Preliminary/incomplete Draft
(Not for circulation or citation)

Labor Market Persistence:
A Cohort-level Analysis

Charles A. Fleischman, Joshua Gallin, and Christopher L. Smith*

Board of Governors of the Federal Reserve System

December 2018

*Division of Research and Statistics, Federal Reserve Board, Washington, DC 20551. Contact author: Fleischman (cfleischman@frb.gov). We would like to thank XX, XX, and XX for helpful comments and discussions, and Nicko Gladstone and Jesse Wedewer for assistance with this research. All errors are our own. The views expressed in this paper are not necessarily those of the Federal Reserve System or its employees. [This is a substantially updated version of an earlier working paper (FEDS 2001-25).]

I. Introduction

Monetary policymakers are keenly interested in whether there is hysteresis in the labor market. That is, can recessions result in a permanently lower long-run sustainable level of employment and/or can long economic expansions characterized by an extended period of tight labor markets result in a permanently higher level of long-run sustainable employment? The answers to these questions can affect optimal monetary policy, because hysteresis dynamics raise the long-run costs of an economic downturn and the benefits of long expansions.

There is a substantial macroeconomics literature on unemployment hysteresis that is largely focused on whether there is a unit root in the aggregate unemployment rate. If the unemployment rate has a unit root, even if it explains only a small fraction of the variation in the unemployment rate, then cyclical fluctuations could result in permanent changes in the natural rate of unemployment. More micro-oriented labor economists have often focused on the relationship between the exit rate from unemployment (a fundamental determinant of the natural rate) and the duration of an unemployment spell using individual-level panel data sets, arguing that true duration dependence implies unemployment rate hysteresis (e.g. Kroft, Lange, Notowidigdo, and Katz 2016).

In this paper, we depart from both the macro and micro literatures along several dimensions by turning the question around and asking whether cyclical fluctuations in current macroeconomic conditions affect an individual's future probability of being employed. Essentially, this boils down to asking whether there is path dependence (hysteresis) in employment; that is, does current employment beget future employment? Our work extends and updates the analysis described in Fleischman and Gallin (2001, hereafter FG).

Following FG, we write down a highly-stylized, simple dynamic model where an individual's probability of being employed in the current period depends on the individual's past employment history as well as macroeconomic conditions.[1] Lacking long time-series

[1] Although we do not formally link our approach to insider-outsider models of employment, the implications are similar.

of individual-level employment records, we estimate the model using synthetic cohorts constructed from the repeated cross-sections of the CPS and defined by birth year, gender, and education (at most a high-school education vs. at least some college education), and report results primarily based on cohorts of less-educated men.

In our main specification, we regress the cohort-level employment-to-population ratio (EPOP) on its lag, the current and lagged output gap[2], a quadratic in age, and cohort fixed effects. Including the lagged output gap in the regression allows for richer dynamics. We include the quadratic in age to capture the regular, life-cycle pattern of EPOPs. We include the cohort fixed effects to allow for variation across cohorts in average EPOPs that may reflect, in part, the composition effects associated with the declining trend in the share of individuals with at most a high-school education.

To assess the persistence of cohort-level EPOPs in response to fluctuations in macroeconomic conditions, we report the dynamic responses of the cohorts' EPOPs (net of the usual life-cycle patterns, which are captured by the quadratic in age) to "shocks" to the output gap. The key parameters governing our estimates of persistence are the coefficients on lagged EPOP, the current output gap, and the lagged output gap.

Our main findings come from a two-stage least-squares (2SLS) estimator that addresses key limitations of using pseudo-panel data constructed from repeated cross sections from the monthly CPS. To ensure that each annual observation is not based on multiple observations of the same people, we estimate cohort-level EPOPs using only one of the eight "rotation groups" sampled each month in the CPS, specifically individuals in their fourth month in sample (MIS-4). Importantly, because of the relatively small number of individual observations used to construct the cohort-level EPOP, both our dependent and lagged dependent variable are measured with error, which would lead to a downward bias in our estimate of the coefficient on the lagged EPOP. To address the attenuation bias, we use a second measure of EPOP based only on those who have been in the sample for eight months (MIS-8) as an instrument; this instrument is correlated with lagged EPOP measured using those in MIS-4, but the sampling errors in the two measures should be uncorrelated.

---

[3] See Roberts and Morin's (1999) summary of the existing macroeconomics literature, and their conclusion that there is little compelling evidence of hysteresis in the unemployment rate.

Indeed, compared with results based on OLS regressions, we found noticeably more persistence with our 2SLS estimator, which highlights the importance of addressing the attenuation bias.

Although our 2SLS estimates imply greater employment persistence than OLS estimates, in most of our 2SLS specifications we nonetheless find at most only modest persistence in the response of the EPOP to an output gap shock. Specifically, we typically find that the coefficient on the current output gap is positive and the coefficient on the lagged output gap is negative, and that the product of the coefficient on lagged EPOP times the coefficient on the current output gap is somewhat larger—in absolute value—than the coefficient on the lagged output gap. Accordingly, the initial response of the EPOP to a positive output gap shock is positive, the next period response remains positive but is notably smaller, and in subsequent periods the response decays at a rate of one minus the coefficient on lagged EPOP.

Our results are generally inconsistent with the presence of substantial and statistically significant labor market hysteresis, as the effects of a one-period "blip" in the output gap have nearly fully dissipated within the first five years after the shock. That said, in response to an output shock that is longer-lasting (which is more similar to what is observed empirically), the effect on the EPOP in can persist well beyond five years. These results point to the possibility of lasting gains from countercyclical policy to dampen or shorten cyclical downturns.

We also find that output gap shocks have large initial effects on the EPOP for younger prime-age men, and our point estimates indicate somewhat more persistence as well. In addition, it appears that the effects on EPOP are amplified (though no more persistent) during periods when the unemployment rate is rising than when it is declining.

The rest of the paper is organized as follows: Section 2 places our approach within the macro and micro literatures on hysteresis. Section 3 discusses our data and our choice of instruments. Section 4 lays out our stylized model, and discusses some econometric issues related to our synthetic cohort approach. Section 5 presents and interprets our regression results. Section 6 concludes and offers directions for future work.

II. How our approach differs from the previous literature

Is there hysteresis in the US labor market? Both macroeconomists and labor economists have attacked this question. The macro literaturehas focused largely on whether there is a unit root in the aggregate unemployment rate, the presence of which is interpreted as evidence of hysteresis.[3] There is no consensus on whether there is a unit root in the U.S. unemployment rate, with some specifications rejecting a unit root (Mitchell, 1993; Song and Wu, 1997) and others failing to reject (Nelson and Plosser, 1982). Roberts and Morin (1999) test for hysteresis in the context of a price-price Phillips curve turned up no evidence of what they call "permanent" hysteresis. However, they did find some indications that past levels of the unemployment rate have a transitory effect on the current level of the NAIRU. Ball (2009) studies the behavior of the NAIRU and inflation over multiple decades for 20 countries and finds many examples of movements in aggregate demand seeming to be related to movements in the NAIRU, consistent with some hysteresis.

Whether there is a unit root in the unemployment rate may provide only a weak test of hysteresis, as this approach ignores the labor force participation margin, as well as other margins such as the share of the employed working part-time for economic reasons. Taking a broader approach, testing for a unit root in the employment-to-population ratio (EPOP) would encompass hysteresis in both the unemployment rate and/or labor force participation rate—and Gustavsson and Österholm (2007), who engage with this exercise using US macro data, fail to reject the hypothesis of a unit root.

Even so, failure to reject the unit-root hypothesis for EPOP (or even for the unemployment rate) should not necessarily be interpreted as evidence of labor market hysteresis, per se. A more useful definition of labor market hysteresis would differentiate between, on the one hand, permanent movements that can be explained by structural changes that are orthogonal to the business cycle—such as changes in the composition of the population, like the aging of the baby boom—and, on the other hand, permanent movements in response to transitory business cycle fluctuations. Monetary policymakers

[3] See Roberts and Morin's (1999) summary of the existing macroeconomics literature, and their conclusion that there is little compelling evidence of hysteresis in the unemployment rate.

should be primarily interested in the latter, although understanding both types of movements is essential to the evaluation of resource utilization in the labor market.

In contrast to this aggregate time-series approach, micro-labor economists have mainly focused their analysis of hysteresis on the relationship between the exit rate from unemployment (a fundamental determinant of the natural rate) and the duration of an unemployment spell using individual-level panel data sets. Beginning at least with Heckman and Borjas (1980), labor economists have asked whether unemployment begets further unemployment. If the probability of exiting unemployment declines as the length of the spell increases—there is negative duration dependence—then hysteresis is possible because changes in the distribution of unemployment spells by duration could affect the natural rate. Wilkinson (1997) finds some evidence of negative duration dependence in spells of unemployment in Canada consistent with hysteresis, though the effect is quite small, and Arulampalam, Booth, and Taylor (2000) find evidence for hysteresis-like effects—what they call scarring—in an individual-level panel of British labor force participants. Heckman and Borjas (1980) found no evidence that previous occurrences of unemployment or their duration affected future labor market outcomes once they controlled for sample selection and heterogeneity bias. A challenge facing the literature that seeks evidence for hysteresis at the individual-level is that it is difficult to separate out duration dependence (which would be related to hysteresis) from unobserved heterogeneity (longer-duration unemployed are less productive)—studies that attempt to separate these factors typically find some role for true duration dependence, albeit often a limited one (e.g. van den Berg and van Ours, 1996).

More recently, one strand of literature related to individual-level hysteresis focuses on estimating the long-term effects of entering a labor market during a boom or a bust. Oreopoulos, von Wachter, and Heisz (2012), and Kahn (2009) demonstrate long-lasting effects (up to 10 years) on the wages and earnings of college graduates who graduate during a recession; when they examine employment outcomes (employment, unemployment), they find less persistent effects, however. Even more recently, Fallick and Krowlikowski (2018) take an approach somewhat similar to the one we use in this paper, and use CPS data to examine the relationship between the EPOPs in a state for non-college prime-age men and the lagged EPOP; these sorts of regressions find only limited persistence in EPOPs, with shocks mostly dissipating within three years.[4] Similarly, Hotchkiss and Moore (2018) use the NLSY to explore whether exposure to a tight labor market mitigates the effects of subsequent downturns on employment (which could be consistent with hysteresis), but find only limited support for this idea. Our findings are consistent with all of these studies, in that we find only limited evidence for significant employment persistence (with the effects on employment of a shock mostly dissipating after a few years). Finally, Yagan (2018) studies the consequences of the Great Recession on subsequent employment outcomes, and finds that a local labor market's employment rate remained depressed five years after the Great Recession despite unemployment rates mostly recovering—which could be interpreted as evidence for hysteresis. One way to reconcile Yagan's findings in support of employment hysteresis with other studies lack of support for hysteresis may be that the extent of employment hysteresis is related to the severity of the recession, and the other studies cited here (as well as our paper) estimate average employment persistence/hysteresis over multiple business cycles.

In this paper, we take a hybrid approach. We estimate persistence in the employment-to-population rates for single birth-year cohorts in response to a shock to the output gap. Our approach side-steps some of the key limitations of using only aggregate data or of focusing only on the unemployment rate. Specifically, we ask whether business cycle fluctuations can leave a permanent imprint on an individual's probability of employment in the future. Previous work has found that an individual who loses a job during a recession faces longer spells of unemployment/non-employment (e.g. Song and von Wachter 2014). If longer spells of non-employment contribute to depreciation of skills (both hard, task-specific skills and soft, workplace-related skills) or influence potential employers' beliefs about such workers' quality and skills, or lead to addiction or other mental-health issues, it is possible that a transitory drop in labor demand could permanently—or at least persistently—reduce a person's future probability of employment.

[4] A key distinction between their work and ours is that Fallick and Krowlikowski (2018) estimate the relationship between the EPOP for prime-age men and its lag, while we estimate the relationship between a *cohort's* EPOP and its lag (for prime-age men). Another distinction is that Fallick and Krowlikoski utilize cross-state variation, whereas we rely on variation across cohorts at the national level.

As noted earlier, there is also evidence that entering the labor market during a recession may lead to persistent employment and wage effects.

Conversely, during a period of very strong labor demand, increased employment opportunities may draw some individuals into the labor force who otherwise would have remained on the sidelines, while helping others improve their job-match quality. Indeed, anecdotes suggest that as the labor market tightens, employers broaden the pools of applicants considered for vacant positions—perhaps by narrowing the range of jobs subject to drug screening, lowering education/experience requirements, and by making greater use of on-the-job training. And, as a result, those who benefit from the opportunities available in a tight labor market may be able to accumulate additional skills, learn improved work habits, and build social and business networks, each of which could enhance future employment probabilities.

III. Stylized model

A. Individual-level model

Our cohort-level analysis is motivated by an individual-level model of employment determination. In the individual-level model, a person (i) in cohort (c) is employed for a fraction of year (t), E_{ict}, that can take on a value between zero and one. We posit that E_{ict} depends, in part, on Z_{ict}, a vector of personal characteristics and X_t, a vector of aggregate factors that can vary over time, but not across cohorts. E_{ict} will also depend on employment experience in earlier years, E_{ict-k}, for k = 1, 2, 3…, because human capital accumulation, in the form of on-the-job training and the acquisition of basic workplace skills may increases an individual's future employability, while periods spent not working could lead to skill depreciation and reduce future employability. In addition, adjustment cost, tenure rules, etc… that limit turnover, will also tend to induce path dependence at the individual level. This can be represented as a reduced-form linear equation:

$$E_{ict} = \alpha_1 Z_{ict} + \alpha_2 X_t + \beta(L) E_{ict-1} + u_{ict} \qquad (1)$$

where u_{ict} captures various unobserved (or unmodeled) influences on E_{ict}. We assume that the unobserved determinants can be written as a combination of a cohort effect, η_c, an aggregate time effect, υ_t, an individual fixed effect, ω_i, and an idiosyncratic error, ε_{ict}:

$$u_{ict} = \eta_c + \omega_i + \upsilon_t + \varepsilon_{ict}$$

where each component is mean zero.

B. Cohort-level model

Although equation (1) above is specified at the level of an individual, Deaton (1985) shows that for a model such as ours that is linear in parameters, though not necessarily linear in variables, it is possible to take cohort means and rewrite (1) as a cohort-level relationship:

$$EPOP_{ct} = \alpha_1 Z_{ct} + \alpha_2 X_t + \beta(L)EPOP_{ct-1} + u_{ct} \qquad (2)$$

where all (i) subscripts have been removed and the cohort-level employment-to-population ratio, $EPOP_{ct}$, is equal to the average share of weeks worked across all cohort members.[5] Our estimating equation, (equation 3), which we present below, is consistent with (2), where we replace Z_{ct} with a cohort fixed effect and a quadratic in age, we replace X_t with the current-period and lagged output gap, and we include only one lag of the cohort's EPOP.

IV. Data

A. Cohort-level EPOP

We constructed annual estimates of the EPOP at the cohort level, using monthly microdata from the Current Population Survey (CPS), as retrieved from IPUMS (Flood et.

[5] The unobserved component, u_{ct}, now includes the cohort fixed effects, time fixed effects, and a cohort-level idiosyncratic component.

al. 2018).[6] Given the interview structure of the CPS (respondents are interviewed for four consecutive months, out of the sample for eight months, and then interviewed for another four months), the monthly CPS data include observations from eight waves of individuals, known as rotation groups defined by how many times they have been interviewed, or "month-in-sample" (MIS). As a result, we cannot include data on the same individual (or group of individuals) for all twelve months and our annual estimates must be based on 12 monthly estimates based on observations of different groups of cohort members. To ensure that a unique set of respondents is used for estimating employment in each year, we only use respondents from a single rotation group—specifically those in their fourth month of being interviewed, MIS4. Had we, instead, included all eight rotation groups, we would had multiple observations per person per year—but not the same number of observations per person—and observations for the same people in consecutive years, but also for varying numbers of months.[7] Our results are robust to constructing our annual averages using observations only from MIS1, MIS2, or MIS3.

One drawback to this approach is that it leaves us with relatively few individual-level observations per cohort, raising the possibility of substantial measurement error that likely attenuates somewhat estimates of the correlation between current and lagged employment. We could have doubled the sample size for each cohort by using individuals in both MIS4 and MIS8 (or any pairs of MIS1 and MIS5, or MIS2 and MIS6, or MIS3 and MIS7) to estimate annual cohort-level EPOPs and reduced the variance of measurement

---

[6] In 2001, for a sample that ended in 1998, FG constructed synthetic cohorts using monthly observations from the March CPS, which are available back to 1964, rather than from the full complement of monthly CPS data, which were available for fewer years, primarily because of the longer sample period. For this paper, we use data through 2017, and opted instead to use the monthly CPS in order to take advantage of the larger sample sizes and full-year coverage rather than limit ourselves to one monthly observation per year.

[7] If we included all rotation groups, the annual averages for year t and year t-1 would be estimated from different numbers of observations per person, depending on that person's rotation group in January or September of year t-1of year t. A person in MIS1 in January of year t-1 would contribute 4 observations to both the year t and year t-1 averages. For members of the other rotation groups in January of year t-1, the contributions of MISX to year t-1 and year t annual averages would be: MIS2 (4, 3); MIS3 (4, 2); MIS4 (4, 1); MIS5 (4, 0), MIS6 (3, 0); MIS7 (2, 0), and MIS8 (1, 0). Individuals in MISX in September of year t-1 cannot be in the January t-1 sample, but can also contribute unequal numbers of observations to the year t-1 and year t averages: MIS1 through MIS4 (4, 4), and MIS5 through MIS8 (4, 0). Finally, individuals whose rotation group is not interviewed in both January t-1 and September t-1 will contribute four observations to the year t-1 average, and either four observations to the year t average (MIS1 in February through May of year t-1) or zero observations (MIS5 in February through May of year t-1).

error by a factor of four. However, with this approach, half of the individuals would be included in consecutive annual estimates, which would have caused the measurement errors to be positively serially correlated and would bias up our estimated coefficient on the lagged EPOP.

To overcome the attenuation bias caused by the substantial, though serially uncorrelated, measurement errors caused by using only one rotation group, we use an instrumental variables approach, where we instrument the lagged EPOP measured from MIS4 respondents using estimates of the lagged EPOP measured from MIS8 respondents. The MIS8 measure should be a powerful instrument for the MIS4 measure because the two are strongly correlated and, because they are constructed from two distinct pools of individuals, their measurement errors should not be correlated. Because the MIS8 respondents are in their final month of being interviewed in the CPS in year t-1, they do not appear in the MIS4 sample used to construct the year t estimate of the dependent variables—hence the measurement error in lagged EPOP from MIS8 should also be uncorrelated with the measurement error in the current period EPOP based on MIS4 respondents, the latter of which will be included in the equation residual.

B. Output gap

We are particularly interested in estimating the persistence in each cohort's EPOP in response to shocks to aggregate macroeconomic conditions, which we proxy for by the output gap.[8] Our measure of the output gap is the (log) ratio of real GDP to the CBO's estimate of potential GDP. As of this writing, the CBO's latest estimate of potential GDP was estimated using data available prior to the BEA's 2018 Comprehensive Update of the National Income and Product accounts (NIPA). Accordingly, we also use estimates of real GDP published prior to the 2018 Comprehensive Update.

V. Results

[8] In later drafts of the paper, we will explore the sensitivity of our results to using alternative proxies for macroeconomic conditions.

A. Cohort and sample selection

We focus on cohorts of less-educated men (those with at most a high-school degree and zero years of college education) because, compared with more-educated men, the cohort-level EPOPs for less-educated men are more cyclically sensitive (see FG) and, compared with women, there has been less variation over time in the life-cycle profile of male EPOPs.[9]

Our sample covers the period from 1976 to 2017, and within this period we follow less-educated men during their prime working years (ages 25 to 54). Given that we subdivide each cohort by educational attainment, we impose the 25-year age minimum because most people who will acquire at least some college education have done so by age 25. Verbeek (2008) argues that cohorts should be defined on the basis of variables that do not vary over time and constructed using samples drawn for each period from of the same population. This rules out defining cohorts by whether individuals have at least some college education that include younger ages when college entrance is still common. We chose the 54-year age maximum because employment rates for older individuals have been rising as retirement has occurred later for more recent cohorts, likely due in part to increases in longevity and changes in pension structure and the availability of Social Security payments. Thus, it seems likely that factors guiding employment decisions for older ages, and hence the degree of persistence in employment, may be substantively different than those for younger ages. Nevertheless, because of late college enrollment and early deaths, any pair of age cut-offs may result in some changes in the population from which the individuals in the cohort have been drawn.

We further restrict our sample to include only the 43 birth-year cohorts for which we have at least 15 observations. We follow our oldest cohort, born in 1936, from age 40 (1976) to 54 (1990) and we follow our youngest cohort, born in 1978, from ages 25 (2003) to 39 (2017). We are mindful of the trade-off between including more cohorts (increasing N) and including more annual observations per cohort (increasing T). Importantly, estimating dynamic pseudo panel models that include both a lagged dependent variable and

[9] In future work we intend to broaden our analysis to include both more- and less-educated men and women.

cohort fixed effects can cause the coefficient on the lagged dependent variable to be biased down. In effect, some of the persistence in the EPOP that should be captured by the coefficient on the lagged dependent variable is instead captured by the cohort fixed effects, and this problem is particularly severe in panels with few observations per cohort. However, there is no established standard for what constitutes an adequate number of observations, as the bias on the lagged dependent variable can depend on the particular empirical exercise. We believe that our sample restriction is fairly conservative, and our estimates are robust to requiring additional observations per cohort and including fewer cohorts.

Given our sample, our panel is "unbalanced," in that we include different numbers of years per cohort *and* different number of cohorts per year. Importantly, we observe some cohorts only during the early part of our sample and others only towards the end of the sample. Thus, in addition to the potential bias described above, the unbalanced panel structure may impose other challenges to our empirical strategy, which we hope to explore in future work.

B. Empirical specification

In our main specification, we regress the cohort-level EPOP on its lag, cohort fixed effects, a quadratic in age, and the current and lagged output gap:

$$EPOP_{ct} = \lambda EPOP_{ct-1} + \eta_c + \gamma_1 age_{ct} + \gamma_2 age_{ct}^2 + \beta_1 gap_t + \beta_2 gap_{t-1} + \varepsilon_{ct} \qquad (3)$$

We include the quadratic in age to capture the regular, life-cycle pattern of EPOPs. We include the cohort fixed effects to allow for variation across cohorts in average EPOPs that may reflect, in part, the composition effects associated with the declining trend in the share of individuals with at most a high-school education.[10] And, we include the lagged output gap to allow for richer EPOP dynamics in response to a cyclical shock.

[10] On balance, there has been a noticeable downward trend by birth year in the share of the male population with at most a high-school education. If an individual's choice of education is not random, and instead reflects sorting along the distribution of latent skills, then the average skill level of less-educated males will be declining, on average, with each subsequent cohort. To the extent that this concept of skills is correlated with employment probability, then we should expect to see downward shifts in the age profile of EPOPs with each subsequent cohort.

To assess the persistence of cohort-level EPOPs in response to fluctuations in macroeconomic conditions, we report the dynamic responses of the cohorts' EPOPs to shocks to the output gap. A key identification assumption is that the life-cycle profile of EPOPs is the same across all cohorts—net of parallel shifts due to changes in cohort composition and other factors that are captured by the cohort fixed effects. If this assumption is false because differences in life-cycle profiles are correlated with the differences in macroeconomic conditions faced by each cohort, but not caused by these differences in macroeconomic conditions, then interpretation of our results may be somewhat problematic. (net of the usual life-cycle patterns, which are captured by the quadratic in age) to "shocks" to the output gap. The key parameters governing our estimates of persistence are the coefficients on lagged EPOP, the current output gap, and the lagged output gap.

C. Baseline results

Table 1 presents estimates for our primary regression specification (equation 3). Note that unless otherwise indicated, the sample used for the following regressions are non-college men age 25-54 and include all data from 1976-2017, with the limitation that only cohorts with at least 15 observations are included. Also, as described in the previous section, another critical sample selection criteria important in our analysis is that we only include respondents who are MIS4 when constructing the dependent variable and lagged dependent variable.

In the top panel of table 1, and in the regression tables that follow, we display coefficients and associated standard errors for the lagged employment rate, the output gap, and the once-lagged output gap. (The standard errors that we present throughout the analysis are always clustered at the year level.) As described in equation 3, the regressions also always include a quadratic in age and cohort fixed effects; in what follows, we do not show estimates of the age quadratic or cohort effects, but estimates are available upon request. The bottom panel uses these estimates to construct a 95 percent confidence interval for the implied impact of a 1 percentage point increase in the output gap on the dependent variable at various points after the shock (0, the year of the impact; and 3 and 5 years later).

The dynamic behavior of the EPOP around the average age profile and cohort-specific intercepts (based on the estimates of α_1, α_2, and the η_c in equation (3)), is governed by the coefficient on the lagged EPOP (λ) and the coefficients on the contemporaneous and lagged output gap (β_1 and β_2). To fix ideas, we describe the response to a unit shock to the output gap in year t. On impact, the EPOP is β_1 above the life-cycle profile. In period t+1, the deviation from the life-cycle profile is $\lambda^* \beta_1 + \beta_2$. Because we typically find that β_2 is negative, λ is positive, and $\lambda^* \beta_1 + \beta_2$ is also positive, the t+1 EPOP deviation is smaller than the period t deviation. In the following periods, the period t+1 deviation decays geometrically at rate λ; thus, the period t+k deviation is $\lambda^{(k-2)} * (\lambda^* \beta_1 + \beta_2)$. Because λ is less than one, we generally find that there is little imprint of a unit shock to the GDP gap beyond three years.

We report our estimates of the dynamic EPOP response to both a one-off shock to the output gap in year t, which we refer to as "no persistence in the GDP shock" and report in the top rows of the lower panels of our tables, and to a shock to the output gap that is modestly serially correlated ($\rho = 0.65$), which we refer to as "persistence in the GDP shock" and report in the bottom rows of the lower panels of our tables.[11] We consider the most persistent case to be a more realistic depiction of how a cohort's EPOP would evolve in response to an innovation in the output gap. However, because the "no persistence in GDP shock" case is a more straightforward translation of the regression coefficients, in the visual depictions that follow of impulse response functions associated with our regressions we focus on the "no GDP persistence" case.[12]

The first column of Table 1 displays the estimated coefficient on lagged EPOP from equation (3), but where we exclude the output gap from the regression. Because we are limiting our sample to respondents from only a single MIS (MIS4) , sample sizes are much smaller than if we use the full sample of respondents. A potentially important drawback of

---

[11] This is based on the estimated coefficient from a regression of the annual average output gap on its lag. Using annual average data, we cannot reject that the coefficient on the second lag of the output gap is 0. As a result, the annual average output gap does not exhibit the familiar hump shape dynamics associated with most quarterly measures of the output gap.

[12] To be clear, the impulse responses that we display are translations of regression coefficients from the same regressions for both the persistence in GDP and no persistence in GDP cases; the only difference is that in the persistence case we feed through the impulse response a shock to GDP that decays gradually over time.

these OLS regressions, therefore, is that measurement error in estimates of lagged employment may attenuate estimates of the relationship between current and lagged EPOPs. We address this concern with an instrumental variables strategy. For our 2SLS estimates, we include only MIS4 respondents when calculating current and lagged EPOPs. Then, we construct a second estimate of lagged EPOP using only MIS8 respondents, which we use as an instrument. Given the interview structure of the CPS (respondents are interviewed for four consecutive months, dropped from the sample for 8 months, and then interviewed for another four consecutive months), there is no overlap in the respondents used for constructing the current EPOP (for year t), the lagged EPOP regressor, or the lagged EPOP used as an instrument. The third column displays the estimated coefficient on lagged EPOP based on this same specification used in column 1, but estimated with 2SLS.

As expected, the OLS results indicate economically and statistically significantly less persistence than do the 2SLS estimates, which we attribute in large part to the measurement error in our estimates of annual EPOPs based only on MIS4 respondents.[13] The OLS estimate (point estimate 0.48, standard error 0.03) indicates that only about 10 percent of a period t innovation to a cohort EPOP remains after three years. In contrast, the 2SLS estimate shown in column 3 (0.80, standard error 0.08) indicates that about half of the period t innovation remains after three years. The 95 percent confidence intervals around these estimated period t+3 effects, which are not reported, are non-overlapping.

Columns 2 (OLS) and 4 (2SLS) of Table 1 display the results when the regression also includes the contemporaneous and lagged output gap, and the impulse responses to the non-persistent and persistent output gap shocks are shown in Figure 1. The OLS coefficients on the output gap and its lag are both positive, but the coefficient on the lagged EPOP is very small. These OLS estimates suggest very little persistence in the EPOP response to an output gap shock beyond period t+1. In contrast, the 2SLS coefficient on the lagged output gap is negative and statistically significant, while the coefficient on the lagged

[13] We repeated the OLS exercise using annual averages constructed from both MIS4 and MIS8 respondents and the estimated coefficient (0.64, standard error of 0.02) is about right in the middle of the range of the OLS estimate using MIS4 and the 2SLS estimate. The larger coefficient when using both MIS respondents could reflect the reduced measurement error and/or the serial correlation of the measurement errors. OLS estimates using only respondents in MIS8 are quantitatively and qualitatively identical to the MIS4 OLS estimates; the point estimate of 0.45 with a standard error of 0.03

EPOP is large (0.75, standard error 0.14). This configuration of coefficients indicates more persistence in the EPOP response to an output gap shock after period t+1. Indeed, as shown in the left panel of Figure 1, the OLS estimate of the period t+1 EPOP response is somewhat larger than the 2SLS estimate (because of the positive coefficient on the lagged output gap), but is much smaller by period t+2 and is nearly gone by period t+1.

Moreover, the 2SLS estimates indicate some modest persistence in the EPOP response to a one-off shock to the GDP gap. After three years, the point estimate shown in the lower panel points to a small effect on the EPOP after three years (0.13, standard error of 0.07), which is statistically different from zero. After five years, the point estimate is small but still noticeable (0.08) and zero is barely within the 95 percent confidence interval. And, as shown in the bottom rows of Table 1 and in the right panel of Figure 1, the 2SLS estimates imply that the EPOP remains significantly elevated for several years in response to a GDP shock of average persistence. After five years, the lower end of the 95 percent confidence interval is still 0.20, and the point estimate is 0.36.

We interpret this limited, but statistically significant, persistence as only weak evidence *against* hysteresis, per se, at least for less-educated prime-age males. However, the absence of statistically significant hysteresis does rule out the potential benefit of active countercyclical policy. Indeed, the output gap measured with the CBO's estimate of potential output fell to about -6½ percent in 2009. If this had been a one-off output gap shock, and the gap had returned to zero by 2010, our point estimates imply that the EPOP for less-educated men would have still been about ½ percentage point lower in 2014 than in the absence of this large shock. And, had the output gap only shown average persistence following the Great Recession, which would have been a less stark outcome than what seems to have occurred, the EPOP for this group would have been depressed by nearly 2½ percentage points in 2014 and by about 1¼ percentage points in 2018.

To more fully understand our findings, we now turn to estimates of persistence in the labor force participation rate and the unemployment rate for less-educated prime-age men. We will then explore heterogeneity in EPOP persistence by labor market outcomes, by demographic characteristics, and by time periods. Because of the apparent attenuation bias in our OLS results, we prefer the 2SLS estimates and report only these throughout the remainder of the analysis.

D. Persistence in the labor force participation rate and unemployment rate

Does the labor force participation rate or unemployment rate exhibit different persistence than the employment rate? Since labor force entry or exit decisions may be somewhat sticky (for example, individuals who have returned to school or retired may be unable or unwilling to quickly re-enter the labor force in periods of strong labor demand), we may expect to see greater persistence in labor force participation than on employment or unemployment rates. To explore this idea, we re-estimate our 2SLS baseline regressions, replacing the dependent variable and lagged dependent variables with the LFPR and unemployment rates.[14]

Table 2 and Figure 2 display results from these regressions. One difference across these measures of labor market outcomes is that the effect on the LFPR from an increase in output lags a year behind the other measures—that is, an increase in the LFPR appears a year after the initial shock to output. Also, the LFPR appears to exhibit somewhat greater persistence than the unemployment rate; the coefficient on lagged LFPR is roughly twice as large as on the lagged unemployment rate, and by year 4 essentially none of the impact of a one-time GDP shock on the unemployment rate remains, while about half of the shock on the LFPR remains. This is consistent with other research finding a significant relationship between the labor force participation rate and measures of the cycle up to three years previously (e.g. Aaronson et. al. 2014).

E. Variation in employment persistence over time

Although we estimate fairly modest employment persistence over our full sample period, it is possible that shocks to employment have had more persistent effects over different periods, for example over the second half of the sample period when the decline in the EPOP for non-college prime-age men has steepened or when the labor market is tightening or loosening.

[14] Analogously to our baseline specifications, we use MIS4 respondents to estimate the dependent and lagged dependent participation rate or unemployment rate, and we use MIS8 respondents to construct our instruments.

In Figure 3 we display 2SLS estimates when the sample is limited to four periods: 1976-2017 (the same as our baseline results), 1976-1996, 1997-2017, and 1976-2007 (dropping observations from the Great Recession and thereafter). Estimated employment persistence is very similar across the sample periods, although the initial impact of an increase in the output gap on the employment rate is somewhat larger in the second half of the sample (column 3, the green line) than in the first half (column 2, the red line). Despite this larger initial impact, however, by three years after the shock to output the effect on the EPOP is pretty similar regardless of which years are used in the estimation.

It also seems plausible that there may be asymmetry in the degree of employment persistence during expansions compared with recessions. There is much discussion in the macroeconomics literature of aggregate hysteresis following downturns, but little discussion of hysteresis following booms (see, e.g. Martin, Munyan, and Wilson (2015), and Blanchard, Cerutti, and Summers (2015)).[15] To test for asymmetries related to the state of the business cycle, in the results shown in Figure 4 we interact the lagged employment rate, output gap, and lagged output gap with an indicator for whether the unemployment rate decreased (first column, blue line) or increased (second column, red line) since the previous year.[16] The coefficients on lagged employment and the lagged output gap are nearly identical; the only difference between the impact of a shock to output during strong labor markets compared with weak labor markets is that the initial effect on employment is larger during weak labor markets. Indeed, during periods when the unemployment rate is falling, the effect of a one-time increase in output essentially dissipated the year after the shock.

F. Variation in employment persistence by demographic characteristics

i. By age

Since employment persistence is affected by an individual's ability and willingness to substitute time spent in the formal labor market with time spent on other activities (e.g. going to school, engaging in home production, relying on non-wage income from SSDI),

[15] We defer to future drafts a discussion of which labor-market features might lead to asymmetric responses.

[16] As instruments, we include the lagged EPOP estimated from MIS8, adapted for this specification by interacting the instrument with our indicator for the state of the business cycle.

employment persistence may vary by age. To test this, we interact the lagged employment rate, the output gap, and the lagged output gap with a quadratic in age and evaluate the effect of these variables (and the corresponding impulse response to an increase in the output gap) at various ages.[17] In the table to Figure 5, we show that although the impact from an output shock on EPOP is larger for younger workers (younger workers' employment and labor force participation is more cyclically sensitive), there is little difference in estimated employment persistence for non-college men age 25, 35, or 45.

ii. By race

While it is well-known that employment rates are more cyclical for non-whites (e.g. Cajner et al)—suggesting relatively greater harm during recessions to the employment of non-whites and relatively greater benefit during expansions—it has been largely left unexplored whether the effect of output shocks on employment rates also have more persistent effects for non-whites. To test for differences in persistence by race, we divide our sample into whites, blacks, and others (non-white, non-black), and estimate our baseline 2SLS on each sample. One difficulty with dividing our sample this way is that observation counts per cell are very small for blacks and non-whites, non-blacks—for most cohorts in most years, between 20 and 40 observations per year! We recognize that our preferred regression specification, which requires limiting the sample used to calculate the dependent and lagged dependent variables to MIS4, will limit possible inference regarding differences in persistence by race due to small sample sizes. In the results we show here, we include the twice-lagged EPOP from MIS8 as a second instrument in hopes that it additionally helps address attenuation bias, but we admit that these results are especially tentative and still a work-in-progress. Comparing columns 1, 2, and 3 of Figure 6, the initial effect of an increase in the output gap on EPOPs is larger for non-whites, consistent with other research finding greater cyclicality in employment for these groups (e.g. Cajner, Radler, Ratner, and Vidangos 2017). The coefficient on lagged employment is larger for non-whites than for whites (0.75 vs 0.50 for whites), but the standard error on this estimate for non-whites is too large to conclude that these estimates differ. In any case, by years 3 and 5 the estimated

[17] As instruments, we include the lagged EPOPs from MIS8 interacted with a quadratic in age.

impact of a shock to output remaining on the EPOP is fairly similar across races (again, possibly larger for non-whites, but standard errors are large).

VI. Summary

Overall, we estimate that persistence in employment rates for non-college men age 25 to 54 is generally small: in response to a one-time increase in the output gap, the effect on the EPOP for this group has mostly dissipated by three to five years later. The EPOP response to a more persistent increase in the output gap is more drawn out, which suggests a possible role for countercyclical policy, especially in response to a prolonged downturn. Employment persistence appears to be very similar for a number of cuts of the data: in the last 20 years relative to 20 years previously, in periods of improving labor markets relative to weakening labor markets, and for younger men relative to older men. A shock to employment may be somewhat more persistent for non-white men than for white men, but our estimates are not precise enough to conclude this with an acceptable degree of precision. Unsurprisingly given previous evidence on the persistence of labor force participation decisions, the LFPR appears to exhibit somewhat greater persistence than the EPOP and especially the unemployment rate.

We believe that much work remains to be done to understand both the extent of persistence in the labor market and its underlying causes. Importantly, we have focused only on employment persistence for less-educated males. In future work, we hope to extend the analysis to both more- and less-educated men *and* women.

In the current paper, a key identifying assumption is that the life-cycle EPOP profile for less-educated prime-age men has been stable over the past 40 years. Given our unbalanced panel, this is to a large extent an untestable hypothesis because we observe some cohorts only when young and others only when relatively older. Nevertheless, because violations of this assumption could bias our results in a number of ways, we think it rates a more thorough investigation that is beyond the scope of this paper.

Another issue that this paper does not adequately address is whether changes in the composition of the less- and more-educated cohorts can affect the key parameters in this study. If, for example, the types of jobs held by men with and without at least some college education has been changing over time, it is also possible that the cyclicality of employment for the two groups has been changing over time. In future work, we would like to exploit the information on the cohort shares in different education groups, including as interaction terms with our cyclical variables.

A third issue left unexplored in this paper is whether labor market conditions at the time when members of a cohort are entering the labor market and/or making education choices can have persistent effects on cohort-level employment rates. In earlier work (FG), found a correlation between estimated cohort fixed effects and conditions upon labor market entry as well as between educational attainment and initial labor market conditions. This bears further work as well.

VII. Citations

Aaronson, Stephanie, Tomaz Cajner, Bruce Fallick, Felix Galbis-Reig, Christopher L. Smith and William Wascher. 2014. “Labor Force Participation: Recent Developments and Future Prospects,” *Brookings Papers on Economic Activity, Fall.*

Arulampalam, Wiji, Alison Booth, and Mark P. Taylor. 2000. “Unemployment Persistence.” *Oxford Economic Papers* 52:23-50.

Ball, Lawrence. 2009. “Hysteresis in Unemployment: Old and New Evidence.” *National Bureau of Economic Research Working Paper* No. 14818, March 2009.

Cajner, Tomaz, Tyler Radler, David Ratner, and Ivan Vidangos. 2017. “Racial Gaps in Labor Market Outcomes in the Last Four Decades and Over the Business Cycle.” *Finance and Economic Discussion Series*, No. 2017-071, Board of Governors of the Federal Reserve System.

Deaton, Angus. 1985. “Panel Data from Time Series of Cross-Sections.” *Journal of Econometrics*, 30:109-126.

Fallick, Bruce, and Pawel Krolikowski. 2018. “Hysteresis in Employment Among Disadvantaged Workers. *Federal Reserve Bank of Cleveland Working Paper* No. 18-01, February 2018.

Fleischman, Charles A., and Joshua Gallin. 2001. "Employment Persistence." *Finance and Economics Discussion Series*, 2001-25. Board of Governors of the Federal Reserve System of the U.S.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren. 2018. Integrated Public Use Microdata Series, Current Population Survey: Version 6.0. Dataset. Minneapolis, MN. *https://doi.org/10.18128/D030.V6.0*

Gustavsson, Mangus and Pär Österholm. 2007. "Does Unemployment Hysteresis equal Employment Hysteresis?" *Economic Record*, 83: 159-173.

Heckman, James and George Borjas. 1980. "Does Unemployment Cause Future Unemployment? Definitions, Questions, and Answers from a Continuous Time Model of Heterogeneity and State Dependence." *Economica*, 47: 247-284.

Hotchkiss, Julie L., and Robert E. Moore. 2018. "Some Like it Hot: Assessing Longer-Term Labor Market Benefits from a High-Pressure Economy." *Federal Reserve Bank of Atlanta Working Paper* No. 2018-1.

Jacobson, Louis, Robert LaLonde, and Daniel Sullivan. 1993. *The Costs of Worker Dislocation*. Kalamazoo, Mich: W.E. Upjohn Institute for Employment Research.

Jaeger, Albert and Martin Parkinson. 1994. "Some Evidence on Hysteresis in Unemployment." *European Economic Review*, 38: 329-342.

Kahn, Lisa. 2009. "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy." *Labour Economics*, 17: 303-318.

Kroft, Kory, Fabian Lange, Matthew Notowidigdo, and Lawrence Katz. 2016. "Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Non-Participation." *Journal of Labor Economics*, 34 (S1).

Nelson, C.R. and C. I. Plosser. 1982. "Trends and Random Walks in Macro-Economic Time Series: Some Evidence and Implications." *Journal of Monetary Economics*, 10:139-162.

Oreopoulos, Phillip, Till von Wachter, and Andrew Heisz. 2012. "The Short- and Long-Term Career Effects of Graduating in a Recession." *American Economic Journal: Applied Economics*, 4(1): 1-29.

Roberts, John M. and Norman J. Morin. 1999. "Is Hysteresis Important for U.S. Unemployment?" *Finance and Economics Discussion Series*, 1999-56. Board of Governors of the Federal Reserve System of the U.S.

Song, F. and Y. Wu. 1997. "Hysteresis in Unemployment: Evidence from 48 U.S. States." *Economic Inquiry*, 35: 235-243.

Song, Jae and Till von Wachter. 2014. "Long-Term Non-Employment and Job Displacement." In *Evaluating Labor Market Dynamics, a Symposium Sponsored by the Federal Reserve Bank of Kansas City,* Jackson Hole Wyoming, August 2014.

van den Berg, Gerland and Jan van Ours. 1996. "Unemployment Dynamics and Duration Dependence." *Journal of Labor Economics*.

Wilkinson, Gordon. 1997. "A Micro Approach to the Issue of Hysteresis in Unemployment: Evidence from the 1988-1990 Labour Market Survey." *Bank of Canada Working Paper*, 97-12.

Yagan, Danny. 2018. "Employment Hysteresis from the Great Recession." Forthcoming, *Journal of Political Economy.*

Table 1. Estimates of persistence in EPOP, non-college men age 25-54

	OLS		2SLS	
	(1)	(2)	(3)	(4)
Lagged employment rate	0.48	0.15	0.80	0.75
	(0.06)	(0.05)	(0.08)	(0.14)
Output gap		0.67		0.74
		(0.10)		(0.09)
Lagged output gap		0.29		-0.32
		(0.10)		(0.16)
Implied impact of a 1 pp increase in the output gap, 95 percent CI				
No persistence in GDP shock				
Year 0		0.47 to 0.87		0.56 to 0.93
Year 3		-0.00 to 0.02		0.06 to 0.20
Year 5		-0.00 to 0.00		-0.01 to 0.16
Persistence in GDP shock				
Year 0		0.47 to 0.87		0.56 to 0.93
Year 3		0.34 to 0.46		0.40 to 0.70
Year 5		0.14 to 0.20		0.20 to 0.53

Note: The top panel of the table displays coefficients (and standarded errors clustered at the year level, in parentheses) from regressions of the employment-to-population ratio for a birth cohort-age-year on the lagged EPOP, the output gap and lagged output gap (in columns 2 and 4), a quadratic in age, and birth cohort dummy variables. For columns 3 and 4, the EPOP and its lag are estimated only using observations from month-in-sample 4, and the lagged EPOP is instrumented with the lagged EPOP estimated with observations from month-in-sample 8. The bottom panel provides the 95% confidence interval of the implied shock at year 0 (the initial shock) and years 3 and 5, assuming no persistence in the GDP shock (top rows) or persistence as described in the text (bottom rows). The sample is limited to men age 25-54 without any college experience; included cohorts are those with 15 or more years of data for the years of the sample, and data for 1976-2017 are included. The total number of cohorts is 43; the total number of observations for each column is 1,034.

Figure 1: Effect of a 1 pp increase in the output gap on the EPOP for non-college men

Comparing impulse responses on the EPOP, by regression specification and perstistence of GDP shock

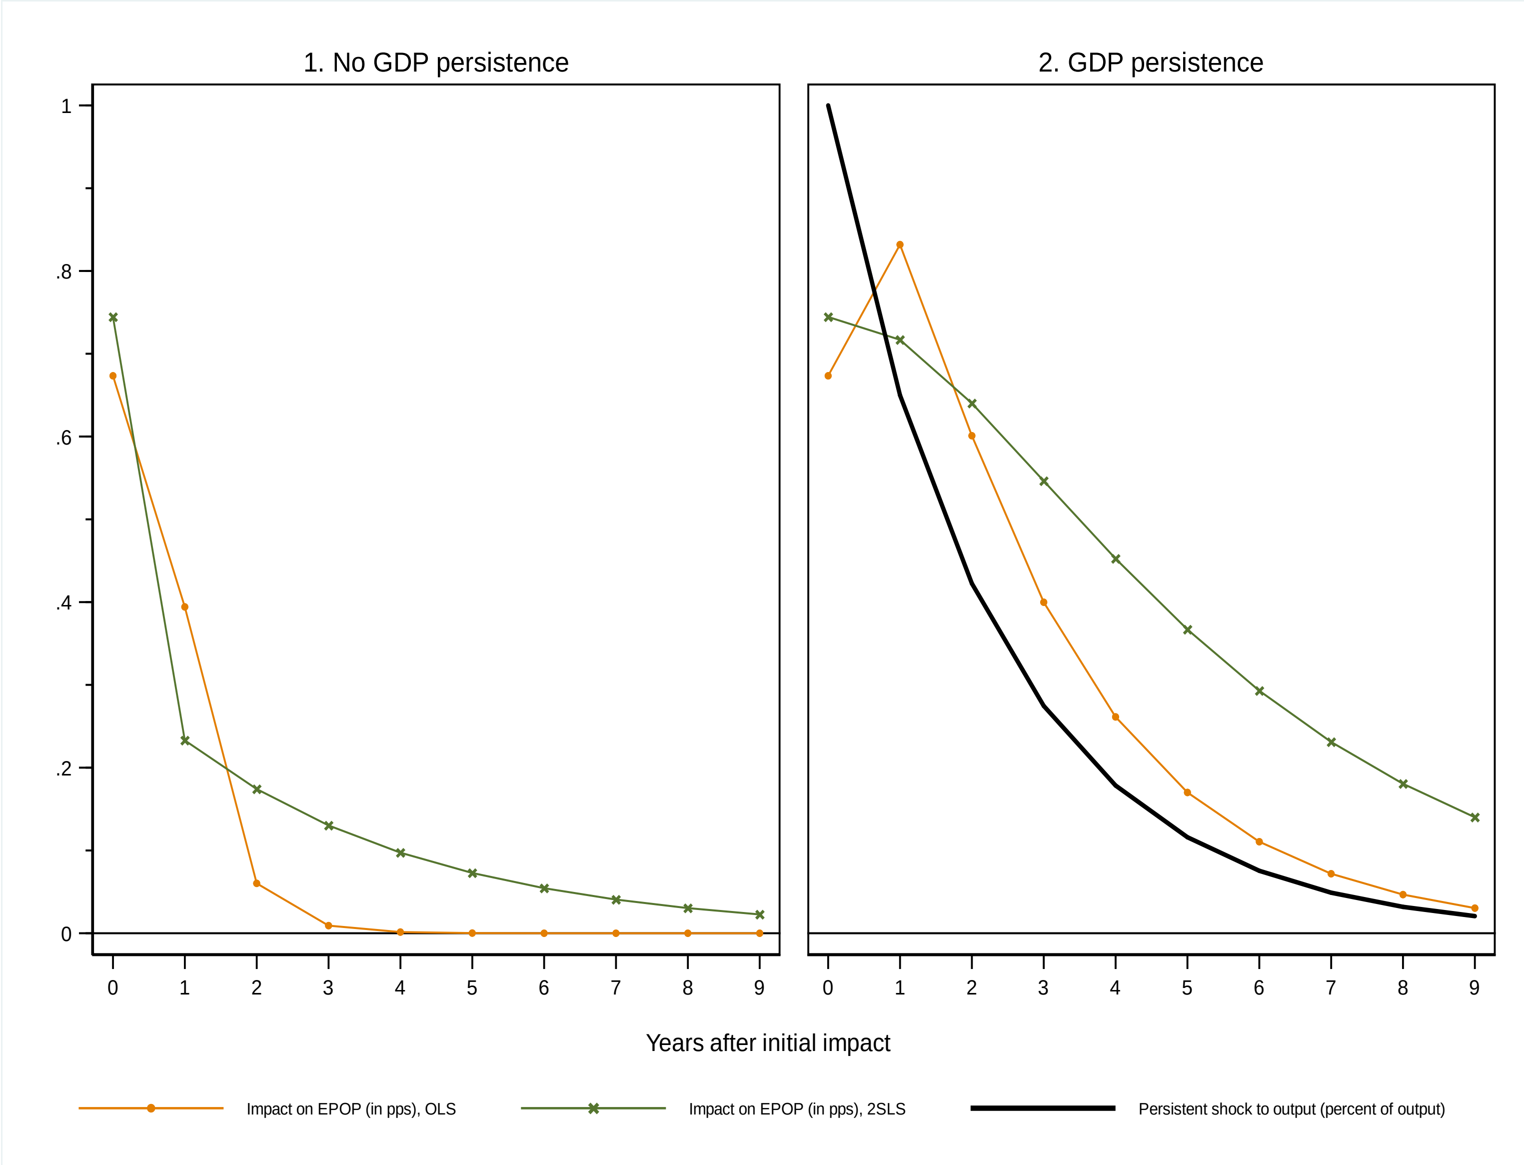

Note: Figure displays the impulse response of the EPOP for non-college men age 25-54 from a 1 percentage point increase in the output gap. Each panel shows the impulse response using regression coefficients from the OLS specification in Table 1 (the blue line, corresponding with column 2 of the table) and the 2SLS specification (the red line, corresponding with column 4 of the table). The left panel shows the impulse responses from a single-period 1 percentage point shock to the GDP gap that decays entirely away by the next period. The right panel shows the impulse response from a 1 percentage point shock to the GDP gap in the first period that decays by 35 percent each period; the black line in this panel shows the evolution of this output shock.

Table 2. Estimates of persistence in EPOP, LFPR, and unemployment rate, for non-college men

	Dependent variable:		
	EPOP	LFPR	Unemp. Rate
	(1)	(2)	(3)
Lagged dependent variable	0.75	0.83	0.47
	(0.11)	(0.16)	(0.07)
Output gap	0.74	0.00	-0.86
	(0.05)	(0.05)	(0.03)
Lagged output gap	-0.32	0.13	0.21
	(0.12)	(0.05)	(0.08)
Implied impact of a 1 pp increase in the output gap, 95 percent CI			
No persistence in GDP shock			
Year 0	0.65 to 0.84	-0.10 to 0.10	-0.92 to -0.80
Year 3	0.07 to 0.20	0.00 to 0.17	-0.07 to -0.02
Year 5	-0.00 to 0.15	-0.04 to 0.16	-0.02 to 0.00
Persistence in GDP shock			
Year 0	0.65 to 0.84	-0.10 to 0.10	-0.92 to -0.80
Year 3	0.42 to 0.67	0.04 to 0.37	-0.47 to -0.38
Year 5	0.21 to 0.53	-0.03 to 0.42	-0.24 to -0.17

Note: Table displays coefficients (and standarded errors clustered at the year level, in parentheses) from regressions of the EPOP, LFPR, and unemployment rate for a birth cohort-age-year on the lagged dependent variable, the output gap, lagged output gap, a quadratic in age, and birth cohort dummy variables. For all regressions the dependent variable and its lag are estimated only using observations from month-in-sample 4, and the MIS4 lag is instrumented with the lag estimate using only observations from month-in-sample 8. The bottom panel of the table provides the 95% confidence interval of the implied shock at year 0 (the initial shock) and years 3 and 5, assuming no persistence in the GDP shock (top rows) or persistence as described in the text (bottom rows). Sample is limited to men age 25-54 without any college experience; included cohorts are those with 15 or more years of data for the years of the sample, and data for 1976-2017 are included. The total number of cohorts is 43, and the total number of observations is 1,043.

Figure 2. Effects of a shock to output on the employment-to-population rate, labor force participation rate, and unemployment rate

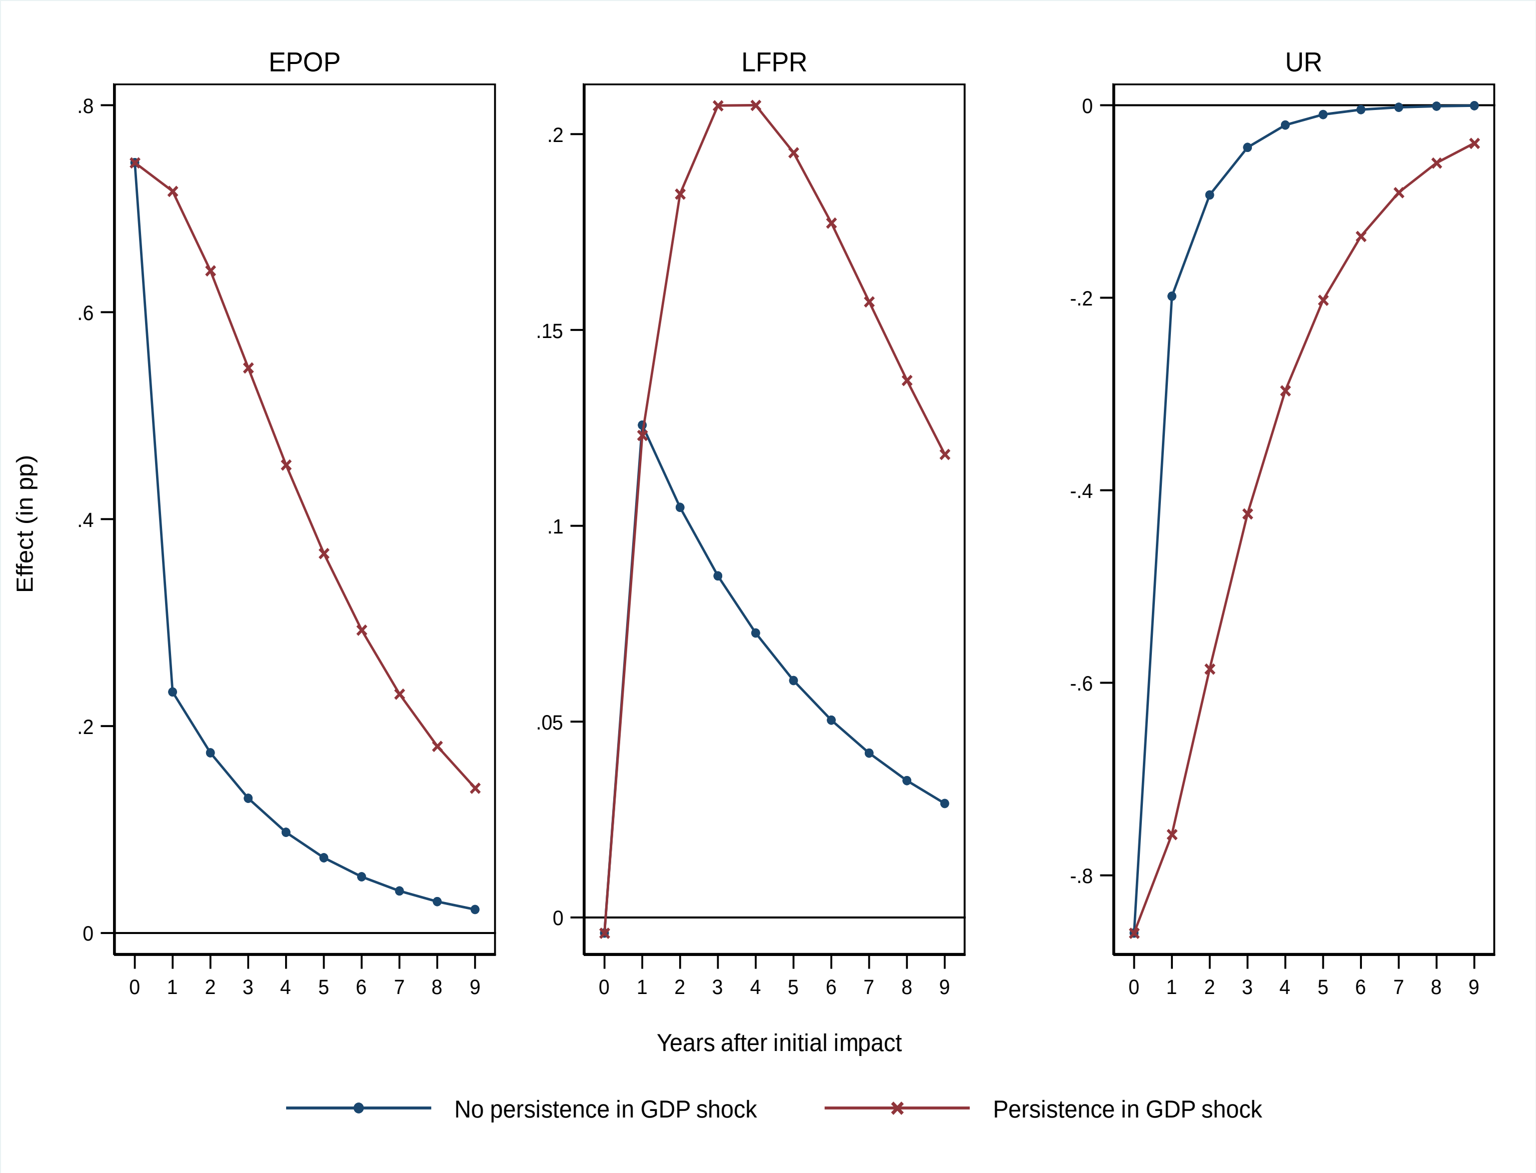

Note: Figure displays the impulse response of the employment-to-population rate, labor force participation rate, and unemployment rate for non-college men age 25-54 from a 1 percentage point increase in the output gap. Impulse responses are based on 2SLS regressions where the dependent and lagged dependent variables are estimated for respondents in month-in-sample 4, and the lagged dependent variable is instrumented by the lagged estimates from respondents in the month-in-sample 8. Estimates correspond with regression coefficients presented in table 2. The blue line shows the impulse responses from a single-period 1 percentage point shock to the GDP gap that decays entirely away by the next period. The red line shows the impulse response from a 1 percentage point shock to the GDP gap in the first period that decays by 35 percent each period.

Figure 3: Effect of a 1 pp increase in the output gap on the EPOP for non-college men
Comparing differences in persistence by years of sample

A. Regression coefficient and standard errors

	1976-2017	1976-1996	1997-2017	1976-2007
	(1)	(2)	(3)	(4)
Lagged employment rate	0.75	0.81	0.72	0.86
	(0.14)	(0.22)	(0.25)	(0.19)
Output gap	0.74	0.60	1.04	0.61
	(0.09)	(0.07)	(0.12)	(0.06)
Lagged output gap	-0.32	-0.26	-0.48	-0.32
	(0.16)	(0.19)	(0.36)	(0.17)
Implied impact of a 1 pp increase in the output gap, 95 percent CI				
No persistence in GDP shock				
Year 0	0.56 to 0.93	0.46 to 0.74	0.79 to 1.28	0.50 to 0.72
Year 3	0.06 to 0.20	-0.02 to 0.31	-0.02 to 0.30	0.01 to 0.30
Year 5	-0.01 to 0.16	-0.09 to 0.28	-0.11 to 0.25	-0.07 to 0.30
Persistence in GDP shock				
Year 0	0.56 to 0.93	0.46 to 0.74	0.79 to 1.28	0.50 to 0.72
Year 3	0.40 to 0.70	0.18 to 0.85	0.39 to 0.93	0.23 to 0.81
Year 5	0.20 to 0.53	-0.02 to 0.80	0.01 to 0.82	0.03 to 0.80

B. Impulse Response Functions (no persistence in GDP shock)

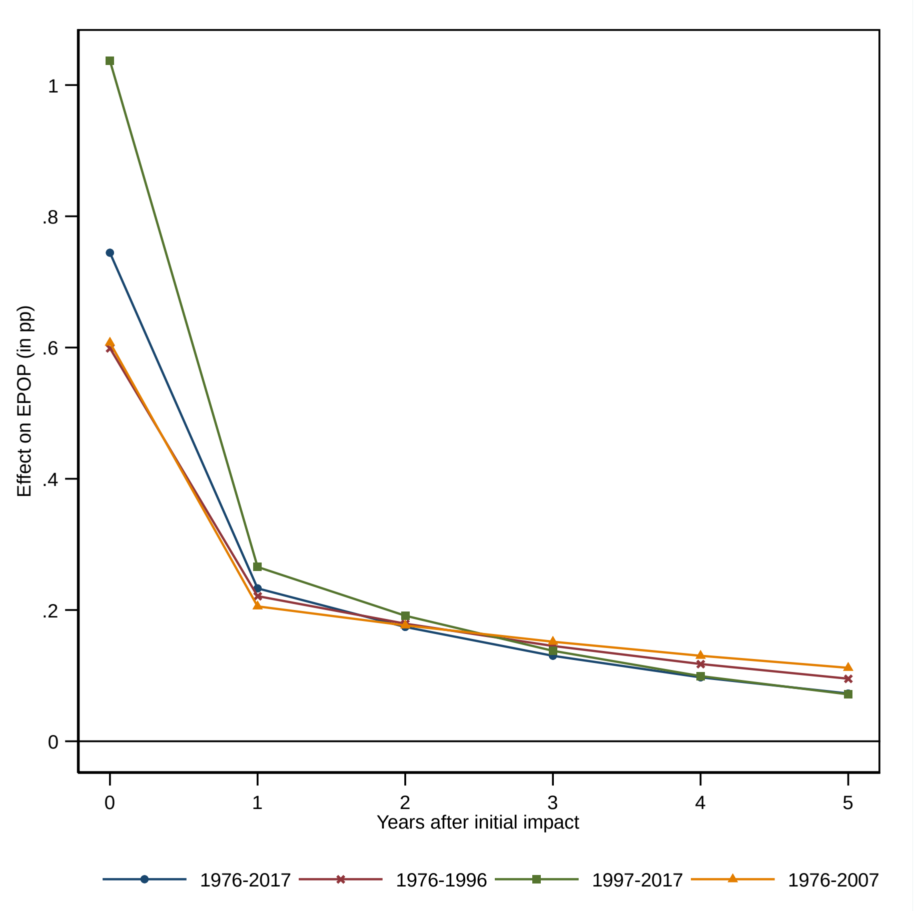

Note: Regression coefficient estimates (and standard errors clustered at the year level, in parentheses) and corresponding impulse responses from a 1 percentage point increase in the output gap, as shown in the table and figure, are derived from 2SLS regressions of the employment-to-population ratio for a birth cohort-age-year on the lagged EPOP, the output gap and lagged output gap (in columns 2 and 4), a quadratic in age, and birth cohort dummy variables. The EPOP and its lag are estimated only using observations from month-in-sample 4, and the lagged EPOP is instrumented with the first- and twice-lagged EPOP estimated with observations from month-in-sample 8. The bottom panel of the table provides the 95% confidence interval of the implied shock at year 0 (the initial shock) and years 3 and 5, assuming no persistence in the GDP shock (top rows) or persistence as described in the text (bottom rows). The sample is limited to men age 25-54 without any college experience; for the 1976-2017 and 1976-2007 periods, included cohorts are those with 15 or more years of data for the years of the sample; for the other year restrictions, included cohorts are those with 10 or more years of data. The total number of cohorts and observations, by column, are: 43 and 1,034 (column 1); 32 and 519 (column 2); 32 and 540 (column 3); 43 and 849 (column 4).

Figure 4: Effect of a 1 pp increase in the output gap on the EPOP for non-college men
Comparing the effects during periods of labor market tightening vs. loosening

A. Regression coefficient and standard errors

	Falling unemp. rate	Rising unemp. rate
	(1)	(2)
RHS variables interacted with indicator variable for rising/falling unemployment rate:		
Lagged employment rate	0.90	0.91
	(0.18)	(0.18)
Output gap	0.38	0.77
	(0.16)	(0.11)
Lagged output gap	-0.30	-0.26
	(0.15)	(0.18)
Implied impact of a 1 pp increase in the output gap, 95 percent CI		
No persistence in GDP shock		
Year 0	0.07 to 0.68	0.55 to 1.00
Year 3	-0.21 to 0.27	0.01 to 0.71
Year 5	-0.15 to 0.20	-0.19 to 0.77
Persistence in GDP shock		
Year 0	0.07 to 0.68	0.55 to 1.00
Year 3	-0.45 to 0.79	0.38 to 1.65
Year 5	-0.46 to 0.69	-0.05 to 1.91

B. Impulse Response Functions (no persistence in GDP shock)

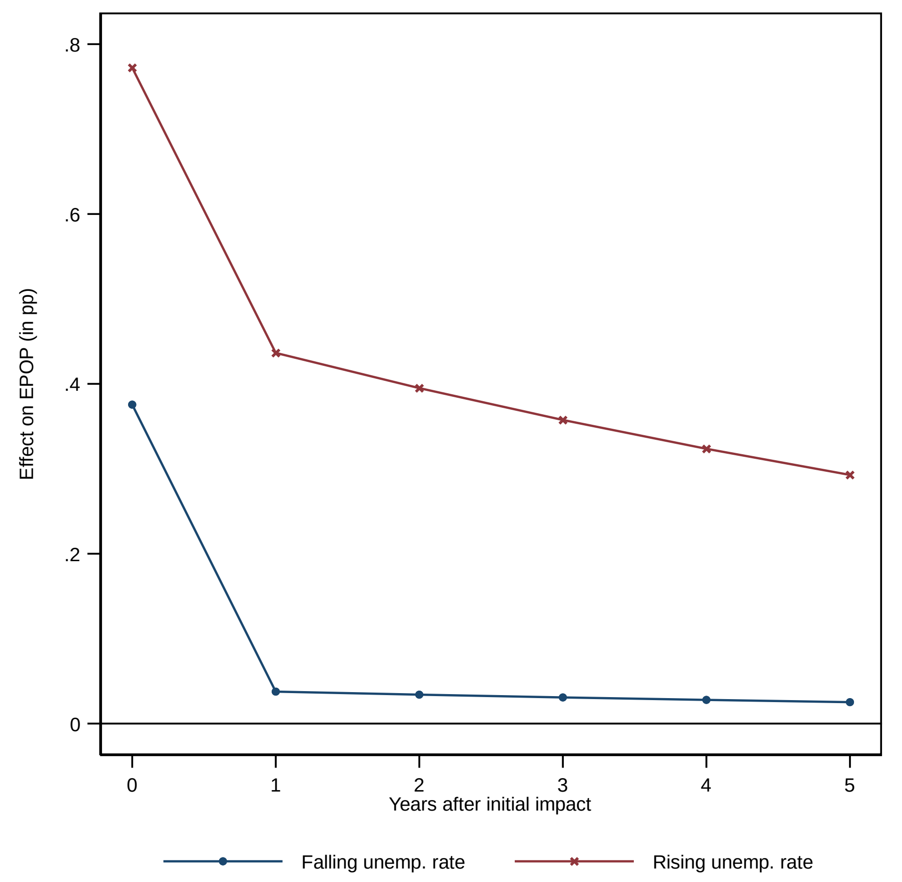

Note: Regression coefficient estimates (and standard errors clustered at the year level, in parentheses) and corresponding impulse responses from a 1 percentage point increase in the output gap, as shown in the table and figure, are derived from 2SLS regressions of the employment-to-population ratio for a birth cohort-age-year on the lagged EPOP, the output gap and lagged output gap, a quadratic in age, and birth cohort dummy variables. The EPOP and its lag are estimated only using observations from month-in-sample 4, and the lagged EPOP is instrumented with the lagged EPOP estimated with observations from month-in-sample 8. The lagged EPOP (RHS variable and instrument), output gap, and lagged output gap are interacted with indicator variables for whether the unemployment rate rose or fell since the previous year; each column displays regression coefficients and implied impulse responses corresponding with the appropriate indicator variables (that is, the regression coefficients in the table correspond with a single regression). The bottom panel of the table provides the 95% confidence interval of the implied shock at year 0 (the initial shock) and years 3 and 5, assuming no persistence in the GDP shock (top rows) or persistence as described in the text (bottom rows). The sample is limited to men age 25-54 without any college experience; included cohorts are those with 10 or more years of data for the years of the sample. The total number of cohorts is 43 and total number of observations is 1,034.

Figure 5: Effect of a 1 pp increase in the output gap on the EPOP for non-college men
Comparing differences in persistence by age

A. Regression coefficient and standard errors

	Age 25 (1)	Age 35 (2)	Age 45 (3)
Implied effect of RHS variables evaluated at corresponding ages:			
Lagged employment rate	0.63	0.74	0.79
	(0.14)	(0.15)	(0.20)
Output gap	0.89	0.80	0.67
	(0.17)	(0.11)	(0.11)
Lagged output gap	-0.29	-0.36	-0.31
	(0.20)	(0.17)	(0.22)
Implied impact of a 1 pp increase in the output gap, 95 percent CI			
No persistence in GDP shock			
Year 0	0.60 to 1.14	0.66 to 0.89	0.52 to 0.76
Year 3	-0.01 to 0.23	0.04 to 0.13	0.03 to 0.13
Year 5	-0.03 to 0.11	-0.00 to 0.05	-0.01 to 0.05
Persistence in GDP shock			
Year 0	0.60 to 1.14	0.66 to 0.89	0.52 to 0.76
Year 3	0.31 to 0.90	0.40 to 0.61	0.37 to 0.58
Year 5	0.09 to 0.58	0.18 to 0.36	0.16 to 0.34

B. Impulse Response Functions (no persistence in GDP shock)

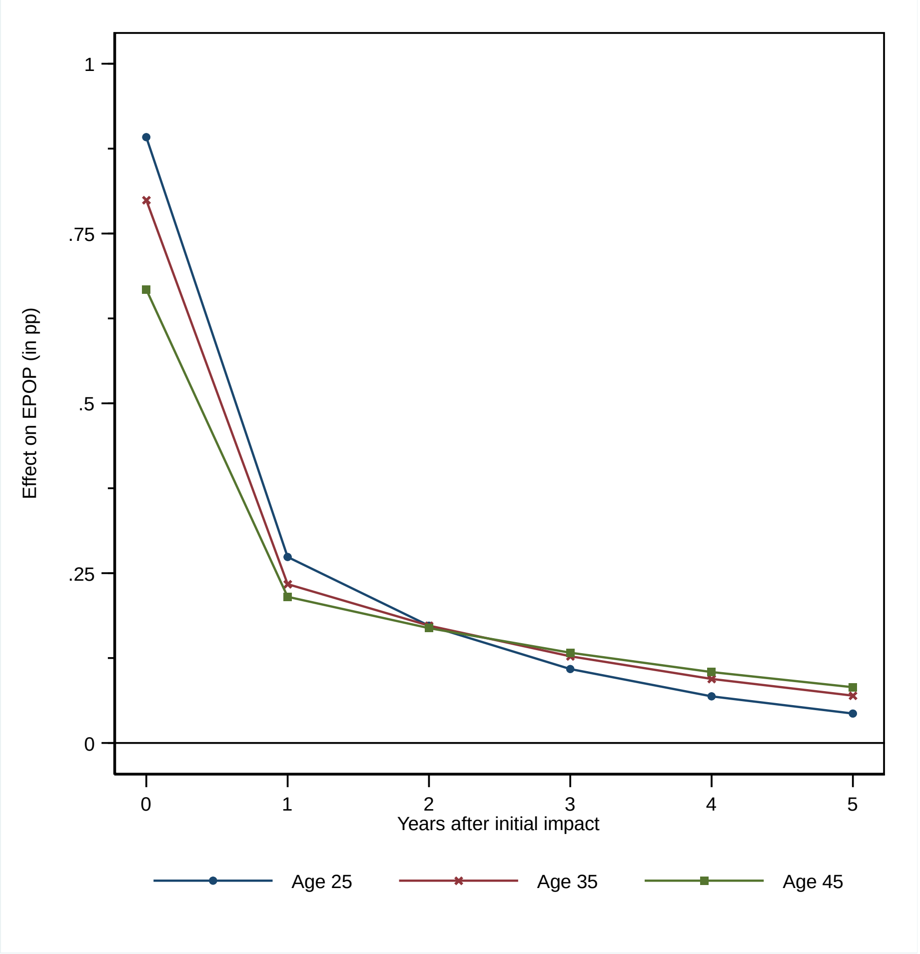

Note: Regression coefficient estimates (and standard errors clustered at the year level, in parentheses) and corresponding impulse responses from a 1 percentage point increase in the output gap, as shown in the table and figure, are derived from 2SLS regressions of the employment-to-population ratio for a birth cohort-age-year on the lagged EPOP, the output gap and lagged output gap, a quadratic in age, and birth cohort dummy variables. The lagged EPOP, output gap, and lagged output gap are interacted with a quadratic in age. The EPOP and its lag are estimated only using observations from month-in-sample 4, and the EPOP is instrumented with once- and twice-lagged EPOP estimated with observations from month-in-sample 8, interacted with a quadratic in age. The estimates in the table and figure correspond with a single regression; age-specific effects are estimated by evaluating the impact of the listed RHS variable (top panel of the table) or impulse response (bottom panel of the table and figure) at ages 25, 35, and 45. The bottom panel of the table provides the 95% confidence interval of the implied shock at year 0 (the initial shock) and years 3 and 5, assuming no persistence in the GDP shock (top rows) or persistence as described in the text (bottom rows). The sample is limited to men age 25-54 without any college experience; included cohorts are those with 15 or more years of data for the years of the sample. The total number of cohorts is 43 and total number of observations is 1,034.

Figure 6: Effect of a 1 pp increase in the output gap on the EPOP for non-college men
Comparing differences in persistence by race

A. Regression coefficient and standard errors

	White	Black	Non-white, non-black
	(1)	(2)	(3)
Lagged employment rate	0.50	0.77	0.75
	(0.10)	(0.36)	(0.37)
Output gap	0.68	0.87	1.25
	(0.10)	(0.14)	(0.31)
Lagged output gap	-0.07	-0.36	-0.78
	(0.12)	(0.46)	(0.53)
Implied impact of a 1 pp increase in the output gap, 95 percent CI			
No persistence in GDP shock			
Year 0	0.48 to 0.88	0.59 to 1.16	0.65 to 1.85
Year 3	0.02 to 0.11	-0.09 to 0.47	-0.19 to 0.38
Year 5	-0.01 to 0.04	-0.24 to 0.46	-0.20 to 0.31
Persistence in GDP shock			
Year 0	0.48 to 0.88	0.59 to 1.16	0.65 to 1.85
Year 3	0.36 to 0.55	0.20 to 1.24	-0.06 to 1.23
Year 5	0.16 to 0.30	-0.21 to 1.23	-0.35 to 1.04

B. Impulse Response Functions (no persistence in GDP shock)

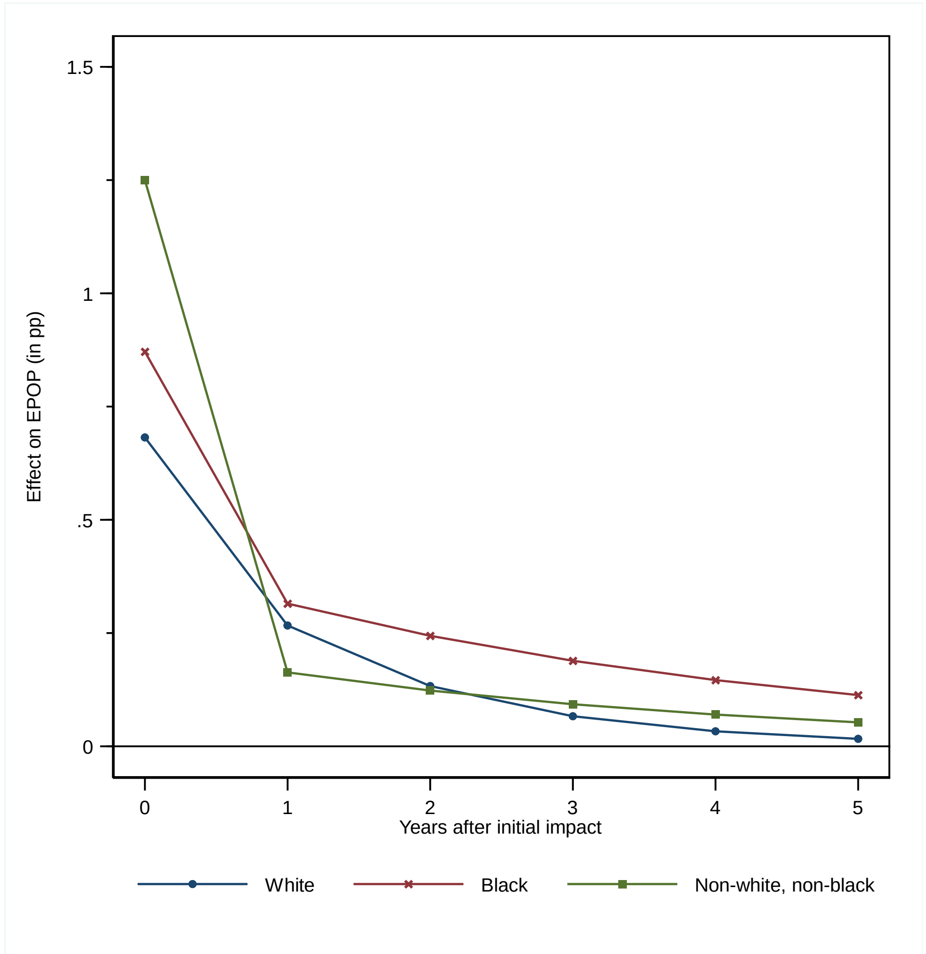

Note: Regression coefficient estimates (and standard errors clustered at the year level, in parentheses) and corresponding impulse responses from a 1 percentage point increase in the output gap, as shown in the table and figure, are derived from 2SLS regressions of the employment-to-population ratio for a birth cohort-age-year on the lagged EPOP, the output gap and lagged output gap, a quadratic in age, and birth cohort dummy variables. The EPOP and its lag are estimated only using observations from month-in-sample 4, and the lagged EPOP is instrumented with once- and twice-lagged EPOP estimated with observations from month-in-sample 8. The regression is estimated separately for each of three races (corresponding to the estimates in each column). The bottom panel of the table provides the 95% confidence interval of the implied shock at year 0 (the initial shock) and years 3 and 5, assuming no persistence in the GDP shock (top rows) or persistence as described in the text (bottom rows). The sample is limited to men age 25-54 without any college experience; included cohorts are those with 15 or more years of data for the years of the sample. The total number of cohorts is 43 and total number of observations is 1,034.