Student Debt and Early-Career Earnings

Emerson Au, Justin Sweeney, Harish Rajagopalan

Repayment of student loans usually starts when graduates begin low-paying jobs, which creates concerns that repayment obligations might lower early-career wages. This study investigates whether changes in borrowers' yearly incomes are related to the beginning of student loan repayment. We track a cohort of individuals born between 1980 and 1984 from 1997 to 2007 using the National Longitudinal Survey of Youth 1997 (NLSY97). We use a staggered difference-in-differences event-study approach to estimate wage dynamics around the first year in which a borrower's outstanding loan total decreases. Pre-treatment estimates violate the parallel trends assumption necessary for a causal interpretation by demonstrating consistent disparities in salary trajectories between earlier and later repayers. Post-treatment effects are negligible and statistically indistinguishable from zero. We find no conclusive evidence that salaries change at the start of repayment within these bounds, and our research is further limited by missing controls on hours worked, occupation, sector, and socioeconomic background.

Introduction

The way US graduates enter the labor market is influenced by their student loan debt. Repayment usually starts when employees start low-paying, unstable jobs in their early careers. Previous research indicates that debt can impact labor supply, sector choice, and job search, but there is little data on whether salaries change after repayment starts, particularly among comparable borrowers.

This paper provides descriptive evidence on wage patterns around the start of repayment from the NLSY97 panel. We track annual wages across early adulthood and examine how they change in the years before and after repayment starts. By using an event-study framework that compares earlier and later repayers, we assess whether there is a difference in wage trajectories after the beginning of repayment.

Literature Review

Recent studies consider links between student loan debt, early-career job search, sector choice, and income, but results differ across disciplines and identification strategies. Ji (2020) develops a dynamic general-equilibrium search model calibrated to NLSY97 data and shows that, because of repayment obligations, indebted graduates search less intensively and accept lower-paying first job offers. In his calibration, borrowers search about 0.9 weeks less and earn roughly $1,479 less in their first year than comparable non-borrowers. Reduced-form OLS estimates indicate that an additional $10,000 of debt is associated with 1.41–1.57 fewer weeks of search and 2.7–4% lower early earnings. Ji also runs counterfactuals that highlight the role of policy design: income-based repayment (IBR) insures job-search risk, raises average wage income, nearly eliminates default, increases college attendance, and generates welfare gains of about 1.8% through higher lifetime income and consumption. Overall, the paper provides evidence for a liquidity-constraint mechanism in which initial wages and search effort are both reduced by student debt.

Complementing this mechanism, Rothstein and Rouse (2007) present causal evidence that debt alters occupational choice. Analyzing a no-loans financial-aid policy reform at an elite university through difference-in-differences and instrumental variables methods, they show that increases in debt push students into the higher-salary private sector, and away from public service. In particular, they find that each additional $10,000 of debt decreases the probability of entering a government, nonprofit, or education career by 5–6 percentage points (from a $\sim 17\%$ baseline), and is associated with an approximately $2,000 higher annual starting salary. In a contrasting study, Daniels and Smythe (2019) utilize the National Longitudinal Survey of Youth cohort 1997 (NLSY97), within a difference-in-differences framework, to find that debt holders earn +8% more income, work +6% more hours, have +1% higher wages, and have a +5 percentage point higher probability of full employment than non-debt holders. Their findings indicate that the majority of income differences are due to labor supply and not a wage premium, suggesting a behavioral adjustment, where debt holders work more hours to fulfill their obligations, despite small differences in hourly earnings.

Policy designs elsewhere around the world help to support these observations. For example, Beyer, Hastings, Neilson, and Zimmerman (2015) examine the loan-cap system in Chile which ties allowable borrowing to potential future earnings by degree level and characteristics about the student. The loan caps shift resources to programs with higher expected returns which can affect the rates of graduation, repayment, and default. While the Chilean analysis does not provide estimates about individual wages as the studies in the U.S. do, the Chilean evidence still points to how design features of the programs can create incentives that drive both borrowing and labor market outcomes based upon expected earnings. Across these studies, a consistent narrative is built: debt shifts

early-career behavior–through search intensity (Ji), sector sorting (Rothstein & Rouse) and labor supply (Daniels Jr. and Smythe)–potentially augmented by policy rules (IBR; loan caps). Our contribution is to extend this research using NLSY97 data from 1997–2007, restricting attention to individuals who ever hold student loans, and applying a staggered difference-in-differences/event-study design to examine wage trajectories around the start of loan repayment. We focus on log annual wages and ask whether there is any detectable change in wages when repayment begins, relative to borrowers who have not yet started repayment. Given data limitations on hours worked, sector, and detailed socioeconomic background, we do not attempt to separate wage effects from labor supply or sectoral mechanisms; instead, we provide descriptive evidence on whether wages themselves move at the onset of repayment.

Data

We use data from the U.S. Bureau of Labor Statistics (NLSY97), a sample of 8,984 people born from 1980–1984 and surveyed annually or biennially since 1997. The panel structure allows us to observe wage trends before and after the repayment of loans has started. We focus our study on the years from 1997–2007 to avoid labor market disruptions from the 2008 financial crisis and because survey participation becomes sparser in later years, with more survey years missing. This time window captures respondents in their early labor market years.

Our analysis is restricted only to those who have received a student loan. Non-borrowers differ systematically from borrowers in ways that affect both wages and repayment behavior. Including them would create a third, structurally different comparison group and would obscure the actual effect of starting repayment.

A limitation of the dataset is that relatively few people have fully repaid their loans within our observation window. As a result, our analysis focuses on the effects of starting repayment rather than the effects of completing repayment. We take the natural log of yearly wages to reduce skewness and heteroskedasticity in the distribution and to interpret estimated effects as approximate percent changes.

Methodology

We estimate the effect of beginning student loan repayment using a staggered difference-in-differences (DiD) approach designed by Callaway and Sant'Anna(2021). A traditional two-period DiD would be inappropriate because individuals begin repayment in different years, so there is no well-defined "before" and "after" period. The staggered DiD framework addresses this by treating each repayment year as its own treatment cohort.

Treatment is defined as the first year in which an individual's outstanding loan balance decreases relative to the previous year, indicating the start of the repayment process. For example, all individuals who begin repayment in 2004 form the 2004 group, those who started repayment in 2005 form the 2005 group, and so on. For each treated cohort, the control group consists of borrowers who have not yet started repayment in that year (and are therefore not yet treated) or never start repayment in the sample period. By restricting the treatment and control group to only borrowers, we ensure that comparisons are made between individuals with similar repayment responsibilities and financial constraints.

The estimator computes the Average Treatment Effect on the Treated (ATT) for each separate cohort and each post-treatment year. We then aggregate each cohort's ATTs using a simple average to obtain an estimate that reflects the average effect of beginning repayment across all treatment cohorts. Parallel trend assumptions are evaluated in the Analysis section.

Results and Analysis

Table 1 reports the event-study estimates from the staggered DiD design. Each event-time coefficient measures the difference in ln(wage) between borrowers who have begun repaying their loans at event time t and borrowers who have not yet started repayment in that same year, averaged across repayment cohorts.

The pre-treatment coefficients (Tm5–Tm1) reveal notable differences between the two groups before repayment starts. In particular, Tm4 is large and highly significant, and the pre-treatment average (Pre_avg) is strongly negative and statistically different from zero. These results indicate that borrowers who repay their loans earlier follow different wage trajectories than those who repay their loans later. Nearly all the post-treatment coefficients (Tp0–Tp4) are statistically insignificant, meaning there is no consistent pattern of wage changes after the repayment process has begun. Although Tp3 is significantly different from zero, the violation of the parallel trends assumption in the pre-treatment period prevents this from being interpreted as a causal effect of repayment.

Figure 1 visualizes these event-study estimates. Each point corresponds to an ATT estimate for a given year relative to the start of repayment, and the vertical bars show 95% confidence intervals. The horizontal axis measures

Table 1: Event-Study Estimates of Aggregate Differences in ln(wage)

Event Time	Coefficient	Std. Error	z-stat	p-value	CI Lower	CI Upper
Pre_avg	-0.4329	0.0655	-6.61	0.000	-0.5612	-0.3046
Post_avg	-0.1492	0.2926	-0.51	0.610	-0.7227	0.4242
Tm5	-0.3993	0.2613	-1.53	0.126	-0.9115	0.1128
Tm4	-2.2567	0.4817	-4.68	0.000	-3.2009	-1.3126
Tm3	0.3266	0.2683	1.22	0.223	-0.1992	0.8524
Tm2	0.1435	0.2252	0.64	0.524	-0.2979	0.5849
Tm1	0.0214	0.1374	0.16	0.876	-0.2478	0.2907
Tp0	0.0605	0.1139	0.53	0.595	-0.1626	0.2837
Tp1	0.2407	0.1878	1.28	0.200	-0.1274	0.6087
Tp2	-0.2774	0.3696	-0.75	0.453	-1.0019	0.4471
Tp3	-0.5664	0.1202	-4.71	0.000	-0.8019	-0.3308
Tp4	-0.2036	1.0956	-0.19	0.853	-2.3510	1.9438

years relative to treatment, while the vertical axis shows the estimated change in ln(wage).

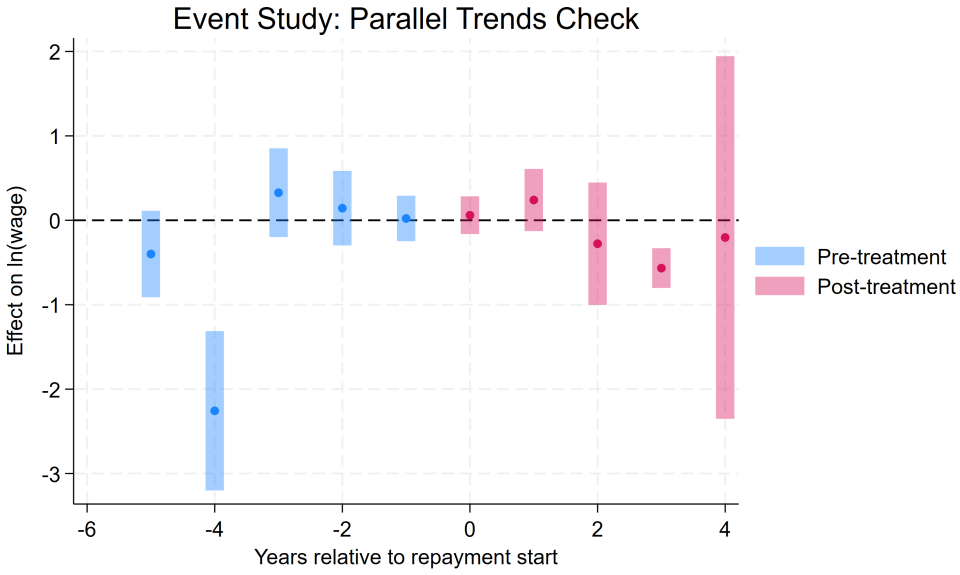

Figure 1: Event-study plot of ATT estimates before and after repayment starts.

Most of the pre-treatment points are close to zero, but the joint test of all pre-treatment coefficients strongly rejects the null that they are all zero ($\chi^2(10) = 261.23$, $p < 0.001$), confirming that early and late repayers differ systematically even before repayment begins.

Table 2 reports the overall Average Treatment Effect on the Treated (ATT) combining post-treatment periods. The point estimate is small in magnitude and very imprecise.

Table 2: Average Treatment Effect on the Treated (ATT)

	Coefficient	Std. Error	z-stat	p-value
ATT	-0.0573	0.3531	-0.16	0.871

With a high p-value (0.871), there is no statistically significant effect of loan repayment on wages detectable in these results. Combined with the strong pre-treatment differences documented in Table 1, this means we cannot interpret the ATT as the causal effect of starting repayment: there is no clear, robust relationship between the onset of student loan repayment and annual wages in this sample.

Conclusion

Within this sample of NLSY97 borrowers we do not see a significant, consistent shift in yearly salaries at the start of student loan repayment. The parallel trends assumption that supports a causal difference-in-differences interpretation is violated by the substantial pre-treatment differences between borrowers who begin repayment earlier and those who begin later, who already have distinct salary trajectories. These pre-existing differences suggest that borrowers differ well before repayment begins, which makes it difficult to isolate any causal effect of repayment itself. We cannot say that repayment itself consistently increases or decreases wages since, after payback begins, the estimated post-treatment effects are insignificant and imprecise, and the aggregate ATT is statistically indistinguishable from zero.

It is important to read these results with several caveats. Our data window is relatively short and includes only the early labor-market years; few respondents fully repay their debts, so we cannot analyze the impacts of being debt-free. We cannot isolate salary impacts from labor-supply or sorting mechanisms and cannot rule out omitted-variable bias since we do not have access to important controls like hours worked, occupation and sector, institutional selectivity, major, and comprehensive socioeconomic background. Furthermore, complexity like postponement, forbearance, or consolidation may be overlooked by our treatment measure, which is the first recorded drop in outstanding balance. The paper's contribution is therefore descriptive: we show that early and late repayers are on distinct salary paths well before payback begins, and we find no discernible wage split at repayment initiation. To determine the causal effect of student debt payment on wages and to better link wage dynamics to search behavior, sector choice, and labor supply, future research utilizing richer administrative loan data or policy-driven variation in repayment terms is required.

References

Beyer, H., Hastings, J., Neilson, C., & Zimmerman, S. (2015). Connecting student loans to labor market outcomes: Policy lessons from chile. *American Economic Review*, *105*(5), 508–513.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.

Daniels Jr., G. E., & Smythe, A. (2019). Student debt and labor market outcomes. *AEA Papers and Proceedings*, *109*, 171–175.

Ji, Y. (2021). Job search under debt: Aggregate implications of student loans. *Journal of Monetary Economics*, *117*, 741–759.

Rothstein, J., & Rouse, C. E. (2011). Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics*, *95*(1–2), 149–163.

U.S. Bureau of Labor Statistics. (2023). *National longitudinal survey of youth 1997 (nlsy97).* National Longitudinal Surveys, U.S. Department of Labor. (`https://www.nlsinfo.org/content/cohorts/nlsy97`)