For Online Publication

Appendix 1: Matching Algorithm and Effects on Sample

I first discuss the general logic of the match of fathers and mothers to children and then document the exact routine employed. Linking parents and children in IRS data for my event study and event-age study designs requires care for several reasons. Marital status and children are only reported by filers, and filing is reduced by layoff. Therefore it is important to use information prior to layoffs to match parents and children, a rule I follow with one exception, discussed below. All matches of fathers with children rely on claims from 1996-1998, giving a buffer of two years before the first layoff can occur, in 2000. I also restrict to claims in years before a child turns 18, because after that age claims depend endogenously on child college outcomes for eligibility reasons. Over 90% of matches occur in the first available year, 1996, while virtually all the rest are made in 1997.[1] About 10% of children are only claimed by mothers and therefore excluded from my sample. An additional 25% of children are either never claimed, or claimed by too many different people for my matching algorithm to assign them a single father in all years 1996-2009 with confidence, and therefore removed from the sample. Multiple claimers are a much bigger problem than no claimers, because most low-income parents file taxes and claim children in order to collect large EITC benefits (Athreya *et al* 2010).

Removing children claimed by multiple fathers before age 18, even when a second father claims the child after a first father is laid-off, violates the rule that only pre-layoff information be used in matching children to parents. My match algorithm errs on the side of strong linkages to assure that children are linked with their primary, contemporary source of parental support. Measurement error in family linkages can be a minor problem when using parents as background controls in the context of some external treatment, because even a non-contemporaneous parent likely contains a lot of information about fixed characteristics of a child's family background. However, erroneous linkages are a major problem when measuring the effects of *changes* in parental circumstances over time on child outcomes. Changes in the time-varying characteristics of a child's multiple claiming fathers are likely far less correlated than time-invariant characteristics of these fathers. Therefore I err on the side of excluding children claimed by more than one father to assure I have strong parent-child linkages. This restriction eliminates nearly 20% of children ever claimed in IRS records. Unsurprisingly, children claimed by multiple fathers or no father have much lower college enrollment than children claimed by one father (note that college enrollment is observed for all children, both matched and unmatched).

To the extent we think income matters more for children in lower-SES families, my estimates may be smaller than estimates for the full population of children. My estimates may also be smaller if layoffs affect children through mechanisms that correlate with divorce and remarriage (the most likely path

[1] The claims data are mostly missing in 1998-2000. It is therefore reassuring that many more kids are claimed for the first time in 1996 than 1997, because this suggests that the missing data only cause a tiny fraction of missed linkages.

to multiple claimers) before a child turns 18. To explore this problem, Figure A1.1 shows two simple validations of my algorithm for matching parents and children. It plots the pre-tax income distribution for a random sample of children in my IRS data who are age 14-16 in 2001, compared with two samples of age 14-16 children drawn from the 2001 American Community Survey (ACS). My sample selection criteria cannot be validated exactly in ACS data because it relies on the time dimension of my panel data. I therefore use two ACS samples with income distributions that I expect to "bracket" that of the IRS sample. The main issue is that children in my data are in households that were headed by men at some point during 1996-1998, several years before the year of observation, 2001. These households are higher-SES than average Census households due to their male-headed status, but lower-SES than Census households headed by men in 2001. Figure A1.1.a confirms this intuition when these three income distributions are normalized into PDF's. The income distribution resulting from my linked sample looks very close to what would be expected from ACS data. Figure A1.1.b makes another point: the Average sample of children is smaller (using appropriate sampling weights) than the Census male-headed sample. This is because I exclude children who are claimed by more than one father before age 18, and children who are never claimed on tax returns.

The algorithm is as follows.

1. Make a list all unique pairs of children and claimers in every year in the sample, 1996-2009. Each individual is indexed by a unique identifier. About 95% of individuals currently in the US who were children during this sample period are linked to at least one claimer. Some of the remaining 5% may have arrived as immigrants after age 18.

2. Restrict this list to rows in which the child is under 18, because claims beginning in 18 are typically only valid conditional on college enrollment.

3. Get the sex of all claimers and the unique identifier of each claimer's spouse, if any, in every year. In each year call the claimers "PE's" for "primary earnings," and their spouses "SE's" for "secondary earners."

4. Case 1: Child has only one PE claimer (63.7% of children)

   -Restrict to SE's who claim child largest number of times.

   -If multiple SE's, break tie by selecting SE who claims child first

5. Case 2: Child has exactly one male PE claimer and one female PE claimer (11.4% of children)

   -Assign child this man and this woman as mother and father

6. Discard remaining children (20% of children)

Appendix 2: Firm Closure Sample: Details

I here briefly discuss the sample restrictions of the closure sample I construct in IRS data, which are based on those in OPS.

Firms can only enter the sample if they employ at least 30 workers at time of closure.

One restriction requires that fewer than 35% of workers experiencing closure at T at a particular firm be working at the same firm in a future year, and is intended to remove re-organizations mistakenly identified as closures. This restriction eliminates 45% of candidate firm closures.

Other restrictions require at least two years of zero UI and two years of tenure at the closing firm. The tenure restriction eliminates 35% of workers and the zero-UI restriction eliminates 15% of the remaining workers. To identify closure at T under these restrictions, we need to confirm zero employment at that firm in T+1, no excess bunching of displaced workers at the same firms in T+2, and the tenure and no-UI restrictions in T-1 and T-2. It is also important to note that many spurious closures arise if less than two years are allowed for late updates of the W2 earnings data in IRS records. Imposing all these restrictions require me to limit my sample to closures to 2001-2007. I also impose a restriction that fathers earn less than $150,000 (2009 dollars) in both of the two years prior to layoff, because there is not enough overlap in this region to adequately reweight Non-Closures to Closures.

The resulting closures sample is a 100% sample of workers displaced by closure who take up UI, combined with a 30% random sample of workers displaced by closure who do not take up UI, with appropriate sampling weights.

Table A2.1 shows number of firms that close and their average size, by year of closure, in my sample. While some fairly large firms do close every year, the vast majority of closing firms are small. This leads to the small average size of closing firms.

Table A2.2 displays summary statistics for the Closure and Non-Closure samples, and is analogous to Table 1. The Closure and Non-Closure samples display similar overall patterns, though smaller declines in child earnings because many fewer cohorts are included in this sample due to the computational demands of identifying closures according to the above restrictions.

Appendix 3: Identification Using Only Layoffs

In this appendix I develop an estimator that relies entirely on the Layoff sample.

The challenge is to estimate period effects when event-year and cohort both have effects on outcomes that are large relative to period effects, given that these three variables are linearly dependent[2]. This linear dependence makes it infeasible to estimate period effects while controlling for cohort and event-year

[2]The problem occurs when linearly dependent covariates enter a conditional expectation function, and the researcher is primarily interested in effects of a subset of these linearly dependent variabless. Leading examples of this are age, year, and cohort effects in labor economics (Hall, Mairesse and Turner 2005) and age, year, and vintage effects in studies of capital goods (Hall 1971, Berndt, Griliches and Rappaport 1995).

fixed effects. I first discuss the approach I take intuitively, and then formalize it using the notation developed above.

Estimating treatment effects requires estimation of potential outcomes under non-layoff for the Layoff children after layoff takes place. Above, I use Survivors for this. Here, I use Layoff children prior to realization of layoffs. This is another form of DD estimator. The first difference is the same: the difference between two moments in the Layoff sample on either side of period $k = 0$. Above, the second difference is between the two corresponding moments in the Survivor sample. Here, the second difference is between two moments in the layoff sample, both of which involve $k < 0$. There are typically a number of candidate pre-layoff differences that can be used to estimate the desired potential outcomes. The approach I develop here pinpoints a particular weighted combination of these differences that addresses the problem of confounding event-year and cohort shocks.

Figure A3.1.a displays average enrollment for the Layoff group at age 19 for three cohorts, each plotted by event-age $a_E \equiv a - (t_O - t_E)$. One option would be to pool all of these cohorts into event-age means, but this throws away a lot of useful information. The key information to exploit is that event-age is collinear with event-year for a fixed cohort. Consider the difference $A - B$ for the 1984 cohort. This is a particular "treatment difference," which can be defined as the outcome at an event-year that takes place after the age of the outcome, minus the outcome at an event-year that takes place before the age of the outcome. For cohort 1984, the points $A - B$ reflect outcomes for children with event-year 2002 minus outcomes for children with event-year 2004. The difference $A - B$ therefore contains both the desired difference in outcomes across children with different event-ages, and a confounding difference across children with event-years 2002 and 2004. We would like to estimate the confounding difference across children with these event-years.

There are many ways to estimate this confounding difference. Any cohort for which both of these event-years occur too late to affect the outcome at age 19 provides an estimate of this event-year effect difference. Define a "control difference" as a difference across outcomes for event-years that occur too late to affect these outcomes. Figure A3.1.a presents two such control differences. The control difference $C - D$ uses the 1986 cohort to estimate the difference across event-years 2002 and 2004. The control difference $E - F$ uses the 1988 cohort to estimate the difference across event-years 2002 and 2004. There are many such control differences. Each control difference yields a different double-difference estimator of the treatment effect. One estimator is $A - B - (C - D)$. Another estimator is $A - B - (E - F)$. The unweighted mean of these double-differences provides one estimate of the difference in event-age effects across event-ages 18 and 20, which is the short-run treatment effect of interest.

This estimate relies on the treatment difference in cohort 1984. There are many other cohorts, each of which offers one treatment difference across event-ages 18 and 20. Each cohort uses a different pair of event-years in its treatment difference, and therefore requires a different set of control differences to remove the confounding event-year variation. Each cohort then yields a different "treatment effect," defined as a treatment difference minus the mean of all available control differences that share the event-years used in the treatment difference. I then take the mean of all these treatment effects.

A similar argument holds for treatment differences that occur within event-years, rather than within cohorts. The analogous graph is presented in Figure A3.1.b. I omit the discussion of these estimators to save space; it is conceptually analogous to that just presented. The surprising fact, however, is that the treatment effects that emerge from these two approaches contain independent information. The amount of independent information decreases in the smaller dimension of the outcome-year by event-year matrix. I therefore calculate the complete set of treatment effects and pool them into a single estimate. This approach does count some information multiple times, and therefore overestimates the precision of the final estimate. I ignore this problem.

Write birth cohort as $t_B \equiv t_O - a$. I now rewrite the model in terms of cohort instead of outcome-year. Write the potential child outcome function in terms of a treatment effect, and age interacted with cohort effects and event-year effects, dropping the g subscript because now all observations set $g = T$:

$$y_{a,t_B,t_E} = \alpha + \sum_{j=k_{\min}}^{k_{\max}} \beta_j \cdot I\left\{k = j\right\} + \sum_{j=t_B^{\min}}^{t_B^{\max}} \theta_{a,j} \cdot I\left\{t_B = j\right\} + \sum_{j=t_E^{\min}}^{t_E^{\max}} \psi_{a,j} \cdot I\left\{t_E = j\right\} + u_{a,t_B,t_E},$$

where β_j is the effect of layoff on the outcome, $\theta_{a,j}$ is the effect of cohort j on the outcome at age a, $\psi_{a,j}$ is the effect of selection into layoff in year j on outcome at age a, and u_{a,t_B,t_E} is an error term. A key restriction here is that cohort effects are constant across event-years, and event-year effects are constant across cohorts, within age groups. Without Survivors we have no way to distinguish these interaction terms from event-age effects. This is why the estimates using only Layoffs are much noisier: event-ages are capturing both treatment effects and random cohort by event-year by age interaction shocks.

It is not possible to identify all of the parameters in this model without further assumptions, due to the collinearity $k = t_O - t_E = a + t_B - t_E$. I therefore make an additional assumption that period effects not driven by treatment effects are linear:

$$A1 : \beta_j = \phi_{a,0} + \phi_{a,1} k + \sum_{j=0}^{k_{\max}} \lambda_j \cdot I\left\{k = j\right\},$$

where $\phi_{a,0}$ and $\phi_{a,1}$ capture the linear trend in period, and λ_k captures treatment effects, assumed to equal zero for outcomes prior to events. This is not a strong additional assumption; it is weaker than the parallel trends assumption $\phi_{a,1} = 0$ used for the main results.

Under these assumptions we identify many different treatment effects using DDs, as previously described. I here characterize the set of these treatment effects. All such DD's consist of one treatment difference that crosses the cutoff where $k = 0$ (e.g.,$a_W = a$, depicted as difference $A - B$), and one control difference that is contained entirely in the untreated region where $k < 0$ (e.g., $a_W > a$, depicted as differences $C - D$ and $E - F$).

Writing event-age in terms of event-year and cohort and fixing age a for simplicity, this set of DD's identifying treatment effects $k = a + t_B^{i_2} - t_E^{i_2}$ years after layoff can be characterized as:

$$\Gamma\left(t_E^{i_1}, t_E^{i_2}, t_E^{i_3}, t_E^{i_4}, t_B^{j_1}, t_B^{j_2}, t_B^{j_3}, t_B^{j_4}\right) \equiv E\left[y_a\left(t_E^{i_2}, t_B^{i_2}\right) - y_a\left(t_E^{i_1}, t_B^{i_1}\right)\right] - E\left[y_a\left(t_E^{i_4}, t_B^{i_4}\right) - y_a\left(t_E^{i_3}, t_B^{i_3}\right)\right] \text{ (DD's in event-age)}$$

such that

1. $t_E^{i_1} - t_B^{i_1} + 1 \quad < \quad a + 1 \leq t_E^{i_2} - t_B^{i_2} \leq t_E^{i_3} - t_B^{i_3} < t_E^{i_4} - t_B^{i_4}$ (1)
(one treatment difference minus one control difference) (2)

and *either* of the following hold:

2A. $t_B^{i_1} \quad = \quad t_B^{i_2}$, $t_B^{i_3} = t_B^{i_4}$, $t_E^{i_1} = t_E^{i_3}$, and $t_E^{i_2} = t_E^{i_4}$
(treatment differences removes cohort, control difference removes event-year)

2B. $t_E^{i_1} \quad = \quad t_E^{i_2}$, $t_E^{i_3} = t_E^{i_4}$, $t_B^{i_1} = t_B^{i_3}$, and $t_B^{i_2} = t_B^{i_4}$
(treatment differences removes event-year, control difference removes cohort).

To avoid clutter we can re-write $\Gamma_a\left(t_E^{i_1}, t_E^{i_2}, t_E^{i_3}, t_E^{i_4}, t_B^{j_1}, t_B^{j_2}, t_B^{j_3}, t_B^{j_4}\right)$ as $\Gamma_a\left(a_E^{i_1}, a_E^{i_2}, a_E^{i_3}, a_E^{i_4}\right)$, where the assumptions embodied in $\Gamma_a(.)$ are implicit. As stated above for the example, under assumption $A1$ selection terms cancel out in DD's of this nature, and we have:

$$\Gamma_a\left(a_E^{i_1}, a_E^{i_2}, a_E^{i_3}, a_E^{i_4}\right) = \lambda_k. \tag{3}$$

Table A3.1 calculates treatment effects as the unweighted average all these DD's, and is analogous to Table 3. Column 1 shows the mean treatment effect for $\pi_a(a-1, a+1)$ by age a for ages $18 - 22$, as well as a total effect that combines all ages in this range. The estimated effects are very similar to those estimated with Survivors as a control group, but much noisier due to the lack of any way to eliminate event-year by cohort interaction terms.

Appendix 4: Estimation with Linear Differential Trends

In this appendix I derive formulas for point estimates and standard errors on a treatment effect estimator that allows for linear differential selection in outcomes with respect to period k for $k < 0$. This is a weaker version of the parallel-trends assumption.

The key parameters are the β_k^T and β_k^C terms; their difference captures the difference in child outcomes around period of father's layoff. These terms are estimated using OLS on Equation (2). I here employ a small amount of new notation for convenience. Write a conditional expectation function for a scalar child outcome Y as

$$E\left[Y|X\right] = \beta X,$$

where Y is the child's outcome, β is a K by 1 vector of parameters, and $\mathbf{X}$ contains the covariates, including the period terms interacted with type of event (layoff or survival) and controls for event-year and cohort. Let $\mathbf{V}_\beta = Var\left(\hat{\beta}\right)$.

First define a 7 by K matrix $\mathbf{L_T}$ such that $\mathbf{L_T}\beta = \left(\beta^{\mathbf{T}}_{-\mathbf{7}}, ..., \beta^{\mathbf{T}}_{-\mathbf{1}}\right)$, and similarly define $\mathbf{L_C}$ such that $\mathbf{L_C}\beta = \left(\beta^{\mathbf{C}}_{-\mathbf{7}}, ..., \beta^{\mathbf{C}}_{-\mathbf{1}}\right)$, where I have here imposed a cutoff of seven years prior to layoff. Now let $\mathbf{L} = \mathbf{L_1} - \mathbf{L_0}$, such that $\mathbf{L}\beta = \left(\beta^{\mathbf{T}}_{-\mathbf{7}} - \beta^{\mathbf{C}}_{-\mathbf{7}}, ..., \beta^{\mathbf{T}}_{-\mathbf{1}} - \beta^{\mathbf{C}}_{-\mathbf{1}}\right)$. This vector contains the points in the pretreatment region of the graph. We want to estimate a line through these points, i.e., we want to regress these points on a constant and on a linear period trend, where period goes from -7 to -1. Therefore define a covariate matrix

$$\mathbf{Z} = \begin{pmatrix} 1 & -7 \\ 1 & -6 \\ ... & ... \\ 1 & -1 \end{pmatrix},$$

and define the parameter vector γ as the least-squares approximation

$$\gamma \equiv \arg\min_a \left(\mathbf{L}\beta - \mathbf{Za}\right)' \left(\mathbf{L}\beta - \mathbf{Za}\right) = \left(\mathbf{Z'Z}\right)^{-1} \mathbf{Z'L}\beta.$$

We can write the estimator of γ as $\left(\mathbf{Z'Z}\right)^{-1} \mathbf{Z'L}\hat{\beta} \equiv \mathbf{\Omega}\hat{\beta}$, and the covariance matrix for $\hat{\gamma}$ as $\mathbf{\Omega V}_\beta\mathbf{\Omega}'$.

The target parameter is the estimated difference between the imputed counterfactual outcome under survival and the realized outcome under layoff in period $k > 0$. Define this scalar parameter as $\theta \equiv (\gamma_0 + k \cdot \gamma_1) - \left(\beta_k^T - \beta_k^C\right)$, where we here focus on the case of $k = 1$, the year after layoff. This can be rewritten by defining two matrices $\mathbf{H_0} = (1, k)$ and $\mathbf{H_1}$ such that $\theta = \mathbf{H_0}\gamma - \mathbf{H_1}\beta$, or

$$\theta = \left(\mathbf{H_0\Omega} - \mathbf{H_1}\right)\beta.$$

This neatly writes the target parameter as a linear combination of the original regression of outcomes on period dummies for each group and other controls. We can therefore write the variance of $\hat{\theta}$ as

$$V_\theta = \left(H_0\mathbf{\Omega} - H_1\right)\mathbf{V}_\beta \left(H_0\mathbf{\Omega} - H_1\right)',$$

yielding a standard error on $\hat{\theta}$ as $V_{\theta}^{1/2}$.

Appendix 5: Institutional Non-Filing of the 1098T Form

My results rely on data contained in 1098T forms filed by all Title IV post-secondary institutions. Title IV institutions contain most four-year, two-year, and professional schools in the U.S.. However, recent work by Cellini and Goldin (2012) suggests that 27% of college students are not enrolled in Title IV institutions, and will therefore not receive 1098Ts. In addition, schools are only required to file 1098T forms for individual students who pay any positive dollar amount for tuition, room, board, or other fees, net of financial aid received from the school or other sources.

My analysis data set defines college enrollment as a non-missing 1098T form, and non-enrollment as a missing 1098T form. I have interpreted reductions in 1098T filing for children with recent paternal layoffs as reductions in college enrollment. However, a decline in non-missing 1098T forms could be generated by increased enrollment in non-Title-IV institutions, or by an increase in financial aid for students that pushes their net payments to zero at a school that does not report 1098Ts when not legally required to do so. Therefore, it is possible in theory that my key enrollment decline results represent a switch from Title IV to non-Title-IV institutions, or an increase in 1098T non-filing. The switch into non-Title-IV institutions seems unlikely because Cellini and Goldin (2012) estimate that these schools are approximately equal in price to Title IV schools, after accounting for subsidies (mainly Pell Grants) at Title IV schools. Therefore the main worry is that layoffs increase 1098T non-filing rather than decrease real college enrollment. While the evidence I present in this appendix suggests 1098T-nonfiling is not driving my results, such a problem would anyway work in favor of my conclusion that layoffs and their associated income declines have at most very small effects on child college outcomes.

To address this concern I first construct an alternative measure of college enrollment in IRS data. I then provide two tests suggesting that the problem is unlikely to drive my main results.

The alternative measure of college enrollment is based on the claiming of children age 19-24. Parents are allowed to claim children ages 19-24 if and only if the child is "permanently and totally disabled" or enrolled full-time at a school[3]. The key features of this rule, for our purposes, are that children can qualify as students if the family pays zero net tuition, and if the school is not a Title IV school. Therefore, conditional on parents filing a tax return, the fraction of parents that claim a child age 19-24 represents a potential alternative measure of college enrollment that can validate findings with 1098T-based enrollment.

Figure A5.1 plots the two measures of college enrollment for children at age 19 in 2002 by mean three-year family income. I restrict to children in 2002 to facilitate comparison with statistics in the NLSY97 for cohorts 1979-1982 in Bailey and Dynarski (2011, Figure 2). I restrict to age 19 because after age

[3] See instructions for the 1040 online at <http://www.irs.gov/instructions/i1040a/ar01.html>.

19 children gradually start to claim themselves as dependents. While parent claims continue to track 1098T-based enrollment for higher-income families, the mechanical decrease in levels makes them less useful as a gauge of total enrollment.

The first pattern is that, for richer children, college enrollment based on claims roughly tracks college enrollment based on 1098Ts. About 20% more high-income children are enrolled in college using the claims measure. This is slightly less than the fraction of students estimated to be enrolled in non-Title-IV colleges in Cellini and Goldin (2012). The smaller figure here could be due to the small fraction of 19-year-olds who claim themselves and hence cannot be claimed by their parents.

The second pattern is that, for poorer children, college enrollment based on claims is much too high to represent actual college enrollment. While 1098T-based enrollment matches enrollment in the NLSY for children in the bottom quartile of family income, the claims-based measure of college enrollment is about twice this level. The most likely explanation for these implausibly-high claiming rates is that low-income families claim children after age 18 in order to claim the EITC[4]. This seems especially likely since the only reason most low-income families file a 1040 at all is to claim EITC benefits[5]. Therefore the claims-based measure of college enrollment cannot help us validate the 1098T-based measure of college enrollment for low-income children: false claims overwhelm genuine non-1098T college enrollment with zero tuition payments or at non-Title IV schools.

I would like to see if the main results in the paper hold up when using this alternative enrollment measure. For this to make sense, I restrict to richer families for two reasons. The first reason, as just discussed, is that this measure appears to approximate enrollment only for richer families. The second reason is that layoffs reduce filing rates, because filing correlates positively with income. This introduces a potential spurious effect by reducing claims by reducing filings. However, note that families who do not file are unlikely to have children in college, because such families can claim EITC benefits. Nonetheless, restricting to higher-income families alleviates most of the filing problem. I use the same definition of "high-income" used in the main results in the text, three-year mean incomes 1996-1999 above $40,000.

Figure A5.2.a plots treatment effects from Equation (2) where the outcome is a dummy for whether the child is claimed by her parents, instead of whether the child receives a 1098T as in Figure 13. The claim-based enrollment measure shows almost exactly the same pattern of treatment effects as the 1098T-based measure for high-income families: about one percentage point lower enrollment for children experiencing a father's job loss one year before college decisions

[4]Note that by using three-year mean family income on the x-axis, families with incomes in the EITC benefit range for at least one of these three years are likely to have much higher three-year mean income, due to mean-reversion, explaining why the excess claims extend "too high" relative to the EITC benefit range, which reaches zero around $30,000.

[5]Only 60% of the poorest parents in this graph file, where as almost 100% of parents file starting at around $50,000

versus one year after college decisions. I conclude that this test supports the main results of the paper: layoffs reduce college enrollment by a small amount.

I also develop a second robustness check on the main 1098T enrollment results. I create a dummy variable that equals one when a child receives a 1098T from a college that often files 1098Ts for students with zero net-tuition payments[6]. I then estimate Equation (2) again for this outcome variable, now for all children, not just high-income children. The resulting event-age effects are much noisier due to the much lower rate of enrollment at this restricted set of institutions, but the overall pattern is reassuring: children experiencing paternal layoffs one year before making enrollment decisions enroll in these colleges about 1% less than children experiencing paternal layoffs one year after making enrollment decisions. The only way this pattern could be generated spuriously is if these schools (1) raise financial aid for students who experience paternal layoff and then (2) selectively decide not to file 1098Ts for students whose greater aid package reduced their net payments to zero, despite filing 1098Ts for many other students with zero net payments. As this seems much less likely than the alternative explanation that paternal layoffs cause a small fraction of students to forego college, I conclude that this test also supports the main results of the paper.

Appendix 6: A Simple Model of Schooling and Parental Income

I here present a stylized two-period model of schooling, parental income, and borrowing constraints. The model clarifies the endogeneity problem, the role of borrowing constraints, and the relation between income and price effects on college outcomes. The model is a simple extension of Autor and Acemoglu (2014).

In Period 1, a parent allocates income W across own consumption c_0, savings s that earn interest at rate r, and a discrete schooling investment $e \in \{0,1\}$ that costs θ. In the second period, children generate earnings and consume all available resources. If parents invest in schooling then children's earnings increase from w_U to $w_S > w_U$. Parents solve

$$\max_{c_0,c_1,e} \ln c_0 + \alpha \ln c_1$$

subject to

$$\begin{aligned} c_0 &\leq W - s - e\theta \\ c_1 &\leq w_U + e\left(w_S - w_U\right) + (1+r)\, s, \end{aligned}$$

where $\alpha > 0$ indicates the relative weight placed on child utility. The parameter α depends on the degree of altruism and the relative bargaining power of parents and children (Browning, Chiappori and Weiss 2011). Without borrowing

[6] I define "often" by ranking colleges by the fraction of their 1098Ts recording zero net tuition payments, and restricting to those in the 75th percentile of this distribution.

constraints, optimizing parents invest in schooling if and only if

$$\theta < \frac{w_S - w_U}{1+r}. \tag{4}$$

This condition only depends on the costs and benefits of schooling, not parental income Wor the Pareto weight α. The key endogeneity problem is that these costs and benefits of schooling may correlate with parental income, potentially generating spurious evidence of parental borrowing constraints. The research design below addresses this problem by isolating variation in parental income that is uncorrelated with unobserved costs and benefits of schooling.

Now introduce a borrowing constraint: $s \geq 0$. Then at low enough incomes, parents invest in schooling if and only if

$$\theta < W \cdot \left(1 - \left(\frac{w_U}{w_S}\right)^{\alpha}\right) \equiv W\Phi, \tag{5}$$

where $\Phi \in [0,1]$. At high enough incomes, borrowing constraints do not bind, and schooling decisions are made without reference to parental income as in equation (4).

It is possible to link the results reported in the paper to the parameters of this model, but only under strong assumptions due to the stylized nature of the model. I show how to do this here merely as a thought experiment. From the NPSAS data discussed in Section VII I am able to observe *realized* parental spending conditional on college attendance, but not maximum *potential* parental spending. Under the strong assumption of constant Φ, we can express parental spending conditional on income and college attendance as:

$$\mu(W) \equiv E[\theta_i | \theta_i < \Phi W] = \frac{\int_0^{\Phi W} \theta_i f(\theta_i)\, d\theta_i}{F(\Phi W)} \tag{6}$$

Note that $F(\Phi W)$ is just the college attendance rate at income level W. Differentiating this conditional expectation with respect to income, we obtain:

$$\frac{d\mu(W)}{dW} = \frac{f(\Phi W)}{F(\Phi W)} \left[\Phi^2 W - \Phi \mu(W)\right] \tag{7}$$

I now solve equations (6) and (7) for Φ as follows. I focus on the case of a family considering a four-year college investment versus no college investment. I view W as all income earned by parents over the ten years prior to college enrollment, and I treat annual income in NPSAS data as a perfect measure of annual lifetime income. I interpret θ as parental spending on four years of college, ignoring other costs paid by children. I then obtain $\mu(W)$ and $\frac{d\mu(W)}{dW}$ $\left(= 0.02 = 0.05 * \frac{4}{10}\right)$ from data underlying Figure 11. If I then assume that college costs θ are distributed uniformly starting from θ_{min}, I obtain $\frac{f(\Phi W)}{F(\Phi W)} = \frac{1}{\Phi W - \theta_{min}}$. I set $\theta_{min} = 0$ based on the median spending of the lowest-income group in Figure 10.

These assumptions yield a quadratic function in Φ at each level of income, the solution to which depends primarily on $\frac{d\mu(W)}{dW}$. This yields solutions for Φ in the range of $0 - 0.15$ for annual incomes in the range of \$0-100,000 (W between \$0 and \$1 million). If I also assume that four years of college attendance raise earnings by 80% so that $\frac{w_U}{w_S} = 0.8$, these solutions imply α in the range of 0.03-0.25. In other words, under this very strong set of assumptions, the data suggest a weak propensity to spend on schooling out of marginal income and a low weight placed on child consumption relative to parental consumption. The data could also be rationalized with lower perceived returns to schooling, and greater parental altruism.

Appendix 7: Unemployment Insurance Takeup by Child Age in SIPP Data

As discussed in Section V, there is a concern that use of UI benefits to proxy for layoff may bias the results up (toward zero for adverse effects) due to differential UI takeup by child age. I can examine this concern in SIPP data, which offers samples large enough to cut the data on presence of children by age among fathers who experience layoff, and allows me to identify layoff and UI takeup separately.

I pool all SIPP data from the 1996, 2001, 2004 and 2008 waves and restrict to men between ages 25 and 65. I construct dummies $\text{flag\_kid}_{12} = 1\{\text{father has child age 12}\}$, $\text{flag\_kid}_{13} = 1\{\text{father has child age 13}\}$, ..., $\text{flag\_kid}_{25} = 1\{\text{father has child age 25}\}$. Note that after age 18, children begin to leave the parental home rapidly, such that the share of men in the sample with children of ages 12-18 in holds steady around 8.6% and falls to 1.9% by age 25.

In order to examine endogenous UI takeup, I run the following regression on a sample restricted to men who experienced a layoff in the prior quarter:

$$\text{flag\_UI}_{i,t} = \alpha + \sum_{k=12}^{25} \beta_k \cdot \text{flag\_kid}_{i,t,k} + \lambda \cdot X_{i,t} + e_{i,t}$$

where i indexes individual men, t indexes time (month and year), and $X_{i,t}$ contains a full set of time dummies, and quadratic in age, and a full set of race dummies. If UI takeup depends heavily on children's college decisions and inclinations, then we would expect to find a discontinuity in the β_k terms around ages of college attendance. Figure A7.1 plots the estimated β_k terms by age of child, k. There is no evidence of any such discontinuity, though even pooling four waves of the SIPP yields fairly noisy estimates. One might worry that this finding merely reflects high correlation between β_k and β_{k+1} due to high overlap between fathers of, say 18 and 17 year-olds. This is not the case: only about 15% of fathers with $\beta_k = 1$ have $\beta_{k'} = 0$ for any particular $k \neq k'$, because many parents have only one child and because birth spacing is quite dispersed.

While the prior test does not find any evidence of endogeneity of UI takeup around ages of college attendance, I now test for evidence of endogeneity of UI takeup with respect to having children in college. Again using the SIPP data, I regress $\text{flag\_UI}_{i,t}$ on a dummy for whether a father has any child in college, controlling for full sets of dummies for father's age, whether the father has a child at each age in the range 12-25, year, month, and state of residence, along with a quartic in monthly earnings prior to layoff.[7] I run this regression restricting to fathers age 25-64 who experience a layoff and with at least one child between the ages of 18-22. The estimated coefficient on the dummy for any child in college is .0046 [.039], providing no evidence that fathers' UI takeup depends endogenously on whether their children attend college.

Because this result is again somewhat imprecise, I also conduct an additional test. I re-run the same regression adding a variable that equals 1 if a father has a child who attends college at *any time* during the four-year panel. For example, this variable will equal 1 for fathers with 15 year-olds who are observed attending college three years later at age 18. This regression yields a similar coefficient of .0046 [.039] on whether a father currently has any child in college, but yields a larger positive coefficient of .085 [.049] on whether a father ever has a child in college during the sample period. While this result is also underpowered, it is more consistent with the idea that having a child in college may proxy for other time-invariant determinants of UI takeup such as knowledge of social programs or skill at navigating bureaucracy. This finding–if it held up under the much richer controls used in the main analysis–would pose no threat to the identification strategy because it would not differentially alter the composition of fathers in the sample by children's age at layoff.

[7]Ideally I would control for identical variables used in the propensity score used to contruct the survivor sample in the main analysis. As this is not feasible, I only control for a small subset of these variables here, for example controlling for state instead of three-digit zipcode, and omitting detailed controls for pre-layoff firm characteristics and for total family income and marital status. This narrow subset of controls biases me toward finding spurious evidence of an endogeneity problem even when none exists.

References

1. Acemoglu, Daron and David Autor. 2014. "Lectures in Labor Economics." Unpublished manuscript.

2. Athreya, Kartik B., Devin Reilly and Nicole B. Simpson. 2010. "Earned Income Tax Credit Recipients: Income, Marginal Tax Rates, Wealth, and Credit Constraints," *Economic Quarterly* 96(3): 229-258

3. Bailey, Martha J. and Susan M. Dynarski. 2011. "Inequality in Postsecondary Education," In G.J. Duncan and R.J. Murnane (eds.), *Whither Opportunity? Rising Inequality, Schools, and Children's Life Chances.* Russell Sage: New York, New York.

4. Berndt, Ernst R., Zvi Griliches and Neal J. Rappaport. 1995. "Econometric Estimates of Price Indexes for Personal Computers in the 1990's." *Journal of Econometrics* 68: 243-268

5. Cellini, Stephanie Riegg, and Claudia Goldin. 2012. "Does Federal Student Aid Raise Tuition? New Evidence on For-Profit Colleges." NBER Working Paper 17827.

6. Hall, Bronwyn H., 1971, "The Measurement of Quality Change from Vintage Price Data," chapter 8 in Zvi Griliches (ed.) Price Indexes and Quality Change, Cambridge, MA: Harvard University Press, 240-271.

7. Hall, Bronwyn H., Jacques Mairesse, and Laure Turner. 2005. "Identifying Age, Cohort and Period Effects in Scientific Research Productivity: Discussion and Illustration Using Simulated and Actual Data on French Physicists," NBER Working Paper 11739.

FIGURE A1.1

IRS Linked Sample vs Census

Family Income Distributions of Children Age 14-16 in 2001

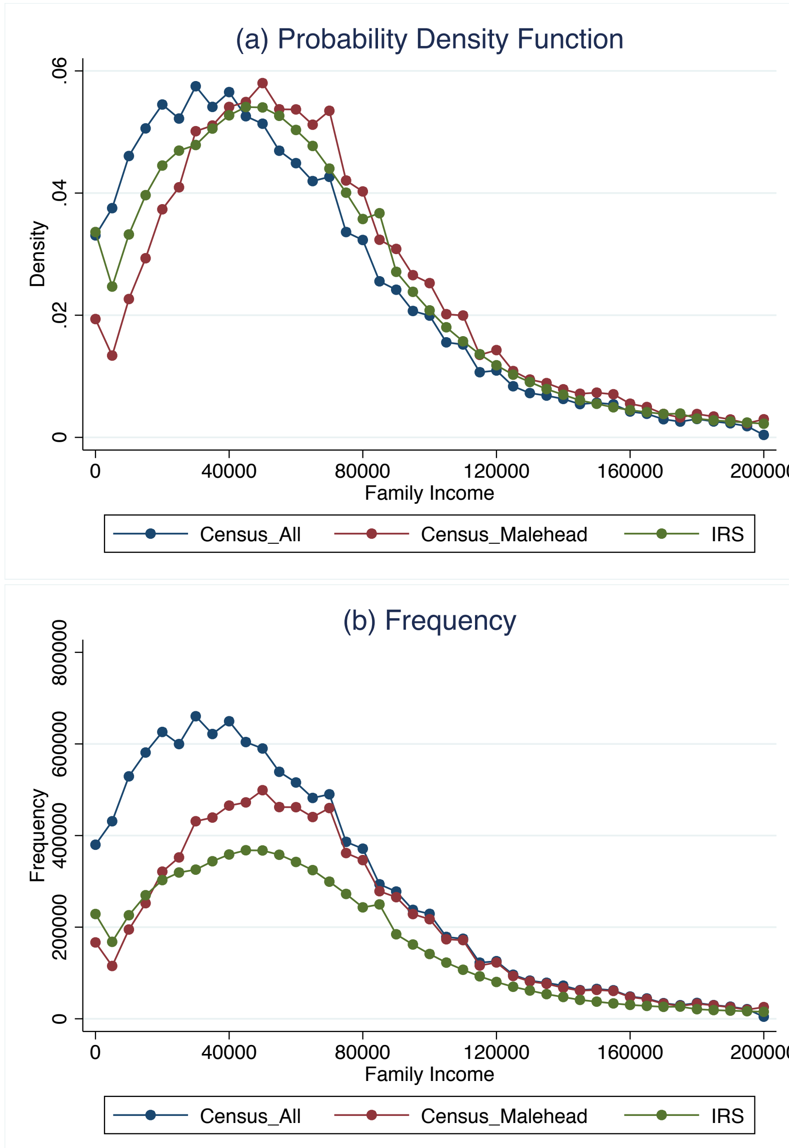

Notes: IRS linked sample is a random sample of the children age 14-16 in 2001 who I match to fathers in 1996-1998. Census_All is a random sample of children from the ACS age 14-16 in 2001. Census_Malehead is a random sample of children from the ACS age 14-16 and residing in male-headed households in 2001. Income is year 2001 pre-tax family nominal income. Panel (a) presents the probability density functions for the income distributions of these three samples. Panel (b) presents the corresponding histograms for these three samples.

FIGURE A3.1

Outcomes at Age 19 by Event-Age, Layoff Sample

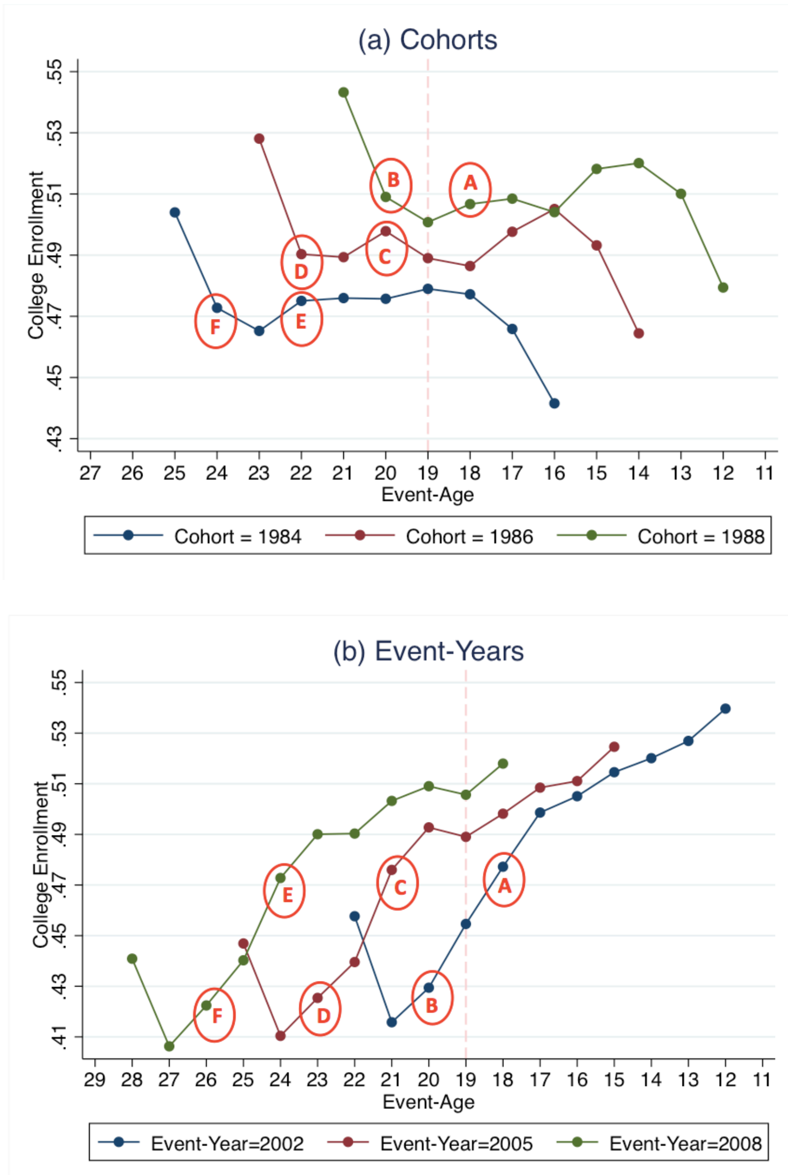

Notes: Figures plot mean college enrollment for children grouped into cohort by event-age bins in the Layoff sample. Panel (a) displays these group means for three cohorts. Panel (b) displays these group means for three event-years.

FIGURE A5.1

College Enrollment at Age 19 in 2002 by Family Income

Two Alternative Measures

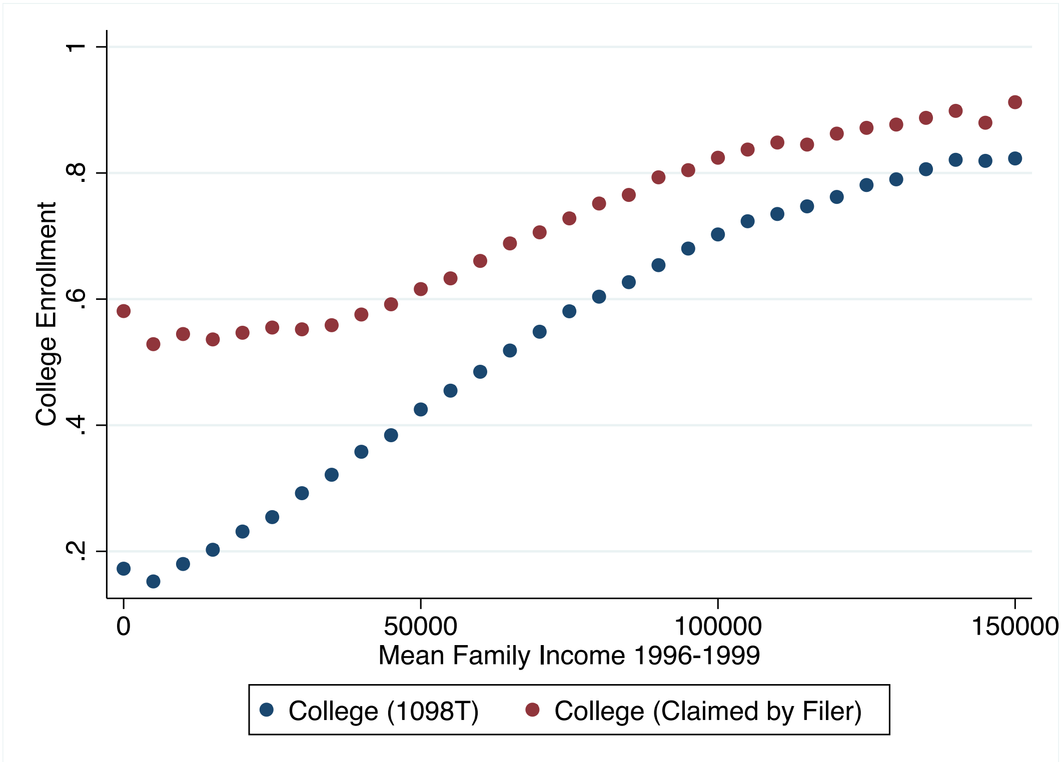

Notes: Figure plots two measures of college enrollment in the Survivor sample. "College (1098T)" uses an indicator variable for receipt of a 1098T form by a student. "College (Claimed by Filer)" uses an indicator for whether a child is claimed by an adult (not necessarily the linked father).

FIGURE A5.2

Event-Age Studies for Alternative College Enrollment Measures

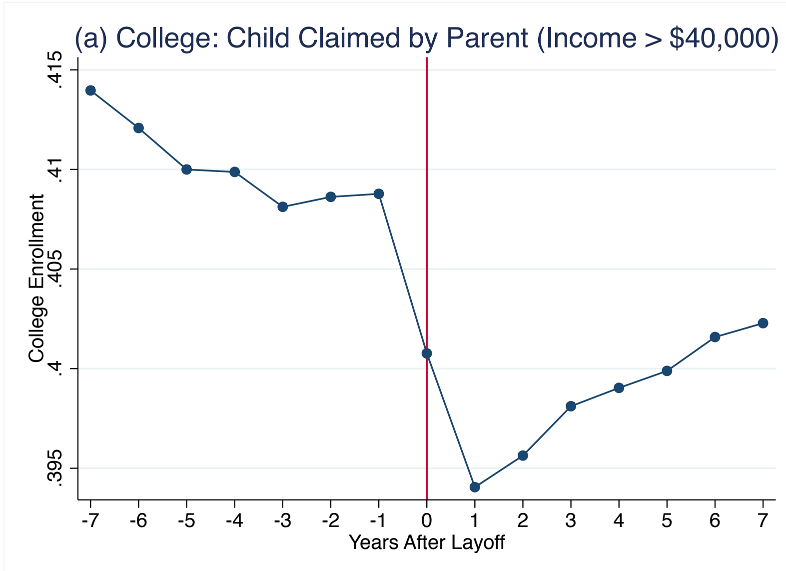

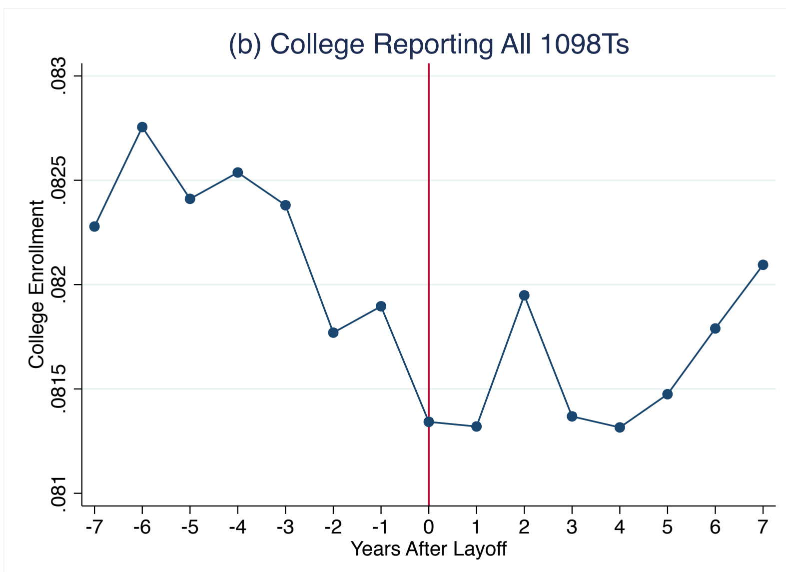

Notes: Panel (a) is analogous to Figure 4.b but replacing the 1098T indicator for college with a claimed-by-adult-filer indicator for college, and adding a restriction to age 19-22 and family incomes in 1996-1999 above $40,000. Panel (b) is also analogous to Figure 4.b but adding a restriction that the 1098T come from a college that appears to report 1098Ts even for children who make no payments for college.

FIGURE A7.1

Unemployment Insurance Take-Up by Child Age

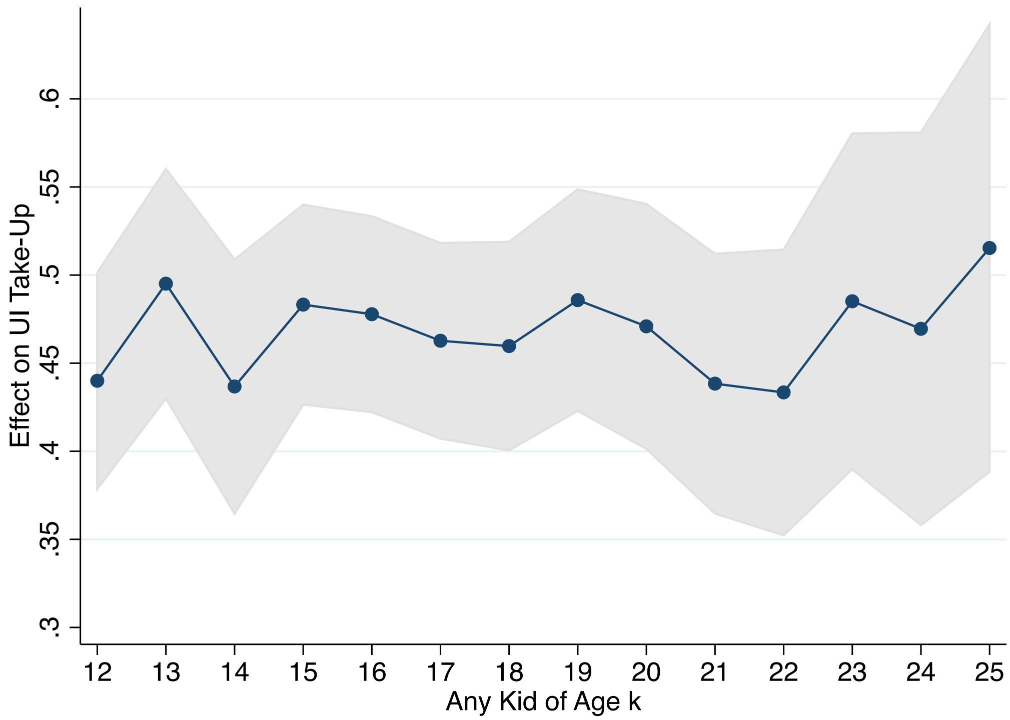

Notes: Plots estimated coefficients from regression in pooled SIPP waves 1996, 2001, 2004, 2008 of UI take-up dummy on dummies for "any child of age k" for k=12, 13, … , 25. Restricts to men age 25-65 experiencing layoff in the prior three months with at least one dependent child, and includes additional controls for year, month, and father's age and race. Regression incorporates SIPP person weights. Confidence interval constructed from standard errors clustered at father level.

Table A2.1: Firm Closures by Year and Size

Year of Closure	Number of Firms	Mean Firm Size
2000	12,894	156
2001	13,164	153
2002	10,679	164
2003	9,652	121
2004	8,966	118
2005	8,979	140
2006	9,688	121
2007	10,910	99
Total	84,932	136

Table A2.2: Summary Statistics 1999-2009 for Children at Age 19

Sample:	Closure			Non-Closure		
	Pre-Shock	Post-Shock	% Diff	Pre-Shock	Post-Shock	% Diff
<u>Parent Outcomes</u>						
Father's Earnings	48,370	41,501	-14.2%	50,761	51,927	2.3%
Father Married	0.787	0.794	0.9%	0.810	0.807	-0.4%
Mother's Earnings	21,944	23,216	5.8%	23,608	24,195	2.5%
<u>Child Outcomes</u>						
Enrollment	0.505	0.522	3.4%	0.564	0.575	2.0%
Earnings	6,882	6,564	-4.6%	6,542	6,523	-0.3%
Freq	23,220	66,895		1,198,446	2,443,755	

Notes: Non-Closure sample is propensity-score reweighted to match Closure sample on observables. "Pre-Shock" includes cells in years before events occur, "Post-Shock" includes cells in years after events occur. "% Diff" is the percent difference between Post-Shock and Pre-Shock columns. Averages pool all available cohorts.