings and Student Loan Repayment," *NBER Working Paper 24863*, 2018.

Long, Bridget Terry and Erin Riley, "Financial aid: A broken bridge to college access?," *Harvard Educational Review*, 2007, *77* (1), 39–63.

Love, David A, "The effects of marital status and children on savings and portfolio choice," *Review of Financial Studies*, 2009, *23* (1), 385–432.

Lovenheim, Michael F, "The effect of liquid housing wealth on college enrollment," *Journal of Labor Economics*, 2011, *29* (4), 741–771.

Lucca, David O, Taylor Nadauld, and Karen Shen, "Credit supply and the rise in college tuition: Evidence from the expansion in federal student aid programs," *Review of Financial Studies*, 2018, *32* (2), 423–466.

Lusardi, Annamaria, Jonathan Skinner, and Steven Venti, "Saving puzzles and saving policies in the United States," *Oxford Review of Economic Policy*, 2001, *17* (1), 95–115.

McPherson, Michael S and Morton Owen Schapiro, "Does student aid affect college enrollment? New evidence on a persistent controversy," *American Economic Review*, 1991, *81* (1), 309–318.

Mezza, Alvaro, Daniel Ringo, Shane Sherlund, and Kamila Sommer, "Student Loans and Homeownership," *Journal of Labor Economics*, 2019. Forthcoming.

Mondragon, John, Janice Eberly, and Gene Amromin, "The Housing Crisis and the Rise in Student Loans," *Northwestern University Kellogg / FRB Chicago Working Paper*, 2017.

Mueller, Holger M and Constantine Yannelis, "The rise in student loan defaults," *Journal of Financial Economics*, 2019, *131* (1), 1–19.

Nielsen, Helena Skyt, Torben Sørensen, and Christopher Taber, "Estimating the effect of student aid on college enrollment: Evidence from a government grant policy reform," *American Economic Journal: Economic Policy*, 2010, *2* (2), 185–215.

Olafsson, Arna and Michaela Pagel, "Family Finances: Intra-Household Bargaining, Spending, and Financial Structure," *Columbia Business School / Copenhagen Business School Working Paper*, 2018.

Parker, Jonathan A, "Spendthrift in America? On two decades of decline in the US saving rate," *NBER macroeconomics annual*, 1999, *14*, 317–370.

Quadrini, Vincenzo, "Entrepreneurship, saving, and social mobility," *Review of Economic Dynamics*, 2000, *3* (1), 1–40.

Rothstein, Jesse and Cecilia Elena Rouse, "Constrained after college: Student loans and early-career occupational choices," *Journal of Public Economics*, 2011, *95* (1-2), 149–163.

Scott-Clayton, Judith and Basit Zafar, "Financial aid, debt management, and socioeconomic outcomes: Post-college effects of merit-based aid," *Journal of Public Economics*, 2019, *170*, 68–82.

Solis, Alex, "Credit access and college enrollment," *Journal of Political Economy*, 2017, *125* (2), 562–622.

Sun, Stephen Teng and Constantine Yannelis, "Credit constraints and demand for higher education: Evidence from financial deregulation," *Review of Economics and Statistics*, 2016, *98* (1), 12–24.

Van der Klaauw, Wilbert, "Estimating the effect of financial aid offers on college enrollment: A regression–discontinuity approach," *International Economic Review*, 2002, *43* (4), 1249–1287.