Effects of confidential oral contraceptive access in late adolescence on work and earnings

Randy Cragun*

August 15, 2023

Abstract

This article uses state-specific timing of legal changes and data from the Current Population Survey and National Longitudinal Survey of Young Women to estimate effects of confidential access to the pill in late adolescence on earnings, labor force participation, and human capital accumulation over the life cycle, finding no evidence of effects. These results are contrary to past research, which imposed a restriction on the regression model that this article shows explains the past results.

*Birmingham-Southern College, Box 549007, 900 Arkadelphia Rd, Birmingham, AL 35254, rc@rcragun.net, rscragun@bsc.edu, +1-205-226-7748, Fax:+1-205-226-4847. This research benefited from a grant from Birmingham-Southern College and from non-material support and feedback from Jason Lindo, Caitlin Knowles Myers, Ishita Chatterjee, James Cragun, Clinton Jenkins, Pamela A. Meyerhofer and other participants in the "Impacts of access to contraception and abortion" session of the 2020 meetings of the Southern Economics Association, Leo Rossi, Mallika Garg, Jacob Burgdorf, Martha Bailey, Brad Hershbein, and Amalia Miller. The author has no conflicts of interest to disclose.

Starting in the 1960s in the US, oral contraceptives ("the pill") provided a safe method to reduce uncertainty about fertility without the high cost of sexual abstinence.[1] Theory suggests that greater certainty about fertility timing would reduce the cost of early human capital accumulation (**???**) and that such investments require foregone earnings early in life that pay off later (**??**).

This article exploits differences in the policy environments faced by young women in different US states to assess evidence for those effects.[2] Young people who wanted the pill needed a prescription, and it was not initially clear that physicians could legally accept the consent of minors (people under age 21 in most states at that time). Around the same time as the introduction of the pill, US states made legal and policy changes that removed barriers for young people who wanted to get the pill without their parents' consent or knowledge. States lowered the age of majority (typically from 21 to 18), allowed minors to consent to treatment if they were mature enough to understand the consequences of their choices (mature minor doctrines), and passed medical consent and family planning laws extending the right to medical or contraceptive choice to minors. I follow **?** (hereafter "Bailey") in calling this state of lower cost of pill access in late adolescence "Early Legal Access" (or "ELA") to the pill.

This research uses a difference-in-differences (DD) framework to estimate age patterns of effects of ELA. State policies turned on ELA for birth cohorts that reached relevant ages only after the policy changed, and the timing of ELA was state-specific. Thus, I compare changes in outcomes between birth cohorts in states where ELA turns on for later cohorts to changes in states where ELA does not turn on. I esti-

[1] The FDA approved Enovid as contraception in 1960. It was in use for regulating menses since 1957.

[2] In discussions of theory, this article defines "women" as people who have or believe they have a non-trivial risk of getting pregnant from sex (or who are perceived by others as having that risk).

mate effects on earnings and labor force participation with person-level observations from the Current Population Survey (CPS).

The results in this article conflict with existing research that estimates effects on the same outcomes (and often with the same data) but with different choices about comparison groups. **?** (hereafter "Bailey") uses CPS data and presents evidence that ELA made women more likely to participate in the labor force in their 20s and 30s but not at later ages. **?** (hereafter "BHM") use data from the National Longitudinal Survey of Young Women (NLSYW) and show evidence that ELA decreased women's earnings in their 20s but increased their earnings starting around age 30, caused young women to invest more in schooling and on-the-job training, and caused women to be more likely to pursue careers that were historically dominated by men. **?** extend estimates to later ages, finding no evidence for effects near retirement age.

This article demonstrates that the effects estimated by Bailey and BHM are the result of a restriction on their regression models that is demonstrably false. They estimate separate effects by age using DD regressions with fixed effects for state and birth cohort but do not allow the state or cohort effects to differ by age. I demonstrate that the restriction on the age×cohort effects is unlikely to hold in the data. I then use a randomization inference procedure to show that randomly-assigned fake ELA produces effects that are indistinguishable from the apparent effects of actual ELA in existing research. Finally, estimates of a version of the model that does not impose the restriction give no evidence of effects.

1 Methods

1.1 Data

1.1.1 CPS

The primary data are person-level observations from the 1963–2009 Annual Social and Economic Supplement to the Current Population Survey (CPS) taken from IPUMS.org (?). Like Bailey, I construct synthetic birth year by subtracting the respondent's age from the survey year.[3] I limit the sample to civilian women aged 14–49 for whom the synthetic birth year is between 1935 and 1960 (inclusive).These persons turned 20 during or a few years surrounding the period when states changed ELA policies. Changes to these cutoffs do not alter estimates meaningfully. I split the sample into seven age groups (14–19, 20–24, 25–29, 30–34, 35–39, 40–44, and 45–49) to estimate effects at different ages. Regressions with CPS data weight observations by the ASEC sample weights from IPUMS.

This sample is nearly identical to Bailey's. The primary differences are that I use a broader range of survey years to observe later cohorts at all ages, that I use a broader range of ages, and that I use age groups with the slightly different boundaries used by BHM. The appendix contains estimates using approximately the same sample as Bailey, and the results are not substantively different. Unfortunately, I am not able to reproduce Bailey's exact sample. The differences may be due to differences between the Unicon files used by Bailey (which Unicon no longer distributes) and the IPUMS files that I use.

The dependent variables are labor force participation and wage and salary earnings

[3] The ASEC data were collected around March (sometimes in February or April), so most respondents would not have had their birthday yet in that sample year. However, the definitions of treatment variables already account for this.

over the previous year deflated by the personal consumption expenditures price index (**?**) with 2000 as the base year. Because of the difficulty in determining labor earnings for the self-employed, I do not include self-employment earnings and exclude from the sample women who report that earnings at the job they worked at the longest over the previous year were not from wages or salary. Estimates of effects on earnings include only observations with positive earnings, but including non-positive earnings does not substantially alter estimates.

1.1.2 ELA

I define ELA policies as those under which any physician in the state (not just those at state-run clinics) could accept the consent of an unmarried (and not emancipated) 19-year-old woman for medical treatment (or specifically for contraception) and under which the woman could legally obtain that pill from a physician or pharmacist.

I define the ELA variable as 1 if an ELA policy was enacted before March in the year that the woman would have reported being age 20 if she were surveyed every year on the same date (in other words, if the law changed before March in the year equal to her synthetic birth year plus 20). March is the cutoff because (most) CPS data are from March. The timing of legal changes comes from **?** and **?** with exact dates added by me when months were unavailable in those sources (and imputed to be July 1 in Mississippi, where I was not able to determine the date). I make a few small corrections and interpret some laws differently. See the online appendix for details. The main body of the article omits extensive discussion because estimates in the appendix using the measures of ELA from Bailey and BHM differ little from from those presented here.

1.1.3 State groups

The CPS groups state of residence in many cases from 1968 through 1976. When using CPS data, I follow Bailey in grouping states into 21 consistently-identifiable groups. For a birth cohort in a state group, ELA is an estimate of the probability of being treated with ELA that equals a weighted average of the states' values of ELA for that cohort, where the weights are the state populations when the cohort was age 20. Population estimates are from the US Census (**????**).

1.1.4 NLSYW

In addition to the CPS data, I follow BHM in using data on various measures of earnings, human capital accumulation, and occupational choice from the National Longitudinal Survey of Young Women (NLSYW). The first wave of the NLSYW surveyed women aged 14 to 24 in 1968. Respondents were interviewed 21 more times through 2003. These women were born from 1943 to 1953. Wherever this study uses NLSYW data, the samples and variables come from replication files from BHM. The sample is limited to the same age range (14–49) and split into the same age groups as the CPS data.[4]

1.2 The unrestricted model

For each age group, I estimate effects with DD regressions with fixed effects for state and birth cohort:

$$Y_{ibs} = \alpha + \delta \times ELA_{bs} + \beta_b + \sigma_s + \varepsilon_{ibs} \tag{1}$$

[4]The regressions in the replication code provided by BHM include women up to age 59. I follow their published description rather than their code and modify their code to drop observations over age 49. Including older ages does not alter conclusions.

where Y_{ibs} is the outcome of person i in birth cohort b in state s, ELA_{bs} is 1 if unmarried women in state s who were born in year b could legally consent to treatment with oral contraceptives in late adolescence and 0 otherwise, β_b is a set of birth year fixed effects, and σ_s is a set of state fixed effects.

This age-specific model (hereafter the "unrestricted model") is equivalent to one that pools observations across age groups and interacts every right-hand-side variable with indicators for age group. The model compares changes in outcomes at one age in a newly-treated state to changes in outcomes at that same age in other states.

Recent research shows that OLS regressions with group and time fixed effects may fail to estimate an economically-meaningful parameter when treatment effects are heterogeneous or vary with time exposed to the treatment (**?**). As a check on results, the appendix estimates average treatment effects on the treated (ATT) following **?**. The estimator compares changes in outcomes between cohorts in newly-treated states to changes in outcomes in states that have not yet enacted ELA and estimates a separate effect for each combination of treatment time (year when the law changed) and year a cohort turned 20. These group-time ATT estimates are then averaged within treatment exposure (how long after the law changed the cohort turned 20) to give event study estimates of effects. The results are not substantively different from the OLS estimates here except that there is weak evidence for negative effects on earnings and labor force participation in people's 40s. However, the estimates do not follow any pattern suggested by theory or past research, and their signs often differ over time exposed to treatment.

The appendix also estimates marginal effects on labor force participation with a probit model, and the results differ little from estimates of the linear probability model here.

1.3 The restricted model

Bailey and BHM estimate a similar model (hereafter the "restricted model") that pools age groups and interacts the treatment variable with age category indicators but does *not* interact the state and time (birth cohort) dummies with those age categories. As a result, the restricted model compares changes in outcomes at one age in a newly-treated state to changes in outcomes at *all* ages in other states. This imposes two important assumptions: (1) pre-treatment differences in the outcome variable (e.g. earnings) between age groups must be equal across states and, more importantly, (2) growth of the outcome in comparison states must not differ by age. These assumptions are about observable age-specific trends in the absence of treatment, so it is straightforward to show that they are inconsistent with the data. Furthermore, violation of either assumption induces bias regardless of whether the difference in trends (or initial levels) by age is exogenous or an effect of previous treatment. For a general treatment of this issue, see **?** and **?**.

The second restriction is violated in the data. The top panel of Figure 1 shows average real earnings in seven age groups for women born from 1935 to 1960 in the CPS data. Later cohorts earn more in middle age than earlier cohorts did, but later cohorts earn about the same as earlier cohorts when they are young (and less than previous cohorts when under age 20). Because the comparison group for the effect of ELA at age 40 in the restricted model is a weighted average of the earnings growth at all ages (including the less-positive growth in earnings at age 20), the estimated effect of ELA at age 40 is biased upward. Similarly, the estimated effect at young ages suffers from negative bias.

The problem is less severe in the case of labor force participation, as seen in the bottom panel of Figure 1. Unlike with earnings, the relationship between growth in labor force participation and age is not monotonic. Labor force participation grew fastest for women in their late 20s and early 30s. Thus, we should expect the restricted model to overestimate effects at those ages and underestimate effects at other (particularly younger) ages.

1.4 Randomization inference

For hypothesis tests and to demonstrate the severity of the bias in estimates of the restricted model, I use a randomization inference (RI) procedure that estimates the models with randomly-assigned fake treatments. I randomly permute the state-specific years of ELA, construct fake ELA with the permuted years, estimate the regressions, and then repeat the process at least 7,200 times.

Because these fake legal changes never actually occurred, they should be unrelated to the outcome variables on average, so the distribution of estimated effects should be centered at 0 for unbiased estimators. The set of RI estimates is an approximation of the sampling distribution of the regression coefficient (under the null hypothesis that there is no effect).

1.5 Other controls

Regressions in Bailey and BHM include additional controls. Bailey includes indicators for year of observation. BHM control for abortion legalization and distance to the nearest abortion clinic. The online appendix presents estimates matching their specifications and with alternative measures and choices of controls. For brevity and to aid in comparison across outcomes and with past research, estimates with CPS data in the main body use baseline models without additional controls. However, due to limitations on data access to protect privacy, estimates with NLSYW data use the specification from BHM.

2 Results

Figure 2 presents estimates of effects of ELA at various ages on annual earnings (top panel) and labor force participation (bottom panel) using CPS data with both the unrestricted model (Eq. 1) and the restricted model used in past research. The thick horizontal tiles are estimates with actual ELA. The violin plots give kernel density estimates of RI effects of fake ELA. Vertical bars are 95% confidence intervals allowing for error correlation within states.

Estimates of the unrestricted model provide no evidence for effects of ELA. The RI estimates are centered at 0, and the estimates with actual ELA are not extreme relative to the RI estimates. Furthermore, the confidence intervals rule out large effects on earnings seen in past research, and the point estimates for labor force participation have opposite signs from those in past work. These estimates do not support theoretical predictions that ELA would cause young women to increase accumulation of human capital and that such investments require foregone earnings early in life that pay off later.

Figure 2 also shows the result of omitting interactions of age group with state and birth cohort effects (the restricted model). Women who could obtain the pill without parental consent in late adolescence appear to earn less in their 20s but thousands more per year after age 30 and appear to participate less in the labor force when young and more after age 25. However, the RI estimates tell a different story. None of the estimates with actual ELA is extreme relative to apparent effects of fake ELA. The RI estimates suggest that the restricted model induces substantial bias and that this bias can explain apparent effects.

It is not simply that the estimates are less precise when including the full set of interactions. Neither the confidence intervals using cluster-robust standard errors (allowing for error correlation at the state level) nor the dispersions of the RI sampling densities based on randomly-permuted laws are much larger with the unrestricted model.

3 Comparison to BHM

BHM estimate effects of ELA on earnings and labor force participation (among many related outcomes) using the restricted model with data from the NLSYW. Like estimates with the restricted model and CPS data in the previous section, they show evidence of lower earnings early in life and higher earnings later. They also find evidence for positive effects on human capital accumulation across multiple outcome measures.

We can again use a randomization inference procedure to show that the restricted model induces substantial bias that can explain these apparent results. Unfortunately, the procedure cannot randomly permute the dates of ELA because geographic (state) identifiers in the NLSYW are limited to use at Census Research Data Centers. Thus, I instead randomly assign each person in the NLSYW to a state. While this is not ideal, randomly assigning persons to states with the CPS produces estimates of the restricted model that are nearly identical to those produced by randomly permuting ELA. Furthermore, adding cohort$\times$age interactions to the restricted model but not including state$\times$age interactions eliminates most of the bias implied by the RI estimates (with CPS data), so preserving state groupings of observations may be unimportant. I report results with equal probability of a person being assigned to each state, but using 1968 state populations as weights produces nearly-identical estimates.

Because I cannot access geographic data with the NLSYW, I must also rely on BHM's published point estimates. These are estimates of coefficients on ELA, which are effects only for those without abortion access. In order to compare the RI estimates to the estimates with actual ELA, all regressions use the definition of ELA from BHM, which cannot time travel to incorporate more recent scholarship.[5] However, the purpose of this exercise is not to estimate effects but to illustrate the bias due to the model restriction.

Figure 3 contains violin plots of RI estimates of effects of ELA with the restricted model using NLSYW data. The horizontal tiles show point estimates from BHM using their base (restricted) specification with abortion controls. Only estimates of effects on career choice (bottom two panels) are extreme relative to apparent effects of fake ELA.[6] In every case, RI estimates suggest that the restricted model induces substantial bias, and, for most outcomes, this bias entirely explains the apparent effects of ELA.

4 Discussion

This article estimates effects on career outcomes of 1960s and 1970s legal changes in the US allowing young people to consent to medical treatments before age 21 and other laws that may have lowered the cost of obtaining oral contraceptives confidentially in late adolescence.

[5]The years of ELA published in their appendix differ in a few cases from the years in their replication package. I use the years from the replication package.

[6]The permutation test estimates still imply much bias in these cases, so the point estimates are still not informative (confidence intervals are not centered at the point estimates). Furthermore, the RI estimates with the NLSYW may understate the bias due to the restrictions because they eliminate one of the sources of bias—differences between states in their pre-treatment differences between earnings by age.

Past research suggests that these legal changes increased earnings, hourly wages, human capital accumulation, and labor force participation for women after their 20s. This article presents evidence that the data do not support those conclusions. Past research uses difference-in-differences regressions to estimate effects at multiple ages but does not allow for age patterns of outcomes to differ between birth cohorts—a restriction that the data suggest is false. The outcome variables grew faster from cohort to cohort at some ages than at other ages, which is exactly what the model restriction assumes is not true. Because the restriction is violated, randomly-assigned fake ELA produces apparent effects that are not centered at 0 and are indistinguishable from estimates using actual ELA.

However, estimates with fake ELA and the unrestricted model are centered at 0, and estimates of the unrestricted model with actual ELA show no evidence of effects. Furthermore, the precision in estimates of the unrestricted model is almost as good as for estimates of the restricted model. These facts suggest that there is no evidence that ELA caused women to earn more or participate more in the labor force at any age. Confidence intervals are also evidence against large effects with magnitudes like those in past research. The results do not constitute strong evidence against *all* effects of ELA. The tests may lack sufficient power to identify effects or measurement error might attenuate estimates. However, these possibilities do not justify using estimates of the restricted model.

Figures

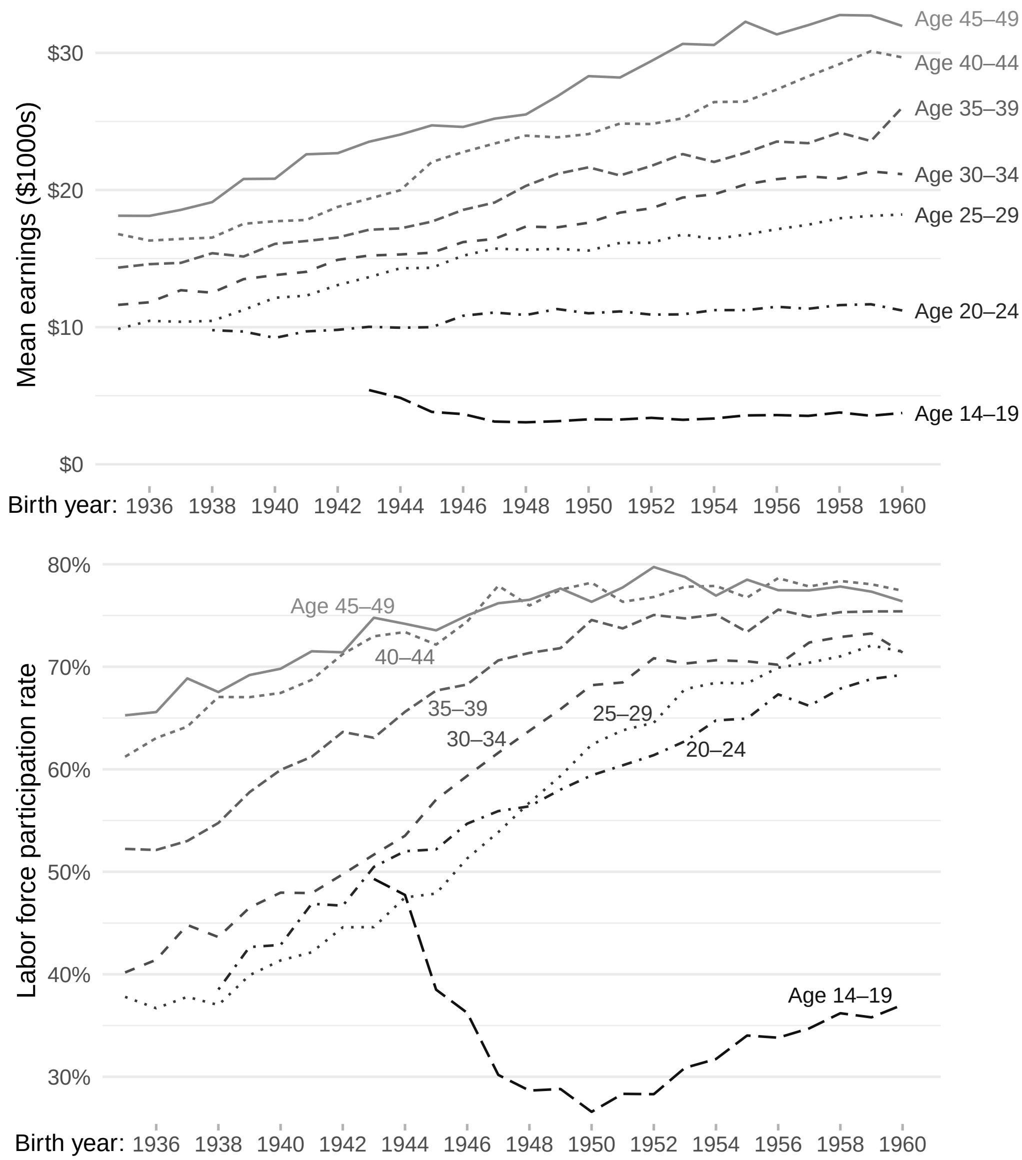

Figure 1: Average earnings and labor force participation by birth cohort and age

Data: women in the 1963–2009 Annual Social and Economic Supplement to the Current Population Survey that were born from 1935 through 1960, were age 14–49 at the time of observation, had positive wage or salary earnings in the last year, and reported that the job they worked at the longest in the previous year paid wages or salary. The data include only wage and salary earnings that are not from self-employment.

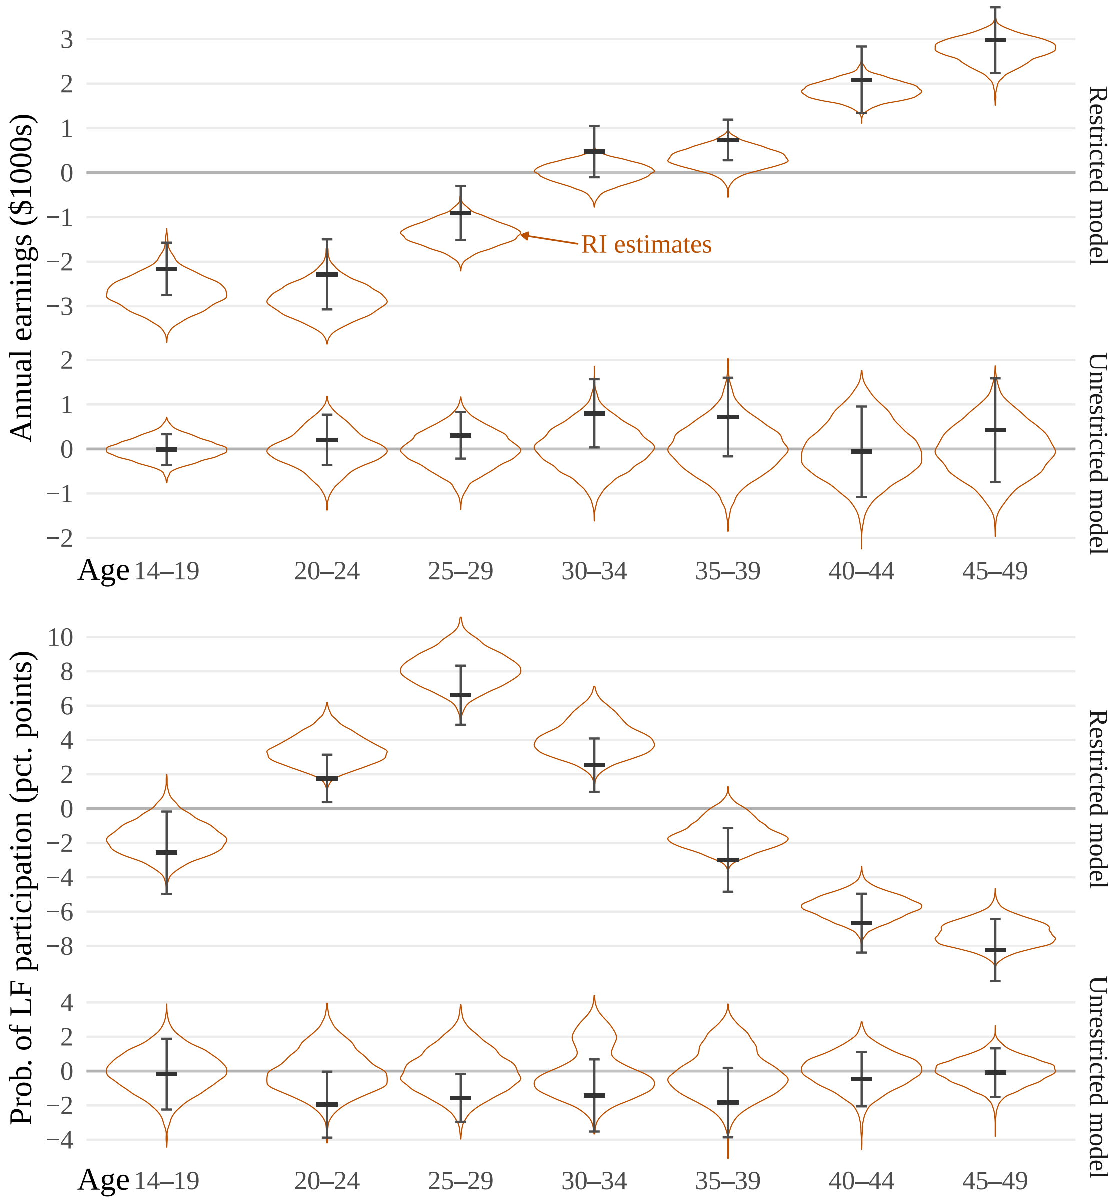

Figure 2: Effects of early legal access to oral contraceptives on annual earnings and labor force participation by age with CPS data

Thick horizontal bars are point estimates. Vertical bars are 95% confidence intervals that correct for heteroskedasticity and error correlation within states. Violin plots are densities of estimates of effects of fake (randomly-permuted) ELA.

Data: 1963–2009 CPS ASEC microdata (from IPUMS) for women aged 14 to 49 who were born from 1935 through 1960. Earnings estimates omit women who did not have positive earnings.

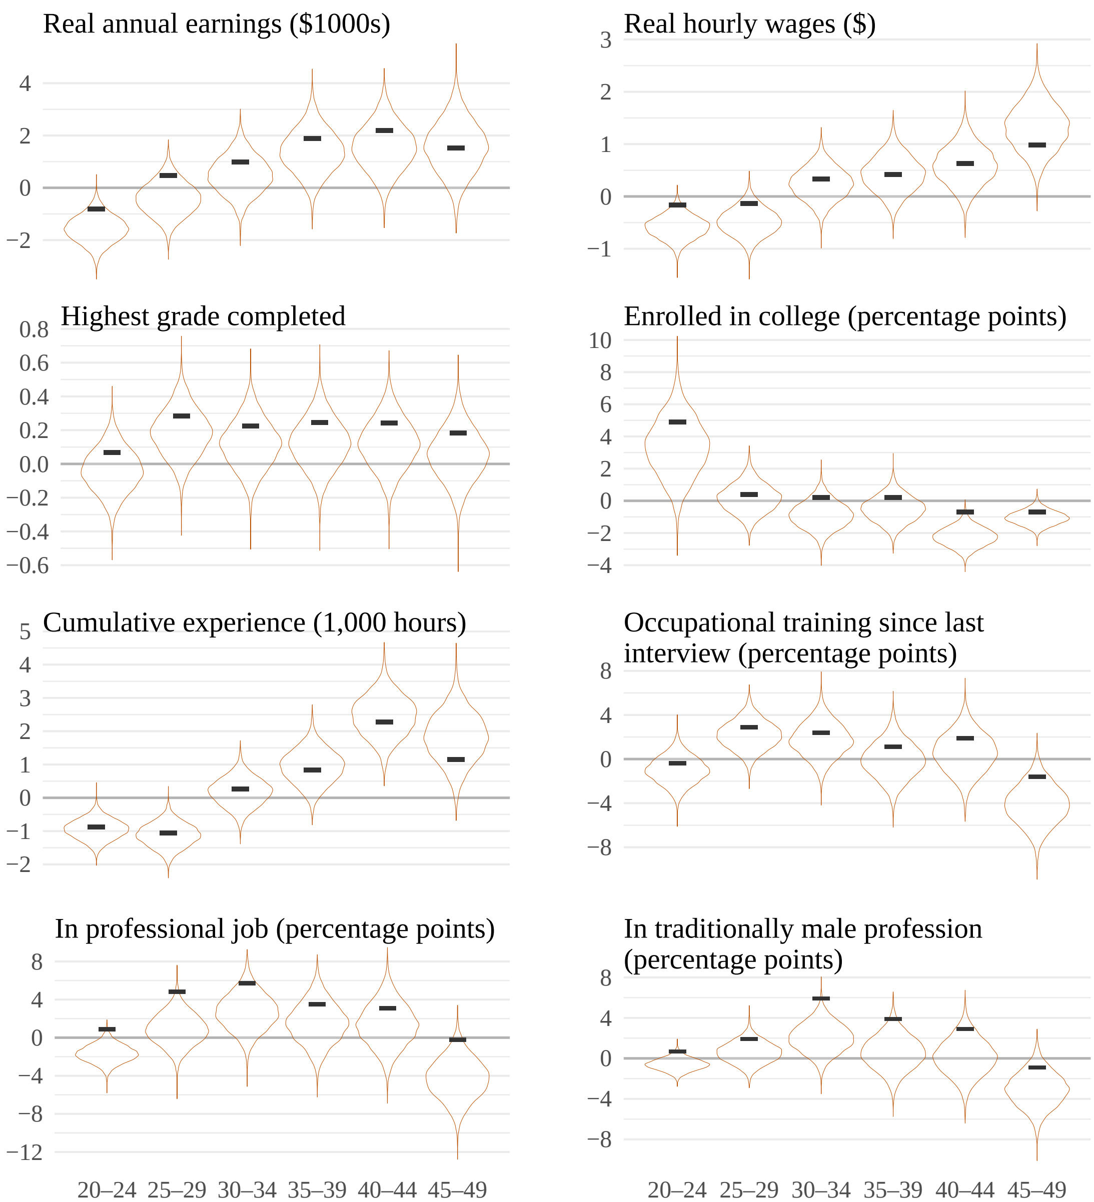

Figure 3: Effects of early legal access to oral contraceptives on various outcomes by age with the restricted model and NLSYW data

Thick horizontal bars are estimates of the restricted model published by BHM. Violin plots are densities of estimates of effects of fake ELA from randomly assigning persons to states. Data: National Longitudinal Survey of Young Women (original cohort). Sample: ages 14–49 with positive earnings (where relevant). Regressions include abortion controls from BHM. For consistency with BHM, effects on binary outcomes are age-specific means of observation-level marginal effects using values of other variables from the data. Effects on earnings, wages, highest grade completed, and cumulative experience are OLS estimates of the coefficient on ELA and so are only for observations without abortion access.