

Public Health Insurance, Labor Supply, and Employment Lock: Effects or Data Artifacts?

ELENA GOUSKOVA*

Populogic Research

July 15, 2016

Using Social Security Administration data on employment, we re-examine the labor supply effect of a large public health insurance disenrollment that took place in Tennessee in 2005. In contrast to previously reported large increases in labor supply obtained with the CPS data, our estimates are small and not statistically insignificant. Furthermore, we find a number of CPS data patterns inconsistent with the "employment lock" interpretation and suggesting data error. Our results narrow the range of existing estimates of the effect of public health insurance on labor supply and highlight importance of considering data error when analyzing survey data.

JEL classification: I1, J22, H75, C81

Keywords: administrative data, difference-in-differences, labor supply, population estimates, public health insurance, survey data

*E-mail: gouskova@populogic.com; mailing address: 3840 Windemere Dr., Ann Arbor, MI 48105; phone: (734)256-4561

1. Introduction

The Affordable Care Act (ACA) of 2010 expanded eligibility for public health insurance (Medicaid) to nearly all low-income adults living in states that opted to participate in the Medicaid expansion. The effect of this policy change on the labor market depends critically on responsiveness of labor supply of those affected by the Medicaid expansion. The major subpopulation benefiting from the expansion are low-income adults without children. Yet empirical studies on labor supply effects of Medicaid among childless adults are few and their results are notably different.¹

The estimated effects of having Medicaid coverage on the probability of being employed range from no effect in Baicker et al. (2013), to a moderate effect of 2 to 10 percentage points in Dague et al. (2014), and to a surprisingly large one of about 60 percentage points in Garthwaite, Gross, and Notowidigdo (2014) (GGN).² How can the estimates of Medicaid effect on labor supply among childless adults differ so much across studies?

GGN suggest labor market conditions as an explanation. The 2008 Oregon and 2009 Wisconsin Medicaid expansions, studied in Baicker et al. (2013) and Dague et al. (2014), took place during the Great Recession when finding

1. Gruber and Madrian (2004) comprehensively review the literature on health insurance and labor supply for other populations, including single mothers, elderly, and married couples. The reviewed studies on Medicaid's effect among low-income single mothers find small or no effects of Medicaid on labor supply.

2. The effect is derived from the estimated increase in employment and decrease in Medicaid enrollment as $\Delta(\text{Probability of being employed})/\Delta(\text{Probability of having Medicaid coverage})$, $0.046/(-0.073)=-0.63$ (see Table II in Garthwaite et. al (2014)). The underlying assumption is that the estimated changes are caused by the disenrollment.

a job could be difficult, let alone a good job with health benefits. Consequently, labor supply appears to be less responsive in those studies. On the other hand, labor supply was not as nearly constrained by job availability around the time of a large public health insurance disenrollment in Tennessee in 2005 — the policy change event investigated in GGN with data from the Current Population Survey (CPS).

In this paper we propose another explanation for the broad range of the labor supply estimates. We examine the 2005 Tennessee disenrollment using the same (as in GGN) difference-in-differences empirical approach but different data. Administrative data from the Social Security Administration (SSA) on employment in combination with Census population estimates show no employment change in Tennessee following the disenrollment. The contrast between no employment change in SSA/Census data and strong employment increases found in CPS data is puzzling.

To shed light on the reason behind this difference, we take a closer look at CPS data. Our replication of GGN using the CPS data yields identical point estimates and confirms that, relative to other Southern states, the share of employed individuals in the population aged 21-62 increased in Tennessee by 2.5 percentage points after 2005. GGN interpret this as evidence of people entering employment in order to secure access to private health insurance after losing Medicaid coverage due to the disenrollment.

However, upon closer inspection, it emerges that the observed increases in the share of workers are driven primarily by declines of Tennessee's civilian

population rather than increases in the number of people being employed, which is clearly at odds with the "employment lock" hypothesis advocated in GGN. Furthermore, we find the labor supply increases estimated in GGN mask dramatic gender differences: remarkably large labor supply increases for men and much smaller and not statistically significant for women. Finally, the alternative employment measure available in the CPS — employment during last year — suggests considerably smaller and not statistically significant labor supply increases in Tennessee following the dis-enrollment. This further evidence suggests another interpretation of GGN's results.

Our explanation is that employment changes estimated in GGN are biased due to data error. During the study period the CPS underwent a number of changes, including changes in the sampling frame, in the weighting procedure, and in the external population controls used in calculation of CPS sample weights. These factors can potentially create shifts in weighted and unweighted sample composition and contribute to the time inconsistency in estimates of labor supply changes.

Thus, we find that estimates of labor supply changes for males are sensitive to controlling for individual characteristics and to shifting of the study period, while estimates for women are not. In particular, when controlling for age, race, education, marital status and health status, the estimates of labor supply increases among males are substantially lower compared to the estimates without these controls.

The rest of the paper is organized as follows. Related literature is reviewed

in Section 2. Section 3 describes three main data sources used in the analysis. Descriptions of the difference-in-differences regression models used in the paper are in Section 4. Section 5 presents results of our replication of key GGN analyses pertaining to estimation of the Medicaid effect on labor supply using CPS and SSA/Census data. In Section 6 we document differences between CPS and SSA/Census data underlying our argument that GGN's estimates of Medicaid effect are biased and data driven. Section 7 discusses two possible sources of bias affecting GGN estimates. Section 8 summarizes and discusses our results.

2. Related Literature

Empirical studies on public health insurance and labor supply of childless adults are few. Research effort in this area was hindered by the lack of data, since this population was not historically eligible for Medicaid coverage. Recently, a number of policy initiatives undertaken in some states and targeted at low-income adults presented an opportunity for such research. Similarly to previous studies focusing on other populations (Gruber and Madrian, 2004), the recent analyses find that availability of public health insurance tends to reduce labor supply of low-income adults. However, the estimated magnitudes are remarkably different, ranging from statistically insignificant to notably large.

Baicker et al. (2013) conduct a study based on the 2008 Oregon Medicaid

expansion in which receipt of Medicaid coverage was determined by a lottery among low-income eligible individuals. The study finds no statistically or economically significant differences in employment and earnings between the lottery participants who received Medicaid eligibility/coverage and who did not.

Dague et al. (2014) estimate employment effects of moderate magnitudes. In 2009, Wisconsin began enrollment of childless adults with household incomes below 200% of the federal poverty line and then within few months had to stop due to lack of funds. Comparing enrollees and eligible applicants on a waiting list, Dague et al. (2014) show that availability of public health insurance leads to statistically significant reduction in labor supply. Their estimated effect ranges from 2 to 10 percentage points depending upon estimation method the authors used.

Among the available estimates, those in Garthwaite, Gross, and Notowidigdo (2014) are conspicuously large. They study a Tennessee's Medicaid disenrollment that, in 2005, left about 170,000 Tennesseans, disproportionately childless adults, without Medicaid coverage. Using the March CPS with difference-in-differences models, GGN provide estimates of the disenrollment effect for two outcomes: public health insurance coverage and employment, measured as share of people having public coverage and being employed correspondingly.

Comparing the outcomes in Tennessee to other Southern states with a difference-in-differences regression model, they find a 4.6 percentage point

decline in public coverage and 2.5 percentage point increase in employment. Additionally, splitting the sample further by childless status and using a triple-difference regression model, GGN estimate a 7.3 percentage points decline in public health insurance coverage and a 4.6 percentage points increase in employment for childless adults in Tennessee. GGN also find a substantial response heterogeneity among childless adults, with those who are older (40-64), less healthy, and less educated being more likely to become employed.

GGN consider the observed changes in the outcome variables as caused by the disenrollment. In particular, the results are interpreted as evidence of people entering employment in order to secure access to private health insurance after losing Medicaid coverage due to the disenrollment. The triple-difference analysis lends support to this interpretation by finding that employment increases were concentrated among childless adults. The response heterogeneity is rationalized by preference heterogeneity, with the older and less healthy individuals placing higher value on health insurance.

In accordance with this interpretation, the public health coverage effect on labor supply is 54 percentage points ($2.5/4.6$), and 63 ($4.6/7.3$) percentage points for childless adults. This means that more than a half of those who lost public coverage entered employment — a notably larger effect compared to other studies.

However, if the observed employment changes do not reflect the causal effect of the disenrollment then GGN’s interpretation and their estimates are no longer valid. In the paper, we argue that the observed employment

increases in Tennessee reflect CPS data error.

3. Data

We employ three main data sources: the Current Population Survey, data from the Social Security Administration on total employment, and the intercensal population estimates.

The CPS is a survey of about 55,000 households conducted monthly by the Census Bureau. The official source of employment statistics for the civilian non-institutional population of the U.S., the survey collects information on demographics and employment situation. We use data from the CPS's Annual Social and Economic Supplement administered during the month of March. In the paper, our focus is on employment.

Following GGN, our main employment measure is derived from the questions on employment status during the survey reference week. An individual is considered employed if he/she is reported 'at work' during the reference week. We refer to this measure as current employment. In addition to this measure, CPS collects information about employment during last year: "Did you work at a job or business at any time during 20...?" For the last year employment we use data from 2001-2008 March CPS to reflect information for the 2000-2007 period.

As an alternative source of employment data we use SSA administrative records. In particular, we use tabulations on the number of persons with

Social Security taxable earnings.³ The SSA tabulations report numbers of workers by state (including District of Columbia), sex, and age intervals: under 20, 20-29, 30-39, 40-49, 50-59, 60-61, 62-64, 65-69, 70 or older.⁴ Conceptually, compared to two CPS employment measures, SSA employment is closer to the CPS last year employment. Covering residential population, SSA employment is, however, a potentially more inclusive concept.

In order to obtain employment rates, we combine the SSA information on the number of workers with the intercensal population estimates provided by the U.S. Census Bureau. The Census Bureau creates two types of population estimates: postcensal and intercensal. Postcensal estimates are calculated each year, by extrapolating estimates since the latest census. Intercensal estimates are postcensal estimates adjusted to smooth the transition between consecutive censuses. They are constructed once every ten years. In this paper we use the 2000-2009 intercensal estimates by single year of age, sex, state and will refer to these estimates as Census population estimates.⁵

There are two sets of data comparability issues that preclude exact repli-

3. The 2000-2007 SSA reports, titled 'Earnings and Employment Data for Workers Covered Under Social Security and Medicare, by State and County,' are available at http://www.ssa.gov/policy/docs/statcomps/eedata_sc/index.html.

4. The numbers of workers in the SSA reports are not the true population values, but estimates derived from a 1 percent sample of W-2 wage reports and from IRS Schedule SE of Form 1040. The sample data are inflated to correspond to SSA estimates for U.S. totals. For comparison, the CPS sample of workers constitutes about 0.07 percent of all workers.

5. For more information on the intercensal estimates see <http://www.census.gov/popest/data/intercensal/index.html> The specific data file we use was obtained from the following location: <http://www.census.gov/popest/data/intercensal/state/files/ST-EST00INT-AGESEX.csv>

cation of GGN using SSA/Census data. One set of differences is related to the composition of analytical sample. Analytical sample in GGN consists of civilian non-institutionalized adults between the ages 21 and 64 with less than a college degree. Our SSA/Census analytical sample includes residential population ages 20 to 64. For the 20-64 age group, the main difference between civilian and residential population are members of the armed forces residing in the US.

Another difference is that the analytical sample in GGN excludes those with more than a college degree. Since SSA/Census data come aggregated by year, state, age, and gender, we cannot select observations based on education. As we show below, some important results in GGN continue to hold in a sample including all educational groups.

For some analyses employing CPS data we use a sample similar to the SSA/Census sample, i.e. individuals between the ages 20 and 64 and including all education groups. In the text, we refer to this sample as full sample. To the analytical sample used in the original GGN analysis we refer as GGN sample.

A more serious comparability issue between CPS and SSA/Census data is that the SSA/Census does not allow separation of childless adults from those living in households with children under 18. To circumvent the issue we exploit the fact of labor supply response heterogeneity found in GGN. In particular, instead of splitting the sample by childless status, we split it by age and compare 20 to 39 year olds with 40 to 64 year olds. In the text below,

we discuss our identification strategy and show that both types of grouping in CPS data produce comparable triple-difference estimates. In particular, the estimates for individuals who are 40 to 64 year old are quite similar to those for childless adults.

Figure 1 demonstrates how total employment using SSA and CPS data compare. It plots SSA and CPS total employment estimates for Tennessee (panel A) and other Southern states (Panel B) for the period between 2000 and 2010. In addition, Figure 1 includes employment estimates from the Local Area Unemployment Statistics (LAUS) data, which we use later in the paper.

As can be seen, the closest to each other are SSA total employment, CPS total last year employment (using the full sample), and LAUS employment measures. These measures also provide the highest levels of employment. CPS current year employment is about 10% lower in the full sample, and considerably lower in the GGN sample. Overall, all employment measures follow similar time trends, and it is particularly so for the combined sample of other Southern states. As expected, for Tennessee alone, CPS employment measures exhibit more volatility.

4. Difference-in-differences regression models

We follow GGN and use difference-in-differences (DD) and triple-difference (DDD) models to estimate the effect of the disenrollment on labor supply.

DD compares before and after the disenrollment changes in outcomes of interest in Tennessee to the corresponding changes in the other Southern states, serving as a control group. The following regression implements the approach:

$$(1) \quad y_{st} = \alpha_s + \delta_t + \beta \cdot I\{s = TN\} \cdot I\{t \geq 2006\} + \varepsilon_{st}$$

where y_{st} is outcome of interest such as employment share for state s in year t , α_s and δ_t are state and year fixed effects, and ε_{st} is an error term. Parameter β measures the effect of the disenrollment under the identifying assumption that the dependent variable in Tennessee and other Southern states would have evolved similarly in absence of the disenrollment.

GGN note that this assumption may not hold in the data. Even prior to the disenrollment, Tennessee appears to follow an employment trend different from other Southern states. Furthermore, the disenrollment policy targeted only a segment of the Tennessee population, childless adults, and therefore its effect should be felt only for that subpopulation. These considerations lead to the following triple-difference regression model:

$$(2) \quad y_{ist} = \gamma_i \cdot \alpha_s + \gamma_i \cdot \delta_t + \alpha_s \cdot \delta_t + \beta \cdot I\{i = 1\} \cdot I\{s = TN\} \cdot I\{t \geq 2006\} + \varepsilon_{ist}$$

where, in addition to the description above, $i \in 0, 1$ indexes subpopulation

and $i = 1$ is the group affected by the policy change (childless adults in GGN analysis), and ε_{ist} is an error term.

Here β measures the effect of the disenrollment on y for childless adults under the condition that the difference in y between childless adults and parents in Tennessee and other Southern states would have followed the same trend in absence of the disenrollment.⁶ Compared to DD model (1), triple-difference model (2) relaxes the requirement of the outcome variable, y , following the same trend in the treatment and control groups prior to the intervention. In model (2), it is replaced with a parallel requirement for the difference in y between childless adults and parents.

As discussed above, a replication of the triple-difference analysis that uses SSA/Census data is not possible because of non-separability of childless adults and parents in SSA/Census data. While it is not feasible to do the triple-difference analysis with grouping by childless status using SSA/Census data, it is still possible to estimate β using grouping by age. The key reason is the response heterogeneity among treated found in GGN. In particular, GGN show that employment rates increases were concentrated among childless adults who are older. We exploit this fact of heterogeneity by age and,

6. The difference between subgroups can be expressed as follows: $y_{1st} - y_{0st} = (\gamma_1 - \gamma_0) \cdot \alpha_s + (\gamma_1 - \gamma_0) \cdot \delta_t + \beta \cdot I\{i = 1\} \cdot I\{s = TN\} \cdot I\{t \geq 2006\} + \varepsilon_{1st} - \varepsilon_{0st}$. Thus, the deterministic difference between subgroups has a state specific intercept and a common (across states in absence of policy change) time trend. Suppose that $\gamma_1 < \gamma_0$, for example labor force participation of childless adults is generally lower than that of parents. Then positive β would indicate narrowing of the difference between childless adults and parents in Tennessee following the dis-enrollment as compared to other Southern states. Since parents were not affected by the policy, positive β would also mean an increase in labor supply of childless adults.

instead of comparing childless adults to parents, we compare 40-64 year olds to 20-39 year olds.

Our approach compares modified treatment and control groups. The modified groups are obtained by swapping young childless adults in the treatment group with old parents in the control group. With both groups not responding to treatment and being similar in size (about 18% each), such exchange should have little effect on the estimate of β . Thus, the identifying assumption in the triple-difference analysis involving grouping by age is analogous to that involving grouping by childless status. We additionally assume that the dis-enrollment had not affected the labor supply decision of young childless adults. The later assumption is justified by the results in GGN analysis.

5. Analysis using CPS and SSA/census data

In this section we use data from the CPS and SSA/Census to replicate key analyses in GGN focusing on the employment effect of the disenrollment. For ease of comparison, results based on CPS and SSA/Census data are presented side by side.

Panel A of Figure 2 reproduces Figure III Panel A in GGN and plots the share of people employed in Tennessee and other Southern states for the 2000-2007 period, with data being aggregated biannually. It shows that after a period of sluggish decline, Tennessee experienced a dramatic increase

in employment rates after 2005. This upturn in employment rates is in a strong contrast to the employment trend in other Southern states where employment rates remain practically flat after 2005.

The next panel of Figure 2 replicates the same analysis using SSA/Census data. Here, clearly different from the previous panel, employment rates in Tennessee are virtually parallel to those in the other Southern states, with no indication of an increase following the disenrollment.

Another noteworthy difference between the plots is the relative position of employment rates in Tennessee and other Southern states. In particular, when using SSA/Census data, employment rates in Tennessee are higher compared to the rest of the South, but when using CPS data they are lower for the most of the period between 2000 and 2007, and particularly so for the early and middle part of the period.

Next we turn to the results of estimation of the difference-in-differences model (1) and the triple-difference model (2), which are reported in Table 1 in panels A and B, respectively. The table has six columns. Column 1 reports GGN's estimates and Column 2 — our replicas using CPS data. As can be seen, the point estimates in both columns are exactly the same.⁷ The next three Columns provide alternative CPS estimates. Column 3 shows GGN estimate using a sample unrestricted by educational level. Columns 4 and 5 report results for the last year employment measure in the GGN and full

7. To keep things simple, we do not adjust standard errors for inter-temporal correlation and instead use OLS standard errors. The OLS standard errors are just slightly larger than the modified bootstrap errors employed in GGN.

samples, respectively. Finally, Column 6 reports SSA/Census estimates.

DD estimates in Panel A in Table 1 echo the graphical evidence above. In particular, using the CPS current employment, Columns 1 and 2 indicate that after 2005, Tennessee experienced an increase of 2.5 percentage points in the share of individuals currently employed, which is statistically different from that of the other Southern states. In contrast, the SSA/Census estimate is small, negative, and not statistically significant (Column 6, Table 1).

Importantly, other DD estimates using CPS data, reported in Columns 3 through 5, provide little support for the original DD estimate as well. Thus, in the sample with all educational levels, the estimate is almost half the size and not statistically significant (see Column 3 in Table 1). Furthermore, the last year employment measure in Columns 4 and 5 show no employment increases following the dis-enrollment. The corresponding DD estimates of 0.012 and 0.006 are not statistically significant.

What makes GGN paper particularly convincing is the estimates of the triple-difference model (2). These estimates suggest that the Tennessee employment increases were concentrated among childless adults. This is exactly what one would expect given that this population was disproportionately affected by the disenrollment. Using the current employment measure, GGN's coefficient of the interaction term of being childless adult and living in Tennessee after 2005 is 4.6 percentage points and statistically significant at the 5% significance level (Panel B, Table 1, Column 1). Our replication yields the same point estimate (Column 2). In the sample with all education groups,

DDD estimate is 0.041 and statistically significant at 10% significance level.

Yet, when we use the last year employment measure, the DDD estimate is not nearly as strong. DDD estimates of 0.023 and 0.006 in the GGN and full samples respectively are both not statistically significant (Table 1, Columns 4 and 5).

Due to non-separability of childless adults and parents in SSA/Census data, there is no SSA/Census estimate in the last Column of Table 1 (Panel B) corresponding to the triple-difference model that compares childless adults and parents. In order to estimate β with SSA/Census data, we run triple-difference model (2) that compares the older and younger age groups corresponding to the modified treatment and control groups. The approach exploits the idea that exchanging non-responding young childless adults from the treatment group with older parents from the control group should not affect the DDD estimate.

Indeed our estimation results using grouping by childless status and by age are quite similar. Thus, for the current employment measure, Column 2 in Table 1, Panel C) shows that the DDD coefficient for 40 to 64 year olds is almost the same as that for childless adults (0.046 vs 0.048). Furthermore, the current employment DDD estimate is robust; it remains large (4.4 percentage points) and statistically significant when respondents of all educational levels are included in the sample (Column 3).

However, despite the strong current employment DDD estimates, the last year employment measure yields considerably smaller DDD estimates of

0.025 (p-value is 0.179) and of 0.032 (p-value is 0.076) in the GGN and full samples, respectively (see Columns 4 and 5).

Importantly, grouping by age allows us to estimate the effect of the disenrollment using SSA/Census data. The last column in Table (2) (Panel C) reports the results. As can be seen, the SSA/Census triple-difference estimate of 0.003 is small and not statistically significant.

6. Differences between CPS and SSA/Census data

6.1. *Different population trends in CPS and Census data*

The results above demonstrate that while SSA/Census data show no increase in employment following the disenrollment, CPS data produce a mixed picture. On the one hand, the current employment measure used by GGN shows large employment increases. On the other, the last year employment measure suggests small or none. In the rest of the paper we argue that the DD and DDD estimates obtained in GGN are likely to be biased due to data error. Since it is the CPS current employment measure that exhibits large population increases after 2005 in Tennessee, in what follows we primarily focus on this measure.

We begin by documenting a puzzling data pattern that exists in CPS data and does not in SSA/Census data. Figure 3 displays the pattern. The figure has three panels and uses CPS data. The first panel plots the share of people employed for Tennessee and other Southern states. It repeats an earlier plot

and provides context for plots in Panels B and C. These two panels retain the structure of the figure in Panel A but, instead of the employment share, plot its components: total employment and population.⁸

Interestingly, there appears no visible break in the trend for total employment in Tennessee compared to other states following the disenrollment (see Panel B in Figure 3). Both total employment measures appear to follow their respective pre-intervention trends. Furthermore, Panel C shows a clear decline in Tennessee’s civilian population after 2005, and no similar decline in the combined group of other Southern states (see Panel C in Figure 3). Hardly consistent with the version of the disenrollment effect advocated in GGN, the graphical evidence above is consistent with presence of data error.

Figure 4 shows analogous graphs using SSA/Census data. Two things to note here. First, according to Census population estimates, Tennessee did not experience population declines after 2005. Second, relative to other Southern states, Tennessee’s employment rates are lower in the CPS data and higher in the SSA/Census data, especially in the early through middle part of the study period. This suggests that Tennessee’s population might be overestimated (relative to the other states) in CPS data in the earlier years.

This indeed appears to be the case. Panel A in Figure 5 plots the ratio of CPS to Census population estimates for Tennessee and other Southern

8. For each state s and year t , these components are obtained for the GGN estimation sample as follows: Total Employment = $\sum_{k=1}^N I(\text{employed}_i = 1)w_i$, Population = $\sum_{k=1}^N w_i$, where N is a number of respondents in the CPS sample for state s in year t , w_i is CPS sample weight for individual i .

states. For the combined population of other Southern states, the CPS population (civilian population) is 4 to 5 percent lower than Census population (residential population). The difference is slightly larger in years 2000-2002, but generally stable throughout the whole period. The CPS-Census differences for Tennessee alone are considerably more volatile, suggesting presence of noise in the population estimates.

Importantly, the CPS-Census differences for Tennessee do not appear random, evolving from relatively small (1-2 percent) to large (about 6-7 percent) over the study period. Relative to other Southern states, the CPS goes from overestimation to underestimation of the Tennessee's civilian population in the years between 2000 and 2007. If this also leads to underestimation in the earlier years and to overestimation in the later years of Tennessee's employment rates, then the difference-in-differences estimate in (1) is likely to pick up this Tennessee-time specific data pattern.

Panels B and C of Figure 5 show CPS-Census population estimates differences separately for the younger and older groups. As can be seen, patterns of the differences are not uniform. For the younger group, the CPS population decreases more or less monotonic throughout the period (year 2004 appears out of line with the general trend). For the older group, the CPS population first increases and then decreases. It increases steadily until 2004, at which point it exceeds the residential population by about 4 percent. After that it rapidly goes down, with the decline being particularly precipitous between 2004 and 2006.

6.2. Quantifying CPS and SSA/Census differences

In this subsection we quantify how much of the observed in the CPS Tennessee’s employment increases in the second half of the study period can be attributed to changes in total employment and how much to changes in population. To do so we use the following identity:

$$\Delta \ln(\text{Share Employed}) = \Delta \ln(\text{Employment}) - \Delta \ln(\text{Population})$$

To estimate these changes, we fit a piecewise linear trend model of the following form:

$$(3) \quad y_{jst} = \begin{cases} \alpha_{js} + \beta_{js} \cdot t + \varepsilon_{jst} & t \leq 2004 \\ (\alpha_{js} - \Delta\beta_{js} \cdot 2004) + (\beta_{js} + \Delta\beta_{js}) \cdot t + \varepsilon_{jst} & t > 2004 \end{cases}$$

where y_{jst} is natural logarithm of one of the three variables: share of people employed, total employment, and population; j indicates age group, 20-39 or 40-64; s is a state index; and t stands for year. Year 2004 has been used as a break point, as it appears to fit CPS data better.

Figure 6 provides an illustration of estimation results of (3) for Tennessee. Six panels depict three dependent variables for two age groups. In particular, each panel juxtaposes four lines corresponding to observed and predicted values of the dependent variable obtained with CPS and SSA/Census data. We focus on changes in growth rates, $\Delta\beta$, which are reported in Table 2.

In line with the previous results, CPS data show that, the growth rate for

the share of people employed in Tennessee increased statistically significantly in the second period for both groups. For the younger group, $\Delta\beta$ is 0.041 (Columns 1). For the older group, the increase is larger and equal to 0.052 (Columns 4). Similarly in line with the evidence above, SSA/Census data show no such increases (0.009 for the younger group and -0.002 for the older group; both are not statistically significant) (Columns 2 and 5). Furthermore, CPS-SSA/Census differences in $\Delta\beta$ are statistically significant (Columns 3 and 6).

Turning to how total employment and population changes contribute to increases in employment rates in the CPS, we see that, for the younger group, the change of 0.041 results from a combination of an increase in total employment growth (0.030) and a decrease in population growth (-0.010). Both estimates are not statistically significant.

For the older group, $\Delta\beta$ of 0.052 in CPS data comes almost exclusively from a population growth decrease of -0.051, which is statistically significant at the 1 percent level. The total employment growth change is small (0.001) and not statistically significant. For comparison, in SSA/Census data, the estimated change in growth of total employment and of population are both small and not statistically significant.

The evidence can be summarized as follows. First of all, the increases in the share of employed people appear to be CPS specific. Second, the increase is mainly driven by a population decline, especially in the older group.

6.3. Which population estimates are correct?

Assuming that CPS estimates are correct and Tennessee's population did indeed decline while total employment did not, the mechanism triggered by the dis-enrollment must be different from the 'job lock' hypothesis. It is possible to imagine that peoples' residential decision is influenced by generosity of social insurance programs in the state. In such case, it is possible that some people might have decided to move out of Tennessee in response to the dis-enrollment.

This line of thinking assumes that CPS population estimates are correct and, consequently, Census intercensal estimates are not. The intercensal estimates are smoothed by construction may possibly miss some temporary population fluctuation. We checked the robustness of our results by using alternative source of population estimates — postcensal population estimates of vintages 2007 and 2008. This has not changed our results in any significant way, casting further doubts about correctness of CPS population estimates.

While biased CPS population estimates do not necessary mean biased employment to population ratio, it is clear, however, that incorrect estimates undermine some important evidence provided in GGN to support their case. Consider Figure VIII in GGN which shows a rapid monthly growth in Tennessee's employment rates around the time of the dis-enrollment and contrasts them with non-increasing rates in other Southern states for the period between 2004 and 2007. The figure uses two types of data: monthly data on

total employment from the LAUS and population estimates from the CPS.⁹

To see the effect of population estimates, we first reproduce the figure, as it was done in GGN, with the CPS population estimates (see Figure 7, Panel A), and then with Census population estimates (see Figure 7, Panel B).¹⁰ The difference between the two panels is dramatic. Panel A shows a vigorously increasing monthly employment rate for Tennessee and practically flat one for the rest of the South. However, with Census population estimates, employment rates in Tennessee and other Southern states move in tandem throughout the period (Panel B Figure 7). With Census population estimates, Tennessee's employment rates are obviously much flatter, with higher values in the earlier years and lower values in the later years, as compared to Tennessee's rates using the CPS population estimates.

7. Possible bias sources

7.1. *CPS sample weights*

In this section we probe further and consider two possible sources of bias in the employment to population ratio: incorrect CPS sample weights and shifting sample composition.

Because CPS population estimates depend on sample weights, the above

9. LAUS data are available from <http://download.bls.gov/pub/time.series/la/>

10. In both cases, monthly population estimates are linear interpolations passing through two points: the March 2004 and the March 2007 for the CPS estimates; and the July 2004 and the July 2007 for the intercensal estimates.

analysis points to a possibility that CPS sample weights are biased. This does not necessarily mean that the employment to population ratio is biased as well. Any bias in the denominator due to weights is likely to be reduced by a bias in the numerator, which is also a function of weights. Yet, in CPS data the sample weights are correlated with labor market outcomes (particularly for males) and the possibility that weights might contribute to a bias in the employment to population ratio should be considered.¹¹

To check the effect of weights, we have explored estimation without weights and re-scaling of weights. It appears that sample weights tend to increase the labor supply estimates in Tennessee after 2005. In particular, compared to weighted DD and DDD estimates, unweighted estimates are smaller, 0.025 vs 0.016 for DD estimate and 0.046 vs 0.044 for DDD estimate (see Table 3 Column 2, Panels A and B). Furthermore, the unweighted DD estimate is no longer statistically significant, while unweighted DDD estimate remains statistically significant at the 5% level.

Naturally, unweighted estimates should be taken with a grain of salt. CPS data are intended to be used with sample weights, designed to account for under-coverage and non-response.¹² Furthermore, to reduce sampling variability of estimates, CPS weights are additionally adjusted to conform

11. For males, the correlation between labor market outcomes and weights is generally negative, suggesting that those who are less likely to work are also less likely to respond to the CPS.

12. The CPS documentation states: "Unweighted counts can be very misleading and should not be used in demographic or labor force analysis" (see <http://www.census.gov/cps/methodology/summarystats.html>).

CPS counts along major demographic dimensions, such as age, gender, race, and Hispanic origin, to independent national and state population controls provided by the U.S. Census Bureau.¹³

To get a better idea of the effect of CPS weights on estimates, we have explored a CPS weights adjustment that ensures that CPS adjusted totals are in line with postcensal estimates for age-gender-state-year cells. We find that such adjustment of CPS weights decreases the DDD estimate from 0.046 to 0.037, which provides further support for the idea that CPS weights account for some of GGN’s results (see Table 3 Column 3, Panel B).

7.2. Sample composition

Another potential and likely related source of bias is the changing composition of the CPS sample during the study period. The CPS sample is based on the rotation system. About half of the the CPS March sample are new respondents while the other half are those who participated in the survey one year prior (U.S. Bureau of Labor Statistics (2006)). Apart from variability due to chance, sample composition changes may also result from

13. Time comparability of CPS labor force estimates can potentially be affected by periodic adjustments of population controls used in benchmarking of CPS weights to reflect the results of Censuses, updated estimates of international migration, and updated vital statistics information. For example, population controls reflecting information of Census 2000 were introduced in January 2003. Such adjustment resulted in increases of about a million to the civilian noninstitutional population and of about 600,000 to the civilian labor force. The overall unemployment rate and other ratios were reported to be not substantially affected by the change (Bowler at al. (2004). Another potential reason for CPS estimates time inconsistency is periodic changes in the weights calculation procedure (Bowler at al. (2004).

the redesign of the CPS sample frame that took place between April 2004 and July 2005.¹⁴

We find that the CPS sample experienced a rise in the fraction of respondents who are white, not married, and more educated in Tennessee following 2004. Some of the compositional shifts appear to affect childless adults and parent differentially. Panels A and C in Figure 8 show that the sample of males without children in Tennessee experienced a noticeable increase from about 14% to 18% in the fraction of those with college degree between 2005 and 2008 and a concurrent decrease from about 18% to 12% in the fraction of those having a work limiting disability. Similar shifts are observed among childless women but one year later (see Panel B and C in Figure 8).

It is important that these patterns differ not only by childless status but by gender as well. This can help shed light on a peculiar feature of the DDD estimates obtained in GGN. In addition to the heterogeneity patterns documented in GGG, the original DDD estimate of 0.046 masks dramatic differences by gender. Thus, for males, the estimate is surprisingly large (in the range between 0.060 and 0.072 depending on estimation method) and statistically significant at the 5% level, while for females it is much smaller (0.25-0.30) and not statistically significant. The stronger preference

14. Shoemaker (2004) describes the scope of the changes and provides details on transition between the sample based on Census 1990 and the sample based on Census 2000. The new sample is similar to old one in terms of size but includes some relocation among the states. The new sample consists of continuing sample areas (90%) and new sample areas (10%). The new sample also incorporates the supplemental sample from the State Children's Health Insurance Program (SCHIP) survey.

for health insurance is a reasonable explanation for the response heterogeneity by age, health and education found in GGN. Yet, it does not seem as convincing in the case of heterogeneity by gender.

We explore the effect of compositional changes with two approaches. One is to control for sample composition in regression framework. Another is to vary the study period. We consider shifting the study period by one year to 2001-2008, which makes it more symmetric around the time of the dis-enrollment and yet avoids the years strongly impacted by the Great Recession.

Table **3** reports the results. Columns 4 through 6 show DD and DDD estimates from the analyses run on state-year means of residuals obtained from an individual-level regression of the current employment indicator on a set of individual controls (see Panels A and B in Table **3**). These controls include age, education, race, marital status, and an indicator of having a work-limiting disability, all interacted with gender. Parallel estimates for the 2001-2008 period are reported in Panel C.

First observation is that compared to the estimates without controls for sample composition, the estimates based on residuals are considerably lower. Thus, the raw DDD estimate of 0.046 declines to 0.038 when using individual controls, but remains statistically significant at the 10% level.¹⁵ Second,

15. This finding is in contrast to GGN who report that their attempt to control for demographic composition, reported in Appendix Table A6, had not affected the magnitude of the triple-difference estimate. Compared to GGN who control for age, education and gender, we additionally control for race, marital status and health status.

the shift of the study period by one year to 2001-2008 further reduces DDD estimates by about 10 percentage points and renders all triple-difference estimates adjusted for individual characteristics not statistically significant.

Importantly, these reductions operate almost exclusively in the sample of males. For males and females separately, Table 4 reports the DDD estimate from three individual-level regressions: without individual controls, with individual controls, and with individual controls and the study period shift. For males, the initial DDD estimate of 0.060 (p-value is 0.026) goes to 0.046 (p-value is 0.024) with individual controls, to 0.026 (p-value is 0.148) with individual controls and the study period shift. Interestingly, for women, individual controls and study period shifts have little effect on the DDD estimate, which are in the 0.024-0.030 range and not statistically significant.

To summarize, the evidence suggests that sample composition is responsible for much of GGN's results. In particular, we find that controlling for individual characteristics substantially reduces the DDD estimate for males. We also find that CPS sampling weights do not appear to mitigate the effect of compositional shifts and may even make the matter worse. The effectiveness of weights is likely to be limited by the division of the CPS sample into fine subsamples in GGN analysis and possibly by changes/errors in CPS weights calculation.

8. Conclusions

Using the CPS data, Garthwaite, Gross, and Notowidigdo (2014) find large employment increases in Tennessee following the 2005 public health insurance disenrollment, which the authors interpret as evidence of people finding employment in order to secure health insurance coverage lost as a result of the disenrollment.

Using SSA employment data and Census population estimates, we find no such employment increases. Our estimate is small and not statistically significant.

To reconcile the differences we take a closer look at CPS data and discover several facts that contradict GGN's interpretation. First of all, employment increases documented in GGN are driven primarily by decreases in civilian population rather than increases in total employment. Second, we find that employment increases documented in GGN are concentrated among males, with the estimates for females being not statistically significant. Finally, using an alternative measure of employment available in the CPS — employment during last year — we find no statistically significant result.

Our interpretation is that employment increases estimated in GGN are biased due to data error. Comparison of CPS and Census data revealed that while CPS estimates show a population decline in Tennessee after 2005, the corresponding Census estimates show no such decline. Furthermore, the CPS-intercensal population differences are not uniform across age groups —

more pronounced in the older group. Finally, CPS-intercensal population differences appear not random. Relative to other states, the CPS population estimates go from overestimation to underestimation over the period between 2000 and 2007.

This suggests two possible sources of bias: error in the CPS sample weights and shifts in CPS sample composition. Exploration of these options revealed importance of both for GGN's results as well as their interconnection. In particular, the CPS weight adjustment, which forces weighted CPS distribution to match Census distribution for age-gender-state cells, leads to smaller estimates in GGN analysis. Even stronger reduction is achieved with controlling for individual characteristics, such as age, education, race, marital status, and health status. Importantly, the reduction affects primarily estimates of males, the subpopulation with unusually large estimated labor supply increases.

We use the term 'data error' inclusively to combine factors such as sampling error (defined under the condition of unchanging repeated sampling), non-sampling error (it is not clear why CPS population estimates differ from Census estimates), and sample/process changes affecting the CPS during the period (e.g. transition from Census 1990 to Census 2000 sampling frame, the expansion of the sample to improve health insurance estimates for children, changes in population controls, and procedural changes in weights adjustment). GGN consider only sampling variability in their analysis, focusing on its effect on standard errors. The last two factors have a potential to produce

biased time differences in the outcome of interest and, therefore, should be paid particular attention in case of difference-in-differences models.

First of all, with difference-in-differences regressions it is intrinsically not clear what effect the interaction term actually measures. Second, data error further complicates model selection. For example, a triple-difference model can be favored as more convincing compared to a difference-in-differences model when only a subpopulation is affected by the policy and also as providing more flexible functional form that allows for better control for unobserved shocks. Yet, an overly detailed model can be counterproductive. In presence of data error, such modeling approach runs the risk of data over-fitting, i.e. introducing bias in the estimates by modeling erroneous data deviations.

Finally, strong "employment lock" effect found in GGN implies potentially large labor market impacts of the Affordable Care Act, which went into effect in January 2014. Based on their results GGN predicted that between 530,000 and 940,000 low-income working adults who, under the law, have become eligible for Medicaid or subsidized coverage would leave their jobs. The empirical evidence in the paper does not justify these predictions. Our results are in line with other empirical studies that predict no or small impact of the ACA on labor supply (Baicker et al. (2013), Dague et al. (2014)). With employment numbers from the first year under the ACA becoming available, these predictions will be ultimately put to test.

References

- Baicker, K., Finkelstein, A., Song, J., Taubman, S., 2014. The Impact of Medicaid on Labor Market Activity and Program Participation: Evidence from the Oregon Health Insurance Experiment. *American Economic Review: Papers & Proceedings* 104(5), 322-328.
- Dague, L., DeLeire, T., Leininger, L., 2014. The Effect of Public Insurance Coverage for Childless Adults on Labor Supply. NBER Working Paper 20111.
- Bowler, M., Ilg, R.E., Miller, S., Robison, E., Polivka, A., 2004. Revisions to the Current Population Survey Effective in January 2003. *Employment and Earnings*. (Bureau of Labor Statistics, February 2003) on the Internet at www.bls.gov/cps/rvcps03.pdf
- Garthwaite, C., Gross, T., Notowidigdo, M.J., 2014. Public Health Insurance, Labor Supply, and Employment Lock. *Quarterly Journal of Economics* 129, 653-696.
- Gruber, J., Madrian, B.C., 2004. Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature. In C. McLaughlin (Ed.), *Health Policy and the Uninsured*. Washington, D.C.: Urban Institute Press.
- Shoemaker, H., 2004. Redesign of the Sample for the Current Population Survey. In *Employment and Earnings*. Bureau of Labor Statistics. 4-6.
- U.S. Bureau of Labor Statistics, 2006. *Design and Methodology: Current Population Survey*. Technical Paper 66. Washington, D.C.: U.S. Department of Labor, Bureau of the Census. (<https://www.census.gov/prod/2006pubs/tp-66.pdf>)

Table 1. The Effect of TennCare Disenrollment on Employment

	CPS Current empl.			CPS Last year empl.		SSA/Census
	GGN	GGN	Full	GGN	Full	Full
	estimate	sample	sample	sample	sample	sample
	(1)	(2)	(3)	(4)	(5)	(6)
A. Difference-in-Difference Estimates						
Tennessee	0.025	0.025	0.013	0.012	0.006	-0.006
× Post 2005	(0.011)	(0.013)	(0.013)	(0.011)	(0.011)	(0.007)
	[0.038]	[0.064]	[0.322]	[0.317]	[0.563]	[0.424]
B. Triple-Difference Estimates (childless adults vs. parents)						
Tennessee	0.046	0.046	0.041	0.023	0.006	
× Post 2005	(0.020)	(0.022)	(0.021)	(0.018)	(0.017)	
× No Children	[0.032]	[0.035]	[0.051]	[0.211]	[0.723]	
C. Triple-Difference Estimates (40-64 year olds vs. 20-39 year olds)						
Tennessee		0.048	0.044	0.025	0.032	0.003
× Post 2005		(0.020)	(0.020)	(0.019)	(0.018)	(0.010)
× Age 40-64		[0.020]	[0.027]	[0.179]	[0.076]	[0.748]

Notes: The table reports estimates of the interaction term coefficient in difference-in-differences regression model and triple-difference model (see Section 4). The dependent variable is the share of people employed, measured in CPS data by employment during the reference week (Columns 1-3) and by employment during last year (Columns 4-5), and in SSA/Census data by the ratio of total number of workers to population (Column 6). Column 1 reports estimates obtained in Garthwaite, Gross, and Notowidigdo (2014); they use modified bootstrap standard errors and a sample of respondents aged 21-64 with college degree or less (GGN sample). Our estimates in Column 2 use GGN sample and OLS standard errors. Estimates in Column 3 are obtained with a sample unrestricted by education (the full sample). Results parallel for those in Columns 2 and 3 for the last year employment measure are in Columns 4 and 5. Standard errors are in parenthesis and p-values are in brackets.

Table 2. Tennessee's changes in growth rates for share employed, total employment and population

Dependent var.	20-39 year olds			40-64 year olds		
	CPS	SSA/Cen.	Diff.	CPS	SSA/Cen.	Diff.
log(Share Empl.)	0.041 (0.015) [0.006]	0.009 (0.008) [0.283]	0.032 (0.017) [0.057]	0.052 (0.013) [0.000]	-0.002 (0.006) [0.787]	0.054 (0.014) [0.000]
log(Employment)	0.030 (0.025) [0.232]	0.019 (0.008) [0.019]	0.011 (0.026) [0.665]	0.001 (0.022) [0.954]	-0.004 (0.006) [0.515]	0.005 (0.023) [0.826]
log(Population)	-0.010 (0.021) [0.615]	0.010 (0.003) [0.000]	-0.020 (0.021) [0.329]	-0.051 (0.019) [0.007]	-0.002 (0.002) [0.215]	-0.049 (0.019) [0.010]

Notes: The table presents CPS and SSA/Census estimates of $\Delta\beta_{js}$ for state Tennessee (s) in the following piecewise regression: $y_{jst} = \alpha_{js} + \beta_{js} \cdot t + \varepsilon_{jst}$ if $t \leq 2004$ and $y_{jst} = (\alpha_{js} - \Delta\beta_{js} \cdot 2004) + (\beta_{js} + \Delta\beta_{js}) \cdot t + \varepsilon_{jst}$ if $t > 2004$, where $2000 \leq t \leq 2007$, y is listed in the first Column, and j indicates age group. Standard errors are in parenthesis, and p-values are in brackets.

Table 3. The Effect of TennCare Disenrollment on Employment, CPS current employment

	No individual controls			With individual controls		
	CPS weight	No weight	Rescaled weight	CPS weight	No weight	Rescaled weight
A. Difference-in-Difference Estimates						
Tennessee	0.025	0.016	0.020	0.014	0.006	0.016
× Post 2005	(0.013)	(0.013)	(0.014)	(0.012)	(0.011)	(0.013)
	[0.064]	[0.218]	[0.152]	[0.235]	[0.601]	[0.194]
B. 2000-2007, Triple-Difference Estimates (childless adults vs. parents)						
Tennessee	0.046	0.044	0.037	0.038	0.038	0.038
× Post 2005	(0.022)	(0.021)	(0.021)	(0.019)	(0.018)	(0.019)
× No Children	[0.035]	[0.042]	[0.087]	[0.050]	[0.035]	[0.052]
C. 2001-2008, Triple-Difference Estimates (childless adults vs. parents)						
Tennessee	0.039	0.032	0.031	0.027	0.024	0.027
× Post 2005	(0.019)	(0.018)	(0.018)	(0.017)	(0.016)	(0.018)
× No Children	[0.039]	[0.076]	[0.093]	[0.128]	[0.135]	[0.121]


Notes: The table reports estimates of the interaction term coefficient in difference-in-differences regression model and triple-difference model (see Section 4) using the 2000-2008 March CPS data. The sample consists of individuals aged 21-64 with college degree or less. The dependent variable is the share of people who report working during the reference week. Columns 1, 2, and 3 show estimation results run on state-year weighted, unweighted, and weighted with adjusted weights (respectively) means of the dependent variable. Columns 4, 5, and 6 show parallel results using means of residuals obtained from an individual-level regression of the current employment indicator on a set of individual controls. These controls include age, education, race, marital status, an indicator of having a work-limiting disability, all interacted with gender. OLS standard errors are in parenthesis and p-values are in brackets.

Table 4. The Effect of TennCare Disenrollment on Employment Using Individual-Level Data by Gender

	Males			Females		
	Baseline model	With indiv. controls	With indiv. control	Baseline model	With indiv. controls	With indiv. control
	2000-2007	2000-2007	2001-2008	2000-2007	2000-2007	2001-2008
Tennessee	0.060	0.046	0.026	0.030	0.027	0.024
× Post 2005	(0.026)	(0.020)	(0.018)	(0.026)	(0.024)	(0.021)
× No Children	[0.023]	[0.024]	[0.148]	[0.240]	[0.246]	[0.239]
R^2	0.03	0.30	0.30	0.01	0.17	0.18
N	115,722	115,722	121,003	132,258	132,258	138,572

Notes: The table reports estimates of the interaction term coefficient in the triple-difference model (2) obtained with CPS individual-level data. The dependent variable is an indicator equal to one for CPS respondents who report working during the reference week. Sample consists of individuals aged 21-64 with college degree or less living in Southern states. Individual controls added to the baseline model (2) include a third-degree polynomial in age, educational level categories, race, marital status, and an indicator for having a work-limiting disability. Standard errors are in parenthesis and p-values are in brackets.

A. Tennessee
(in 1,000,000s)



B. Other Southern States
(in 1,000,000s)


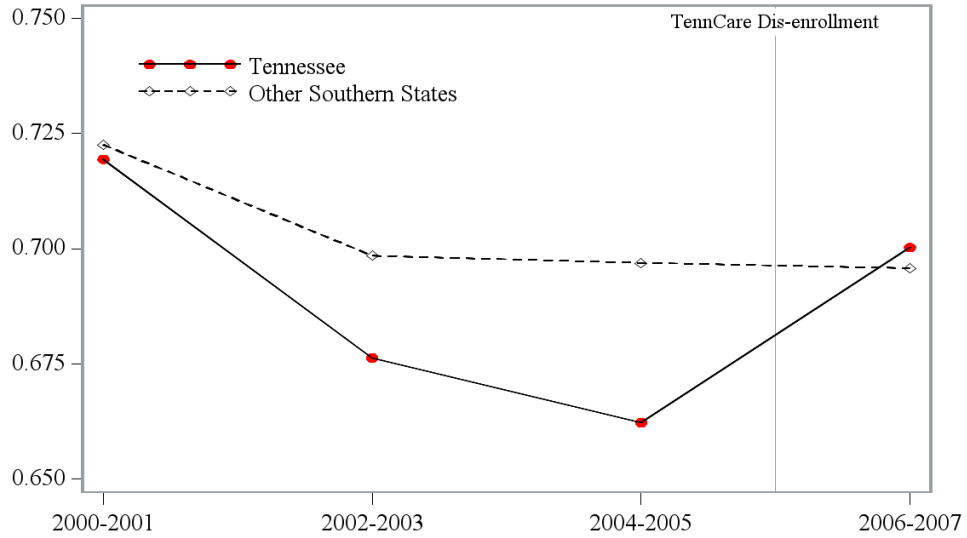


Fig. 1. Total employment

Notes: Panels A and B plot total and average total number of workers (respectively) derived from three data sources: CPS, SSA, and LAUS. CPS employment estimates use information about employment situation during the reference week (current employment) and during last year (last year employment). CPS current employment is shown for two samples of workers: a) those aged 20-64 and b) those aged 21-64 with college degree or less (GGN sample).

A. CPS data



B. SSA/Census data




Fig. 2. Share employed

Notes: The figure plots two-year averages for the share of people employed. The CPS sample consists of civilian individuals aged 21-64 with college degree or less. The CPS employment measure used in the plot is employment during the reference week. The SSA/Census sample consists of individuals aged 20 - 64.

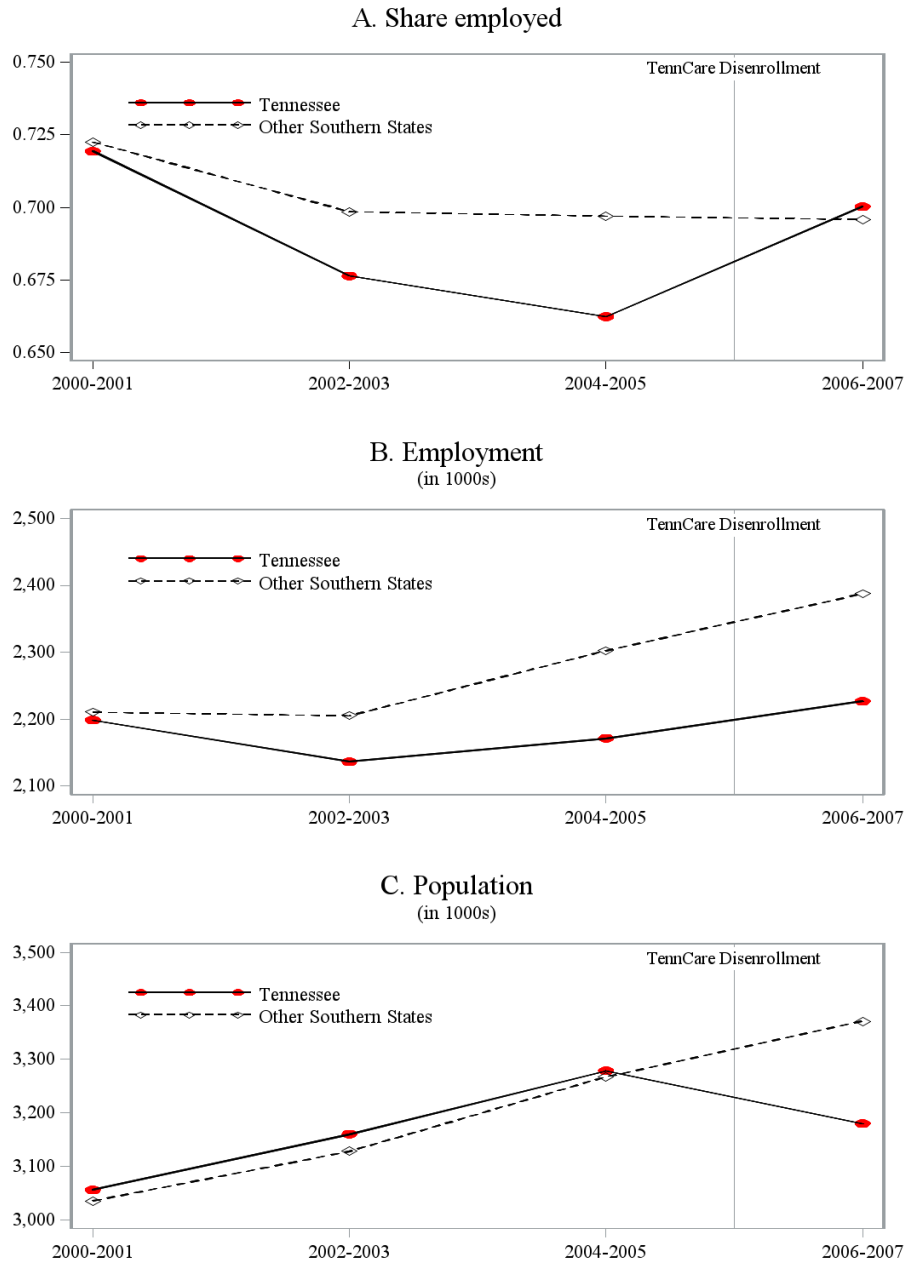


Fig. 3. Share employed and its components, CPS

Notes: Using CPS data for individuals aged 21-64 with college degree or less, the figure plots two-year averages of employment to population ratio, employment, and population. Employment is obtained as $\sum_{i=1}^{n_{st}} (I(\text{employed}_{sti} = 1)w_{sti})$, and population as $\sum_{i=1}^{n_{st}} w_{sti}$, where w is the CPS individual sample weights, and s, t, i represent state, year, and individual indicators.

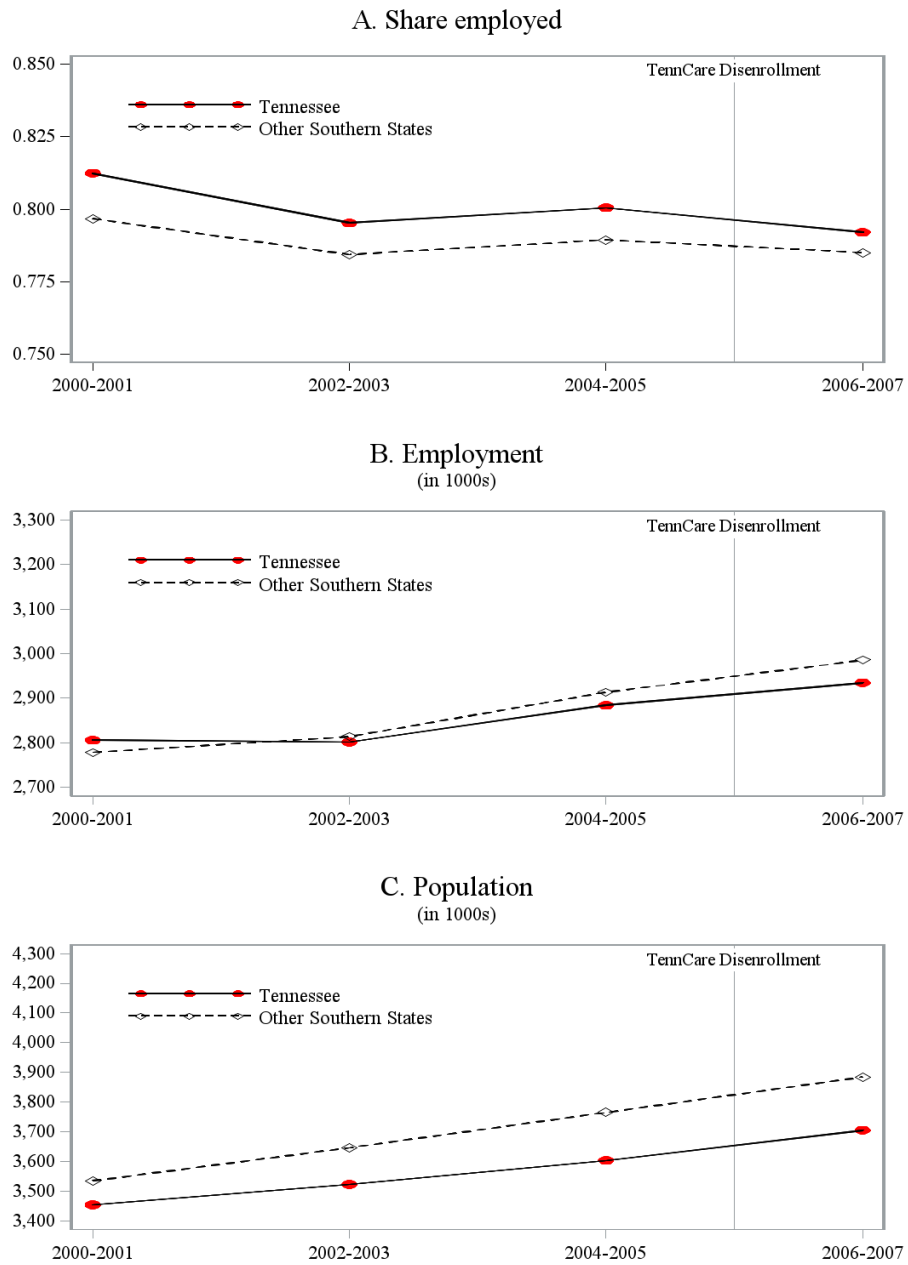


Fig. 4. Share employed and its components, SSA/Census

Notes: Using SSA employment data and Census postcensal population estimates for individuals aged 20-64, the figure plots two-year averages of employment to population ratio, employment, and population.

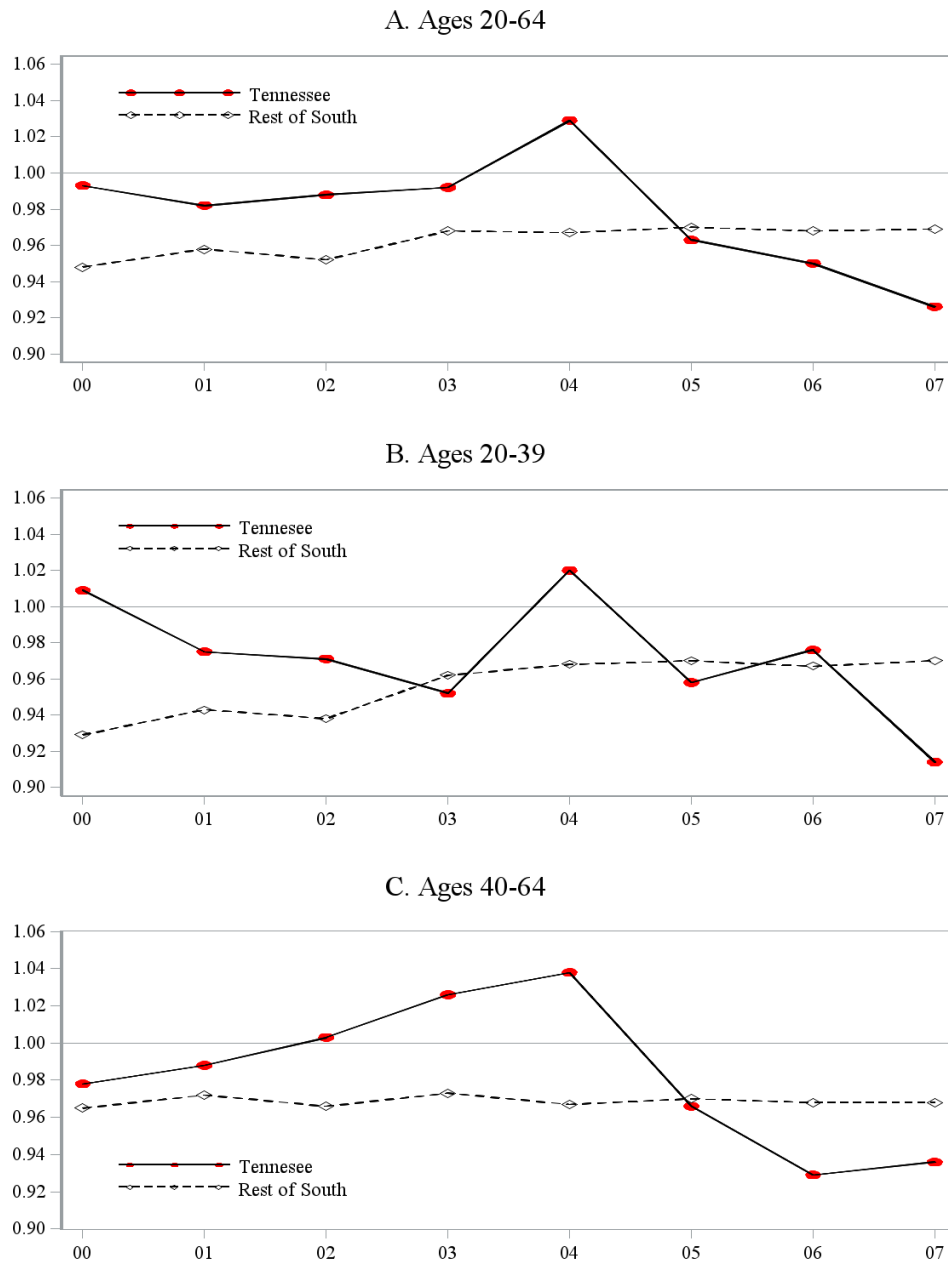


Fig. 5. Ratio of CPS to Census population estimates

Notes: The figure plots the ratio of CPS population estimates to Census population estimates. The Census estimates represent the residential population, while the CPS estimates pertain to the civilian population.

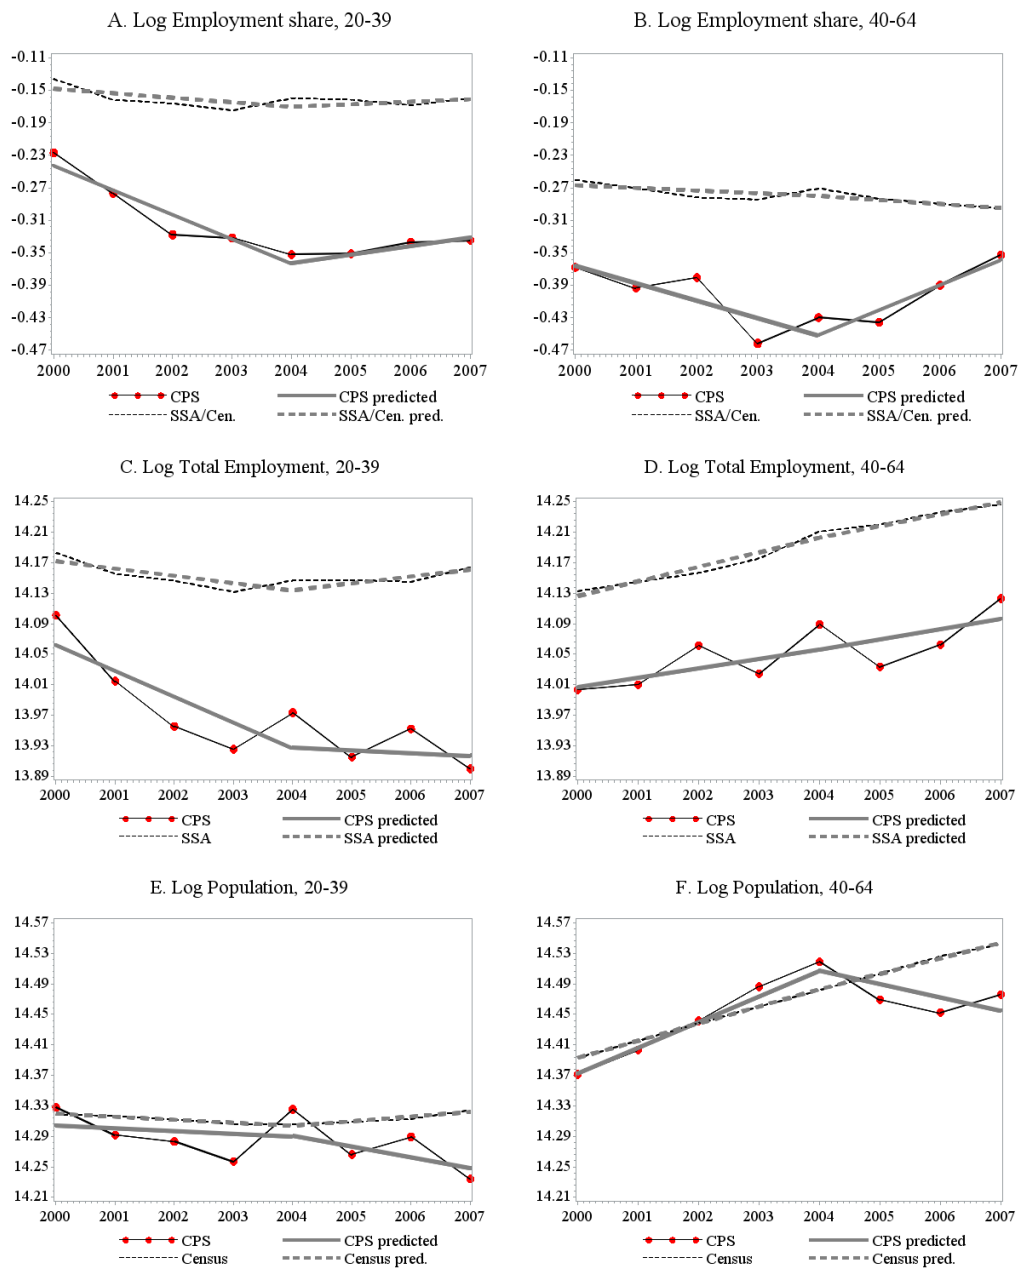
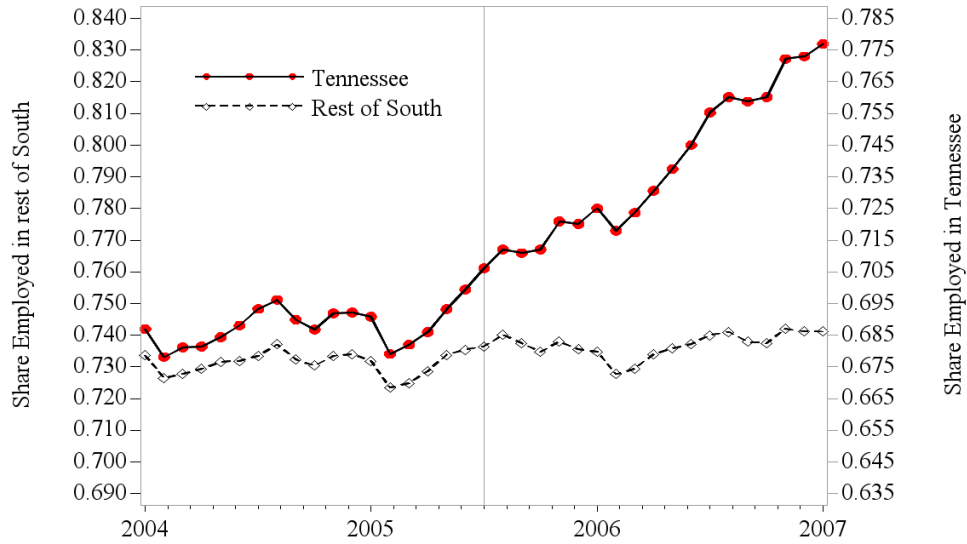


Fig. 6. Piecewise regression, Tennessee

Notes: For three dependent variables, the figure plots reported values and predicted values of the piecewise regression (3) for state Tennessee and two age groups. Each panel juxtaposes results based on the March CPS data with those based on the SSA/Census data.

A. Using CPS Population Estimates



B. Using Census Population Estimates



Fig. 7. Share employed by month, LAUS

Notes: The figure presents monthly employment rates for Tennessee and other Southern states. The rates are estimated as a ratio of total employment from the Local Area Unemployment Statistics (LAUS) to a) civilian population aged 16-64 from the CPS, b) residential population aged 16-64 from Census population estimates.

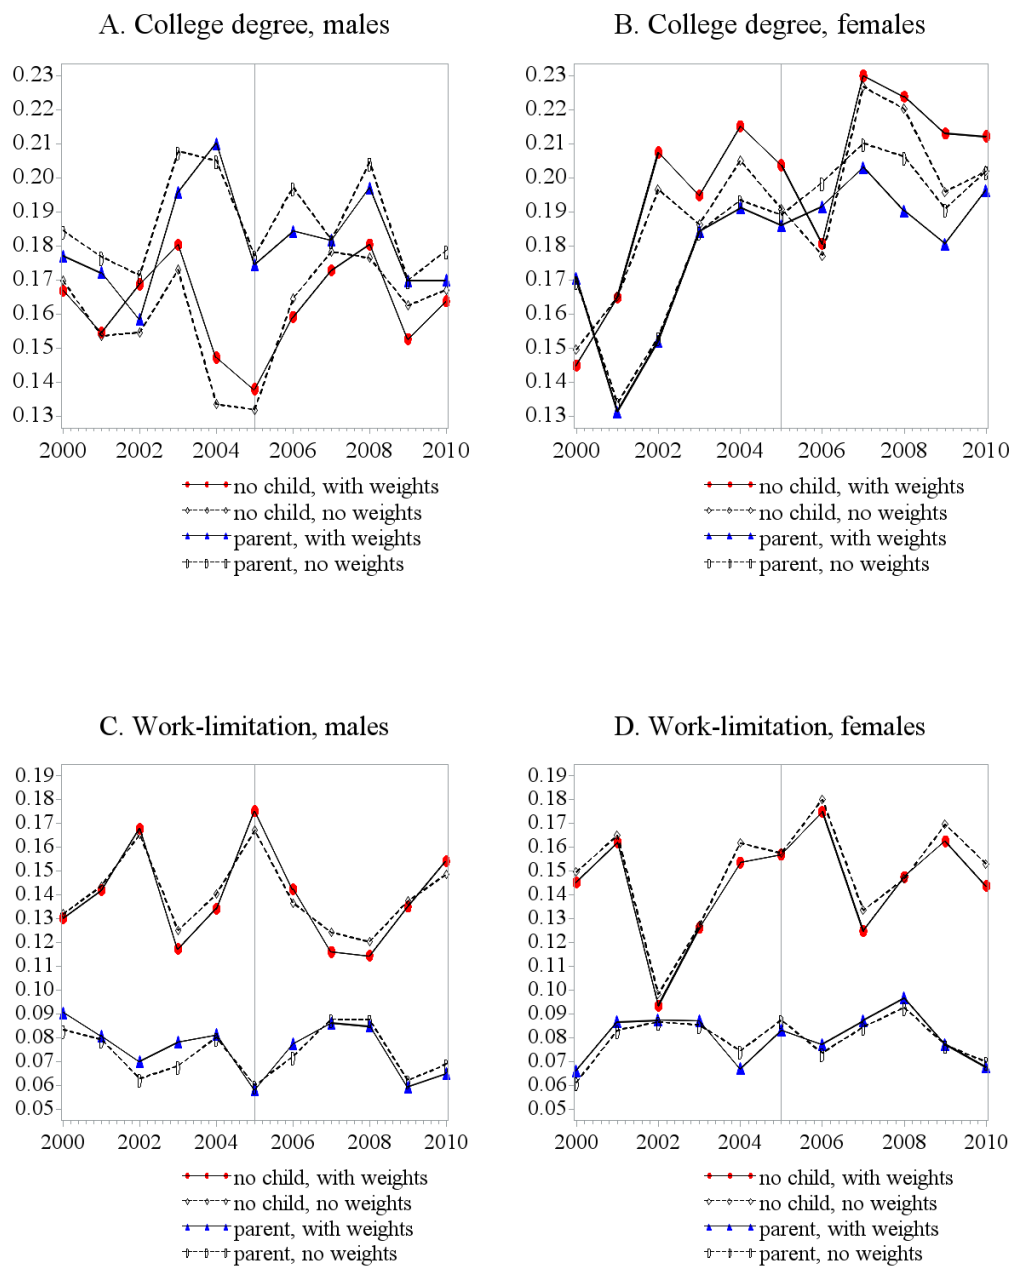


Fig. 8. Share with college degree and work-limitations by childless status, Tennessee

Notes: The figure uses CPS March data. The sample includes Tennesseans aged 21-64 with college degree or less. For males and females, Panels A and B show the fraction of those with college degree by childless status. Similarly, Panels C and D plot the fraction of those with a work-limiting disability. 44