14.662 — Spring 2018

Lecture Note #2: Educational Production and Wage Structure

David Autor, MIT and NBER

February 9, 2018

1 Introduction

A key observation emphasized by Katz and Murphy (1992)—and expanded upon by Card and Lemieux in their 2001 chapter in the Gruber volume *Risky Behavior Among Youth*—is that the rate of increase in educational attainment in the United States slowed dramatically in the early 1970s, particularly for males. This can be seen in Figures 3 and 4 of the CL-G 2001 chapter. Their chapter explores a number of explanations for this pattern including:

1. Leveling of improvements in family background which had increased the rate of college going of children.
2. Rising costs of college
3. Falling returns to college
4. A rise in interest rates (same as a decline in the return to college)
5. Changes in college going associated with the Vietnam war
6. A decline in unemployment
7. Changes in cohort sizes

While I don't discuss the details of this paper, the evidence suggests that only two of these hypotheses are particularly relevant. One is the termination of the Vietnam war. Draft deferments for college students may have artificially expanded male college enrollment in the late 1960s and early 1970s—and this dropped precipitously with the war's end. The second, and at least equally important, is 'cohort crowding.' Very large baby boom cohorts competed for a relatively inelastic supply of college slots, and this may have reduced the rate of college attendance (see Figure 7). Bound and Turner 2007 ("Cohort Crowding" *JPubEc*) provide a more detailed exploration of this hypothesis.

2 Cohort supplies and wages

Following on the analysis of the factors affecting cohort trends in supplies, Card and Lemieux's 2001 *QJE* paper considers the implications of a slowdown in the supply of new college production for wage inequality. Given the work by Katz and Murphy, this was thought to be well-trodden ground. But C-L's paper offers a subtle interpretation of wage trends that advances Katz-Murphy thesis. The C-L paper also probably provides some of the best evidence on how a parsimonious, parametric supply and demand framework can capture many of the important features of the data—and, notably, in several countries simultaneously.

Card and Lemieux (C-L), Figure 1:

- College premium for older men ages 46-50 relatively flat 1959 - 1995 (though does rise somewhat in the 1980s).
- College premium for younger men, ages 26 - 30, varies dramatically by decade.
- These patterns are apparent for US, UK and Canada

Why are these patterns surprising? One reason is that a famous 1993 paper by Juhn, Murphy and Pierce in the *JPE* argued *against* the existence of strong cohort effects in the skill premium. One would therefore expect the College/HS wage gap to roughly move proportionately across cohorts with the 'skill premium.' But this is clearly not the case.

We've so far written the College/HS wage premium as a function of three parameters: 1) H/L relative supply; 2) σ elasticity of substitution; 3) A_H/A_L state of technology.

C-L hypothesize that in addition to these parameters, different age groups within an education group may be imperfect substitutes. Perhaps young college graduates tend to manage McDonalds and Walmart stores while older college graduates take more sedentary management positions.

If this hypothesis is correct, it would imply that the level of wages (or wage inequality) by education may also differ by age. More subtly, it also implies that in a period of decelerating educational attainment, educational premia are likely to twist so that inequality among young workers expands relative to the old. The reason is that a slowdown in educational attainment does not create scarcity of older educated workers; it creates a scarcity of younger educated workers. This observation highlights that it's may not just be the *level* of educational supply that affects inequality but also the *rate of change* of supply.

To see the relevance of this hypothesis, bear in mind that there are many birth cohorts in the labor market in any given time. When education levels are rising, younger cohorts are relatively more educated than older cohorts. In general throughout the 20th century, older cohorts were less educated than young, giving rise to a downward sloping cohort-education profile. When younger cohorts stopped increasing their education levels relative to their elders in the 1980s (i.e., which occurred, with slightly different timing in the U.S., UK and Canada), the inter-cohort pattern of education flattened.

In a model with perfect substitution across cohorts (which we have implicitly assumed so far), the *deceleration* is irrelevant to inequality—it is only the *aggregate* supply of college and high school equivalents that matters. But in a model with imperfect cohort substitutability, a deceleration in the rate of new college graduate production will reduce the effective supply of educated workers by more for younger than older workers. When young college graduates become relatively more scarce than older college graduates, the return to education will rise more for the young than the old.

C-L Table 1 and Figure 2. Several points are visible:

- During 1959 - 1975, the age profile of the college/HS premium was basically concave.
- From 1975 - 1981, the entire profile shifts downward as earnings differentials fell (potentially due to outward supply shifts, as noted by Katz-Murphy).
- From 1982 - 1986, the profile rises again for younger cohorts but not by much for older cohorts. This rise continues in subsequent years so that by 1994 - 1996, the education premium for the youngest 4 cohorts is above that of the oldest 2.
- A similar twisting is evident for the UK, and college premiums for older workers actually fell in the 1990s.
- Canada also experiences some twisting, though it is not nearly as pronounced.

Note that these comparisons are based on age cohorts, not experience cohorts, which could be problematic. Based on the human capital model, you would expect relative earnings of college educated workers to rise relative to workers of the same age (but with less education) over the life cycle since they are on the steeper part of the age-earnings profile later in life (recall: college grads enter the labor market about 4 years later than HS grads of the same birth cohort). We can absorb this difference with age effects in a regression, however. More critically, workers with different levels of education enter the labor market at different points in their life-cycle. Hence, comparing a HS Grad and College grad of the same age 5 years after high school graduate, the HS Grad will typically have 5 years of work experience while the college grad will have 1. You might therefore expect that there would be less substitutability among HS and College grads of the same age cohort than there would be among HS versus College grads of the same experience cohort. (In the extreme case, College grads are still in college while HS grads are acquiring their first 4 years of experience, so it's hard to imagine that they could be close substitutes between the ages of 18 - 22).

2.1 Formalization

Card and Lemieux formalize the imperfectly substitutability hypothesis by writing down a nested, two-level CES model. At the upper level, this model is identical to our simplified, two-factor (high and low educated labor) CES model (without 'extensive' technical change) from the previous lecture. Beneath the upper level CES, the supplies of each education group are themselves CES aggregates of the labor supply of different age groups of members of the education group.

Versions of the multi-level CES appear in a number of papers on your syllabus. These include Krusell 2000, Borjas 2003, Acemoglu, Autor and Lyle 2004, and Carneiro and Lee 2011. The ability to nest CES functions make this function versatile for theory and estimation. You should master this model for your general(s) toolkit.

Here's the Card and Lemieux version. If different age groups are imperfect substitutes in production, a natural way to combine them is as a CES aggregate:

$$H_t = \left(\sum_j \alpha_j H_{jt}^{\eta} \right)^{1/\eta}, \tag{1}$$

and

$$L_t = \left(\sum_j \beta_j L_{jt}^{\eta} \right)^{1/\eta}, \tag{2}$$

where $\sigma_A = 1/(1-\eta)$ is the elasticity of substitution across different age groups j, the parameters α_j and β_j are efficiency parameters, which are assumed fixed by age group (that is, they do not vary across cohorts or over time), and H_{jt}, L_{jt} are age group specific supplies of high and low educated workers in each period t. Note that in the limiting case where $\eta = 1$, cohorts are perfect substitutes (though they may have different 'efficiencies' given by α_j, β_j).

Aggregate output is a function of *total* college and HS supply (i.e., it does not depend on age groups once properly aggregated) and the technological efficiency parameters A_{Ht}, A_{Lt}, which are time varying:

$$Y_t = f(H_t, L_t, A_{Ht}, A_{Lt}). \tag{3}$$

Assume this aggregate production function is also CES:

$$Y_t = (A_{Ht} H_t^{\rho} + A_{Lt} L_t^{\rho})^{1/\rho}, \tag{4}$$

with $\sigma = 1/(1-\rho)$, where σ is the aggregate elasticity of substitution between high school and college workers.

Assuming wages are set competitively and the economy operates on the demand curve, wages will equal marginal products. So, the wage of low skill workers in age group j is:

$$\begin{aligned} \frac{\partial Y_t}{\partial L_{jt}} &= \frac{\partial Y_t}{\partial L_t} \times \frac{\partial L_t}{\partial L_{jt}} \\ &= A_{Lt} L_t^{\rho-\eta} \pi_t \times \beta_j L_{jt}^{\eta-1} \end{aligned} \tag{5}$$

where

$$\pi_t = (A_{Lt} L_t^{\rho} + A_{Ht} H_t^{\rho})^{1/\rho - 1},$$

and similarly for the wages of college graduates. Notice in (5) that provided that $\eta < 1$, the age-specific wage (by education) is declining in age-specific supply.

Efficient utilization of skill groups further requires that relative wages across skill groups are equated with relative marginal products. Writing the relative wages of H versus L workers in the same cohort gives:

$$\begin{aligned} \ln\left(\frac{w_{jt}^H}{w_{jt}^L}\right) \equiv r_{jt} = \ln(A_{Ht}/A_{Lt}) + \ln(\alpha_j/\beta_j) - \frac{1}{\sigma}\ln\left(H_t/L_t\right) \\ - \frac{1}{\sigma_A}\left[\ln(H_{jt}/L_{jt}) - \ln(H_t/L_t)\right] + e_{jt} \end{aligned} \tag{6}$$

(Note that the α_j and β_j terms are switched in these notes relative to the original paper.) Hence, the relative H/L wage ratio for cohort j depends on four factors (in order of above):

- The technology parameters A_{Ht}/A_{Lt},
- The age-specific efficiency parameters α_j/β_j,
- The aggregate supply of H_t/L_t,
- And the gap between the relative supply of H_{jt}/L_{jt} (in a cohort) and aggregate overall supply $\ln(H_t/L_t)$, written as $\ln(H_{jt}/L_{jt}) - \ln(H_t/L_t)$).

Notice if $\eta = 1$ (implying that $\sigma_A = \infty$), this equation is algebraically the Katz-Murphy model (with an extra β_j/α_j floating around, but this would be absorbed into H_t/L_t).

To see how this could generate 'cohort effects' in wages, suppose that the log supply ratio (*not wage ratio*) for workers who are age j in year t consists of a cohort effect for the group λ_{t-j}, and an age effect ϕ_j that is common across cohorts (note that $t-j$ is a cohort's year of birth):

$$\ln\left(H_{jt}/L_{jt}\right) = \lambda_{t-j} + \phi_j.$$

Here, the λ_{t-j} refers to cohort specific relative supply of H versus L labor, and ϕ_j is an age effect. An operative assumption here is that λ_{t-j} is fixed for a cohort because cohorts do not obtain (much) additional education after labor market entry. Of course, the age effect ϕ_j allows for this relative supply term to vary with age, but this age profile must be constant across cohorts.

We can now rewrite equation (6) in estimable form as

$$\ln\left(\frac{w_{jt}^H}{w_{jt}^L}\right) \equiv r_{jt} = \ln(A_{Ht}/A_{Lt}) + \ln(\alpha_j/\beta_j) - \frac{1}{\sigma_A}\phi_j + \left(\frac{1}{\sigma_A} - \frac{1}{\sigma}\right)\ln\left(H_t/L_t\right) - \frac{1}{\sigma_A}\lambda_{t-j} + e_{jt}. \tag{7}$$

This equation says that if we have observations on the college high school wage premium for a set of age groups $\{J\}$ and a set of years $\{T\}$, these wages will depend on:

1. A set of year-specific factors that are common across age groups:

   $$\ln(A_{Ht}/A_{Lt}) + \left(\frac{1}{\sigma_A} - \frac{1}{\sigma}\right)\ln\left(H_t/L_t\right).$$

   These are the 'time effects.'

2. A set of age-group specific factors that are common across years:

   $$\ln(\alpha_j/\beta_j) - \frac{1}{\sigma_A}\phi_j.$$

   These are the 'age effects.'

3. A set of cohort-specific constants:

   $$-\frac{1}{\sigma_A}\lambda_{t-j}$$

4. And a residual, e_{jt}.

Two important special cases:

1. When $1/\sigma_A \approx 0$, $(\sigma_A \to \infty)$ the cohort effects will be ignorable (as above_. This occurs if cohorts are perfect substitutes, as is commonly assumed.

2. When $\ln(H_{jt}/L_{jt}) - \ln(H_t/L_t)$ is approximately constant, meaning that the proportionate growth in cohort supplies relative to average supplies is steady. In this special case, there will be no 'twist' in the education premium by age because cohort effects will be *equal* for all cohorts (meaning that they are present but not identifiable). This is easy to see in equation (6) where this proportionate growth condition implies that

$\frac{1}{\sigma_A}\left[\ln(H_{jt}/L_{jt}) - \ln(H_t/L_t)\right]$ is constant. Conversely, if $\ln(H_{jt}/L_{jt}) - \ln(H_t/L_t)$ varies with time, as will occur when trends in educational attainment accelerate or decelerate, then r_{jt} will exhibit 'cohort effects' (assuming relative education in a cohort is roughly constant over time). This is an astute observation, all the more so because this 'special' case appears surprisingly relevant.

2.2 Existence of cohort effects?

As a first check on the existence of cohort effects, C-L wish to estimate:

$$r_{jt} = b_j + C_{t-j} + d_t + e_{jt}. \tag{8}$$

The problem with this equation as written is that *it is not estimable*—cohort effects are a linear combination of the $b_{j's}$ and the $d_{t's}$. Even though we may believe that age, cohort and time effects exist, we cannot identify them.

What C-L do to solve this problem is to restrict the cohort effects to be the same for the 10 oldest cohorts, which then allows identification of the age effects along with the cohort effects for younger ages. This is identification by assumption, but it may be reasonable.

Results are visible in Table II:

- Large increases in the year effects after 1980.
- But the year effects are not large for the oldest cohorts.
- By implication, cohort effects are detected for youngest cohorts.
- Similar patterns in UK and Canada, though the cohort effects start later in both countries.

These results suggest (consistent with the pictures) that a model with cohort effects has good potential to fit the data. See Figure III, which gives a clear picture (from the supply side) of why these cohort effects would be evident in the data if cohorts are indeed imperfect substitutes.

2.3 Estimating substitutability among cohorts

Now that C-L have demonstrated some evidence for cohort effects, it is time to take this parametric model more seriously—or more literally, if you like. C-L estimate age-group specific elasticities of substitution with the equation:

$$r_{jt} = b_j + d_t - (1/\sigma_A)\ln(H_{jt}/L_{jt}) + e_{jt}. \tag{9}$$

Note that in this equation:

- $\widehat{b}_j = \ln(\alpha_j/\beta_j)$
- $\widehat{d}_t = \ln(A_{Ht}/A_{Lt}) - (\frac{1}{\sigma} - \frac{1}{\sigma_A})\ln(H_t/L_t)$.

Hence, the structure of the CES model allows us to estimate σ_A while absorbing the main effects of σ and H/L.

Results are in Table III. $\widehat{\sigma}_A \approx 4 - 6$.

It is remarkable that this estimation works as well as it does for the US, UK and Canada and that substituting a time trend for year dummies in equation (9) has only a minimal impact on estimates of σ_A.

2.4 Estimating models for aggregate and cohort supply simultaneously

The step above provides an estimate of σ_A. But if we want to fit the full-blown model to the data, we need to construct a measure of aggregate and cohort specific supply. One other ingredient is still missing: an estimate of the age-specific efficiency parameters α_j, β_j (assumed constant by age across cohorts).

C-L estimate these parameters by fitting the following equations (these are adapted from equations (5) and (6) of their paper):

$$\ln(w^L_{jt}) + \frac{1}{\widehat{\sigma}_A} L_{jt} = \ln(A_{Lt} L_t^{\rho-\eta} \psi_t) + \ln\beta_j + e_j \tag{10}$$

and

$$\ln(w^H_{jt}) + \frac{1}{\widehat{\sigma}_A} H_{jt} = \ln(A_{Ht} H_t^{\rho-\eta} \psi_t) + \ln\alpha_j + e'_j.$$

These equations are estimated for each skill group, pooling across all time periods t. The messy term immediately to the right of the equal sign, $\ln(A_{Ht} H_t^{\rho-\eta}\psi_t)$, is absorbed by a set of year dummies, and the efficiency parameters α_j and β_j are estimated with age dummies.

So, we have now recovered estimates of σ_A and the α_j and β_j in the first stage of the estimation. This does not give us σ, the overall elasticity of substitution between college and high school grads, nor does it tell us how important the cohort specific supplies are to inferences about overall changes in relative skill demand (the SBTC term), equal to A_{Ht}/A_{Lt} in the upper level of the CES function. But we're close!

2.5 Stage 2 estimation results

Given estimates of $\alpha_j, \beta_j, \sigma_A$ we are ready to do the grand estimate of the following equation:

$$r_{jt} = \ln(A_{Ht}/A_{Lt}) + \ln(\alpha_j/\beta_j) - \frac{1}{\widehat{\sigma}_A}\phi_j + \left[\frac{1}{\widehat{\sigma}_A} - \frac{1}{\sigma}\right]\ln(H_t/L_t) - \frac{1}{\widehat{\sigma}_A}\lambda_{t-j} + e_{jt}, \tag{11}$$

where

$$\ln(H_{jt}/L_{jt}) = \lambda_{t-j} + \phi_j, \tag{12}$$

and, as above, λ_{t-j} is a cohort effect at market entry, and ϕ_j is an age profile of relative labor supply that is assumed constant across cohorts. Adding ϕ_j to the model allows labor supply or efficiencies by education group to differ over the life cycle and so $\ln(H_{jt}/L_{jt})$ does not have to be exactly constant by assumption (but to be identified, the profile must be fixed over age groups across cohorts). Note that in estimating (11), C-L do not impose that the $\hat{\sigma}_A$ that appears in equation (11) is the same $\hat{\sigma}_A$ that is used to construct the aggregate supply index in (12), which was obtained from estimating (9). By not imposing this restriction, C-L improve the credibility of the exercise. If they obtained wildly different estimates of σ_A from different components of the estimation, this would be worrisome.

Notice that in estimating this equation, C-L include two supply measures as regressors: 1) the aggregate supply measure $\ln(H_t/L_t)$; and 2) the deviation of the cohort supply measure from the aggregate measure, $\ln(H_{jt}/L_{jt}) - \ln(H_t/L_t)$. The coefficient on the former provides an estimate of $1/\sigma$ and the coefficient on the latter provides an estimate of $1/\sigma_A$. In addition, C-L estimate 2nd stage models where the aggregate relative supply index incorporates estimates of σ_A and $\alpha_j's, \beta_j's$ to construct age-specific relative supplies.

Some key points from this table:

- The coefficient on the aggregate supply index (which does not incorporate age-specific supplies) is highly significant and implies an H/L elasticity of substitution of approximately 2.5. (Note, when males and females are combined, this estimate is close to 1.5, which is comparable to Katz-Murphy '92.)

- Estimates of σ_A from this procedure are very similar to the previous exercise that only used age-group and not aggregate supplies. This did not have to work as well as it does; there is no maintained restriction that $\widehat{\sigma}_A$ in stage 2 equals $\widehat{\sigma}_A$ in stage 1.

- As Katz and Murphy noted, the model predicts a large increase in inequality in the late 70s which did not occur. Hence, C-L put in a dummy for this 1975 - 1980 period to soak up this variation. Why this is needed remains a puzzle. (And of course, this is a completely ad hoc fix-up. There's no glory in using a dummy variable to remove this messy discrepancy.)

- Using a more sophisticated version of the relative supply index that implicitly incorporates the imperfect substitution across age cohorts (see equation 4) has surprisingly little quantitative impact on the estimates. Hence, this exercise does *not* change any substantive conclusions from Katz-Murphy '92 (or Autor-Katz-Krueger '98). C-L could not have known this, however, without building the richer measure and running the regression.

- It would have been helpful if C-L had allowed for a less restrictive (i.e., non-linear) time trend in A_{Ht}/A_{Lt} to test whether the assumed smooth trend in demand was really the best fit. But note that they could not include unrestricted year dummies since this would preclude estimation of $1/\sigma$.

2.6 Caveats

1. As is visible in Table V, the importance of cohort specific supplies is vastly reduced when C-L use experience rather than age cohorts ($\widehat{\sigma}_A \simeq 10$). Many would argue that experience cohorts are the right empirical construct (since it is clearly true that *latent* college grads who are *still attending college* are a relatively poor substitute for college or high school grads already in the labor market).

2. When male and females are pooled in age cohorts in Table VI (the first three columns of which are comparable to Table III), the year effects are much larger than in the male only models when cohort effects are included. This indicates that the cohort supply approach is not as effective in explaining the rise in the college/high school gap for men and women jointly.

3. On the other hand, the estimated elasticity of inter-cohort substitution is extremely similar in the pooled gender model.

4. Note also that in the combined-gender specification, $\widehat{\sigma} = 1/0.865 \simeq 1.2$, similar to Katz-Murphy.

5. Again, it would have been useful to test a less restrictive specification of A_{Ht}/A_{Lt}.

2.7 Conclusions

- The C-L paper uses a simple theoretical model with remarkable empirical success. This example testifies to the potential value of parametric (I hesitate to use the word structural) modeling; in this setting, much can be explained with little for three different countries with different time patterns of wages and supplies.

- The paper raises an important historical and policy question as to why the cohort trend in educational attainment slowed. Possibilities include: falling returns to education, the Vietnam war (which caused a boom then bust), physical crowding at universities, and perhaps some 'natural' maximum education for males (note that female educational attainment keeps rising relative to males).

- We appear to be at an important historical turning point at which women have overtaken and surpassed men in educational attainment. This raises numerous, interesting research opportunities. We'll discuss some of these in a few lectures when we consider gender norms, the design of jobs, and the operation of marriage and childrearing 'markets.' We'll only touch on these topics, but hopefully you will find inspiration.

3 Has there been a decline in the quality of college graduates? (Carneiro and Lee, 2011)

As has been noted repeatedly in lecture, the fraction of students attaining a college degree has risen enormously over at least six decades. More than 70 percent of women finishing high school in 1996 entered college, and about 60 percent of males. For cohorts born in 1945, 33 percent of males and 21 percent of females received a four year college degree by age 40.

These secular changes in college attainment create marked differences in the cohort supply of education. But one might also speculate that they create differences in the cohort *quality* of education. This might occur because 'lower-quality' individuals go on to college, or because the quality of education deteriorates when there is a large influx of students (or both). How this might affect the wage structure is non-obvious. When the 'marginal' high school student completes a college degree, this student lowers the average quality of *both* college and high school graduates. Since we are typically measuring the earnings *gap* between college and high-school grads, it's not clear (without distributional assumptions) how a diminution in the average quality of *both* groups affects the *relative* quality of college versus high-school graduates.

But this observation does have implications for the *absolute* quality of cohorts. Under this hypothesis, when college attendance is rising the quality of college cohorts is falling. When attendance stabilizes, quality plateaus. As we saw in Card-Lemeiux, college attendance rose secularly for cohorts born between 1920 and 1945, then stabilized for cohorts born after 1965. Thus, it's conceivable that the cohort quality was falling for cohorts entering the labor market up through the early 1980s, afterwards quality decline might have ceased. This might give rise to a relative increase in the earnings of young college grads, not because there were relatively fewer but also because the intercohort trend of quality decline would have ceased, suggesting (naively) a flattening of the age-earnings profile. This hypothesis could therefore be complementary to, or substitutable for the Card-Lemieux hypothesis.

Before investigating this point further, one needs to fix intuitions on the relationship between college going and skills. The analysis of inequality is most typically measured in terms of earnings ratios of college versus high-school workers. Assume that the best high-school grads go to college and the worst do not. When the fraction of high-schoolers transitioning to college rises, the average quality of *both* college and high-schoolers falls. It's therefore not self-evident that an increase in college-going should reduce the college/high-school premium within a cohort, and it could well raise it. (Assume first that the distribution of ability is uniform; now assume that it is normal. You will see that your conclusions depend greatly on the shape of the ability distribution and the point of truncation.) The Carneiro and Lee paper takes this issue seriously. Interestingly, their results imply that college and high-school 'skills' are distinct, so it's possible for college quality to decline as enrollment increases, even while high school quality remains flat.

A challenge for the Carneiro-Lee approach is that at its basic from, their hypothesis is isomorphic to the Card-Lemieux hypothesis. Both hypotheses predict that education 'returns' rise differentially for young college grads when the supply of young college grads decelerates. But this shift stems from different forces. In the Card-Lemeiux view, it stems from a slowdown in the supply of young college adults. In the Carneiro-Lee view, it could alternatively stem from an *improvement* in the quality of college relative to non-college adults.

How can these ideas be distinguished? Carneiro-Lee draw on a technique first employed by Card and Krueger in a 1992 *JPE* paper. They identify the effect of cohort quality on earnings by regressing the wages of workers living outside of their home region (9 Census regions) on the educational composition of their cohort. Assuming cross-region moves are exogenous to wages (a big if), they can potentially test whether college workers from cohorts with high college-going relative to their region's average level earn relatively lower wages in other regional labor markets.

In particular, write the wage of college worker i as $\ln w^c_{iatrb}$, where a is current age, t is year, r is the region of work, b is the region of birth, and the super-script c refers indicates that i is a college worker. We can write an econometric model for this wage as:

$$\ln w_{iatrb}^{c} = \gamma_{atr} + \gamma_{ab} + \gamma_{tb} + \phi\left(\tilde{P}_{t-a,b}\right) + e_{iatrb}^{c}.$$

The coefficient of interest is ϕ, which Carneiro-Lee parameterize as the odds of the proportion of the cohort that attended college, $\tilde{P}_{t-a,b} = P_{t-a,b}/\left(1 - P_{t-a,b}\right)$. In this model, γ_{atr} is a full set of interactions between age, year and region of work. These dummies absorb average wage levels of all college workers by age in year t in each region. The dummies γ_{ab} take out average wage levels of workers by each region-of-birth by age group. Note that these age by birth-region effects are not time varying, so these dummies allow for cross-cohort within region of birth variation in wages. The dummies γ_{tb} take out average wages of workers by each region of birth by year. These time by birth-region dummies are not allowed to differ across cohorts, so again, cross-cohort within birth-region variation is preserved. Thus, what is left is cohort by birth-region variation in wage levels. This wage variation is identified by individuals born in different regions but working in the same labor markets. (Workers who are born in and working in the same region of birth should not serve to identify the coefficient of interest, ϕ).

So to summarize:

- Wages vary across schooling $\times$ year $\times$ age $\times$ residence-region $\times$ birth-region cells.
- Weeks worked (for labor supply models) vary across schooling $\times$ year $\times$ age $\times$ residence-region cells.
- Composition varies across schooling $\times$ year $\times$ age $\times$ birth-region.

These choices accord with the hypotheses that skill prices vary by region of residence and are impacted by total labor supply in the region of residence, while composition is constant within region of birth (though this is subject to selective migration issues).

Carneiro-Lee attempt to address the selective migration issues with a variety of parametric corrections, but none are ideal. (Card and Krueger also do not fully address these issues, but then again, their paper is 15 years older.) They also do not address the factors that drive cohort-region variation in college-going. One could probably instrument for this variable using cohort size. The migration issue is much harder.

Table 1 of the paper makes the main statistical argument. To interpret the coefficient of -0.085 (Panel A of Table 1) on $\tilde{P}$ consider the following: if college enrollment increases from 50% to 60%, then $\tilde{P}$ increases from 1 to 1.5, implying a fall in college wages of 0.0425. If we want to interpret this effect as occurring through a decline in the quality of the marginal college-goers (meaning that there is no dilution effect on inframarginal goers) Let w_1 be the average log wage of inframarginal students. If marginal students are $17\% = 10/60$ of college goers, then the observed college wage is:

$$\begin{aligned}\Delta\bar{w} &= -(0.084 \times 0.5) = 0.83w_1 + 0.17w_2 - w_1 \\ -(0.084 \times 0.5) &= 0.17(w_2 - w_1) \\ w_1 - w_2 &= 0.247\end{aligned}$$

So, the estimate implies that the marginal college-goers are 25% less skilled than the inframarginal college-goers.

The estimates do not imply any effect of increased college-going on the wages of high-school grads. This may be seen as suggestive evidence in favor of the congestion externality argument for why college-grad quality falls when college-going rises. Since high-school going is universal, there is no congestion component. For college-going, this is clearly not true. Under this hypothesis, it is easy for a rise in college-going to reduce college quality without also reducing high-school grad quality. Carneiro-Lee do not address the issue of selection into post-college schooling. It may be that some of the apparent college-quality dilution is coming from an increase in the proportion of the best college-grads who go onto grad school. (That is, because Carneiro-Lee are contrasting those with exactly a BA with those with exactly a HS diploma, an increase in the fraction of college grads going on for a higher degree could dilute the quality of the BA-only pool.)

I will not 'deconstruct' this paper in exhaustive detail/ I believe it should be accessible based on your understanding of Card and Lemieux '01. The exercise is extremely interesting and potentially highly important. Figure 3 is crucial, though also a bit ill-conceived. Here, the paper plots counterfactual college premium series, holding composition, supply and both composition and supply constant. The problem with this exercise is that one cannot hold composition constant without implicitly changing supply and vice versa (since supply includes composition as an argument). Moreover, the supply-constant college premium would better be titled "Demand." That is, holding supply (and implicitly composition) constant, any movement in the college premium can only reflect movements in demand. Thus, this figure implies that relative demand for college workers has grown more rapidly in recent decades than what would conclude from simply adjusting for 'supply' without adjusting for composition. It also suggest a smaller deceleration of demand growth in the 1990s than is commonly believed—which is potentially a key finding.

To wrap up, this paper appears a valuable step forward for this literature. It both establishes the importance of economically large quality changes in recent cohorts of college graduates and demonstrates (or suggests) that underlying shifts in demand have in part been masked by declines in quality. It is of substantive importance to ask if college quality dilution is due to a reduction in the latent ability of marginal college-goers or instead reflects a congestion externality (or both). These alternative hypotheses have distinct implications for the long-term social returns to raising the educational stock of advanced economies.