Family Resources and Human Capital in the Great Recession*

Garrett Anstreicher
University of Nebraska-Lincoln

January 23, 2025

Abstract

I study how the Great Recession impacted the human capital of young adults. Using a novel dataset from U.S. Census, I document that the negative effects of worse local economic shocks on educational attainment are concentrated among middle-class children. I then develop a model of selection into college and life-cycle earnings. Simulating a recession in the model produces a "hollowing out the middle" in lifecycle earnings — lifetime earnings gaps between poor and wealthy children decrease, but at the same time such gaps between middle-class and wealthy children widen. An expansion to financial aid that more completely insures against parental wealth shocks can mitigate these effects while being approximately revenue-neutral, and the presence of a recession makes progressive financial aid policy more cost-effective in general.

JEL Codes: D15, D64, E24, E32, I2, J2, J24

Keywords: Human capital, college choice, impacts of recessions, borrowing constraints

*Email: ganstreicher2@unl.edu. I am heavily indebted to John Kennan for his insights and patience. Thanks also to Dean Corbae, Diego Daruich, Joel McMurry, Jo Mullins, Mike Ricks, Ananth Seshadri, Jeff Smith, Brittany Street, Joanna Venator, Chris Taber, Sarah Turner, Matt Wiswall, Anson Zhou, and audiences at SOLE, SED, NASMES, and the USC Macro Day for many helpful comments. The editor, Eric French, and three anonymous referees provided feedback that substantially improved the paper. Thanks to Bob Thomas for every step of applying for and using the restricted-use ACS and Decennial data. Additionally, this material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1747503. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2413. (CBDRB-FY20-P2413-R8725, CBDRB-FY22-P2413-R9572, and CBDRB-FY25-P2413-R12040). All remaining errors are my own.

1 Introduction

How do recessions impact human capital accumulation among young adults, and how do these effects vary based on the resources of one's family? Economic downturns decrease the opportunity cost of spending more time building human capital as opposed to working, such as by going to college, but also often feature decreases in family resources through events such as layoffs or asset depreciation. Entries in the economics literature have documented that early exposure to recessions can have long-term negative effects on earnings (Kahn, 2010; Oreopoulos et al., 2012; Stuart, 2022), as well as that the role of parental resources and educational borrowing constraints in determining college attendance has increased in recent decades (Lochner and Monge-Naranjo, 2011). How these forces interact, however, is not well known and has important implications for understanding the incidence of recessions and informing policy responses to them.

This paper studies the role of family resources in human capital accumulation and earnings in the context of the financial crisis of 2007-08 and the subsequent Great Recession. The financial crash induced severe shocks to employment and housing wealth across the country, with the intensity of these shocks varying considerably across labor markets. This spatial heterogeneity has allowed researchers to use US local labor markets as a laboratory in which to study the effects of the Recession (Yagan, 2018). I first conduct an empirical investigation that uses a novel dataset comprised of linked confidential data from U.S. Census. New linkage rules and methods allow for individuals in the yearly American Community Surveys to be linked to themselves in the 2000 Census — in addition to providing the likely local labor market of these individuals at the time of the financial crash, this also allows me to obtain information about their parents and family background if they were a child in 2000. I focus my analysis on the birth cohorts of 1990 through 1994, as these individuals would have been age 13-17 in 2007 and thus were at considerable risk of entering the labor force in a bad economy.

Exploiting the aforementioned spatial variation in shock severity across commuting zones (CZs) and individual variation in parental resources, I find evidence that college enrollment and attainment is considerably less countercyclical for middle-class children than it is for either their poor or rich counterparts. Other things held equal, exposure to a one percentage point larger 2007-2009 CZ unemployment shock in the recession is associated with a 7% long-run reduction in college attainment for children with parents in either the second or third quartile in the parental income distribution. The effects for children with bottom

and top-quartile parents are a 1% gain or 2% reduction respectively, with neither effect being statistically significant. Higher rates of college non-completion appear to be more important in generating this finding than substitution to two-year degrees, and additional evidence points to losses in parental home equity being a particularly important mechanism for reducing college attainment among middle-class children.

The data, however, are limited in their ability to speak to lifecycle and aggregate implications of the recession. To probe these issues and study policy implications, I develop and estimate a model of selection into college and life-cycle earnings. The model most closely resembles that of Colas et al. (2021) in featuring endogenous parental transfers, multiple schooling outcomes, and post-graduation labor market uncertainty. The novel features of my model include multiple schooling options (allowing agents to potentially substitute into cheaper but lower-quality options in response to family resource shocks) and post-graduation uncertainty in *employment* as well as earnings. The model also includes realistic student debt repayment schemes and salient features of U.S. need-based financial aid, including that the assets most relevant for middle-class households — home equity and retirement savings — are excluded from consideration in financial aid calculation.

Simulating a shock resembling the Great Recession results in a "hollowing out of the middle" in lifecycle earnings in the aggregate that mirrors the partial effect found in the reduced-form section: while the recession in fact *reduces* the gap in lifetime earnings between top and bottom-decile children in the parent income distribution, the gap between middle and top-decile children rises by nearly $30,000. Educational borrowing constraints as well as changes made to them in the financial crash play a notable role in generating this result, but the model predicts that this hollowing effect would still have been observed to a substantial extent had the recession occurred in a world with no student loan limits whatsoever. However, a counterfactual expansion to need-based financial aid that considers all family assets, including home equity, in award calculation can mitigate this effect while being approximately revenue-neutral if enacted during the downturn. The presence of the recession enhances the cost-effectiveness of financial aid more broadly, with this being especially true for progressive financial aid policy and the changes to student loan limits and Pell grant generosity that were made in the real world in response to the financial crisis.

Related Literature

This paper contributes to several strands of economic literature, the first being the one that studies the effects of recessions on individual outcomes. Exposure to recessions has long-lasting consequences for earnings and human capital (Stuart, 2022), with particularly large

effects for workers who were displaced during an economic downturn (Davis and Wachter, 2011). Yagan (2018) finds that the local shocks induced by the Great Recession had large effects on the employment and income of working-aged adults, and Rinz (2022) finds that adolescents in the United States were even more strongly affected by the crash in percentage terms. However, neither study is able to explore heterogeneity in effects across socioeconomic background in their samples.

Past literature has also established that entering the labor force in a recession in particular has long-lasting negative impacts on individual earnings (Kahn, 2010; Oreopoulos et al., 2012). However, the extent to which individuals attempt to avoid or mitigate the adverse effects of recessions is not well understood. Yagan (2014) finds that working-age adults did not appear to move to less-affected labor markets in response to local economic shocks. Long (2014) reports that college attendance increased more in states that were more affected by the recession, indicating some evasive action being taken on the part of adolescents; (Charles et al., 2018) also find that college enrollment moved opposite with the housing boom and bust in the 2000s. Explicitly modeling these decisions and focusing the heterogeneity of effects of recessions over the parent resource gradient are my main contributions to this literature.

In doing this, this paper also speaks to the large literature that has studied the role of family resources and credit constraints in college enrollment and completion and extends it through studying the interaction of these factors with the business cycle. While older entries in this literature suggest a limited role of family income in deciding whether children go to college (Cameron and Taber, 2004), other papers have documented that family resources may have become more important in recent years (Belley and Lochner, 2007; Lochner and Monge-Naranjo, 2011). This empirical results in this paper provide support for the latter interpretation and suggest different roles for family income and family wealth in financing college enrollment, offering a potential resolution between papers that have separately documented small parental income effects on college enrollment (Hilger, 2016) but large effects for parental home equity (Lovenheim, 2011).

The third literature I contribute to is the one that models human capital formation in earlier parts of the lifecycle. Parent inputs and resources are vital in early (Becker and Tomes, 1979; Cunha and Heckman, 2007; Lee and Seshadri, 2019) and, increasingly, later (Caucutt and Lochner, 2020) stages of child human capital development — this paper complements these papers by studying the importance of parental inputs and resources on the transition to adulthood, which Danziger and Ratner (2010) argue has become considerably more difficult over recent years. The model I develop most closely resembles the models in Heckman et al.

(1998) and Colas et al. (2021), with the main contribution being that I enhance the realism of educational decisions by explicitly including a richer menu of college enrollment decisions and include employment risk after graduation. Multiple college options is a particularly salient feature, since substitution to lower-quality options can have marked impacts on later-life outcomes (Mountjoy, 2022). In evaluating the incidence of counterfactual tuition and student loan policies, I also contribute to the literature that has studied the optimal design of such policies (Johnson, 2013; Abbott et al., 2019) by providing estimates regarding how it varies over the business cycle.

Methodologically, this paper is one of the first to form linkages *across* mandatory 21st-century Census surveys. Individuals in older vintages of Census data have been linked across time via probabilistic matching algorithms in past research [e.g. Ferrie (1996); Feigenbaum (2016)], but issues of match quality produced by these algorithms can be severe (Bailey et al., 2020). The use of administrative records and the finer geographic information available in restricted-use Census survey data allows for cross-survey linkages to be made considerably more confidently and frequently, which may make the data I use a valuable resource for researchers going forward. A supplemental report that details the availability of intergenerational linkages between the ACS and 2000 Long-Form by birth cohort, geography, and race can be found on the author's website.

The remainder of the paper proceeds as follows. Section 2 introduces the data I use in this paper along with my empirical strategy, while Section 3 details my empirical results. Section 4 lays out my model, and Sections 5 describes my estimation strategy and evaluates the model's fit. Sections 6 studies the effects of the Recession in the model environment and its interaction with counterfactual policy. Section 7 concludes.

2 Data

2.1 Sample Description

To estimate how the effects of the Great Recession on human capital accumulation among young adults varied over family resources, I use newly available linked confidential data consisting of the 2000 Decennial Census and the 2005-2018 waves of the American Community Surveys. In addition to containing considerably richer geographic information for the respondents, individuals in both these datasets have been linked to the Social Security Administration NUMIDENT file using the U.S. Census Bureau's Person Identification Val-

idation System, which allows me to assign Protected Identification Keys (PIKs; essentially scrambled SSNs) to the large majority of individuals in the surveys[1].

While other papers have made use of these data as repeated cross-sections (see for instance Stuart (2022)), recent regulation changes regarding linking mandatory Census surveys allow for respondents in these surveys to be *linked to themselves* over time[2]. With PIKs being stable over time and the large majority of U.S. citizens appearing in either the short or long-form 2000 Census, this allows me to link individuals in the ACS to themselves in the 2000 Census, effectively forming a two-year panel for a 1% sample of the U.S. over the years 2005 to 2018. This linkage allows me to observe detailed family background information among adults in the ACS who no longer live with their parents, overcoming a key limitation in the public-use version of the data.

My baseline sample consists of individuals born between 1990 and 1994 in any ACS year. These birth years translate to these cohorts being aged between 13 and 17 in 2007 — thus, these cohorts would have been at considerable risk of entering the labor force during the recession. Further, because these birth cohorts were children in the year 2000, linking these individuals in the ACS to themselves in the 2000 Census allows me to infer their family background by observing information about household heads and their spouses[3]. For individuals linked to themselves in the 2000 short form, I can see their parents' race, marital status, age, and location down to the census block. For those linked to the long-form Census, I can see the same information along with a rich vector of parent demographic and socioeconomic covariates. I limit the sample to individuals who were observed in the long-form Census and were coded as either the child, adopted child, or stepchild of the household head. Individuals living in Group Quarters in the 2000 Census are dropped, as are any groups of individuals in any survey year who are linked to the same PIK.

After performing the linkages and applying the sample restrictions listed above, I am left with slightly over 500,000 observations of the 1990-1994 birth cohorts over the 2005-2018 period. Table A.1 reports the rates at which I can both assign individuals in the ACS PIKs

[1]The PVS is Census Bureau's system for searching and verifying Social Security Numbers or Protected Identification Keys for person records in demographic surveys and/or censuses. This system first attempts to match individuals in surveys to SSA records by using address information — if this fails, individuals are matched via name and date of birth. A more comprehensive review of this procedure and its performance can be found in Mulrow et al. (2011)

[2]Conversations with Census Research Data Center administrators suggest that this change happened sometime in the 2019 calendar year.

[3]Over 90% of individuals in these birth cohorts are coded to be either the child, stepchild, or adopted child of their household head in 2000.

Table 1: Sample Summary Statistics

Variable	ACS	ACS (PIK Linked)	ACS (Census Linked)
Age	19.07	19.08	19.03
	(4.36)	(4.36)	(4.32)
Male	50.92	50.74	50.86
	(49.99)	(49.99)	(49.99)
White	59.64	61.47	71.92
	(49.06)	(48.67)	(44.94)
Black	15.54	15.37	11.31
	(36.23)	(36.07)	(31.67)
Asian	4.85	4.58	2.78
	(21.49)	(20.89)	(16.44)
Hispanic	21.10	19.70	14.71
	(40.80)	(39.78)	(35.42)
Married	6.99	6.95	6.81
	(25.50)	(25.43)	(25.20)
Employed	41.34	41.98	43.15
	(49.24)	(49.35)	(49.53)
Wage Income ($1,000)	9.85	9.97	10.21
	(16.27)	(16.33)	(16.41)
College Degree	9.47	9.71	10.14
	(29.28)	(29.60)	(30.18)
Associate Degree	13.12	13.48	14.17
	(33.76)	(34.15)	(34.88)
High School Degree	52.85	53.22	53.21
	(49.92)	(49.90)	(49.90)
In Northeast	17.11	17.23	17.35
	(37.66)	(37.76)	(37.87)
In Midwest	21.33	21.95	26.83
	(40.96)	(41.39)	(44.31)
In Southeast	36.99	36.91	34.52
	(48.28)	(48.26)	(47.54)
In West	24.57	23.91	21.29
	(43.05)	(42.65)	(40.94)
N	4053000	3702000	539000

Notes: Data from 2000 Census and 2005-2018 American Community Surveys. Statistics are for 1990-1994 birth cohorts. Standard deviations in parenthesis. Sample sizes rounded for disclosure avoidance purposes. Column 2 reports variable means and standard deviations for entire ACS sample. Column 3 reports variable means and standard deviations for ACS sample linked to PIKs. Column 4 reports variable means and standard deviations for ACS sample linked to 2000 Long-Form Census. CBDRB-FY20-P2413-R8725.

and link them to themselves in the 2000 Census. Broadly speaking, the linkage is a success: I am able to obtain family background information for more than 12% of the relevant individuals in most of my ACS years, which is reasonably high given that only one sixth of households filled out the 2000 Long-Form Census in the first place. Table 1 reports how basic demographic and socioeconomic characteristics of my ACS sample change after removing unlinked individuals in either step. The table suggests that the linkage process does induce some selection, with the economic indicators of the sample improving somewhat after the linkage procedure; in particular, many blacks are dropped in the linkage procedure. This is not surprising, as the PIK system relies most heavily on individual addresses to make linkages; thus, individuals who do not have stable addresses are more likely to be unmatched and subsequently dropped. To address this selection issue, I calculate the ex-ante probability of being included in the baseline sample by state, gender, and race (white, black, Hispanic or Asian) and modify individual sample weights by dividing them by these probabilities[4].

2.2 Empirical Strategy

I estimate the effects of exposure to the Great Recession using an event-study specification that is similar to Yagan (2018) and Rinz (2022). Specifically, I estimate the following regression:

$$y_{iact} = \boldsymbol{\alpha} \boldsymbol{X}_i + \theta_c + \gamma_a + \delta_t + \beta_a S_c + \varepsilon_{iact}, \tag{1}$$

where y_{iact} is some outcome variable of interest (such as earnings or educational attainment) for individual i at age a who was in commuting zone (CZ))[5] c in 2007 and surveyed in year t. Fixed effects for age, CZ, and calendar year $\gamma_a, \theta_c, \delta_t$ control for life-cycle patterns, geographic mean differences, and aggregate yearly trends in the outcome variable, and $\boldsymbol{X}_i$ is a rich vector of demographic and background covariates including parental marital status, parental income, parental homeownership, parental house value, parental education, and a dummy for the year-2000 household head working in either construction or manufacturing. These controls are not included in the baseline specification but are added in later as a robustness check. Age fixed effects γ_a are also interacted with Census division so that comparisons

[4]In practice, using these modified weights vs. the original sample weights does not change the sign or significance of the empirical results.

[5]Commuting Zones are a frequently-used notion of local labor markets and groups of counties characterized by strong economic ties. There are 722 CZs in the United States. CZ shocks are computed by taking county-level unemployment rates from the monthly BLS Local Area Unemployment Statistics series and then aggregating to CZs with population weights before averaging across months to form a yearly CZ measure. For a visualization of these shocks, refer to Figure A.1.

are made within reasonably similar geographic units[6]. The error term is represented by ε_{iact}. Following Yagan (2018), S_c is the 2007-2009 change in unemployment rate observed in commuting zone c. Thus, the β_a terms are the coefficients of interest and have the interpretation of the effect on some outcome at age a of living in 2007[7] in a labor market that experienced a one percentage point larger unemployment shock in the Great Recession.

To facilitate the inclusion of the CZ fixed effects θ_c, I normalize one of the β_a terms to be zero. The identifying assumption in this specification is parallel trends, or children in CZs that received larger shocks in the recession were not systematically trending differently in age profiles in the outcome variable compared to children in CZs that received smaller shocks. However, the age of the cohorts studied result in pre-recession earnings and college attainment being largely missing, thereby preventing a direct validation of this assumption — as such, the identifying assumption is tantamount to selection on observables. I aim to make the identifying assumption credible by including a comprehensive set of controls in my regressions in a robustness check[8] — additionally, I verify that the empirical specification applied to the entire sample can reproduce estimates of Great Recession earnings impacts reported in other entries of the literature as a face validity test. I cluster standard errors by year-2000 CZ to allow for arbitrary serial correlation of error terms within CZs.

3 Empirical Results

3.1 Results from Full Sample

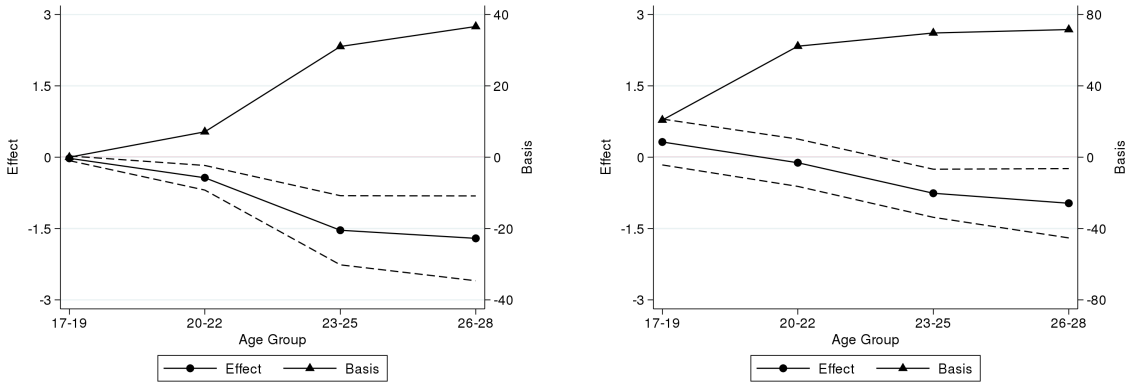
I first estimate the effects of the recession on human capital for my entire sample to assess the face validity of the output of my specification. When estimating (1), I group ages into five bins ($< 17, 17-19, 20-22, 23-25, 26-28$) to increase precision and to reduce disclosure risks. The first age group is used as the omitted category, so all estimated coefficients for a

[6]There are nine Census divisions, listed in Table A.2. I prefer division-by-age fixed effects to state-by-age fixed effects because within-state variation of CZ shock intensity is quite limited: state fixed effects can explain approximately 70% of the variation in CZ unemployment shocks, while the corresponding statistic for division fixed effects is 40%.

[7]One limitation of my data is that the 2007 CZ of residence for survey respondents is not directly observable, so I use their year-2000 CZ of residence as a proxy. Looking at the 2007 ACS allows a direct evaluation of how good this approximation is: I find that 87% percent of my sample in 2007 lived in the same CZ as in 2000, with minimal heterogeneity over the parent income gradient.

[8]I verify that these controls are well-balanced across CZs in Table A.3 by regressing CZ shocks on them. After controlling for division fixed effects, the additional inclusion of the control variables does little to change the explanatory power of the regression, and the coefficients on the variables, while statistically distinguishable from zero, are generally small in magnitude.

Figure 1: Full-Sample Effects of the Recession on Educational Attainment



(a) Four-Year Degree Attainment

(b) Some College Attainment

Notes: Figures display outcome variable means over age when surveyed and estimates of β_a from equation (1), with age grouped into five bins and effects for the bottom age group normalized to zero. Outcome variables include indicator variables for attainment of at least a four-year college degree or at least one year of college attendance. Standard errors are clustered at the year-2000 CZ level, and the dotted lines indicate 95% confidence intervals for coefficient estimates. Coefficient estimates correspond to the circular dots and are measured by the left-hand axes; bases correspond to triangular dots and are measured by the right-hand axes. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY22-P2413-R9572.

given outcome variable are interpretable as the effects of the Recession relative to its impacts on individuals aged younger than 17. Effects on binary variables are converted to percentage points terms to ease interpretability.

Figure 1 displays the estimates of the effects of the Great Recession on educational attainment over age groups. The estimates suggest that exposure to a 1 pp larger unemployment shock is associated with a 1.7 pp lower probability of college attainment or higher by age 26-28. The estimate attenuates to approximately 1 pp when using an indicator variable of having attained at least some college as the dependent variable, suggesting that both non-attendance and non-*completion* are at work in driving the four-year degree results.

Other papers that have used the same CZ unemployment shock variation allow me to assess the face validity of these estimates. Rinz (2022) conducts subgroup analyses of the effects of the Recession on earnings and human capital over different generations (Millenials, Gen-X, Baby Boomers, and the Silent Generation). He estimates a long-run 1.3pp reduction of completion of any post-secondary degree in response to a 1 pp larger CZ unemployment

shock, which is quite similar to my estimate for the effects of the Recession on college attainment[9] That my estimate is somewhat larger is reasonable, given that I focus on college attainment instead of any postsecondary degree attainment — moreover, his Millenial sample contains individuals who were born at 1981 at the earliest and thus would have contained a substantial number of individuals who already completed their education by the time the Great Recession occurred. The results for college attainment are also similar in magnitude to those reported in Stuart (2022), who finds that a one standard deviation[10] increase in exposure to the 1980 recession is associated with a 3 pp decrease in four-year college attainment.

3.2 Parent Resources and the Effects of the Recession

The results using the full sample provide evidence for the validity of my empirical specification. To study how the effects of the Great Recession on human capital accumulation and earnings varied over the parent income gradient, I conduct subgroup analyses by estimating (1) over parent income quartiles[11]. I prefer subgroup analyses to including interaction terms in the regression, as the former allows for age/year fixed effects and the demographic controls in $\boldsymbol{X}_{it}$ to differentially impact the outcome variable of interest based on socioeconomic background in a fully flexible way.

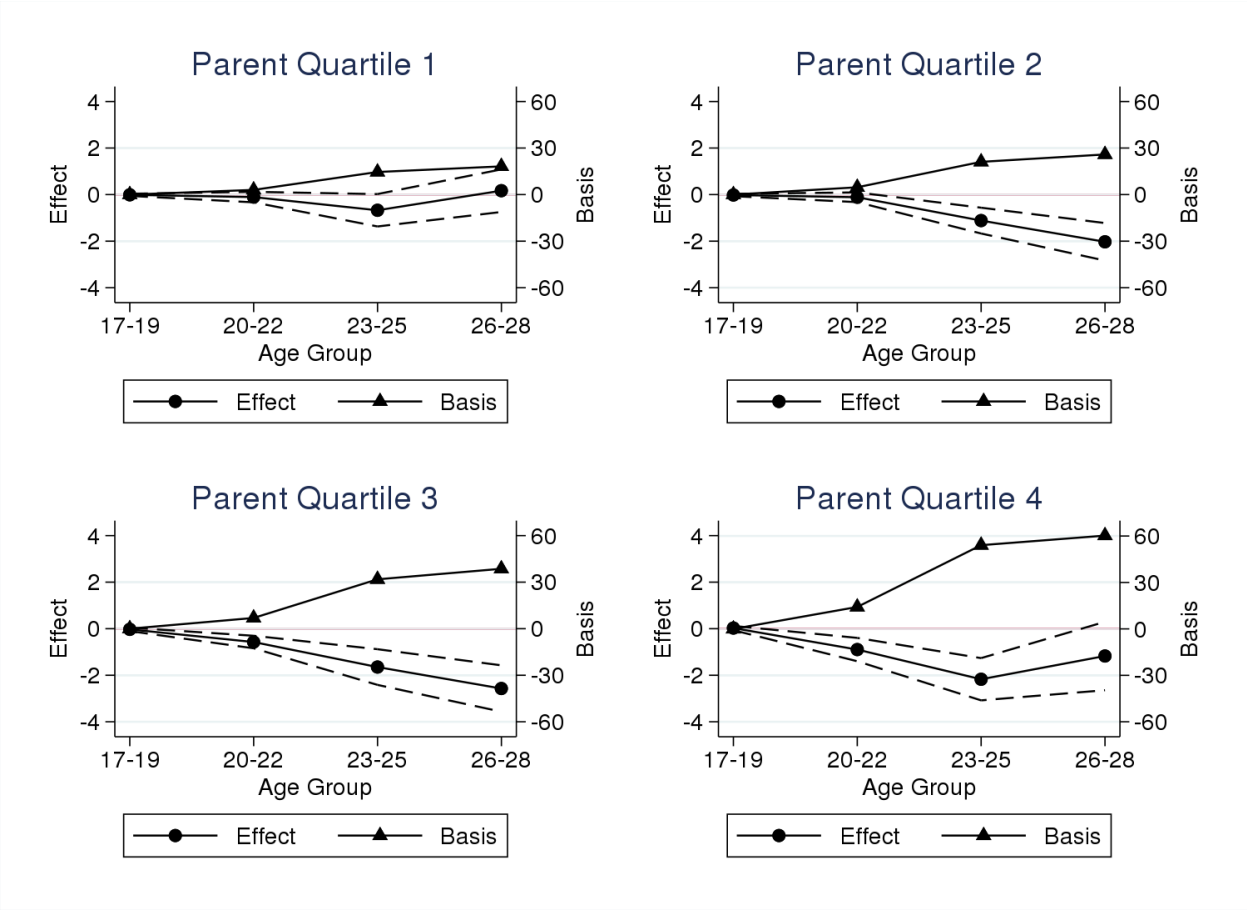
Figure 2 presents the results of this exercise for four-year college attainment or higher, with individual plots structured similarly to those in Figure 1. By age 26-28, at which point college attainment rates flatten considerably in the data, exposure to a 1pp larger unemployment shock in the Great Recession is associated with a 2pp or larger reduction in four-year college degree attainment for children who had parents in the second or third income quartiles. The comparable estimate for children with top-quartile parents is smaller (slightly over 1pp) and not statistically distinguishable from zero, though equivalence of the coefficients for second/third-quartile children is difficult to reject. However, the estimate age 26-28 coefficient for bottom-quartile children is essentially zero and is statistically distinguishable from the second/third quartile effects.

[9] I also estimate the impact of larger unemployment shocks on annual wage earnings in Figure A.2 and find a long-run effect of $600, which is identical to the Millenial earnings estimate from Rinz (2022).

[10] A 1pp larger CZ unemployment shock is around a 0.65 SD increase in my specification.

[11] Parent income is defined as the sum of household head and spouse of household head (if present) wage income in 2000. Quartiles are computed nationally and within-cohort to account for potential age differences of parents for individuals who were born in earlier or later years. While different commuting zones may be more likely than others to have residents in certain parts of U.S. income distributions, quartiles are coarse enough a grouping to guarantee that all CZs contain at least some individuals from each income category.

Figure 2: Effects of Recession Shocks on College Attainment by Parent Income Quartile



Notes: Figures display outcome variable means over age when surveyed and estimates of β_a from equation (1) for the outcome variable of at attainment of at least a four-year college degree, with age grouped into five bins and effects for the bottom age group normalized to zero. Effects estimated separately over year-2000 parent income quartile. Standard errors are clustered at the year-2000 CZ level, and the dotted lines indicate 95% confidence intervals for coefficient estimates. Coefficient estimates correspond to the circular dots and are measured by the left-hand axes; bases correspond to triangular dots and are measured by the right-hand axes. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY22-P2413-R9572.

Also of interest is that children from all types of socioeconomic backgrounds see drops in college attainment at age 23-25. The null effects for top and bottom children come from a recovery in college attainment as opposed to null effects over the entire life cycle. This suggests that the recession may have *delayed* four-year college degree attainment for all types of children but that these delays became permanent primarily for middle-class children while wealthier children became more likely to pursue non-traditional education paths[12].

An important caveat to consider is that the effects displayed in Figure 2 are in levels as opposed to percentage terms, which is relevant because the base rates of college completion vary substantially over the parent income gradient. Roughly 30% of children with second or third-quartile parents had completed a four-year college degree or higher by age 26-28, while the corresponding figure for top-quartile children was over 60%. Figure 3a takes the estimates and standard errors of β_{26-28}[13] for the four types of children and divides them by the bases to obtain percentage effects before plotting these effects over parent quartiles. A clear pattern emerges: in percentage terms, long-term human capital losses from the Great Recession were sharply concentrated among children from the middle class. Exposure to a 1pp larger shock is associated with a roughly 7% decline in college attainment for these children, while top/bottom-quartile children saw an insignificant 2% loss and 1% gain[14].

The remaining panels of Figure 3 conduct the same exercise for other outcome variables, including attainment of at least some college (Figure 3b), attainment of an associate degree or higher (Figure 3c), and annual wage earnings (Figure 3d)[15]. A very similar U-shaped pattern of heterogeneity emerges when evaluating the effects on associate degree attainment, though the impact on middle-class children attenuates to be around 5%. When considering some college attainment or higher, the effects attenuate markedly further, and the U-shape pattern flattens considerably, though the percentage effects for top/bottom-quartile children are still different from those of middle-quartile children at a $p < 0.05$ level. This suggests that college non-completion may have been a more important mechanism than substitution into two-year degrees in generating the results plotted in Figure 3a, consistent with evidence that enrollment expansions associated with recessions are concentrated among lower-quality colleges (Bound et al., 2010; Barr and Turner, 2013).

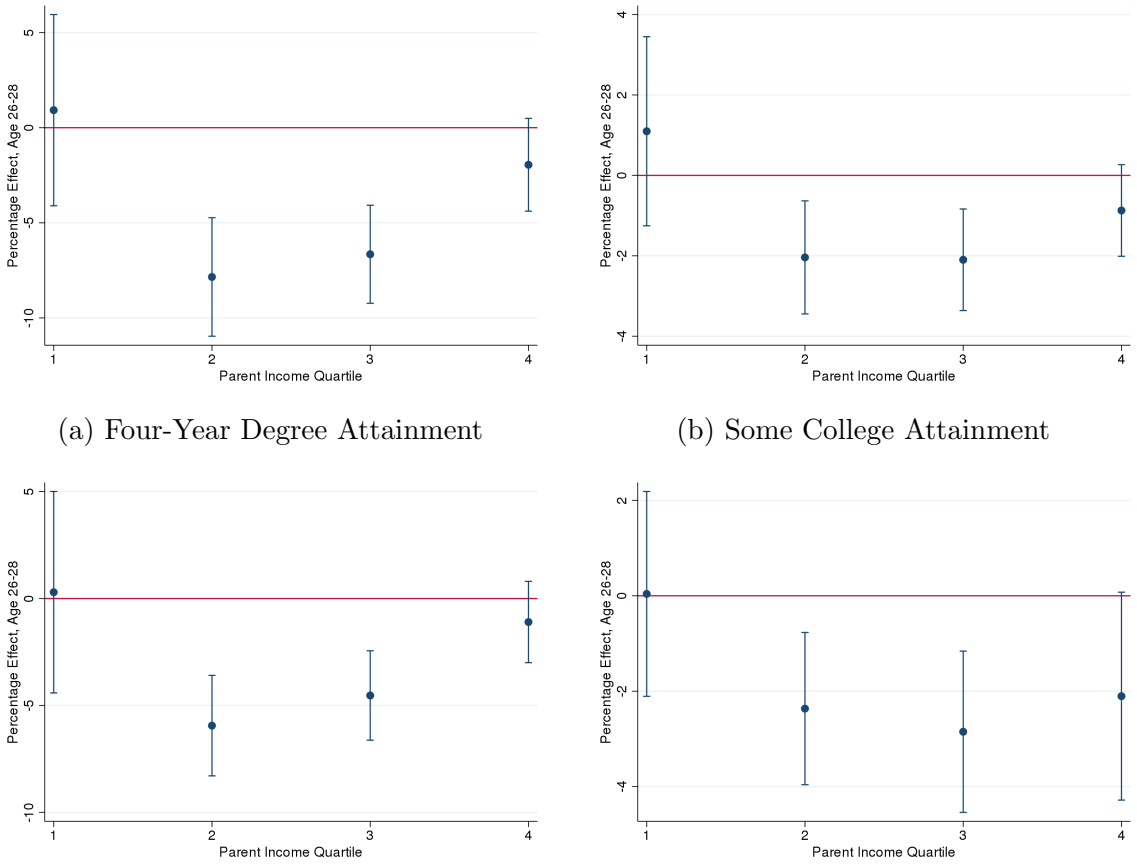
[12]Figure A.3 appears to corroborate this interpretation, with top-income children being more likely than others to be enrolled in college at age 23-25.

[13]One can interact the CZ shocks with time indicators instead and obtain similar results; see Figure A.4. Figure A.5 further demonstrates that the main results are not sensitive to the inclusion of the additional control variables described in Section 2.2.

[14]Discretizing these impacts over income deciles reveals a comparable U-shape; see Figure A.6.

[15]Level plots for these outcome variables in the style of Figure 2 are available upon request.

Figure 3: Percentage Effects at Age 26-28 Over Parent Income Quartile



(a) Four-Year Degree Attainment

(b) Some College Attainment

(c) Associate's Degree Attainment

(d) Annual Wage Earnings

Notes: Figures display estimates of β_{26-28} from equation (1) for the indicated outcome variable over parent income quartile. Coefficient and standard error divided by outcome variable mean (except for log wage earnings) for relevant parent quartile group at age 26-28 (when surveyed) to form percentage effects. Standard errors are clustered at the year-2000 CZ level, and 95% confidence intervals are displayed. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY22-P2413-R9572.

The U-shape heterogeneity persists when looking at wage earnings, though the pattern does become considerably noisier. Figure 3d demonstrates that long-run percentage effects on wage earnings from the Great Recession were only statistically distinguishable from zero for middle-class children, though equivalence of the coefficients for middle-class and top-quartile children clearly cannot be rejected. However, the effects for bottom-quartile children, in

addition to being almost exactly zero, are statistically distinguishable from those of other types of children. This suggests that shocks to human capital accumulation via schooling were a more important mechanism for decreasing long-term earnings than was potential human capital depreciation during unemployment, as children of all types saw decreased employment following the Recession initially (see Figure A.7). If human capital depreciation during unemployment in the Recession was severe, we may expect to see long-term earnings reductions for low-income children as well. At the same time, the young age of the sample may mean that the earnings impacts have yet to fully crystallize, and enabling the study of the full life-cycle implications of these findings is a key motivation for developing a model.

3.3 Mechanisms

The previous results suggest that the Great Recession decreased human capital among middle-class children. In this section, I investigate the mechanisms that drove these impacts. I consider two potential mechanisms in particular that may have made college attendance prohibitively expensive for these children: losses in parental income following layoffs in the Recession, or losses in parental house equity following the housing market crash. To study these possible explanations, I estimate (1) for children with second or third-quartile parents, broken down by whether their parents owned or rented their residence[16] in 2000 or whether the household head worked in either construction or manufacturing in 2000. The latter delineation serves as a proxy for studying the effects of layoffs, as construction and manufacturing were the two most-affected industries in the Recession by a considerable margin.

Table 2 presents the results of this exercise. Two results quickly emerge: the long-run effects of the Recession on college attainment for middle-income children do not differ at all by whether their parents worked in heavily-affected industries, but heterogeneity of effects over home ownership is sharp. Middle-class children whose parents were renting as opposed to owning do not observe any long-term effects of the Recession on college attainment that can be distinguished from zero, suggesting that the declines in their educational attainment seen earlier may be almost *entirely* attributable to losses in parental home equity. The fourth and fifth columns of Table 2 further split up the middle-class children with home-owning parents by below and above-median home value, finding stronger but also noisier effects for

[16] I detect ownership by whether respondents report positive house value (that is, how much the house/residence would sell for) in the 2000 Census, which is only asked to respondents who report owning their residence. While home equity held by parents would admittedly be the more appropriate metric, it is not measured in the survey.

Table 2: Investigation of Mechanisms for Middle-Class Children

VARIABLES	(1) College	(2) College	(3) College	(4) College	(5) College	(6) College	(7) College
Age 17 to 19	-0.022	-0.012	-0.009	-0.004	-0.066	-0.034	0.000
	(0.024)	(0.031)	(0.044)	(0.030)	(0.049)	(0.027)	(0.029)
Age 20 to 22	-0.354	-0.400	-0.126	-0.392	-0.729	-0.407	-0.293
	(0.101)	(0.108)	(0.189)	(0.118)	(0.231)	(0.137)	(0.115)
Age 23 to 25	-1.434	-1.504	-0.970	-1.426	-2.882	-1.047	-1.482
	(0.317)	(0.297)	(0.478)	(0.260)	(0.635)	(0.342)	(0.397)
Age 26 to 28	-2.373	-2.801	-0.838	-2.639	-4.515	-2.229	-2.317
	(0.371)	(0.404)	(0.547)	(0.464)	(0.836)	(0.506)	(0.466)
Sample	Baseline	HV > 0	HV $= 0$	HV<Med	HV>Med	C/M	Not C/M
Observations	303000	234000	69000	131000	103000	106000	196000
R^2	0.219	0.236	0.173	0.176	0.294	0.180	0.230

Notes: Standard errors clustered at the year-2000 CZ level and are in parentheses. Coefficients and sample counts rounded for disclosure avoidance purposes. The table displays estimates of β_a from equation (1) for the outcome variable of attainment of at least a four-year college degree, with age when surveyed grouped into five bins and effects for the bottom age group normalized to zero. HV: parent house value in 2000. Med: median home value among all parents in 2000. HV < Med indicates parents with house value below median but above zero. C/M: household head worked in construction or manufacturing in 2000. Not C/M: household head did not work in construction or manufacturing in 2000. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY20-P2413-R8725 and CBDRB-FY25-P2413-R12040.

the children of parents with more valuable homes.

These results are consistent with a model wherein parent income and wealth (here captured by home equity) factor differently in financing higher education. Such a scenario may happen if middle-class parents consume their income in a hand-to-mouth manner but allow shocks to wealth to influence how much they finance the educational attainment of their child. This story is consistent with multiple entries of the economics literature: Hilger (2016) finds only modest evidence that parental layoffs between 2000 and 2009 depressed the college attainment of their children, but Lovenheim (2011) finds that house equity increases following the housing boom of the early 2000s substantially increased the college attendance of affected families, with particular strong effects for families with relatively low household

income[17]. Kirk (2023) also studies a similar question to mine while focusing more specifically on house value shocks in the Great Recession, finding that larger house shocks net of changing labor market conditions are associated with decreased bachelor degree attainment due to substitution into two-year programs or dropout. Indeed, using 2007-2009 HPI shocks from the Federal Housing Finance Agency as the source of variation instead of unemployment shocks in my specification results in a very similar pattern to the baseline results; see Figure A.8. Allowing parent income and wealth to factor in differently in influencing the educational attainment of their children will be a priority in the subsequent section.

4 Model

The empirical results described in Section 3 suggest that decreases to parental resources may have depressed the human capital accumulation of middle-class children in the Great Recession, with depreciations in parental home equity being an important mechanism. However, the analysis above is limited in several ways. The ACS sample allows me to observe individuals up through age 28 and therefore cannot speak to lifetime impacts of recession exposure. Moreover, the regression specification does not allow me to identify the effect of nationwide changes to labor markets and higher education that occurred concurrently with the local economic shocks it exploits. This "missing intercept" problem may be salient in this context, since individuals could be quite sensitive to macroeconomic as well as local conditions when making college enrollment decisions—indeed, if the former are more salient in informing college attendance, then the design-based analysis may primarily be detecting the effects of reduced parent resources on enrollment, and the aggregate impact of the recession on enrollment may be positive despite the negative partial effects.

Studying the lifecycle impacts of the recession and assessing its interaction with counterfactual educational policies will require a model. To this end, I develop a model that tracks the education decisions and life-cycle wages and earnings of adolescents. A period is one year. Children start the model at age 18 after having graduated high school and are endowed with a level of ability. The agent's parents start the model with a certain level of income and wealth, which may include both housing and non-housing assets. Parents provide an endogenous college transfer to their children who in turn select whether to pursue

[17]Additional survey evidence also suggests that home equity may be a salient resource for financing higher education among middle-class families: in 2010, home equity loans were the second-largest source of borrowing behind PLUS loans among parents with annual income between $35,000 and $100,000 (Sallie Mae, 2010).

any postsecondary education, weighing the potential economic benefits of doing so against psychic costs and idiosyncratic preferences. The agent faces a menu of college attendance options, and the likelihood of actually obtaining a college degree depends on both the agent's ability and the college option selected. After education outcomes are realized, the agent enters the labor force and gradually pays off any student debt while making labor supply and savings decisions every period until retirement at age 65. Agents derive utility from consumption and leisure, with the latter only being consumed if the agent does not participate in the labor force.

This sequence of decisions closely follows those in Colas et al. (2021), with several key distinctions. First, agents have a menu of college options to choose from, allowing for potential substitution to cheaper but lower-quality schools in the face of parent resource shocks. Second, agents face both income and *employment* risk over the life cycle: choosing to participate in the labor market does not guarantee employment, which is critical to consider when evaluating the impacts of the Great Recession, as it featured a cratering of employment with fairly stable (or even slightly rising) real wages. The probability of employment will depend on educational attainment and accumulated experience, providing additional motivations for human capital investment in my framework. Third, wages are determined by accumulated experience instead of age as in Colas et al. (2021), allowing agents to endogenously adjust their labor force participation behavior in the face of economic contractions with the knowledge that doing so affects their future earnings trajectory. Finally, I model a realistic debt repayment process and include psychic costs associated with taking on student loans to better enable the model to capture the importance of borrowing constraints in influencing the effects of an economic downturn.

The following sections lay out the decision problems solved by agents in the model in more detail before turning to the model's solution and estimation.

4.1 Decision Problems

4.1.1 Parent Transfer

Agents begin the model at age 18 with a continuous level of ability θ, the range of which is normalized to be 0 to 1, that influences college tastes and completion probabilities as well as earnings later in life. Parents start with income I^p and wealth W^p and choose a transfer τ to give to their child to finance education. Importantly, W^p may include both housing and

non-housing wealth. Specifically, the parent maximizes:

$$V^p(I^P, W^p, \theta) = \max_{\tau_I, \tau_H} \mathbb{E}[\log(I^p - \tau_I') + \omega \log(W^p - \tau_W') + \alpha V^c(\theta, \tau, I^p)]; \tag{2}$$

$$\tau = \tau_I + \tau_W. \tag{3}$$

The parent trades off utility from helping their child with their own consumption. The parent also chooses how to split the total transfer between transfers from their income and wealth. The child's enrollment decision and educational outcomes will be subject to idiosyncratic shocks that I assume parents are ignorant of before committing to a transfer schedule[18] — as a result, the parent gains utility from the expected utility of their child and their expected *realized* transfers τ_I', τ_W'. Parents are altruistic and place weight α on the utility of their child, allowing for the possibility of crowd-out from policies that expand financial aid. The parameter ω allows for parents to be differentially inclined to draw from income or assets when constructing the transfer, perhaps due to frictions in liquidating particular assets or due to income and wealth factoring in differently when determining need-based financial aid[19]. This parameter will additionally enable the model to reproduce different college enrollment elasticities following parental income and wealth shocks.

While parental transfers are endogenous in that they are sensitive to their child's future utility, I do not model the savings decisions of parents that transpire beforehand. I adopt this simplification for both computational tractability and because such behavior is not germane to my research question, given that the central experiment in the model will be to induce a shock that parents are unable to plan for ahead of time. Allowing for richer parent behavior is important if one is concerned about how parent savings behavior may respond to permanent changes to financial aid policies (Colas et al., 2021), but the unanticipated nature of the recession makes the parsimonious specification above preferable.

4.1.2 Schooling Decision

Having received the transfer from their parents, the agent chooses whether to attend college. If the agent chooses not to go to college at all, they enter the working stage with zero savings. While college is typically modeled in the literature as a binary decision, the United States features a wide array of schooling options that differ considerably in their quality and cost.

[18] This simplifies the analysis greatly through allowing these idiosyncratic shocks to be left out of the parent state space. Colas et al. (2021) adopt a similar structure.

[19] In particular, retirement plans and equity held in house of residence are not included as assets in FAFSA applications.

I allow for a broader menu of college-going options to reflect this and allow for potential substitution towards cheaper options in the face of parent resource shocks. I adopt a menu similar to the one used in Kennan (2021): agents may choose to not attend college at all or may attend a college of one of four types corresponding to community colleges, other public universities, and private colleges of two levels of quality. The agent chooses their preferred schooling option s, solving

$$V^c(\theta, \tau, I^p) = \max_s \Big[\mathbb{E}[V_t^w(\theta, e, x, a, d)|s, \tau, I^p] + \psi_0 + \psi_1\theta\mathbb{1}(s \neq 1) + \psi_2\theta\mathbb{1}(s \in HQ) + \psi_3\mathbb{E}[d|s, \theta, \tau, I^p] + \zeta_s\Big], \quad (4)$$

where ζ_s is a vector of utility shocks distributed type 1 extreme value with scale parameter σ_ζ[20]. This functional form assumption allows for straightforward attendance probability calculations and for the child's expected utility in the parent transfer stage of the model to be solved for closed-form. In addition to idiosyncratic shocks, agents face fixed psychic costs for college enrollment that depends on their enrollment decision and ability: ψ_0 captures an average psychic cost for college attendance, while ψ_1 and ψ_2 capture varying psychic costs by ability for attending college at all and attending high-quality schools in particular. The final psychic cost ψ_3 captures a potential distaste for taking on student loans that goes beyond their impacts on consumption in the post-schooling phase of the model.

The agent's continuation value V^w is formed over their ability θ, completed education e, experience x (which will always start at zero), assets a, student debt d, and age t at the point they enter the labor market. Potential educational outcomes include school/college dropout ($e = 1$), some college/associate degree ($e = 2$), or bachelor's degree ($e = 3$). The probability of an agent achieving a certain level of schooling depends on both their selected schooling option and underlying ability, with the probability of reaching a bachelor's degree being larger for higher-quality schooling options and higher-ability individuals.

Education outcomes also influence the assets agents hold and the age at which they enter the working stage of the model. If agents do not pursue schooling whatsoever or drop out of college, they enter the model at age 18 with assets and debt $a = d = 0$. If the agent obtains a bachelor's degree, they enter the working stage of the model at age 22 with starting assets

[20] Another option is to normalize $\sigma_\zeta = 1$ and allow there to be an additional coefficient in front of V^w. The distinction is immaterial in terms of the model's results, and the current setup more resembles the college choice problem's structure in Heckman et al. (1998).

or debt equal to

$$NC_s = \tau - (C_s - G(I^p, \theta, s)),$$

where NC denotes the net cost of four years of education choice s. If NC_s is positive, then starting assets $a = NC_s$ and start debt $d = 0$. If NC_s is negative, the opposite is true. C_s denotes the sticker price of schooling option s, and $G(I^p, \theta, s)$ represents grants awarded to the child, which varies over parent income, ability, and choice of school[21]. If the agent drops out of college, they start working at age 19 with assets/debt equal to $NC_s/4$ and are otherwise indistinguishable from a high-school graduate. If the agent achieves an associate's degree before dropping out, they start the model at age 20 with the degree and assets/debt equal to $NC_s/2$. Importantly, agents face an educational borrowing constraint $\bar{d}$ and may not choose any schooling option that *could* result in their exceeding this limit. Following Colas et al. (2021), I assume that if parental income I^p is below the median that the student is eligible for subsidized student loans, which do not accrue interest while the student is in college. Otherwise, I assume that the student is eligible only for unsubsidized loans, which do accumulate interest during schooling years. Discounting of V^w is done back to the beginning of the model to account for different ages of completing school.

4.1.3 Post-Graduation Labor Supply and Saving

In each period following their schooling decision, the agent chooses whether to supply labor. In addition to the state variables described above, I also track as a new state variable past labor force participation status (henceforth p) after the first working period to account for potential frictions in moving in and out of the labor market. Denoting $\Omega = \{\theta, e, x, a, p, d\}$, an agent who has made their labor supply decision $\ell \in \{0, 1\}$ and realized employment outcome (should they have participated) $m \in \{0, 1\}$ maximizes utility

$$V_t^w(\Omega; \ell, m) = \max_{c \geq 0} \left\{ \frac{c^{1-\eta}}{1-\eta} + \Lambda(\Omega, t)(1-\ell) + \beta \mathbb{E}[\tilde{V}_{t+1}^w(\Omega' | \Omega, \ell, m, c)] \right\} \tag{5}$$

subject to the budget constraint

$$c + a' = I(\Omega, m, \ell) + a(1+r) - P(d, s, t, m).$$

Consumption depends on how much the individual chooses to save for the upcoming period as

[21] For details on the calibration of the financial aid schedule and a discussion of the broader features of U.S. financial aid, refer to Appendix B.2

well as their flow income $I(\Omega, m)$ and any required student loan payment $P(d, s, t, m)$ (more on this below). If the individual does not participate in the labor market or is unemployed, the government provides transfers so as to guarantee a consumption floor $\underline{c}$. If the individual supplies labor and is employed, their flow income is $I(\Omega, 1) = w - \mathcal{T}(w)$, where w is their wage and $\mathcal{T}$ is the tax-transfer system. The wage is given by:

$$\log(w) = \gamma_0^e + \gamma_1^e \theta + \gamma_2^e x + \gamma_3^e x^2,$$

so wages depend on ability θ, work experience x, and experience, with separate values of γ for each possible education level (high school/dropout, some college/associate, bachelor's degree).

The second term of (5) captures the utility of leisure, which the agent only enjoys if they do not participate in the labor force (i.e. unemployed agents do not consume leisure). Utility from leisure is given by

$$\Lambda(\Omega, t) = \lambda_0 + \lambda_1 \theta + \lambda_2 t + \lambda_3 t \mathbb{1}(t > 50),$$

so leisure utility is allowed to vary by agent ability and age. The final term of (5) is the continuation value: next-period utility is discounted by β, and the state space Ω evolves according to the choices made in the current period. Ability θ and education e are static, while experience x increments by 1 if the agent supplies labor and is employed. Assets a evolve according to $a' = \ell \cdot I(\Omega, m) + a(1 + r) - c$, with the restrictions that individuals may not hold debt greater than the consumption floor $\underline{c}$. Student debt d' evolves according to a process detailed below. The previous-period participation status evolves according to $p' = 0$ if the agent does not participate, $p = 1$ if the agent is unemployed, and $p = 2$ if the agent is employed.

At the start of a period, the agent decides whether to supply labor, solving:

$$\tilde{V}_t^w(\Omega) = \max_{\ell} \left\{ \ell \cdot \left[\sum_{m=0}^{1} \boldsymbol{Pr}(m|\Omega) \cdot V_t^w(\Omega; 1, m) \right] + (1 - \ell) \cdot V_t^w(\Omega; 0, 0) + \xi S(\Omega, \ell) + \varepsilon_\ell \right\},$$

where ξ is a switching cost the agent must bear if they change their labor force participation status (i.e. $p \in \{1, 2\}$ and $\ell = 0$, or $p = 0$ and $\ell = 1$) after their initial working period, and ε_ℓ is a transient preference shock with mean 0 and standard deviation σ_ℓ. Finally, the

probability of employment is given by

$$\boldsymbol{Pr}(m=1|\Omega)=\Phi\Big(\pi_0+\pi_1\mathbb{1}(p=2)+\pi_2\mathbb{1}(e=2)+\pi_3\mathbb{1}(e=3)+\pi_4 x+\pi_5 x\mathbb{1}(e=2)+\pi_6 x\mathbb{1}(e=3)\Big),$$

so the probability of employment varies based on experience, educational attainment, and previous employment. Finally, the agent retires at age 65, so $\tilde{V}^w_{66}(\Omega)=0$.

Student Loan Repayment

The process of student loan repayment in my model closely follows the approximation used in Luo and Mongey (2024). If the agent is participating in the labor force and employed, they are expected to make a standard repayment

$$P^*(d,T)=\frac{d\cdot r_d}{1-(1+r_d)^{-T}},$$

where d is debt held, r_d is the interest rate for student loans, and T denotes the number of payment periods remaining for the borrower. The number of payments remaining starts at $T=10$ upon schooling completion and declines by one after each payment made. Keeping track of both debt held and payments remaining, however, is computationally intractable, so I instead approximate the number of payments remaining by comparing the agent's age t to their schooling level, which will indicate how many years have passed since they left the schooling phase of the model. If this time exceeds 10 years, they will pay off the debt as quickly as possible (i.e. as though payments remaining $T=1$), with annual payments also being bounded below by \$600 and above by \$6,000. Debt in this case evolves according to $d'=d\cdot(1+r_d)-P^*(d,T)$.

Relief is also available to distressed borrowers to ease repayments. In particular, if an individual is unemployed or not participating in the labor force, they are eligible for *Forbearance*, in which case their payment is suspended ($P(d,s,t,m)=0$), and interest accrues on the entire loan. Luo and Mongey (2024) find that the share of eligible borrowers who actually enroll in debt forgiveness programs varies very little by eligibility channels, type of loans, or amount of debt, and I follow their approach in modeling forbearance enrollment as random with probability 0.36. Debt here evolves according to $d'=d\cdot(1+r_d)$.

Finally, if an individual is not in forbearance but does not have enough resources to make the standard payment while attaining consumption above the consumption floor, they enter the state of *Delinquency*. In this state, the borrower makes as large a payment as possible,

conditional on consuming $\underline{c}$:

$$P_{del}(d, s, t, m) = (1 + r)a + I(\Omega, m, \ell) + \underline{a} - \underline{c},$$

where $\underline{a}$ is the lower bound of assets the agent can hold. Debt accrues according to interest as well as a penalty from the unpaid component of the standard repayment:

$$d' = d \cdot (1 + r_d) - P_{del}(d, s, t, m) + \rho \cdot [P^*(d, T) - P_{del}(d, s, t, m)],$$

with the penalty ρ set to 0.185 as in Luo and Mongey (2024) to approximate the effect of payment default.

4.2 Model Solution and Discussion

The model presented attempts to include several parsimoniously specified decisions surrounding early-life human capital investments and subsequent labor market outcomes that, when combined, allow for a rich set of behavioral responses to counterfactual scenarios and policies. If a recessionary shock is introduced to the model that worsens employment prospects, agents can attempt to supply more labor later in the life cycle to compensate for lost experience earlier on. At the same time, frictions to switching labor force participation status and worse employment prospects for individuals with less experience may result in economic contractions having long-lasting impacts on those who begin working during them. Parents can increase their transfers to their children to provide intergenerational insurance for a recession, though their ability to do so may also be affected if the shock impacts parental resources as well. Even if parental transfers are reduced, agents can elect to attend cheaper schooling options in an attempt to delay labor force entry and build more human capital. Thus, agents can employ a wide set of responses to various shocks in the model, and which options they choose may vary systematically over the parental income gradient.

Solutions to the agent's problems are computed via dynamic programming. The model is solved via backward induction, where policy functions and value functions are computed first for age 65, then age 64, and so on. When evaluating the objective function, I discretize the ability level into ten points and approximate continuation values via linear interpolation over the asset state. Policy functions for savings are computed for Brent's method for univariate optimization over a bounded interval. Expectations of $\tilde{V}_t^w(\Omega)$ over the ε_ℓ shock can be solved for closed-form by determining threshold levels of ε_ℓ for any state in the state space

Table 3: Parameters Estimated Externally

Parameter		Value	Source
Tax schedule	$\mathcal{T}(w)$	Various	Heathcote et al. (2017)
CRRA coefficient	η	2	Standard
Interest rate	r	0.03	Standard
Student loan interest rate	r_d	0.07	Standard
Discount rate	β	0.97	Standard
Student loan limit	$\bar{d}$	2.972	Stafford limit
Consumption floor	$\underline{c}$	0.44	2000 Census

Notes: Table presents model parameters that are estimated externally from the model. The first two columns describe the parameter and give its symbolic representation. Column 3 reports values, when possible, and Column 4 reports the source of the value. Monetary units are scaled by 10,000 2012 dollars.

that govern labor force participation and applying the inverse Mills ratio. After solving value functions for the working phase, closed form probabilities of school selections and child expected utility can be formed given a parent transfer t and child starting ability θ, allowing parent transfer policy functions to be computed as well.

5 Estimation

Estimation of the model proceeds in two steps: first, a subset of model parameters is calibrated outside of the model. The remainder of the model's parameters are estimated via the simulated method of moments. All monetary units in the model are scaled by 10,000 2012 dollars.

5.1 Parameters Calibrated Outside the Model

To begin external calibration, I use a variety of data sources to calculate targets to use in model estimation as well as to parameterize graduation probabilities conditional on college choice and ability, the costs of schooling options, and financial aid schedules. Details on this procedure can be found in Appendix B.

A summary of additional parameters calibrated outside the model may be found in Table 3. Taxes $\mathcal{T}(w)$ are calculated using the approximation in Heathcote et al. (2017), which has been shown to do well in representing U.S. tax policies. Utility over consumption is assumed to be CRRA with relative risk aversion parameter $\eta = 2$, which is roughly in the

middle of values used in the empirical human capital theory literature. I set the risk-free interest rate to be $r = 0.03$ and assume $\beta = 1/(1+r)$. The student loan interest rate r_d is set to 0.07, which is roughly the average interest rate used in the repayment period for student loans in the late 1990s and early 2000s.

The borrowing limit $\bar{d}$ is set to \$29,720, which corresponds to the nominal Stafford borrowing limit of \$23,000 in 2000, which is the same limit used in Colas et al. (2021)[22]. The consumption floor $\underline{c}$ is set to roughly \$4,500, which comes from average income from welfare, social security, supplemental security, and "other" sources among unemployed individuals in the 2000 Census who provided 0 hours of work in the previous year and is quite similar to the value used in Fan et al. (2024).

5.2 Method of Moments

After the calibration discussed above, I am left with 31 parameters to estimate via the simulated method of moments. These parameters are collected as

$$\Theta = [\omega, \alpha, \psi_0, \psi_1, \psi_2, \psi_3, \sigma_\zeta, \gamma_0^e, \gamma_1^e, \gamma_2^e, \gamma_3^e, \lambda_0, \lambda_1, \lambda_2, \lambda_3, \xi, \varepsilon_\ell, \pi_0, \pi_1, \pi_2, \pi_3, \pi_4, \pi_5, \pi_6].$$

The first two parameters $[\omega, \alpha]$ represent the utility weight parent's place on wealth as well as the altruism parameter for the utility of their child. One challenge is that because parent income and house equity are strongly correlated, the parameter ω that represents the parent's utility weight on wealth is somewhat difficult to identify credibly. I overcome this by fitting a moment reported by Lovenheim (2011), who leverages plausibly exogenous housing wealth shocks to study the impact of family wealth on college attendance. He finds that an additional \$10,000 in parent house equity increases the probability of college attendance by the child by 0.7 pp. This provides a credible target for identifying ω via indirect inference, as I can easily simulate the probability that an individual attends college in a counterfactual world where their parents have an additional \$10,000 in housing wealth and compare this probability to the baseline simulation. I additionally target rates of college attendance by parent income and wealth equity quartile to target α as well as an elasticity reported by

[22]This abstracts away from considering the possibility of private loans. Conditional on positive debt, two thirds of borrowers in the graduating class of 2008/2009 only held federal loans. Moreover, 100 percent of federal borrowers held Stafford loans, and nearly 80 percent of them held Stafford loans as their only form of federal support (Luo and Mongey, 2024)

Hilger (2016) that a decline in parent income of approximately \$8,000 decreases the college enrollment of their children by approximately 0.43 percentage points.

The set of parameters $[\psi_0, \psi_1, \psi_2, \psi_3, \sigma_\zeta]$ correspond to psychic costs associated with college attendance and debt accumulation as well as the scale of college preference shocks. These parameters are identified by targeting a rich set of college enrollment profiles from the NLSY97. For each ability tercile and parent income/wealth quartile, I compute the share of children who select each educational option (non-enrollment, public/private low/high quality) and use these moments to discipline the parameters. The psychic costs are identified by targeting overall college attendance rates as well as differences in enrollment and high/low quality selections by ability while conditioning on parent resources. I target moments concerning average student loans held at 25 and the share of individuals who borrow in the NLSY97 to discipline ψ_3, and variation in selections made *within* ability/parent resource combinations identifies σ_ζ.

The next set of parameters $[\gamma_0^e, \gamma_1^e, \gamma_2^e, \gamma_3^e]$ govern earnings over the life cycle. These are identified by targeting life-cycle wage profiles by educational status. The effects of ability on earnings γ_1^e, conditional on education, are identified by matching the early-life ratios (as opposed to levels, since the sample composition of the NLSY97 considerably from the 2000 Census, from which other income moments are drawn) of wages of individuals in the top and bottom ability terciles within educational attainment categories in the NLSY97.

The parameters $[\lambda_0, \lambda_1, \lambda_2, \lambda_3, \xi, \varepsilon_\ell]$ control labor force participation in the model. I discipline these parameters to match labor force participation rates over the life cycle and by educational attainment. Uncertainty of labor force participation within a given age and educational attainment is useful in identifying the scale of LFP preference shocks ε_ℓ. Finally, to identify the LFP switching cost ξ, I construct a panel sample of labor force participation and employment over the life cycle from the Panel Study of Income Dynamics (PSID), following a data cleaning procedure that closely follows my treatment of the Census data in Section B.1. Using the PSID sample, I compute the rate at which individuals either supply labor when they did not do so in the previous period, or vice-versa. This rate of change for labor force participation status from year to year is useful in disciplining ξ.

Finally, the parameters $[\pi_0, \pi_1, \pi_2, \pi_3, \pi_4, \pi_5, \pi_6]$ govern the probability of employment for individuals who choose to supply labor. These parameters are targeted to match employment rates over the life cycle by educational attainment. Finally, from the PSID sample described before I compute the probability of being unemployed conditional on supplying labor and either being employed or not employed in the previous period. To again account for the

PSID having a different sample composition than the 2000 Census, I take the ratio of these two probabilities to create a moment that is highly informative in identifying π_1.

I simulate the college attendance and life-cycle earnings for 50,000 individuals. When drawing parental resources for each individual, I separate parent income into ventiles and estimate the distribution of child ability conditional on parent income ventile as a truncated normal distribution in the [0, 1] interval. I then estimate the probability that the parents own a house (as opposed to rent) for each parent income ventile in the NLSY97: conditional on owning a house, I assume that parental house equity follows a log-normal distribution specific to each income ventile. Non-housing wealth is parameterized similarly and is drawn independently of house equity[23]. This allows me to draw parental homeownership and house equity realizations for the child, after which I may compute parental transfers, draw schooling utility shocks and education realizations, and simulate subsequent life-cycle earnings while approximating the savings policy function via linear interpolation.

Denoting $M = [M_1, M_2, ..., M_N]$ as the vector of empirical moments I target in the simulation procedure, denote $g(\Theta)$ as the vector of percentage errors between the data moments and the simulated moments, so:

$$g(\Theta)_i = \left(\frac{M(\Theta)_i - M_i}{M_i} \right), \quad i = 1, ..., N,$$

where $M(\Theta)_i$ is the ith moment simulated from the data given the parameter guess Θ. I find the point estimate $\hat{\Theta}$ by numerically solving

$$\hat{\Theta} = \arg\min g(\Theta)' W g(\Theta), \tag{6}$$

where $M(\Theta)$ are the simulated model moments and W is the identity matrix, with the exception that specific moments that are especially helpful for identifying specific parameters are given more weight. Minimization of the objective function proceeds via the Sbplx[24] routine. I compute standard errors by evaluating the numerical gradient of the objective function and applying the standard indirect inference formula (Gourieroux et al., 1993).

[23] While positive, the correlation between house equity and non-housing wealth is fairly weak in the data. The distributions are computed for parents with equity/non-equity wealth falling between the 25th and 75th percentile (within each income decile) to prevent wealth outliers from raising the spread of the distributions to unreasonable levels.

[24] A variant of the Subplex routine, which itself executes Nelder-Mead on a sequence of subspaces. See https://nlopt.readthedocs.io/en/latest/NLopt_Algorithm/.

Table 4: Parameters Estimated via SMM

Parameter		$\hat{\theta}$	$\hat{\sigma}$	Targeted Moment
Parent Preferences				
Wealth weight	ω	1.680	0.528	Lovenheim (2011) elasticities
Altruism	α	0.099	0.004	Hilger (2016) elasticity
College Preferences				
Fixed cost	ψ_0	-19.06	1.007	Overall enrollment rate
Ability effect	ψ_1	33.03	0.869	Enrollment profiles by ability
Ability-HQ effect	ψ_2	15.17	1.217	Enrollment profiles by ability
Debt psychic cost	ψ_3	-9.312	0.881	Debt holdings
Scale of utility shock	σ_ζ	13.89	0.611	Within-ability enrollment profiles
Earnings Determinants				
Constant, HS	γ_0^1	-0.430	0.169	Lifecycle earnings profiles
Ability effect, HS	γ_1^1	0.337	0.249	Lifecycle earnings profiles
Experience (linear), HS	γ_2^1	0.046	0.011	Lifecycle earnings profiles
Experience (quadratic), HS	γ_3^1	-0.001	4e-7	Lifecycle earnings profiles
Constant, SC	γ_0^2	-0.321	0.144	Lifecycle earnings profiles
Ability effect, SC	γ_1^2	0.266	0.224	Lifecycle earnings profiles
Experience (linear), SC	γ_2^2	0.053	0.009	Lifecycle earnings profiles
Experience (quadratic), SC	γ_3^2	-0.001	6e-7	Lifecycle earnings profiles
Constant, Coll	γ_0^3	-0.303	0.163	Lifecycle earnings profiles
Ability effect, Coll	γ_1^3	0.323	0.215	Lifecycle earnings profiles
Experience (linear), Coll	γ_2^3	0.060	0.009	Lifecycle earnings profiles
Experience (quadratic), Coll	γ_3^3	-0.001	1e-6	Lifecycle earnings profiles
Leisure Preferences				
Constant	λ_0	3.774	0.274	Overall LFP rate
Ability effect	λ_1	-1.588	0.367	LFP by education level
Age effect	λ_2	-0.086	0.009	Lifecycle LFP profiles
Age effect ($t > 50$)	λ_3	0.453	0.064	Lifecycle LFP profiles
LFP switch cost	ξ	17.55	0.243	LFP switching rate
Scale of LFP shock	σ_ℓ	10.41	0.259	Lifecycle LFP profiles
Employment Determinants				
Constant	π_0	0.983	0.123	Overall employment rate
Previously employed	π_1	0.758	0.112	Employment by previous status
Some college effect	π_2	0.339	0.148	Employment rate by education
College effect	π_3	0.564	0.154	Employment rate by education
Experience effect	π_4	0.003	0.006	Lifecycle employment profiles
Experience x Some college	π_5	-0.008	0.008	Lifecycle employment profiles
Experience x College	π_6	-0.013	0.009	Lifecycle employment profiles

Notes: Table presents model parameters estimated via indirect inference. Columns 1 and 2 describe parameters and present symbolic representation. Columns 3 and 4 present parameter estimates and standard errors, obtained via Gourieroux et al. (1993). Column 5 describes data moments used to discipline parameters. See text for details.

5.3 Goodness of Fit

The values of the parameters estimated via SMM are reported in Table 4. The weight on wealth in the parent utility function is greater than one, which will result in parents being less willing to draw from savings than flow income when providing transfers to their children. As a result, shocks to income and wealth can have different implications for college attendance propensities, particularly given that need-base grants are more sensitive to the former than the latter (especially for middle-class families; see Appendix B.2.). The preferences of the child over college attendance results in parental transfers having a large impact on child utility, especially if they move the child under borrowing constraints for particular schooling options — hence, even a low level of parental altruism can result in generous parent behavior.

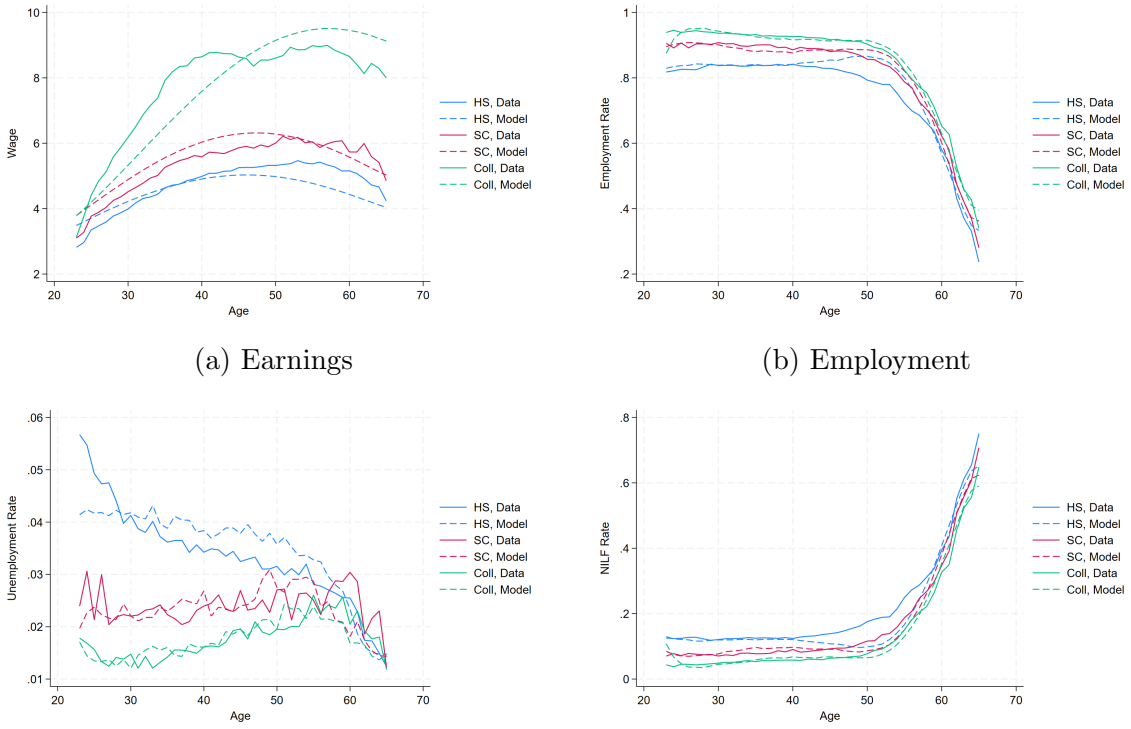
The psychic costs associated with college and the scale of the college preference shocks are large, which is a similar result observed in other entries of the literature (Heckman et al., 1998), and preferences for college attendance (especially high-quality college attendance) are strongly increasing in ability. The wage estimates feature comparable intercepts across educational levels but different returns to experience, allowing for the wages of college graduates to grow faster over the life cycle despite starting off similar to the wages of non-college graduates. Notably, the wage returns to ability are highest for college graduates and high school graduates and lowest for individuals who exactly obtain some college, indicating that the latter category may be especially likely to put individuals on relatively low-variance wage paths.

Utility for leisure is negative in ability, which is needed to match lower rates of non-participation of college graduates even with their higher wages. Switching costs and the scale of LFP shocks are both large, suggesting considerable non-pecuniary factors going into labor force participation decisions. Being previously employed and having either some college or a college degree are all associated with substantially improved employment prospects, and the returns to experience in employability suggest a role for experience that, while positive, eventually fades.

Figure 4 reports the fit of the model with regard to age profiles of earnings and employment. The model obtains a reasonable fit of wage profiles, though it understates the difference between college graduates and non-graduates in labor force non-participation. Additionally, high school graduates and those with more education have different unemployment profiles in that the former exhibits a clear decline over the life cycle, while the unemployment rates of individuals with some college or more hold fairly steady as they age.

The model's fit with regard to college attendance rates may be found in Figure 5. Across

Figure 4: Model Fit — Earnings and Labor Force Profiles



(a) Earnings

(b) Employment

(c) Unemployment

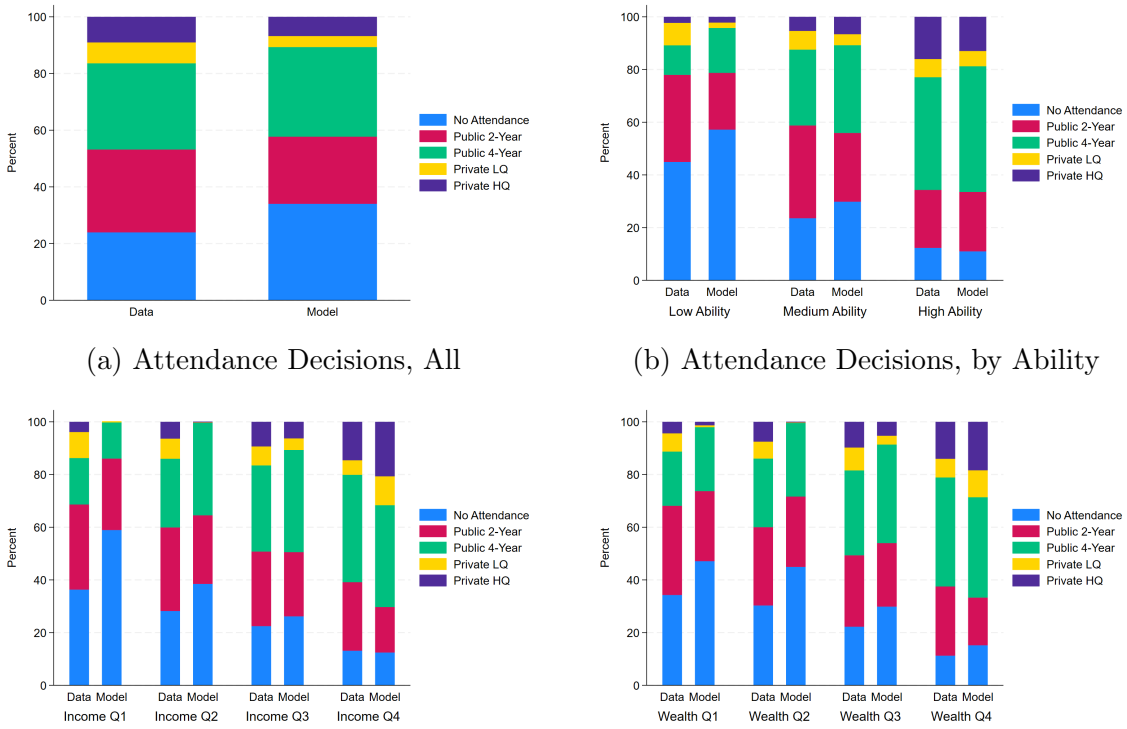
(d) Labor Force Non-Participation

Notes: Figure presents age profiles of wages and employment outcomes from data and model simulation. Moments from 2000 Census. Earnings normalized by $10,000 2012 dollars. See text for further details.

all attendance decisions, the model's fit is reasonable, though it does slightly overstate rates of non-attendance. Breaking the fit of college attendance down by ability level and parental resources reveals that the model especially overstates college non-attendance for the lowest-ability and lowest-resource children, while the model's fit of college attendance profiles for middle to upper-class children is more successful.

Finally, Table 5 reports the model's fit of other moments used in estimation. The model generates reasonable levels of labor force participation switching and the relative effect of previous employment on current employment prospects. The model can also reproduce that most labor force switches are made by individuals entering as opposed to leaving the labor market despite these statistics not being explicitly targeted in the estimation procedure. The model also produces relative wages between high and low-ability individuals, conditional on

Figure 5: Model Fit — College Attendance Decisions



(a) Attendance Decisions, All

(b) Attendance Decisions, by Ability

(c) Attendance Decisions, by Parent Income

(d) Attendance Decisions, by Parent Wealth

Notes: Figure presents shares of college attendance decisions (no enrollment, public 2-year, public 4-year, or private low and high-quality) from data and model simulation. Moments from NLSY97 geocode and IPEDS. See text for details on sample construction.

schooling, compared to the NLSY97.

The model also is successful in fitting elasticities of college enrollment in response to parental resource shocks reported in other literature entries, which I can evaluate by directly computing the change in likelihood of college enrollment for individuals in the model under counterfactual parent resource scenarios. The model generates an effect of a $10,000 increase in housing wealth on college enrollment similar to the effect reported in Lovenheim (2011). Consistent with the paper, the model also predicts that these effects are localized for relatively poor children: an additional $10,000 in home equity increases the likelihood of college attendance for children with parental income less than $70,000 by 0.94 percentage points, while the corresponding statistic for children with parental income above $70,000 is 0.11 points. The effect for poor children is quite lower than Lovenheim (2011)'s estimate

Table 5: Model Fit — Other Moments

Moment	Data	Model	Source
Labor Market Frictions			
Labor force switching rate	0.032	0.039	PSID
Switch rate, entering*	0.19	0.15	PSID
Switch rate, leaving*	0.02	0.03	PSID
$P[\text{m=1}|p=2]/P[\text{m=1}|p\neq 2]$	0.21	0.22	PSID
College Enrollment Elasticities			
Effect of \$10k↑ house equity on enrollment	0.71	0.69	Lovenheim (2011)
House equity elasticity, income<\$70k	5.67	0.90	Lovenheim (2011)
Effect of layoff on enrollment	-0.43	-0.43	Hilger (2016)
Loans held at 25, all	0.15	0.13	NLSY97
Loans held at 25, borrowers	0.59	0.36	NLSY97
Share who borrow	0.26	0.37	NLSY97
Ability Wage Effects			
High/low earnings ratio, High School	1.26	1.24	NLSY97
High/low earnings ratio, Some College	1.17	1.18	NLSY97
High/low earnings ratio, College	1.23	1.23	NLSY97

Notes: Table presents model fit of additional moments used in estimation. Moments from PSID, Lovenheim (2011), and Hilger (2016). A * indicates a moment not explicitly used when evaluating Equation 6. Switch rates refer to probabilities of changing labor force participation status. Loans held measured in \$10,000 2012 dollars. Ability effects computed as the ratio of wages between top-ability-tercile and bottom-ability-tercile individuals, conditional on educational status and age. See text for sample construction.

of 5.67 points, but the latter figure was estimated with considerable noise, and my model elasticity still falls within the confidence interval reported in the paper. I also introduce parental layoffs in the model by reducing parental income by 13.74% and can match their impact on college enrollment pooled over the ages 18 to 22, as estimated by Hilger (2016). The model also generates sensible predictions of student loans held at age 25 compared to individuals in the NLSY97, though it underestimates loans held conditional on borrowing and overstates the share of individuals who opt to take any debt at all. Because of the possibility of forbearance and delinquency, not all debt is paid off within ten years of graduation, though by age 40 nearly all agents have successfully amortized their loans.

6 Counterfactual Scenarios

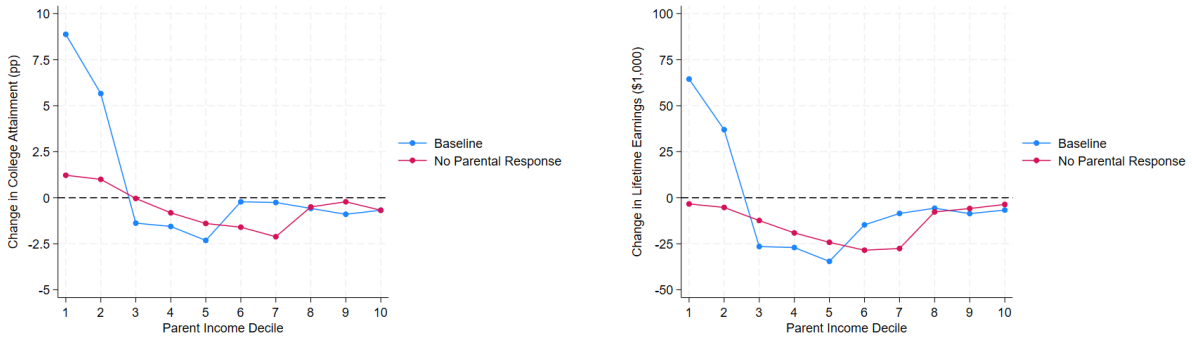
6.1 Impacts of a Recession

Having verified the model's ability to fit salient features of the data, I next turn to using the model to analyze the impacts of a recession. I simulate a contraction resembling the Great Recession in my model by introducing a five-year downturn with three main features. First, initial employment probabilities and employment transitions change for the worse. Agents may avoid some of this by enrolling in college and so forgoing the nadir of the recession, but college graduates will still face worsened prospects in the first year of their working life. Second, parental resources decline, and third, financial aid policy changes.

Modified employment probabilities come from the 2008-2013 waves of the American Community Survey. Specifically, I estimate employment probit coefficients for fresh graduates at each education level in the pre-recession (2008) wave of the ACS and then estimate the extent to which these coefficients changed for fresh graduates in subsequent years. I then modify values of π_0, π_2, and π_3 (the unemployment constant and education-specific effects) in the model by the same fraction as observed in the ACS time series for the first five years of the lifecycle. I calibrate the size of the housing wealth shock to -25.5%, which comes from population-weighted county-level real house price changes between 2007 and 2012 as measured in the Federal Housing Finance Agency House Price Index series. I also allow for potential parental layoffs by estimating, for each income ventile, the change in likelihood that adults in the NLSY79 experience an unemployment spell in the 2009-2013 period compared to the preceding four years and interpret this change as the additional layoff risk induced by the recession for parents of various income levels. I then express layoff probability as a cubic function of parent income ventile to smooth out layoff likelihoods over the parent income distribution and model layoffs as before as a 13.74% reduction in parental income following Hilger (2016).

Finally, I change financial aid policy by raising the Stafford Loan limit to $31,000 and by increasing the existing Pell grant schedule by roughly 25%, following the changes that were actually implemented in the wake of the financial crash (Barr and Turner, 2013). I also allow educational costs to rise in the wake of the recession, which I compute by observing how posted tuition, room and board costs changed across different types of universities in the 2008-2012 period, relative to trend, in the 2013 *Trends in College Pricing* report. After introducing these shocks, I re-simulate the outcomes of individuals in the new economy, allowing me to directly compare simulated individuals in the baseline model to their outcomes

Figure 6: Recession Lifetime Impacts



(a) College Attainment

(b) Lifetime Earnings

Notes: Figures present effects of counterfactual recession on educational attainment, and lifetime earnings by parent income decile. Recession simulated in model via shocks to educational policies and costs, unemployment probabilities, and parental income/wealth. Agents are assumed to have perfect foresight in baseline simulation. See text for details.

in the counterfactual world.

Figure 6 displays the results of this simulation. In the aggregate, college enrollment increases by roughly 1 percentage point in the recession, from 66 to 67 percent. This enrollment elasticity is consistent with the literature — Barr and Turner (2013) estimate that detrended college enrollment for 18-19 year old adolescents rose by a similar proportion in the 2007-2009 period[25]. However, this aggregate measure masks considerable heterogeneity over the parent income gradient: Figure 6a demonstrates a heterogeneity similar qualitatively to the results in the design-based analysis, with college attainment increasing or holding steady for children in top and bottom deciles and decreasing for children in the middle of the income distribution.

Low-income families frequently have low housing wealth or none at all, so parent housing wealth shocks have little impact on the direct financial cost a poor child has to bear in order to attend college. At the same time, low-income parents in the model often modestly increase their transfers to their children to move them over the borrowing constraint for attendance at 2-year or 4-year public institutions in response to the recession, which concurrent expansions to Stafford loan limits and the Pell grant also help facilitate. The same applies to children

[25]Moreover, the strong rise in low-income college attainment also fits with time series evidence — the number of Pell recipients nearly doubled in the same time period (Barr and Turner, 2013)

from high-income families: even with wealth and income shocks, children from rich families can obtain schooling while accruing minimal debt, and even parents in the sixth and seventh income deciles are willing to maintain transfers roughly equal to baseline levels in order to provide intergenerational insurance to their children. As a result, children in the lower-middle of the income distribution experience the worst losses in educational attainment, which translates to greater losses in lifetime earnings (Figure 6b). Fifth-decile children suffer losses in lifetime earnings of over $30,000, while children in the top and bottom parent income quartiles experience only a fraction of these losses or even gains in lifetime income due to increased college attainment.

Intensive-margin adjustments in schooling decisions (i.e. attending lower-quality schools as opposed to not enrolling at all) and college non-completion are also important factors to consider. For instance, the enrollment of children in the sixth and seventh parental income deciles slightly increases in response to the recession, but the college attainment rate for these children declines due to their substituting into lower-quality options with worse graduation probabilities, in particular low-quality private schools. Thus, focusing only on a binary college enrollment decision could substantially underestimate the impacts of the recession on completed education. Table A.4 presents percentage frequencies of schooling decisions in the recession simulation conditional on schooling choices made in the baseline state of the world. The most salient adjustments come from individuals substituting away from private schools toward public options, which is especially important when considering impacts at the upper half of the parental income distribution.

Figure A.9 conducts a further decomposition by introducing the three key ingredients of the recessionary shock (parent resource shocks, worsened labor market prospects, and expansions in financial aid generosity) individually. The key result from this exercise is that changes made to financial aid policy were the most impactful in encouraging low-income and, interestingly, upper-middle children to obtain a college degree. Expansions to the Stafford loan limit in particular made it much more feasible for low-income students to attend public four-year institutions, while the same change also made private schools more accessible to some children at higher points in the parental income distribution.

The results of this simulation compare most directly to the design-based analysis in Figure A.6a, which plots level effects of an exposure to worse unemployment shocks on age 26-28 college attainment by parent income decile. Both feature small or positive impacts on college attainment for children at the bottom and the top of the income distribution as well as negative effects in the second income quartile. In contrast to the data, the model predicts

that parents in the sixth and seventh income decile largely swallow their losses in income and wealth in order to maintain their transfers to their children, while in the data this sort of intergenerational insurance appears to come into effect around the eighth decile.

However, directly comparing the aggregate impacts into the model to the partial effects estimated in the reduced form is difficult. As mentioned previously, the missing intercept in the design-based specification means that aggregate effects may be meaningfully different than what the partial effects may imply alone. Moreover, the baseline simulation makes the strong assumption that parents and children have perfect foresight regarding the path of the recession and fully internalize changes to educational policy and parental resources when making their transfer and enrollment decisions. This may result in some departure from the reduced-form results due to the model overstating behavioral responses to the recession. As a point of comparison, I also simulate a recession with a limited parental behavioral response, where parents do not take into account changed employment probabilities or financial aid when making their transfers. When parental behavior is modeled this way, the effects of the recession on college attainment take a smoother U-shape, with intergenerational insurance only appearing meaningfully at the eighth decile as in the data. Nonetheless, I use the full-behavior model as my baseline simulation so as to be conservative when estimating the impact of the recession on lifetime earnings.

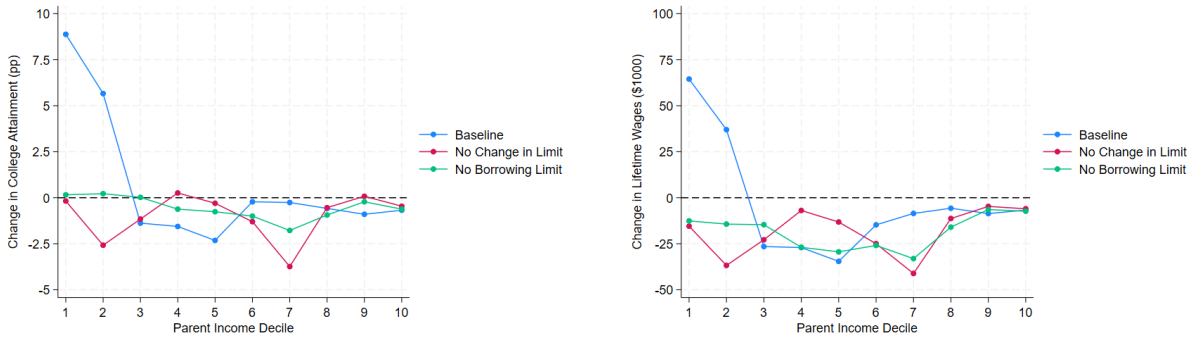
The Role of Borrowing Constraints

To assess the role of student loan limits in influencing the distributional impacts of the recession, I consider alternate scenarios where, in response to the recession, student loan limits were either not changed at all or removed entirely. In each scenario, I re-simulate the recession in the model with the modified student loan limit and re-compare agent outcomes to the baseline world.

While intuition may suggest that expanding borrowing limits should uniformly increase college enrollment and completion, the presence of endogenous parental transfers actually renders the impact of such policies ambiguous. Crucially, parents in the model make transfers to their children in part to move them over borrowing constraints for schooling options and in doing so make those options available to them. Slackening borrowing constraints decreases the contributions necessary from the parents to make college "affordable" for their children, and this decreased parental generosity means that certain schooling options may in fact become more costly (in terms of necessary loans) for students than they were in the baseline world.

Figure 7 displays the impacts of these alternate scenarios on college attainment and

Figure 7: Recession Lifetime Impacts with Alternate Loan Limits



(a) College Attainment

(b) Lifetime Earnings

Notes: Figures present effects of counterfactual recession on educational attainment and lifetime earnings by baseline parent income decile when student loan borrowing limits are not present in the model. Recession simulated in model via shocks to educational policies and costs, unemployment probabilities, and parental income/wealth. Agents are assumed to have perfect foresight. See text for details.

lifetime earnings relative to the model baseline. Modifications to borrowing limits do little to change the outcomes of children on the upper end of the income distribution, but the effects for children in the first three income quartiles are markedly heterogeneous. For instance, the model suggest that the recession's impact on the college attainment of second-quartile children would have been *lower* in magnitude had the borrowing limit not been changed at all in response. This outcome is due to decreased parental generosity in response to the borrowing limit expansion — at the same time, this increase in the borrowing loan limit benefited children in the bottom and third income quartile, as it motivated parents to provide larger transfers to their children to move them over the borrowing limits for high-quality public and private options, respectively.

On the other hand, removing borrowing limits outright during the recession results in a relative flattening of its effect on educational attainment, though widening of lifetime earnings gaps between middle and top-decile children largely remains. Taken together, these results suggest that while borrowing constraints and expansions to them were important in influencing the educational attainment of low-income children in the recession, the binding of borrowing constraints was not the primary reason behind the losses of human capital and earnings among middle-class children.

6.2 Policy Interventions

To what extent could other government interventions mitigate losses of human capital during a recession, and how might the presence of a recession affect the cost-effectiveness of financial aid policy? To probe these questions, I consider three modifications to financial aid policy and assess their impacts in and out of a recession. The first modification expands the generosity of the Pell grant but eliminates its progressive nature, providing a flat transfer of $3,750 per year in college (so, up to $15,000 total) for all individuals with parental income of $125,000 or less. The second counterfactual simply doubles financial aid awards while keeping the same progressive structure as in the baseline model. The final experiment expands the Pell grant to consider all parental assets (notably, including home equity) in award calculation. I do this by adding an additional spline function of parental assets[26] to model's baseline income-based approximation of U.S. need-based financial aid — this makes all financial aid awards weakly more generous while also having the feature that financial aid increases if an individual experiences any loss in parental assets.

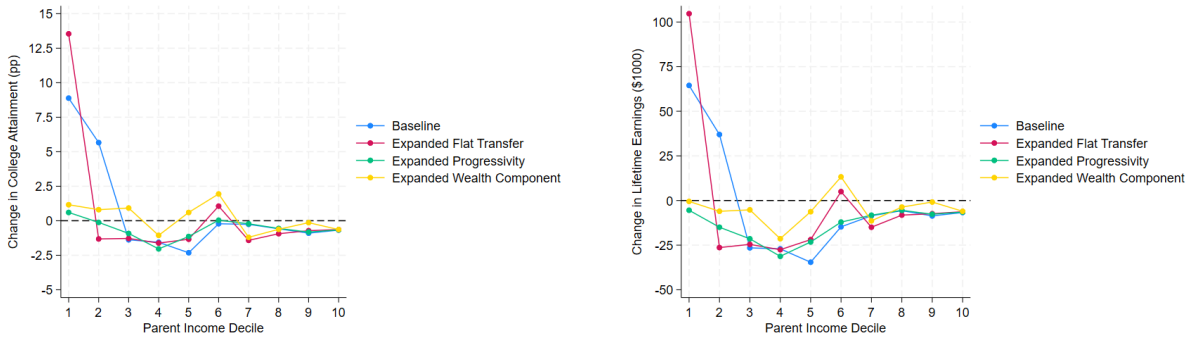
In each counterfactual, I first simulate the college decisions and earnings of the same individuals with the new policies enacted before introducing the recessionary shock to the economy to see how the policies might blunt its impact. Figure 8 presents the results of this exercise for college attainment and lifetime earnings. The recession's impacts under income-based financial aid expansions generally look similar to its impacts in the baseline world, though in some cases its impacts on low-income children are worse[27]. These expansions can mitigate the impacts of the recession on middle-decile children to an extent, but the lack of home equity considerations means that these children remain substantially exposed to the recessionary parental asset shocks. In contrast, the policy that adds an asset-contingent expansion to financial aid is more successful, by and large flattening out the U-shaped impact of the recession on human capital and in particular almost completely eliminating the widening of the gap in lifetime earnings between middle and top-decile children observed in the baseline simulation.

While counterfactual policies may be effective in mitigating the adverse effects of recessions, they may not be cost-effective for governments to implement. To evaluate this, I

[26]Specifically, I include an additional asset-based award that begins at $5,000 per year in college and declines to zero when parental assets reach $100,000, which roughly corresponds to the third quartile in the data.

[27]This happens because income-based aid expansions move certain low-income children over the borrowing constraint for four-year public universities in the baseline world, meaning that expansions to Stafford loans in the recession no longer have the same constraint-relaxing effect.

Figure 8: Impacts of Recession with Different Policy Responses



(a) College Attainment

(b) Lifetime Earnings

Notes: Figures present effects of counterfactual recession on educational attainment and lifetime earnings by baseline parent income decile under different financial aid schemes. Recession simulated in model via shocks to educational policies and costs, unemployment probabilities, and parental income/wealth. Flat transfer provides $15,000 for all individuals with parent income up to $125,000. Expanded progressivity doubles baseline Pell awards. Expanded wealth component adds an additional asset-based spline function to financial aid. Agents are assumed to have perfect foresight. See text for details.

compute the net present value of government expenditure on children in the model as

$$NPV_i = \left(\sum_{t=1}^{T} \beta^t \mathcal{T}(w_{it}) \right) - G_i,$$

where i indexes individual model agents, t iterates over the life cycle, and G_i represents government expenditure on individual i, computed as the sum of income-based or asset-based grants. For each policy, I re-simulate outcomes and calculate the average change in per-capita NPV in thousands of dollars compared to the baseline world.

These statistics are reported in Table 6, with separate entries for whether the policies are introduced in the recession version of the model and whether parents respond to the policies when making their transfer decisions[28]. Three features of the results merit discussion. The first is that regardless of macroeconomic conditions or parental responses, the more progressive policies such as the doubled Pell or the additional wealth component tend to

[28]Parents make transfers in the model solely out of altruism for their children. To the extent that parents may also make transfers for prestige effects or warm-glow motivations, the numbers presented here represent upper bounds of potential parental crowd-out effects.

Table 6: Counterfactual Cost-Benefit Analysis

Policy	Baseline, No PR	Baseline	Recession, No PR	Recession
Expanded Pell Flat Transfer	-3.66	-3.36	-0.39	-0.67
Expanded Pell Progressivity	0.99	0.85	2.63	2.27
Expanded Pell Wealth Component	-1.11	-1.35	-0.39	-0.01
Recession policy	1.87	1.69	1.73	2.85

Notes: PR = Parent response. The table reports changes in average per-capita NPV in thousands of dollars compared to baseline model. Recession simulated in model via shocks to educational policies and costs, unemployment probabilities, and parental income/wealth. Flat transfer provides $15,000 for all individuals with parent income up to $125,000. Expanded progressivity doubles baseline Pell awards. Expanded wealth component adds an additional asset-based spline function to financial aid. Agents are assumed to have perfect foresight. Recession policy includes expansions to Stafford loan limits and Pell grant generosity made in the Great recession. See text for details on simulation and NPV calculation.

be more cost-effective than the policy that eliminates the Pell grant's progressive nature, in keeping with the findings of papers such as Abbott et al. (2019) and Colas et al. (2021). Second, the role of parents in influencing the cost-effectiveness of financial aid policy is mixed. While models with a single college option find substantial parental crowd-out effects (Abbott et al., 2019), the presence of multiple college options and associated costs means that the behavior of altruistic parents and financial aid expansions can potentially complement each other in moving children over the borrowing constraint for higher-quality schooling options. As such, while parental responses decrease the cost-effectiveness of financial aid in some scenarios, they can in fact augment it in others.

The final key result from this exercise is that financial aid generally becomes more cost-effective when implemented in the recessionary state of the world, despite potential negative selection into college induced by worsened labor market conditions. The cost effectiveness of the policy that doubles the Pell grant more than doubles itself when implemented in the recession, and moreover the expanded wealth component policy nearly achieves cost parity when implemented in the recession with responsive parents. I also analyze the cost effectiveness of the actual policy combination that was introduced in response to the recession (namely, the expansion to the Pell program and Stafford loan limits) and find that, rather serendipitously, it performs the best when enacted during the downturn. These results suggest that, in the same way that an economic downturn decreases the opportunity cost of human capital investment for young adults, it may also make government support for such endeavors a better investment.

7 Discussion and Conclusion

Despite large individual literatures on the scarring effects of recessions and the role of family resources and borrowing constraints in education decisions, there have been few attempts to study their interaction. In this paper, I aim to advance this literature by studying how the Great Recession impacted the human capital of vulnerable adolescents and how these effects varied over the parental income distribution. Both a design-based exercise using a novel linked dataset from U.S. Census and a structural model point to a "hollowing out of the middle" in human capital and lifetime earnings, with middle-class children experiencing the worst losses following the downturn. Low-income children see little change in college affordability following parental resource shocks since they were typically receiving little support from their parents beforehand. In the same vein, high-income children can expect large enough transfers from their parents so that student loan limits do not bind, even in the face of an economic contraction. Children from the middle class, as a result, are the most vulnerable to borrowing constraints binding and subsequent losses in human capital following depreciations in parental resources.

In addition to adding to our understanding of the impacts of the Great Recession, these results point to potential drawbacks of current financial aid policies that exclude certain assets from award calculation. In particular, omitting home equity when determining financial aid certainly encourages middle-class college enrollment relative to a policy that penalizes such assets, but this same omission can leave certain children especially vulnerable to parental home equity shocks in an economic downturn. In my model, a simple counterfactual expansion to financial aid policy that adds a wealth-contingent stipend that considers all household assets on top of existing income-based awards can mitigate the effects of the recession on middle-class children while being nearly revenue-neutral when enacted during the downturn.

More broadly, my model also suggests that optimal financial aid policy may be considerably sensitive to the macroeconomic environment. The importance of borrowing constraints and family resources in financing higher education means that governments may have a role to play in ensuring their long-run supply of human capital when the economy experiences a contraction. Extending a more general framework of financial aid policy (Colas et al., 2021) to explore how optimal financial aid generosity and progressivity responds to the business cycle would likely be worthwhile and is left to future work.

References

Abbott, B., G. Gallipoli, C. Meghir, and G. L. Violante (2019). Education policy and intergenerational transfers in equilibrium. *Journal of Political Economy 127*(6), 2569–2624.

Bailey, M. J., C. Cole, M. Henderson, and C. Massey (2020). How well do automated linking methods perform? lessons from us historical data. *Journal of Economic Literature 58*(4), 997–1044.

Barr, A. and S. E. Turner (2013). Expanding Enrollments and Contracting State Budgets: The Effect of the Great Recession on Higher Education. *Annals of the American Academy of Political and Social Science 650*(1), 168–193.

Becker, G. S. and N. Tomes (1979, dec). An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility. *Journal of Political Economy 87*(6), 1153–1189.

Belley, P. and L. Lochner (2007). The Changing Role of Family Income and Ability in Determining Educational Achievement. *Journal of Human Capital 1*(1), 37–89.

Bound, J., M. F. Lovenheim, and S. Turner (2010). Why have college completion rates declined? an analysis of changing student preparation and collegiate resources. *American Economic Journal: Applied Economics 2*(3), 129–157.

Cameron, S. V. and C. Taber (2004). Estimation of educational borrowing constraints using returns to schooling. *Journal of Political Economy 112*(1), 132–182.

Caucutt, E. M. and L. Lochner (2020). Early and late human capital investments, borrowing constraints, and the family. *Journal of Political Economy 128*, 1065–1147.

Charles, K. K., E. Hurst, and M. J. Notowidigdo (2018). Housing booms and busts, labor market opportunities, and college attendance. *American Economic Review 108*(10), 2947–2994.

Colas, M., S. Findeisen, and D. Sachs (2021). Optimal need-based financial aid. *Journal of Political Economy*.

Cunha, F. and J. Heckman (2007). The technology of skill formation. *American Economic Review: Papers and Proceedings 97*(2), 31–47.

Danziger, S. and D. Ratner (2010). Labor market outcomes and the transition to adulthood. *Future of Children 20*(1), 133–158.

Davis, S. J. and T. V. Wachter (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity* (Fall 2011), 1–72.

Dillon, E. and J. Smith (2017). The Determinants of Mismatch Between Students and Colleges. *Journal of Labor Economics 35*(1), 45–66.

Fan, X., A. Seshadri, and C. Taber (2024). Estimation of a life-cycle model with human capital, labor supply, and retirement. *Journal of Political Economy 132*(1), 48–95.

Feigenbaum, J. J. (2016). A Machine Learning Approach to Census Record Linking. *Working Paper*.

Ferrie, J. P. (1996). A New Sample of Males Linked from the Public Use Microdata Sample of the 1850 U.S. Federal Census of Population to the 1860 U.S. Federal Census Manuscript Schedules. *Historical Methods 29*(4), 141–156.

Gourieroux, C., A. Monfort, and E. Renault (1993). Indirect inference. *Journal of Applied Econometrics 8*, S85–S118.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal tax progressivity: An analytical framework. *Quarterly Journal of Economics 132*(4), 1693–1754.

Heckman, J., L. Lochner, and C. Taber (1998). Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents. *Review of Economic Dynamics 1*, 1–58.

Hilger, N. G. (2016). Parental job loss and children's long-term outcomes: Evidence from 7 million fathers' layoffs. *American Economic Journal: Applied Economics 8*(3), 247–283.

Huggett, M., G. Ventura, and A. Yaron (2011). Sources of lifetime inequality. *American Economic Review 101*(7), 2923–2954.

Ishimaru, S. (2023). Geographic mobility of youth and spatial gaps in local college and labor market opportunities.

Johnson, M. T. (2013). Borrowing constraints, college enrollment, and delayed entry. *Journal of Labor Economics 31*(4), 669–725.

Kahn, L. B. (2010). The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics 17*(2), 303–316.

Kennan, J. (2021). Spatial variation in higher education financing and the supply of college graduates.

Kirk, J. (2023). The impact of parental resources on human capital investment and labor market outcomes: Evidence from the great recession.

Lee, S. Y. T. and A. Seshadri (2019). On the intergenerational transmission of economic status. *Journal of Political Economy 127*, 855–921.

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

Long, B. T. (2014). The Financial Crisis and Declining College Affordability. How Have Students and Their Families Responded? In *How the Financial Crisis and Great Recession Affected Higher Education*, pp. 209–233.

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

Luo, M. and S. Mongey (2024). Assets and job choice: Student debt, wages and amenities.

Mountjoy, J. (2022, 8). Community colleges and upward mobility. *American Economic Review 112*, 2580–2630.

Mulrow, E., A. Mushtaq, S. Pramanik, and A. Fontes (2011). Assessment of the U.S. Census Bureau's Person Identification Validation System. Technical report, NORC.

Oreopoulos, P., T. von Wachter, and A. Heisz (2012). The short- and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics 4*(1), 1–29.

Rinz, K. (2022). Did timing matter? life cycle differences in effects of exposure to the great recession. *Journal of Labor Economics 40*(3), 703–735.

Sallie Mae (2010). How america pays for college. Technical report.

Stuart, B. A. (2022). The long-run effects of recessions on education and income. *American Economic Journal: Applied Economics 14*(1), 42–74.

Yagan, D. (2014). Migratory Insurance Over the Great Recession. *Working Paper*.

Yagan, D. (2018). Employment Hysteresis from the Great Recession. *Journal of Political Economy 127*(5).

A Supplementary Tables and Figures

Table A.1: ACS Linkage Rates

Dataset	PIK Assignment Rate	Census Link Rate
2005 ACS	90.5	10.0
2006 ACS	90.5	16.6
2007 ACS	90.3	15.4
2008 ACS	90.2	14.4
2009 ACS	89.6	14.0
2010 ACS	93.1	10.7
2011 ACS	91.9	15.2
2012 ACS	91.6	13.9
2013 ACS	91.6	13.1
2014 ACS	91.9	12.8
2015 ACS	91.9	10.9
2016 ACS	91.5	13.3
2017 ACS	91.9	12.7
2018 ACS	93.0	12.6

Notes: Data from 2000 Census and 2005-2018 American Community Surveys. Statistics are for 1990-1994 birth cohorts. Linkage rates rounded for disclosure avoidance purposes. Column 2 reports the rates at which individuals in the relevant year of ACS data are assigned PIKs. Column 3 reports the rate at which individuals in the relevant year of ACS data are linked to themselves in the 2000 Long-Form Census. CBDRB-FY20-P2413-R8725.

Table A.2: Divisional Groupings of States

Division	States Included
New England (NE)	CT, ME, MA, NH, RI, VT
Mid-Atlantic (MA)	NJ, NY, PA
East North Central (ENC)	IL, IN, MI OH, WI
West North Central (WNC)	IA, KS, MN, MO, NE, ND, SD
South Atlantic (SA)	DE, FL, GA, MD, NC, SC, VA, DC, WV
East South Central (ESC)	AL, KY, MS, TN
West South Central (WSC)	AK, LA, OK, TX
Mountain (MO)	AZ, CO, ID, MT, NE, NM, UT, WY
Pacific (PA)	AK, CA, HI, OR, WA

Notes: This table lists which states are contained in each Census division, which are interacted with age fixed effects in the regression described in Equation 1.

Table A.3: Regression Covariate Balance Test

VARIABLES	(1) CZ Shock	(2) CZ Shock	(3) CZ Shock	(4) CZ Shock
Male		-0.004	-0.004	-0.005
		(0.005)	(0.005)	(0.006)
Black		0.057	0.059	0.144
		(0.008)	(0.008)	(0.011)
Hispanic		0.290	0.281	0.285
		(0.006)	(0.007)	(0.009)
HH Head College Attainment, 2000			-0.061	-0.099
			(0.006)	(0.007)
HH Head Married, 2000			-0.029	-0.026
			(0.006)	(0.008)
Total Parental Income, 2000			0.001	-0.004
			(0.000)	(0.000)
Parental Homeownership, 2000			-0.000	-0.137
			(0.006)	(0.008)
Parental House Value, 2000			0.001	0.008
			(0.000)	(0.000)
Constant	3.536	3.493	3.487	4.600
	(0.013)	(0.013)	(0.014)	(0.009)
Division Fixed Effects	YES	YES	YES	NO
Observations	539000	539000	539000	539000
R^2	0.418	0.424	0.425	0.014

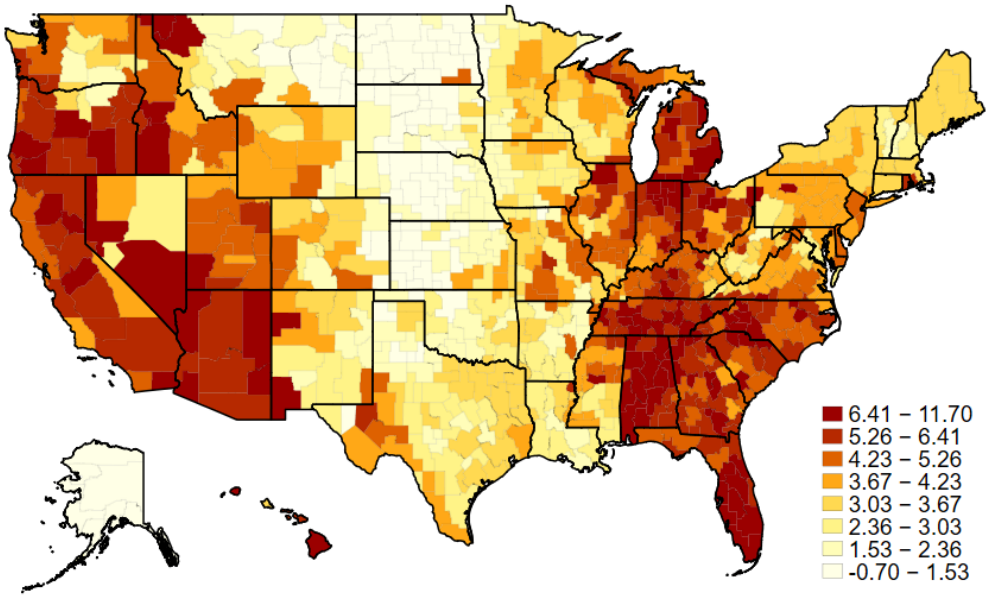
Notes: Standard errors in parentheses. HH Head = Head of household in 2000 Long-Form Census. Table presents results of regression of Commuting Zone shock individual exposed to in Great Recession on individual demographic and socioeconomic characteristics. The first column controls only for Census Division fixed effects; see Table A.2 for Division definitions. Division fixed effects omitted in final column. CBDRB-FY25-P2413-R12040

Table A.4: Recession Schooling Adjustments

	Choice in Recession				
Baseline Choice	Non-Attendance	Public 2-Year	Public 4-Year	Private LQ	Private HQ
Non-Attendance	93.72	0.75	5.07	0.25	0.22
Public 2-Year	0.68	94.68	4.04	0.29	0.32
Public 4-Year	2.16	1.73	95.30	0.27	0.54
Private LQ	1.96	3.36	4.44	90.19	0.05
Private HQ	2.68	4.37	6.67	3.77	85.66

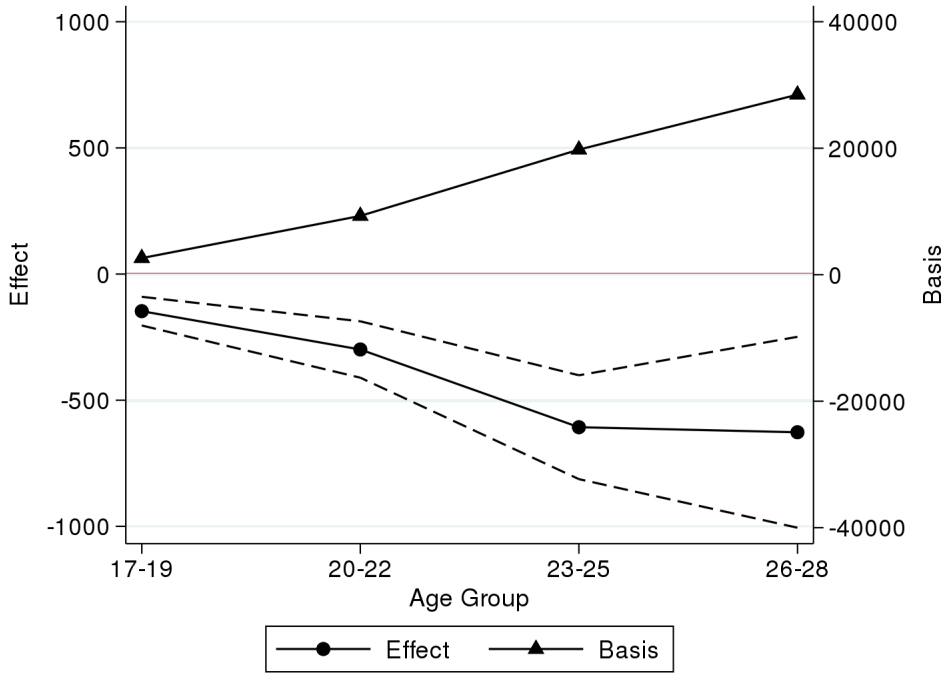
Notes: Table presents percentage likelihoods of schooling decisions in recession simulation conditional on schooling decisions in baseline simulation. Recession simulated in model via shocks to educational policies and costs, unemployment probabilities, and parental income/wealth. Agents are assumed to have perfect foresight. See text for details.

Figure A.1: Great Recession CZ Shocks



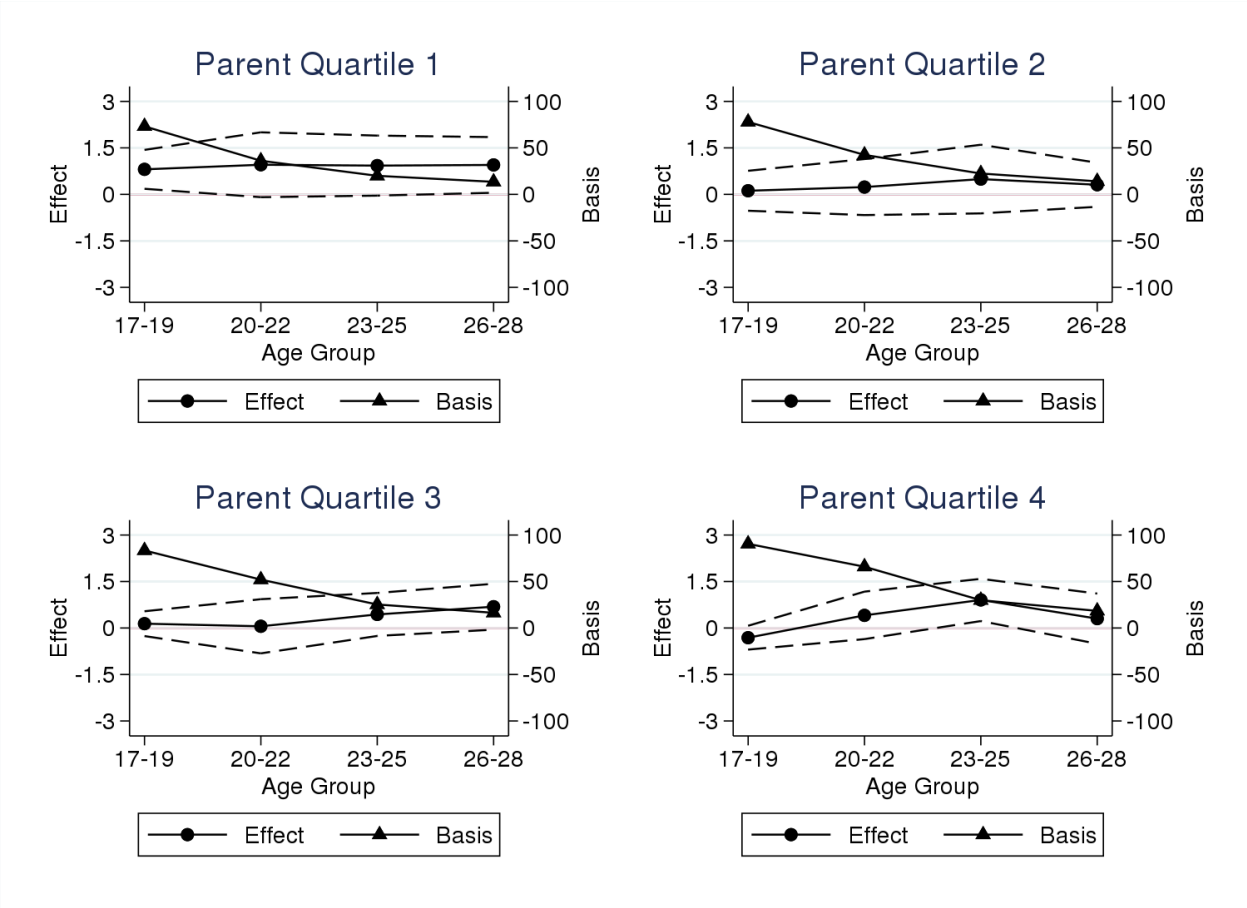
Notes: Figure shows percentage point change in unemployment rate between 2007 and 2009 at the Commuting Zone level. Data from BLS LAUS series.

Figure A.2: Full-Sample Effects of the Recession on Wages



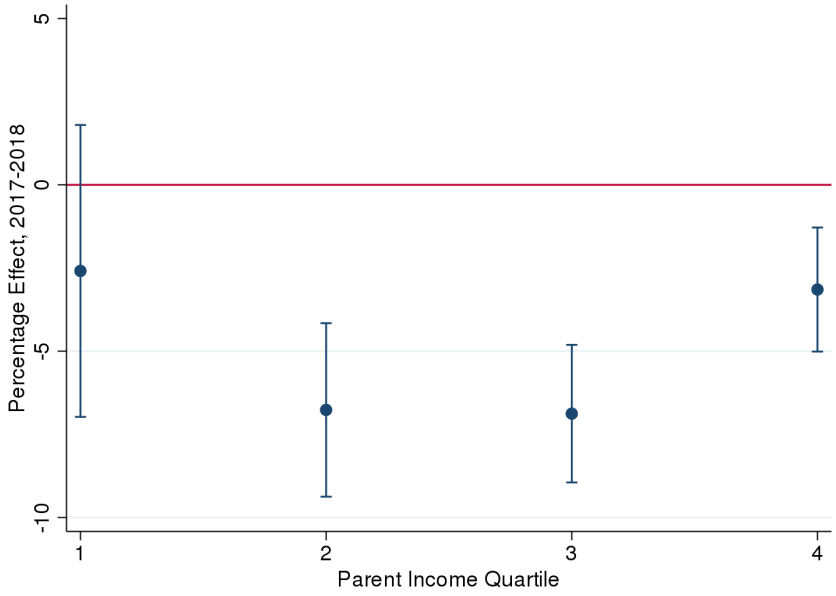
Notes: Figures display outcome variable means over age when surveyed and estimates of β_a from equation (1), with age grouped into five bins and effects for the bottom age group normalized to zero. Outcome variables is annual wage earnings in \$2012 dollars. Standard errors are clustered at the year-2000 CZ level, and the dotted lines indicate 95% confidence intervals for coefficient estimates. Coefficient estimates correspond to the circular dots and are measured by the left-hand axes; bases correspond to triangular dots and are measured by the right-hand axes. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY22-P2413-R9572.

Figure A.3: Effects of Recession Shocks on School Enrollment by Parent Income Quartile



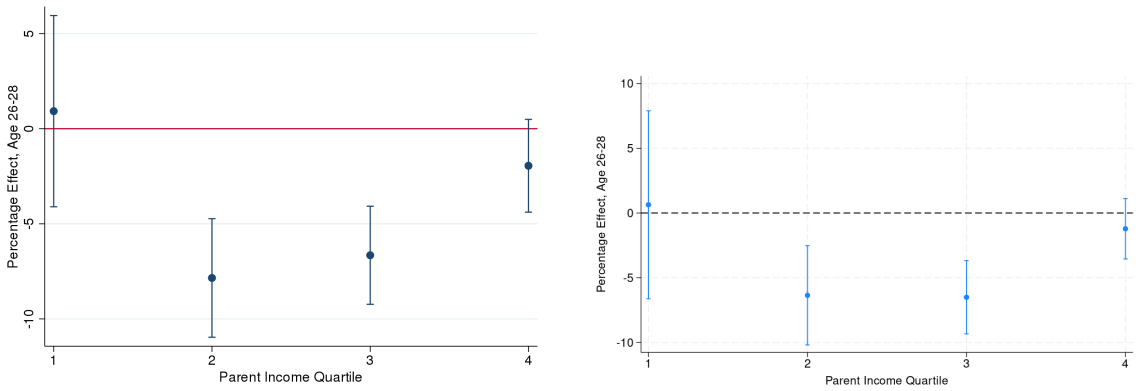
Notes: Figures display outcome variable means over age when surveyed and estimates of β_a from equation (1) for the outcome variable of being enrolled in school at time of survey, with age grouped into five bins and effects for the bottom age group normalized to zero. Effects estimated separately over year-2000 parent income quartile. Standard errors are clustered at the year-2000 CZ level, and the dotted lines indicate 95% confidence intervals for coefficient estimates. Coefficient estimates correspond to the circular dots and are measured by the left-hand axes; bases correspond to triangular dots and are measured by the right-hand axes. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY22-P2413-R9572.

Figure A.4: Percentage Effects in year 2017-2018 over Parent Income Quartile for College Attainment



Notes: Figures display estimates of $\beta_{2017-2018}$ for the outcome variable of attainment of at least a four-year college degree over parent income decile. Coefficient and standard error divided by outcome variable mean for relevant parent decile group in survey year 2017-2018 to form percentage effects. Standard errors are clustered at the year-2000 CZ level, and 95% confidence intervals are displayed. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY22-P2413-R9572.

Figure A.5: Robustness of Results to Inclusion of Control Variables

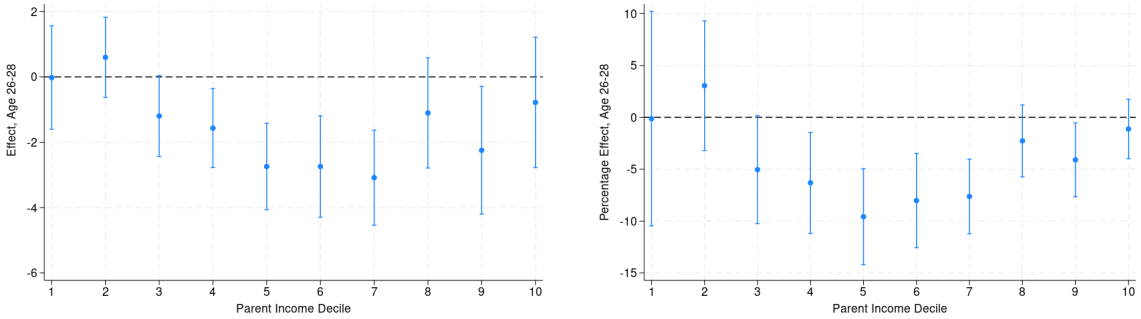


(a) Baseline

(b) With Controls

Notes: Figures display estimates of β_{26-28} for the outcome variable of attainment of at least a four-year college degree over parent income decile. Coefficient and standard error divided by outcome variable mean for relevant parent decile group in survey year 2017-2018 to form percentage effects. Standard errors are clustered at the year-2000 CZ level, and 95% confidence intervals are displayed. The figure displays main results with and without supplementary control variables described in Section 2.2. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY22-P2413-R9572 and CBDRB-FY25-P2413-R12040.

Figure A.6: Effects at Age 26-28 over Parent Income Decile for College Attainment

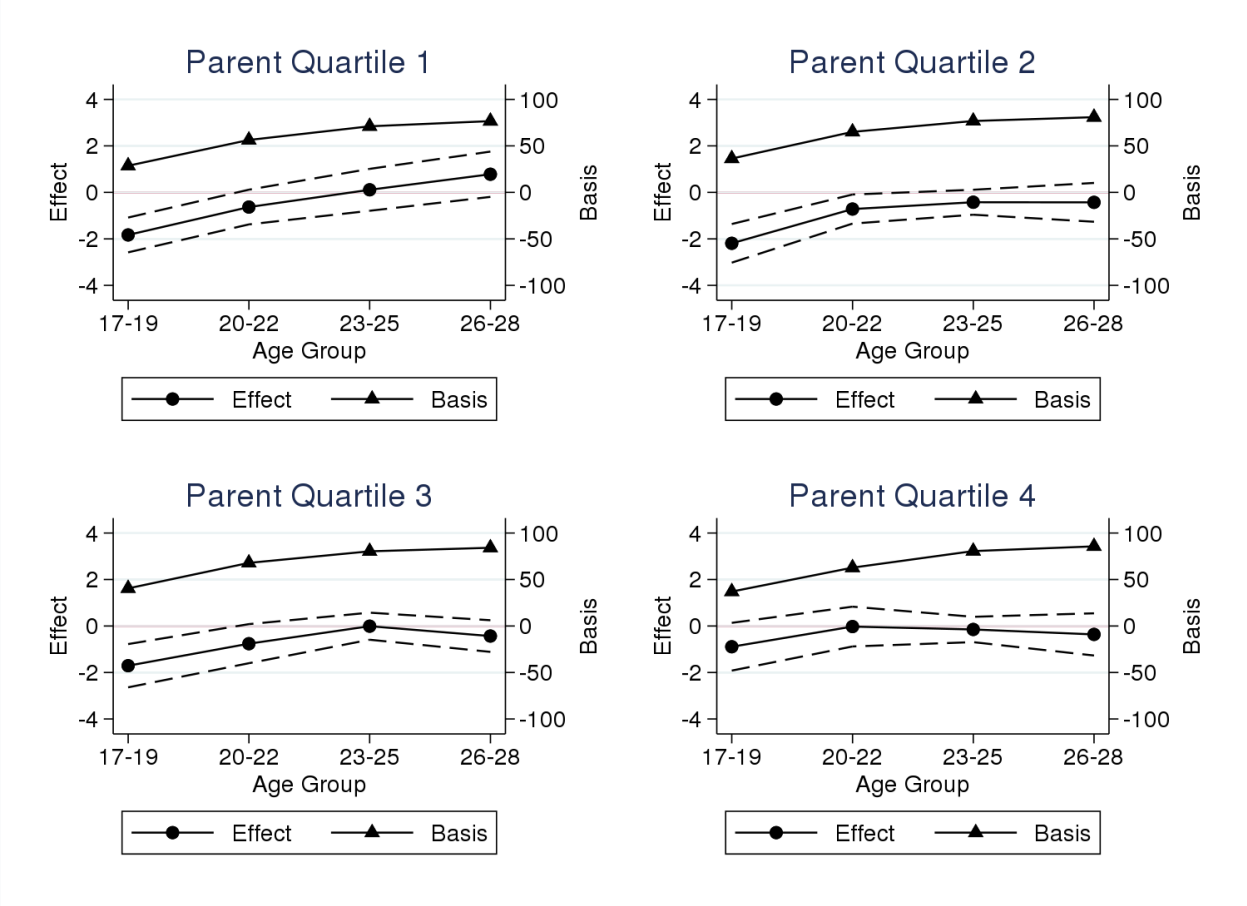


(a) Level Effects

(b) Percentage Effects

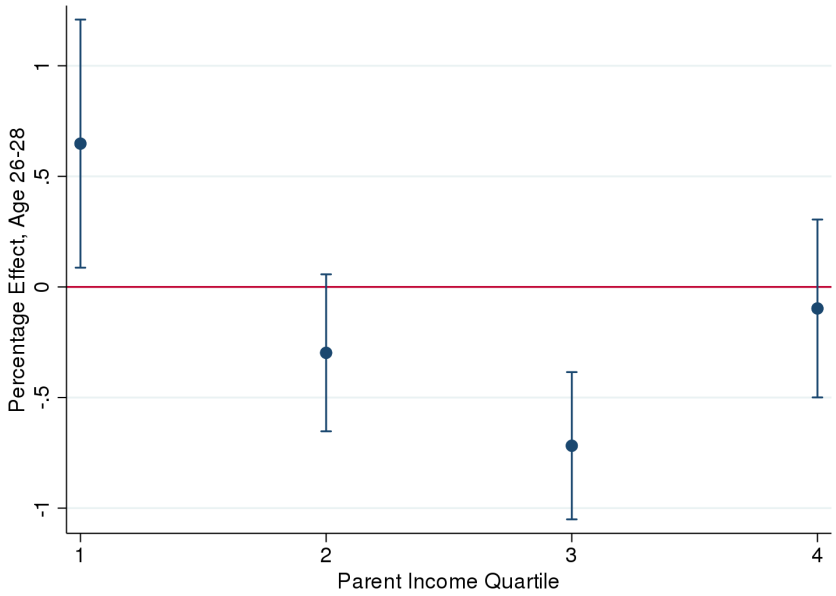
Notes: Figures display estimates of β_{26-28} from equation (1) for the outcome variable of attainment of at least a four-year college degree over parent income decile. Coefficient and standard error divided by outcome variable mean for relevant parent decile group at age 26-28 (when surveyed) to form percentage effects. Standard errors are clustered at the year-2000 CZ level, and 95% confidence intervals are displayed. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY25-P2413-R12040.

Figure A.7: Effects of Recession Shocks on Employment by Parent Income Quartile



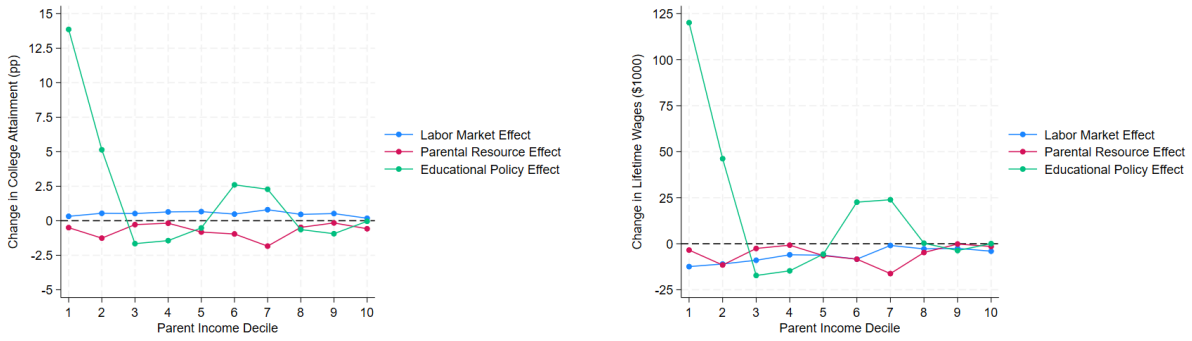
Notes: Figures display outcome variable means over age when surveyed and estimates of β_a from equation (1) for the outcome variable of being employed at time of survey, with age grouped into five bins and effects for the bottom age group normalized to zero. Effects estimated separately over year-2000 parent income quartile. Standard errors are clustered at the year-2000 CZ level, and the dotted lines indicate 95% confidence intervals for coefficient estimates. Coefficient estimates correspond to the circular dots and are measured by the left-hand axes; bases correspond to triangular dots and are measured by the right-hand axes. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY22-P2413-R9572.

Figure A.8: Percentage Effects at age 26-28 over Parent Income Quartile for College Attainment using HPI Shocks



Notes: Figures display estimates of $\beta_{26-28}$from equation (1) for the outcome variable of attainment of at least a four-year college degree over parent income decile and CZ-level 2007-2009 HPI shocks used as variation instead of unemployment shocks. Coefficient and standard error divided by outcome variable mean for relevant parent decile group at age 26-28 (when surveyed) to form percentage effects. Standard errors are clustered at the year-2000 CZ level, and 95% confidence intervals are displayed. The sample includes individuals of the 1990-1994 birth cohorts observed in the 2005-2018 ACS and linked to themselves in the 2000 Long-Form Census; see text for details. CBDRB-FY22-P2413-R9572.

Figure A.9: Decomposition of Recession Components on Human Capital and Earnings



(a) College Attainment

(b) Lifetime Earnings

Notes: Figures present effects of counterfactual recession on educational attainment and lifetime earnings by baseline parent income decile while separately introducing individual recession components. Recession components include shocks to educational policies and costs, unemployment probabilities, and parental income/wealth. Agents are assumed to have perfect foresight. See text for details.

B Estimation Appendix

B.1 Sources of Estimation Targets

The main source of data I use to estimate the model is the restricted-use NLSY97 geocode, a dataset that surveys 8,984 youths aged 12-16 as of December 31, 1996. The NLSY97 is divided into two subsamples: a nationally representative sample of 6,748 youths and an oversample of 2,236 minorities. A crucial feature of the NLSY97 is that it contains both measures of ability (an individual's ASVAB/AFTQT score) and parental income and wealth. An additional important feature of the NLSY97 geocode is that it contains information on the exact institution a respondent is studying at should they choose to pursue any post-secondary schooling.

My estimation sample uses all NLSY97 individuals who completed a high school degree by age 20 in the data and have non-missing ASVAB scores, parental income, and parental wealth[29]. Late college-going and return college-going are salient features of these data but are not modeled explicitly in my framework Kennan (2021) — to account for these, I code the first institution an individual is observed to be enrolled at up to age 25 as their choice of schooling and their educational attainment at up to age 29 as their final completed schooling. Individuals for whom final educational attainment is missing are dropped, resulting in a sample of 3,959 individuals. Longitudinal sample weights are used when computing all moments to target.

To assign ability levels to individuals in the NLSY97, I follow Dillon and Smith (2017) and use results from the ASVAB, a test designed for applicants to the U.S. military that most NLSY97 respondents took in 1997. The ASVAB has twelve separate component scores — I convert these scores to standard deviations within category and birth cohort (to account for individuals taking the test at different ages) before using the first principal component of the scores as my measure of ability. Consistent with Dillon and Smith (2017), I find that the first component explains roughly 60% of the total variance across the 12 sections and is strongly correlated with other test scores such as the ACT. I then compute percentiles of this measure and divide them by 100 to assign ability levels between 0 and 1 to individuals in the sample.

[29]Measured as the sum of parental house equity and wealth held in prepaid tuition accounts, checking accounts, bonds and stocks, vehicles, real estate, businesses, and other, minus any money owed on vehicles, child payments, or other. Wealth in retirement accounts and pensions is not included as regulation regarding taking loans out of such accounts makes using them to finance college highly difficult.

Since the NLSY97 sample is relatively small and only contains individuals born between 1980 and 1984, it does not contain adequate information about lifetime earnings profiles. To obtain lifetime profiles of earnings and labor force participation, I use data from the 2000 U.S. Census. I limit my sample to household heads aged 19 to 65 and restrict the sample in a manner similar to Huggett et al. (2011) when computing wage profiles to target in estimation. For household heads aged 30 or above, I require that individuals worked at least 520 hours in the previous year and made at least $6,000 2000 dollars in wage earnings. For household heads aged younger than 30, I require at least 260 hours worked and $4,000 2000 dollars. From these data, I then compute average earnings and the variance in log earnings by educational status and age as targets for my model to match. From the U.S. Census, I also obtain labor force participation and unemployment rates for household heads not enrolled in any higher education, again broken up by educational attainment and age, as additional targets in estimation.

B.2 Educational Setting

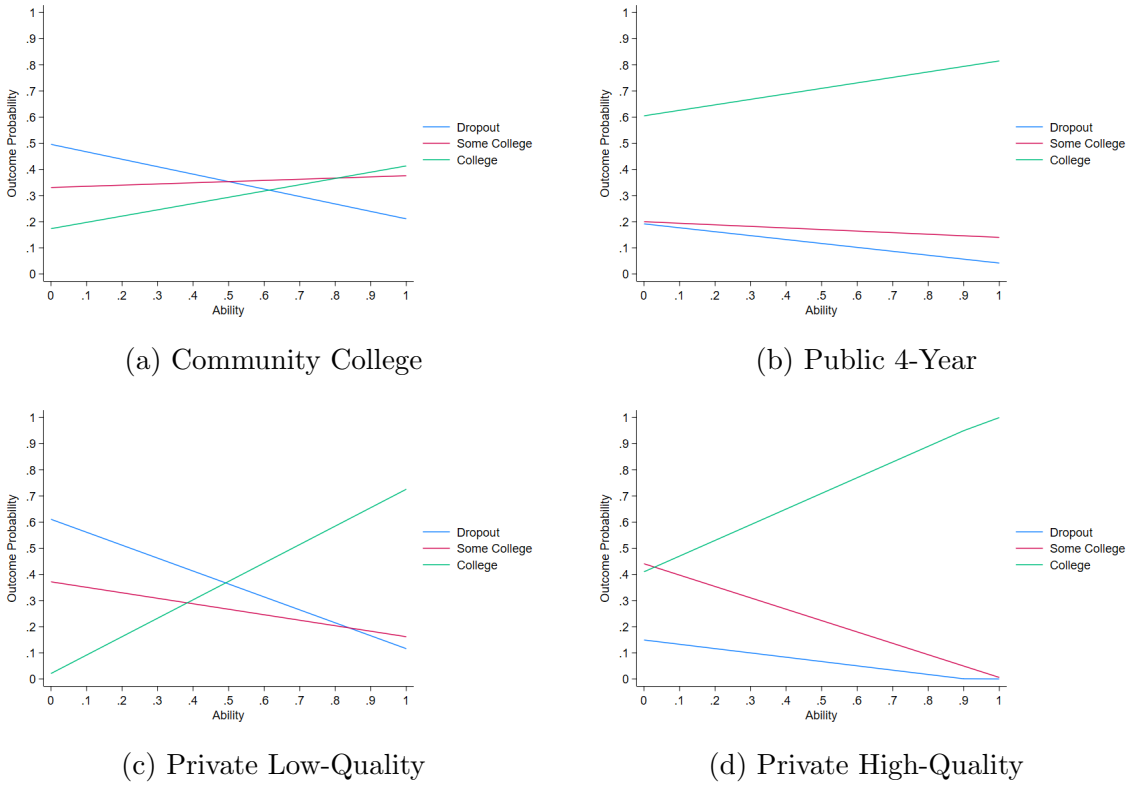
College Costs and Transitions

I use data from the Integrated Postsecondary Education Data System (IPEDS) to sort higher education institutions into the categories used in the model, which then allows me to calibrate the costs of schooling options. Additionally, since exact institutions of attendance are available in the NLSY geocode, these data can then be merged to the NLSY97 sample to compute moments of school attendance decisions by ability and parental resources as well as calibrate probabilities of educational outcomes conditional on agent ability and schooling selection.

In the model, agents may choose to attend community colleges, public four-year colleges, or private colleges of two levels of quality. College options differ in their cost and probability of graduation, which also depends on the agent's ability. To obtain costs for each option, I use data from the 2000-2004 rounds of IPEDS data. I categorize private as high quality if they are both non-for-profit 4-year institutions and either 1) have a 25th percentile SAT math score of admitted students above 380, 2) have a 25th percentile ACT math score of admitted students above 16, or 3) have an admittance rate of less than 50%[30]. All other private institutions are coded as low-quality. The full cost of each schooling option is taken as the full-time-enrollment-weighted average sticker price (including tuition, room and board, and

[30] Ishimaru (2023) categorizes institutions as selective or not using similar criteria.

Figure B.1: Probabilities of Educational Outcomes over Ability by Enrollment Decision



(a) Community College

(b) Public 4-Year

(c) Private Low-Quality

(d) Private High-Quality

Notes: Figures present probabilities of educational outcomes based on ability level and enrollment decision. Data from IPEDS and NLSY97 geocode. See text for details.

required fees) for schools in each category, multiplied by four[31]. Costs for public institutions differ substantially based on whether students are attending in-state or out-of-state; when computing the costs of community colleges and other public universities, I weight the costs according to the proportion of individuals in the NLSY97 who attend such institutions in-state vs. out-of-state. This procedure results in calibrated schooling option 4-year costs of 3.36, 6.08, 10.24 and 13.656 ($33,600, $60,800, $102,400, and $136,560 in 2012 dollars).

Next, I merge the IPEDS data to data on enrollment choices and educational outcomes in the NLSY97 geocode. I calibrate the probabilities of dropout, obtaining some college, or

[31]Given that students not enrolled in college must also pay living expenses, one may worry that this modeling choice overstates the opportunity cost of college. However, the college preference parameters will prevent the model from mechanically underestimating college-going, and it is important for the model to consider as many potential sources of student debt as possible, which certainly includes room and board.

obtaining a college degree by fitting a linear function of the outcome probability over ability terciles conditional on enrollment decision, and I impose that non-attendance results in the individual having a high school degree with certainty[32]. The transitions obtained from this process can be viewed in Figure B.1. The probabilities of favorable educational outcomes conditional on college attendance choice are increasing in ability in all cases. Notably, low-ability children have a lower chance of obtaining a college degree at a high-quality private college than a four-year public university, while the reverse is true for high-ability individuals.

Financial Aid

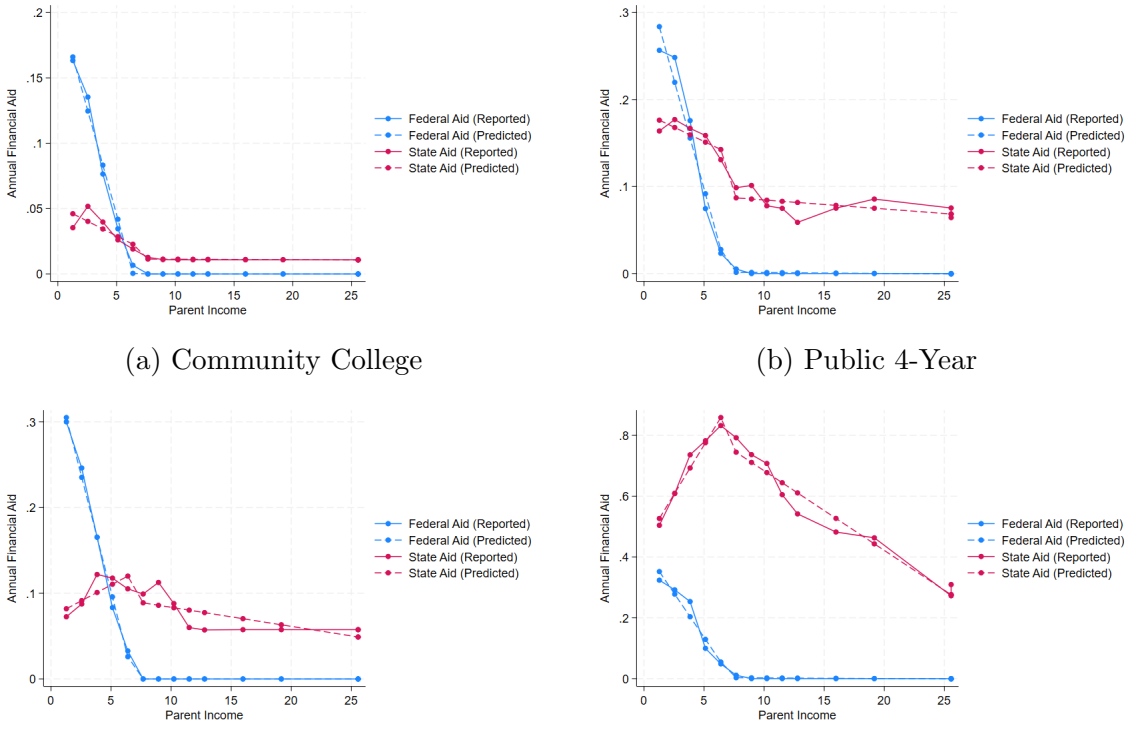
Financial aid in the model $G(I^p, \Theta, s)$ is assumed to depend on parental income, the ability of the child, and the selected schooling option. In practice, financial aid awards depend on myriad factors calculated from the Free Application for Federal Student Aid (FAFSA) form, including parental and student income and assets, enrollment levels (e.g. part-time vs. full-time), and the presence of siblings or spouses. This presents a challenge, as including the full set of these considerations in the model would be computationally intractable. Additionally, these considerations determine financial aid *eligibility*, but incomplete program take-up would result in these figures being overestimates of how much financial aid the typical student actually receives.

Approximating financial aid as a function of parental income, child ability, and schooling selection addresses the tractability concern. In particular, I abstract away from considering parental assets in calculating financial aid because, notably, the FAFSA formula does *not* include home equity or retirement accounts when calculating parental assets. Accounting for this in the parental transfer policy functions would demand that I keep track of different asset types as separate state variables, which again raises tractability concerns — moreover, the exclusions of home equity and retirement savings from FAFSA consideration justifies focusing on parental income in determining financial aid because these two assets are, by far, the most relevant for the large majority of U.S. families[33]. Among the parents of children in my NLSY97 sample, the average parents at the bottom income ventile hold less than $6,000 in non-FAFSA-excluded assets in 2012 dollars. This number rises slowly over the parent income distribution: 5th-ventile parents hold approximately $14,000 in non-excluded assets on average, compared to $28,000 for 10th-ventile parents (by which point need-based financial aid is already quite sparse) and $64,000 for 15th-ventile parents. It is only at the

[32] A very small proportion of individuals in the NLSY97 report having some college or a college degree despite never seeming to enroll at a postsecondary institution.

[33] Colas et al. (2021) also model financial aid as a function of parental income and not assets.

Figure B.2: Need-Based Financial Aid Schedules



(a) Community College

(b) Public 4-Year

(c) Private Low-Quality

(d) Private High-Quality

Notes: Figures present reported and predicted values of financial aid receipt by schooling choice over different levels of parental income. Units in $10,000 2012. Financial aid modeled as a spline function of parental income with a kink at $50,000 2000, estimated separately for school type and source of aid. Data from the 2000 National Postsecondary Student Aid Survey.

second-highest ventile that average non-excluded asset holdings exceed $100,000.

To make progress on the issue of financial aid eligibility vs. receipt, I use data on *ex-post* grant aid from the 2000 wave of the National Postsecondary Student Aid Survey (NPSAS)[34]. These data contain precise information on the amount of aid that students received from federal and non-federal (such as state or institutional) sources. Additionally, the data record background covariates such as parental income (reported in $10,000 increments and larger increments past $100,000) and high school GPA, as well as categories of colleges attended that can be mapped to the schooling options used in the model.

From these data, I obtain measures of federal and non-federal financial aid receipt over

[34] https://nces.ed.gov/surveys/npsas/

the parent income distribution for students attending the four schooling options in my model. After inspection, I parameterize need-based financial aid as a linear spline function with a kink at \$50,000 in 2012 dollars, separately for each schooling option and for federal and non-federal aid. Figure B.2 displays how this approximated financial aid function tracks with reported aid received in the NPSAS — outside of some erratic behavior for non-federal aid at low-quality private institutions, the fit of this approximation is quite strong.

Finally, merit aid is calibrated non-parametrically by directly observing average merit aid received among NPSAS respondents, broken down by schooling option and ability tercile, here approximated crudely by ranges of high school GPA ($\text{GPA} < 2, 2 \leq \text{GPA} < 3, \text{GPA} \geq 3$). Even for high-achieving students at the costliest schooling options, however, average levels of merit-based aid are marginal.