

NBER WORKING PAPER SERIES

THE IMPACT OF YOUNG WORKERS
ON THE AGGREGATE LABOR MARKET

Robert Shimer

Working Paper 7306

<http://www.nber.org/papers/w7306>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

August 1999

I have received helpful comments from Daron Acemoglu, Andrew Caplin, Hank Farber, Chris Foote, Bo Honore, Alan Krueger, Christina Paxson, Harvey Rosen, and seminar participants at George Washington University, NYU, and Princeton. Thanks to Susan Gore for providing me with cross-state unemployment and participation rates for different age and sex groups, to David Neumark for providing me with his cross-country time series of unemployment rates, to Joydeep Roy for research assistance, and to the National Science Foundation grant SBR-9709881 for financial support. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 1999 by Robert Shimer. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Impact of Young Workers on the Aggregate Labor Market

Robert Shimer

NBER Working Paper No. 7306

August 1999

JEL No. E24, J41, J64

ABSTRACT

This paper estimates the response of the unemployment rate and labor force participation rate to exogenous variation in the youth share of the working age population, using cross-state variation in lagged birth rates as an instrumental variable. A one percent increase in the youth share reduces the unemployment rate of young workers by more than one percent, and of older workers by more than two percent, holding conditions in other states constant. It raises the labor force participation rate by about a third of a percent for young workers, and by much less for older workers, again *ceteris paribus*. These results are consistent with increasing returns to scale (“thick market externalities”) in the labor market. Young workers are frequently mismatched in their employment, and firms create jobs to take advantage of this mismatch. Data on gross job creation and destruction in manufacturing support this theory. I also reconcile these results with existing evidence on the labor market impact of young workers.

Robert Shimer

Department of Economics

Princeton University

001 Fisher Hall

Princeton, NJ 08544-1021

and NBER

shimer@princeton.edu

1 Introduction

The baby boom has profoundly altered the demographic structure of the U.S. population during the past fifty years. A number of authors have argued that this anticipated supply shock can explain part or all of the secular changes in the unemployment rate during this period. First Perry (1970) predicted that the entrance of the baby boom cohort into the labor force would push up the unemployment rate during the 1970s. Later, Flaim (1979) and Gordon (1982) confirmed the increase, and predicted declines during the 1980s, which were in turn confirmed by Flaim (1990) and Shimer (1998).

The effects of the baby boom on unemployment can be grouped in two categories. First, since the aggregate unemployment rate is a weighted average of the unemployment rates of different age workers, demographic changes may alter the weights and thus the aggregate unemployment rate without affecting the age-specific unemployment rates. Shimer (1998) finds that this ‘direct’ effect of the baby boom can account for about an eighty four basis point (0.84 percentage point) increase in the aggregate unemployment rate from 1954 to 1978, and an eighty one basis point decline from 1978 to 1998.

Second, changes in demographics may have ‘indirect’ effects on age-specific unemployment rates. For example, the conventional neoclassical growth model predicts that an increase in the labor force growth rate will reduce the capital-labor ratio, raising interest rates and lowering wages. Augmenting such a model with labor market frictions, low wages may lead to high unemployment, for example if unemployed workers reduce their search effort. In addition, if different age workers are not perfect substitutes, an increase in the youth labor supply may have a differential impact on young and old workers.¹ Shimer (1998) estimates that the indirect effects of the baby boom were about as large as the direct effect, so that the baby boom caused a 180 basis point increase in the aggregate unemployment rate until 1980, and a 145 basis point decrease in the subsequent years.

Unfortunately, while the direct effect of this supply shock can be precisely es-

¹This possibility has received considerable attention from labor economists. See the literature review in Korenman and Neumark (1997).

timated, most previous calculations of the indirect effects of the baby boom use conjectures based on time series evidence. For example, Shimer (1998) attributes an increase in the youth unemployment rate relative to the prime age unemployment rate to the baby boom. Although the timing of this increase coincides with the entrance of the baby boom into the labor market, at least two criticisms can be levied at this attribution. First, there may have been coincident macroeconomic fluctuations that raised the youth unemployment rate relative to the prime age unemployment rate in the 1970s and 1980s. Second, if young and old workers are complements in production, the baby boom may have simply *reduced* the unemployment rate of prime age workers. If that is the case, the indirect effects of the baby boom at least partially offset the direct effects. It is impossible to reject these possibilities using time series evidence from a single baby boom.

Recent work by Korenman and Neumark (1997) partially addresses these issues by using time series data on unemployment rates for 15 countries. They look at the relative unemployment rate of young and old workers in an effort to muffle the noise introduced by country-specific macroeconomic shocks, and find that an increase in the youth share of the working age population raises the youth unemployment rate relative to the prime age unemployment rate, with an elasticity of approximately 0.5. The use of data from multiple countries allows them to include time dummies in their regression, thereby addressing the first criticism. However, the use of relative unemployment rates exposes it to the second: they cannot tell whether an increase in the youth share of the labor force lowers the prime age unemployment rate or raises the youth unemployment rate.

This paper's innovation is to focus on data from within the U.S.. I use annual observations of unemployment rates from all fifty states and the District of Columbia, from 1978 to 1996. Because of the relatively large sample size and the relative irrelevance of state-specific macroeconomic shocks, I can tightly estimate the impact of changes in the youth share of the population on youth and prime age unemployment rates. Contrary to the existing literature, I find that an increase in the youth share of the working age population *reduces* the youth unemployment rate, with an elasticity of about -1; and that the effect on the prime age unemployment rate is even

larger in magnitude. A one percent increase in the youth share of the population reduces the prime age unemployment rate by more than two percent. Not only are the signs of these estimates surprising, but the magnitudes are enormous, with an (out of sample) implication that the entry of the baby boom cohort into the labor market should have halved the prime age unemployment rate!

One possible explanation for these results is that young workers migrate to states with low unemployment rates. I control for the endogeneity of the youth share using instrumental variables techniques. When I instrument the youth share with appropriately lagged birth rates, I find that the estimates do not differ significantly from those found by ordinary least squares (OLS). Indeed, a specification test fails to reject the exogeneity of the youth share of the population. A number of robustness checks confirm the basic results.

Finally, I find that an increase in the youth share of the population raises the labor force participation rate rate for young workers, with an elasticity of about a third. It has a smaller effect on the participation rate of older workers, with an elasticity of around 0.05. This implies that the employment-population ratio rises for each group of workers.

The second task of this paper is to understand why an exogenous increase in the youth share of the working age population leads to such a dramatic reduction in unemployment rates and increase in participation. The change in the youth share of the working age population represents an anticipated supply shock. Standard theories predict that in response to an increase in the supply of an input, its price and utilization rate will decline. Yet the data indicate that labor utilization rates increase in response to an increase in the youth labor supply.

I propose that in the presence of search frictions and increasing returns to scale, labor markets with many young workers will be more fluid. There is no inherent difference between young and old workers; however, it takes time to find good matches, and so younger labor markets will tend to have more workers who are mismatched in their current employment. Firms will find it more profitable to participate in such labor markets, boosting job creation and reducing the unemployment rate of both young and old workers. In contrast, a labor market with mostly older workers is

more rigid. By increasing the fluidity of the labor market, an increase in the youth share of the population also induces more workers to participate in the market, consistent with the evidence.

A testable implication of this theory, is that labor markets with more young workers should have more turnover. Since reliable worker flow data does not exist on a state level, I instead use state level job flow data for the manufacturing sector (Davis, Haltiwanger, and Schuh, 1996). I find that a one percent increase in the youth share of the population raises the job creation rate by 0.8 percent and the job destruction rate by 0.7 percent, which supports the theory.

The third task is to reconcile the findings of this paper with the existing ‘cohort crowding’ literature on the relationship between the youth share of the population and the youth unemployment rate. I confirm that one cannot find the relationship in the 15-country time series assembled by Korenman and Neumark (1997). But this is not surprising, because cross-state data ignores other channels that lead to an increase in the unemployment rate. As previously noted, an increase in the growth rate of the labor force reduces the capital-labor ratio, raises unemployment and reduces wages. If capital flows freely across states, however, this effect will not appear in cross-state data; a large increase in the population growth rate in one state will cause a modest reduction in the capital-labor ratio in all states, which I will pick up as a year fixed effect and interpret as a macroeconomic shock. In contrast, to the extent that capital flows imperfectly across national borders, this effect will appear in the cross-country panel regression using Korenman and Neumark’s (1997) data. General equilibrium interactions may also rationalize why the entry of the baby boom into the labor force did not lead to a fifty percent decline in the U.S. unemployment rate.

Section 2 describes the data used in the main empirical analysis, whose results are presented in Section 3. Section 4 argues that the endogeneity of fertility decisions is unlikely to explain these results. Instead, I develop a simple model in Section 5 that illustrates how an increase in the youth share of the population can reduce the unemployment rate by increasing the fluidity of the labor market. Section 6 offers a simple test of this theory using job creation and destruction data from

manufacturing. Section 7 reconciles these results with existing evidence on the relationship between the youth share of the population and labor market conditions. Section 8 concludes by exploring the broader implications of the findings.

2 Data

The main empirical analysis draws on cross-state differences in birth rates within the U.S., and the consequent impact on the youth share of the population and on unemployment and participation rates. The basic source of unemployment and participation rate data is the Current Population Survey (CPS), which is designed to yield an accurate description of the national labor market. The Bureau of Labor Statistics (BLS) has estimated state unemployment and participation rates since 1970 by augmenting the CPS with information from the unemployment insurance system and ‘time series modeling’.² This yields an official series for the state rates, which is publically available from the BLS web site <http://stats.bls.gov/>. In addition the BLS computes (but does not make generally available) data on state unemployment and participation rates for different age cohorts from 1978 to 1996, my primary sample period.

The Census Bureau produces annual estimates of the number of workers in each state in many different age cohorts, supplementing the decennial census. Although the BLS also produces analogous numbers, BLS and Census data are surprisingly different. To avoid any measurement error that might be correlated with measurement error in the BLS unemployment rate estimates, I use Census data for the age structure of the population.

The third source of data is birth rate data for each state from 1954 to 1980 (16 to 24 years before my sample period). This data comes from various years of the Statistical Abstract of the United States, and is measured in births per thousand people. Whenever possible, I use birth rates corrected for undercounting, rather than the official birth census.

²Time series models are used for 41 small states. Unemployment rates are calculated directly from the CPS in the ten largest states (California, Florida, Illinois, Massachusetts, Michigan, New Jersey, New York, Ohio, Pennsylvania, and Texas). I return to the importance of the estimation procedure in the robustness checks.

3 Empirical Analysis

3.1 Empirical Model

The unemployment rate varies substantially over time. At a national level, it averaged 7.5% during the first half of the sample period, and fell to 6.1% during the second half. The age structure of the population shows similar temporal variation, with the share of youths, 16 to 24 year olds, in the working age population, 16 to 64 year olds, declining from 26.3% at the beginning of the sample to 18.9% at the end. Given the importance of macroeconomic shocks during this time period, it would be naïve to interpret this correlation causally.

Similarly, unemployment and demographics show considerable cross-sectional variation. The youth share of the working age population averaged 27.1% in Mississippi and 22.3% in Connecticut during the sample period, while the state unemployment rates averaged 7.9% and 6.0% respectively. It is not obvious whether there is a causal relationship, and if so, in which direction the causation goes.

To avoid these issues, I do difference-in-difference estimation. The empirical model looks at how the unemployment rate and labor force participation rate in state i and year t depend on state and year fixed effects and on the youth share of the working age population, share_{it} , defined as the number of 16–24 year olds divided by the number of 16–64 year olds:

$$\log \text{rate}_{it} = \alpha_i + \beta_t + \gamma \log \text{share}_{it} + \varepsilon_{it} \tag{1}$$

where the dependent variable rate_{it} is either the unemployment or participation rate in state i and year t , and ε represents other sources of variation in the dependent variable, such as state-specific economic shocks, which are orthogonal to the youth share of the population.³

³Other papers studying the behavior of state employment and unemployment rates do not include any other exogenous variables in the regression (Blanchard and Katz, 1992). There are a few papers that study the determinants of regional fluctuations, and find that variables like state-specific military spending and fluctuations in the price of oil interacted with cross-state differences in the mix of oil-sensitive industries have significant explanatory power (Davis, Loungani, and Mahidhara, 1997). If, as seems reasonable, these variables are uncorrelated with the de-measured youth share of the population, they will not bias the estimates of the variable of interest, the

The null hypothesis is that the elasticity γ is zero. Because I allow for both state and year fixed effects, I can only estimate γ to the extent that $\log\text{share}_{it}$ cannot be predicted by state and year fixed effects. Since about 94.9% of the variation in $\log\text{share}_{it}$ is predicted by the two fixed effects, only the remaining 5.1% of the variation can be used to estimate γ . However, given the fairly long sample period of nineteen years and large cross section of 51 ‘states’ (including the District of Columbia), this is enough to obtain tight estimates.

3.2 OLS Results

Panel A of Table 1 shows the results from estimating equation (1) using aggregate unemployment and participation rate data from 1970 to 1996. The estimated elasticity of the unemployment rate with respect to the youth share of the population is approximately -1, significantly different than zero at any standard confidence level. This is best interpreted as a partial equilibrium correlation. A one percent increase in the youth share of the working age population in one state is correlated with a one percent reduction in the unemployment rate in that state. The elasticity of the participation rate has the opposite sign and is much smaller, although it is still statistically significantly different than zero. A one percent increase in the youth share of the working age population is correlated with a 0.05 percent increase in the participation rate in that state.

For both regressions, both the state and year fixed effects are highly significant, justifying the panel data analysis. One can also compare the results of the fixed effects OLS regression with a random effects feasible generalized least squares estimate. The latter is more efficient if the random effects are uncorrelated with any omitted variables, but will otherwise be inconsistent; while the former is consistent in either case. A Hausman specification test cannot reject the null hypothesis that the random effects model is consistent in the unemployment rate equation; however, the coefficient estimates are not significantly changed compared with the fixed effects estimates. On the other hand, in the participation rate equation, the Hausman test strongly rejects random year effects. To avoid these inconsistent estimates, I use

elasticity γ . Finally, I have verified that the racial composition of a state has no effect on the youth share coefficient, and so can safely be omitted.

fixed effects estimates throughout the remainder of the paper, possibly sacrificing some efficiency. Since almost all of the estimates are statistically significant, this should not affect my conclusions.

The estimates in Panel A suffer from autocorrelated residuals, because an unexpected increase in the unemployment rate only gradually disappears as the laid-off workers find new jobs. A regression of the OLS residual $\hat{\varepsilon}_{it}$ on itself lagged one period, $\hat{\varepsilon}_{it-1}$, yields a coefficient of 0.73. A similar regression of the residuals in the participation equation yields a coefficient of 0.72. There is no evidence of higher ordered autocorrelation, so Panel B performs an AR(1) correction. The point estimates of the elasticities are statistically unchanged. Although the standard errors are somewhat larger, the results remain significant. All the results in the remainder of the paper include an AR(1) correction.

There is also a mechanical bias in the estimates in Table 1. Since young workers have a higher unemployment rate, an increase in the youth share of the population will raise the aggregate unemployment rate. On the other hand, young workers generally have low participation rates, so an increase in the youth share of the population will lower the aggregate participation rate. Similarly, male and female labor market participation rates systematically differ; and the relationship between them changes with age. This too may bias the estimates in Table 1.

To address these concerns, I estimate equation (1) separately for seven different age groups and both sexes. Column I in Table 2 shows the results. For the data availability reasons described in Section 2, I use the sample period 1978–96. In each case, the independent variable of interest is the youth share of the working age population. The dependent variable is the unemployment or participation rate of a particular age and sex group in the state at that point in time. The first row also shows the results from a regression of the aggregate unemployment or participation rate on the youth share using this shorter sample period.

The results are qualitatively similar for almost all age and sex groups. An increase in the youth share of the population is correlated with a statistically significant reduction in the unemployment rate of all workers contingent on their age and sex (Panel A). Quantitatively, an increase in the youth share of the population has twice as large an effect on the unemployment rate of older workers as on younger

workers. The elasticity of the teenage unemployment rate is -1.3, and the elasticity of the unemployment rate for workers age 45–64 is -2.5. As a result, the relative unemployment rate of young workers rises in response to an increase in the youth share of the population, consistent with the ‘cohort crowding’ hypothesis (Korenman and Neumark, 1997). However, the absolute decline in the youth unemployment rate is not consistent with the standard hypothesis. The response of male and female unemployment to changes in the youth share is the same.

An increase in the youth share of the population is correlated with a much larger increase in the participation rate of young workers than of older workers (Panel B). Since the results in Panel A suggest that labor market conditions improve more for older workers than for younger workers, one might expect the opposite relationship. However, the participation rate results make sense if young workers have a more elastic labor supply. Again, the response of male and female unemployment is about the same for each age group.

Finally, a reduction in the unemployment rate and an increase in the participation rate imply an increase in the employment-population ratio. One can show that the employment-population ratio is higher for each age and sex group when the youth share of the population is higher. Thus these are generally periods of tight labor markets.

3.3 IV Results

One possible explanation for the results in Column I of Table 2, is that young workers who are eager to participate in the labor market flock to states with low unemployment rates. Note that the story is not as simple as saying that state i always has low unemployment rates, so young workers move to state i . A persistently low unemployment rate would be captured by state i ’s fixed effect. The concern is more subtle: a temporary reduction in the unemployment rate in i might temporarily attract more young workers to i .

I control for this possibility using instrumental variables (IV). I look for exogenous variation in the youth share of the working age population, share_{it} , caused by the birth rate in that state 16 to 24 years before. More precisely, birth_{it} in year

$t = 1978$ is equal to the sum of the number of births per person in state i from 1954 to 1962. A simple regression of $\log \text{share}_{it}$ on $\log \text{birth}_{it}$ and a constant yields an elasticity of 0.647, with a standard error of 0.007. This single variable explains 86.0% of the variation in the youth share. Including state and year fixed effects lowers the elasticity estimate to 0.595 with a standard error of 0.016, and the R^2 rises to 0.977. Since the elasticity is less than one, an increase in the birth rate leads to a smaller increase in the youth share 16–24 years later, a tendency towards mean reversion. However, the instrument is a remarkably good predictor of future youth shares.

Column II of Table 2 then estimates the basic regression using IV instead of OLS. None of the results are significantly changed, and the magnitudes are generally larger, not smaller. A one percent exogenous increase in the youth share of the population, caused by an increase in the birth rate 16 to 24 years before, reduces the unemployment rate of teenagers by one percent and of older workers by two or three percent. It raises the participation rate of teenage workers by about half a percent, and slightly raises the participation rate of older workers. Because the standard errors are somewhat larger using IV, a few of the elasticity estimates are not significantly different from zero.

Since lagged birth rates and state and year fixed effects predict 97.7% of the variation in the youth share of the population, the youth share of the population cannot fluctuate too much in response to short term economic conditions. Thus it should not be too surprising that IV and OLS estimates are similar. Indeed, the predictability of future youth population shares suggests that instrumental variables may be inappropriate. If the youth share of the population is exogenous, OLS is a consistent and efficient estimator, while IV is inefficient. I examine whether this is the case using a two-stage procedure proposed by Davidson and MacKinnon (1993). In the first stage, predict the youth share of the population from a regression on all the exogenous variables, here the lagged birth rates and state and year dummies. Then in the second stage, regress the unemployment or participation rate on the youth share of the population, the *predicted value* of the youth share, and state and year dummies. If the predicted value of the youth share enters significantly into this regression, then we can reject the null hypothesis that the youth share is exogenous.

The final column in Table 2 shows the p-values from this exogeneity test, which is just the p-value for a t-test of whether the coefficient on the predicted youth share is different from zero. One can reject exogeneity at the ten percent confidence level in only eleven of 44 cases, and at the five percent level in only four cases. Moreover, in about half the cases where one can reject exogeneity of the youth share at conventional significance levels, IV has the ‘wrong’ effect on the elasticity estimates. Exogenous variation in the youth share has a more negative effect on the unemployment rate, and a more positive effect on the participation rate, than does endogenous variation. I conclude that OLS is a consistent and efficient estimator, and use it throughout the remainder of the paper.⁴

3.4 Robustness

This section describes several robustness checks on the results in Table 2. All support the preceding results.⁵

Regional Fluctuations

Shocks are likely to be correlated according to economic, cultural, and geographic proximity. This may bias down the standard errors reported in Table 2. Ideally, one could estimate the complete cross-sectional variance-covariance matrix of shocks, but this is impossible because I have fewer time observations than states. Instead, I try several alternative controls for regional economic fluctuations.

⁴This may appear to contradict the conclusion of Blanchard and Katz (1992), that regional labor markets adjust to shocks via migration. There are two possible reconciliations. First, their findings concern relatively long term adjustment, while with my methodology, long term adjustment is dumped into state fixed effects. Second, the population variables that I use are estimated by the Census Bureau in years between decennial censuses. To the extent that the Bureau relies on lagged birth rates to construct the population variables, I will be unable to observe temporary population fluctuations. As a result, the *measured* youth share of the population is already somewhat exogenous to current labor market conditions.

⁵I have performed, but do not report, a number of other robustness checks. These include population weighted regressions, quantile regressions to control for outliers, and a variety of sub-sample estimates: the 41 small states; within census regions; during recessions or expansions; and in each half of the sample period. None of the results contradict the main findings, although in some cases (e.g. within the East and Midwest census regions) the standard errors explode and the point estimates are imprecise. This reflects the fact that much of the usable variation in population shares is at the regional level, as pointed out in Columns I and II of Table 3.

First, I include a full set of census division/year dummies in the basic regression. The Census Bureau divides the country into nine geographic divisions, ranging in size from three states (New York, New Jersey and Pennsylvania) to nine (South Atlantic division). Let $d(i)$ denote the census division of state i . I include a separate year fixed effect for each census region, $\beta_{d(i),t}$, constant for any observation in census division d and year t and zero otherwise. This will pick up regional fluctuations.⁶

Column I of Table 3 shows the results from estimating this equation.⁷ Only one of the point estimates changes significantly (at the five percent level) from the point estimates in Table 2. This is what we would expect, since regional fluctuations should be uncorrelated with the youth population share.

More important is what happens to the standard errors. Including all these extra dummy variables comes at some cost in terms of degrees of freedom. If, in fact, shocks are independent across states within census divisions, this should raise the estimated standard errors by a little over 9%. On the other hand, to the extent the results are driven exclusively by fluctuations within census divisions, the standard errors should blow up, which would indicate that the standard errors reported in Table 2 are too small. We find that in the unemployment rate equation, the standard errors increase by about 15% on average, while in the participation rate equation they increase by 34%, evidence for a weak correlation of shocks within census divisions. However, this correlation is not so large that it alters the statistical significance of any of the results.

A complementary way to control for regional economic fluctuations is to aggregate all the data to the regional level and rerun the regression. That is, I regress the log unemployment rate in region d and year t , on region and year dummies and the log youth share of the working age population in region d and year t . This aggregation comes at an even larger cost in terms of degrees of freedom, since we are throwing away over 82% of the observations. However, this should completely eliminate any concern about correlation of shocks.

The results are presented in Column II of Table 3. An exogenous increase in

⁶Allowing for state-specific time trends yields similar results.

⁷I omit the results for sex-specific unemployment and participation rates to avoid overwhelming the reader with numbers. The results are again similar.

the youth share of the population in a census division causes a decrease in the unemployment rate, possibly larger for older workers than for teenagers. It causes an increase in the youth participation rate, and we cannot discern a significant effect on prime age participation. This confirms the robustness of the main results.

Estimates Using Only Large States

Recall that the BLS uses ‘time series methods’ to calculate the unemployment and participation rate in many states. Is it possible that this estimation procedure is somehow driving the results? One way to answer that question, is by running the regression using only data from the ten large states where the rates are computed directly from the CPS. Column III of Table 3 reports those results. They are consistent with the full sample estimates, eliminating this cause for concern.

Sectoral Composition of Employment Growth

An increase in the youth share of the population causes an increase in participation and a sharp decline in the unemployment rate, and so it must lead to an increase in total employment. I test this using state level data on total employment from 1970–97, also available from the BLS website. Regressing log total employment on the log youth share and year and state dummies yields an elasticity of 0.36 (Table 4, Panel A); controlling for endogeneity of the youth share raises the estimate slightly. Moreover, upon breaking employment down by sector, one sees that the response is widespread, with the strongest effect on construction, and a slightly weaker impact on manufacturing, wholesale and retail trade, and services. This conclusion holds whether we estimate the elasticity of the level of employment (Panel A) or the sectoral employment-working age population ratio (Panel B).

The large response of construction employment might suggest that the strong labor market conditions are due to a construction boom, as residential and commercial real estate are built to accommodate a large youth cohort. However, if that were the case, the price of labor should be bid up, causing substitution away from employment in other sectors. This is inconsistent with the evidence in Table 4. Instead, the economy-wide increase in employment suggests a growth in labor de-

mand, which then leads to a construction boom. The large response of construction employment is simply due to the fact that that sector is highly cyclical.

3.5 Summary

Cross-state evidence suggests that an exogenous one percent increase in the youth share of the population in a state will reduce the unemployment rate of workers in that state by over one percent, with the strongest effect on prime age unemployment and a somewhat weaker effect on young workers. It will also raise the participation rate of young workers by at least one-third of a percent, with a smaller and less significant effect on prime age workers.

These elasticities not only have an unexpected sign, but they are very large. Consider the regression of the log youth share of the population, $share_{it}$, on state and year dummies. The residual has a standard error of 0.033, so a one standard deviation increase in the youth share of the population, relative to that state's history and to the youth share in other states at that point in time, will reduce the unemployment rate in that state by about six percent and raise the participation rate of young workers by about one percent.

The implied impact of the baby boom on the aggregate unemployment rate is simply enormous. The youth population share bottomed out at 18.0% in 1953, rose to 26.7% by 1976, and has since fallen back to 19.0%. Roughly speaking, this change should have first halved and then doubled the prime age unemployment rate! However, one must be cautious with this calculation, since the estimated elasticities concern changes in the youth population share in one state relative to others, while the baby boom was an international phenomenon that may have induced general equilibrium effects.

4 The Endogenous Fertility Hypothesis

The empirical findings contradict the existing literature on cohort crowding. Why might an increase in the youth share of the population caused by an increase in the birth rate twenty years earlier lead to a reduction in the unemployment rate and an

increase in the participation rate? One possibility is that a third factor affects the birth rate now and will later alter unemployment and participation rates.

Birth rates are endogenously determined by the collective decisions of parents. Suppose parents expect a strong labor market in the future. This might lead to an increase in fertility today, as parents anticipate being able to support more children, or as they have more children in an effort for the family to take advantage of the economy. I label this the endogenous fertility hypothesis.

While endogeneity of fertility is undoubtedly important in many situations, it seems unlikely to be driving the empirical findings. These are ‘difference-in-difference’ estimates. Parents in state i from year $t - 24$ to year $t - 16$ must anticipate low unemployment in state i in year t , relative to the norm in state i and relative to other states in year t . Such precise beliefs seem implausible.

To get at this more concretely, I try other demographic variables on the right hand side of the regression. I start with share_{it}^{5-15} , the number of 5–15 year olds in state i and year t as a fraction of the working age population 16–64. The first line in each block of Table 5 shows that an increase in the share of ‘school children’ in a state raises the unemployment rate and reduces the participation rate for each working age group, opposite to the previous results. Next I try $\text{share}_{it}^{25-34}$, the share of 25–34 year olds in the working age population. This yields qualitatively similar results, although they are quantitatively and statistically less significant. The timing of births implied by these results is so precise and the foresight required is so incredible, that I reject the endogenous fertility hypothesis.

The question remains as to why the share of school children and 25–34 year olds enters significantly into these regressions. One possibility is the mechanical negative correlation between these share variables and the share of 16–24 year olds. To see whether this is the case, I include the share of 16–24 year olds in the two regressions, and report the results in the remaining lines of each block. The effect of the share of 25–34 year olds largely goes away. However, the effect of school children appears to be robust, and slightly dampens the effect of 16–24 year olds. Again, this result is quite surprising. One would expect the parents of school children to be more attached to the labor market, but an increase in the number of school children raises the unemployment rate and reduces the participation rate in a state for all

age groups. Although a detailed exploration of this finding goes beyond the scope of this paper, I will briefly return to it at the end of the theory developed in the next section.

5 The Fluid Labor Market Hypothesis

If the birth rate is exogenous, then the relationship between the youth population share and the unemployment and participation rates must be interpreted causally. The most obvious causal connection is that an exogenous increase in the youth share of the population leads to more job creation, reducing the unemployment rate and raising the participation rate. However, simple theories of equilibrium unemployment do not admit such behavior (Pissarides, 1990). An increase in the youth share of the population directly raises the unemployment rate, because young workers are unemployed more frequently. Even if job creation responds more than proportionately, as would be the case if there is some form of increasing returns to scale in the labor market (Diamond, 1982), this cannot lead to a *decline* in the aggregate unemployment rate, or it would choke off the new job creation.

Instead, there must both be increasing returns to scale and a reason for firms to create jobs even when unemployment falls.⁸ This section uses a very simple model to illustrate a plausible mechanism. Younger workers are more likely to be mismatched with their current employer. Having many mismatched workers encourages job creation, since firms that locate in such markets find it relatively easy to attract employees.

⁸Some readers will strongly oppose the increasing returns to scale assumption. Indeed, the standard theory of equilibrium unemployment (Pissarides, 1990) assumes constant returns to scale, typically justifying this by the empirical work of Blanchard and Diamond (1989). However, this finding is controversial. Hall (1989) questions it in his discussion of their paper. And in a later paper, Blanchard and Diamond (1992) recognize that “the process through which workers and jobs find each other, surely has increasing returns over some range.” The empirical evidence in this paper supports this view.

5.1 Model

There are two types of agents, workers and firms. Let $L(t)$ denote the measure of workers in the economy at time t and $\theta(t)L(t)$ denote the measure of firms. Both types of agents are risk neutral, infinitely lived, and discount the future at rate $r > 0$. Each agent may be in one of three states: unmatched, mismatched, or well matched. While unmatched, an agent earns nothing. A mismatched agent (worker or firm) gets $x_1 > 0$. A well matched agent gets $x_2 > x_1$.⁹ In addition, new workers are born, unmatched, at rate $n(t) > 0$, so $L(t) = L_0 e^{\int_0^t n(s) ds}$. New firms enter the market by paying a one-time fixed cost c . The assumption that the supply of jobs is perfectly elastic is reasonable in this setting, since the empirical evidence looks at cross-state variation in birth rates. A change in the birth rate in one state will lead to a flow of capital from other states, but will not change the cost of capital.

Let $\alpha_1(t)$ denote the fraction of workers who are mismatched and $\alpha_2(t)$ denote the fraction who are well matched in steady state, so $1 - \alpha_1(t) - \alpha_2(t)$ is the fraction who are unmatched. Simple accounting shows that at time t , a fraction of firms $(\theta(t) - \alpha_1(t) - \alpha_2(t))/\theta(t)$ are unmatched, while fractions $\alpha_1(t)/\theta(t)$ and $\alpha_2(t)/\theta(t)$ of the firms are in bad and good matches, respectively.

These fractions are limited by search frictions. Workers and firms periodically meet and have an opportunity to match. A firm meets a worker at a rate $\eta(\theta)$, decreasing in the contemporaneous firm-worker ratio $\theta(t)$. A worker meets a firm at an increasing rate $\mu(\theta) \equiv \theta\eta(\theta)$. In each case, the potential partner is drawn randomly from the appropriate population, regardless of the partner's current match quality. The two agents then realize the quality of their match; it is good with probability p , independent across workers and firms and over time.

Finally, potential partners match if it improves both agents' state. That is, a good match is accepted unless either agent is already in a good match. A mismatch is accepted only if both agents are unmatched. Thus it is easier to match when the fraction of unmatched or mismatched agents is high. As in Diamond (1982), this 'thick-market externality' is the source of increasing returns to scale.

Whenever a matched agent accepts a new partner, her old partner becomes

⁹Note that these payoffs are not transferable. This assumption simplifies the analysis by eliminating wage determination issues.

unmatched.¹⁰ There is no recall of past partners. One can think of a worker leaving a mismatch for a good match as a quit; a firm leaving a mismatch for a good match as a layoff; unmatched workers as unemployed; and unmatched firms as vacant.

Equilibrium demands two things. First, the two state variables $\alpha_1(t)$ and $\alpha_2(t)$, the fraction of mismatched and well matched workers, are determined by worker and firm flows and initial conditions $\alpha_i(0) = \alpha_{i0}$, $i \in \{1, 2\}$. And second, the value of creating a new unmatched firm must always equal the startup cost c .

This model builds on work by Burdett and Mortensen (1998), who introduce on-the-job-search into an otherwise standard search framework. There are two significant differences between the models. First, their model is more ambitious, in that it endogenizes the division of match surplus, and demonstrates that wage dispersion is an equilibrium phenomenon. I treat the division of surplus as exogenous, and take wage dispersion ($x_1 < x_2$) as a primitive. Second, I only allow firms to hire one worker, and thus develop a theory of temporary employment. In the Burdett-Mortensen model, the marginal product of labor is constant, so firms never fire workers. One can show that without temporary employment, this framework cannot explain a decline in the aggregate unemployment rate in response to an increase in the birth rate.

5.2 State Variables

A good match is an absorbing state in this economy. The measure of good matches increases when unmatched or mismatched firms meet unmatched or mismatched workers and the pair matches well, and never declines.

$$\frac{d\alpha_2(t)L(t)}{dt} = \mu(\theta(t))L(t)(1 - \alpha_2(t))\frac{\theta(t) - \alpha_2(t)}{\theta(t)}p$$

The right hand side of this equation describes the flow of new good matches at time t . There is a flow $\mu(\theta(t))L(t)$ meetings by workers, a fraction $1 - \alpha_2(t)$ of which are not well matched; and a fraction $(\theta(t) - \alpha_2(t))/\theta(t)$ of those meetings are with firms

¹⁰These matching patterns are efficient if $x_2 > 2x_1$. Otherwise, output might be higher if mismatched agents did not match with other mismatched agents. However, since utility is non-transferable, a spurned partner can still not prevent the termination of her match.

that are not well matched; and a fraction p of those meetings result in a good match. Since $dL(t)/dt = n(t)L(t)$, the labor force cancels from the above equation, yielding a differential equation for α_2 in terms of model parameters and the firm-worker ratio θ .

$$\dot{\alpha}_2(t) = \eta(\theta(t))(1 - \alpha_2(t))(\theta(t) - \alpha_2(t))p - n(t)\alpha_2(t) \quad (2)$$

The expression for the measure of mismatches is more complicated. The share of mismatches decreases due to new births, quits, and layoffs. It increases when previously unmatched workers and firms meet. The difference is the growth rate of the share of mismatches.

$$\begin{aligned} \dot{\alpha}_1(t) = & \eta(\theta(t))(1 - \alpha_1(t) - \alpha_2(t))(\theta(t) - \alpha_1(t) - \alpha_2(t))(1 - p) \\ & - \left(n(t) + \mu(\theta(t))\frac{\theta(t) - \alpha_2(t)}{\theta(t)}p + \eta(\theta(t))(1 - \alpha_2(t))p \right) \alpha_1(t) \end{aligned} \quad (3)$$

The first term on the right hand side reflects the fact that new mismatches occur only if both partners are unmatched. The first term on the second line is the adjustment for new births. The second term is the probability that the worker quits for a good match. The third term is the probability that the firm lays off the worker for a good match. Equations (2) and (3) can be solved numerically for α_1 and α_2 in terms of the firm-worker ratio θ and initial conditions.

Now consider the effect of a permanent increase in the birth rate $n(t)$, holding the firm-worker ratio constant. From equation (2), this will push down the steady state share good matches, as young workers have not yet had a chance to find them. It will raise the share of mismatches, since there are more workers passing through that intermediate stage. And it will push up the unemployment rate, as many workers will not yet have even found a first job, or will be suffering from layoffs.

5.3 Entry

For a firm, the reduction in the share of good matches is advantageous. When firms meet well matched workers, their overtures are always rejected. Instead, firms prefer

to operate in environments with more mismatched or unmatched workers. To show this formally, calculate the value of a firm in different states, letting $W_0(t)$, $W_1(t)$, and $W_2(t)$ denote the expected present value of profits for an unmatched firm, a mismatched firm, and a well-matched firm, respectively.

A well-matched firm earns x_2 forever, yielding a payoff $W_2(t) \equiv x_2/r$. A mismatched firm earns x_1 , but may suffer a quit or be able to create a good job.

$$rW_1(t) - \dot{W}_1(t) = x_1 + \mu(\theta(t)) \frac{\theta(t) - \alpha_2(t)}{\theta(t)} p(W_0(t) - W_1(t)) + \eta(\theta(t))(1 - \alpha_2(t))p(W_2(t) - W_1(t)) \quad (4)$$

The flow value of a firm in a mismatch is equal to the sum of the current payoff x_1 , plus the probability that its employee realizes a good match with a firm not already in a good match, times the capital loss associated with becoming unmatched, plus the probability that the firm realizes a good match with a worker not already well-matched, times that capital gain. Similarly, the value of an unmatched firm satisfies:

$$rW_0(t) - \dot{W}_0(t) = \eta(\theta(t))(1 - \alpha_2(t))p(W_2(t) - W_0(t)) + \eta(\theta(t))(1 - \alpha_1(t) - \alpha_2(t))(1 - p)(W_1(t) - W_0(t)) \quad (5)$$

The first term is the probability of realizing a good match with a worker not already in a good match, times the resulting capital gain. The second term is the probability of realizing a bad match with an unmatched worker, times that capital gain.

For a given time-path of θ , α_1 , and α_2 , equations (4) and (5) can be solved for W_0 and W_1 . For this to be an equilibrium, it must be the case that $W_0(t) = c$ for all t , the free entry condition. Thus an equilibrium is a tuple of functions $\{\alpha_1, \alpha_2, W_1, \theta\}$ that satisfies equations (2)–(5) with $W_0(t) \equiv c$ and $W_2(t) \equiv x_2/r$.

In practice, I solve this system numerically using a simple algorithm. Conjecture a path for the firm-worker ratio, perhaps from a steady-state solution. Use this to calculate a candidate path for the state variables from equations (2) and (3). Then under the assumption that $W_0(t) \equiv c$ and $W_2(t) \equiv x_2/r$, solve equation (4) for $W_1(t)$ and invert equation (5) to solve for $\theta(t)$. If the solution coincides with the

initial guess, this is the equilibrium. Otherwise, perturb the initial guess towards the firm-worker ratio that comes out of the algorithm, and restart. In practice, this algorithm rapidly converges to a unique equilibrium for any initial conditions.

5.4 Simulation

Clearly this model is too simple to be taken quantitatively seriously. However, a simulation shows how an increase in the population growth rate can lead to a sufficient increase in job creation, so as to reduce unemployment.

Let $\mu(\theta) = 10\sqrt{\theta}$ and $\eta(\theta) = 10/\sqrt{\theta}$. Take $x_1 = 1$ and $x_2 = 2$, with $p = 0.04$, so few potential matches are good. Also let the interest rate $r = 0.05$ and set the population growth rate $n = 0.02$, reasonable numbers for annual data. Finally, set the entry cost $c = 30$. Run the model for many periods, so the system converges to a steady state, independent of the initial conditions. In this steady state, $\theta = 1.00$, while 5.2% of workers are unmatched; 14.8% are mismatched; and 80.1% have found good matches.

Now consider an anticipated permanent increase in the population growth rate to $n = 0.03$ in year T , starting from the $n = 0.02$ steady state. If the firm-worker ratio did not change, the resulting increase in the youth population would eventually raise the unemployment rate by 1.1 percentage points and the share of mismatches by 2.7 percentage points. However, this creates profit opportunities for firms. In response, θ increases to 1.05, which reduces the steady state unemployment rate to 4.5% and leaves the share of mismatches at 17.7%. All workers benefit from this increase in θ , as it makes it easier to find jobs when unmatched, to find good jobs when mismatched, and it reduces the chance that an employer will lay them off to hire another worker.

This is a dynamic model, so I can look not only at the new steady state, but at the dynamic adjustment path (Figure 1). Since firms enter the market in anticipation of the high population growth rate, about 9% of the decline in unemployment occurs *before* the increase in population growth in year T . The remaining adjustment takes some time, although about 95% of the decline in unemployment happens by year $T + 10$.

I can also check whether the decline in unemployment rates is larger for young or old workers. Since by assumption all newborn workers are unemployed regardless of the birth rate or firm-worker ratio, it must in some sense be bigger for old workers. Indeed, in this example the change in the birth rate reduces the steady state unemployment rate of workers with ten years of labor market experience by 49%; and the unemployment rate of workers with thirty years of experience by 69%; and is generally a monotonic function of experience. This demonstrates that the model is consistent with the empirical evidence that a change in the youth share of the population has a bigger effect on the prime age unemployment rate than the youth unemployment rate (Table 2, Panel A).

5.5 Discussion

The simplicity of this model highlights its main elements. In a labor market with many older workers in good matches, firms are reluctant to create jobs, because they meet mismatched young workers too infrequently. The conclusion might seem sensitive to the assumption that well-matched workers continue to meet firms. If firms could focus their search on unmatched or mismatched workers, the mechanism in this paper would disappear. Moreover, in this model, well matched workers have nothing to gain by meeting firms, so it might be feasible for firms to avoid meeting such workers.

Such an interpretation would take the model too literally. A straightforward extension allows for many job qualities. Fairly well matched workers continue to search because there is always a chance of finding a better job. Even if they can reduce the time and effort devoted to search as they climb the quality ladder, they will still clog the labor market to some degree.

Another critique of the model is that it fails to address the relationship between the youth share of the population and the participation rate. However, a simple extension to the model endogenizes workers' participation decision. Assume that more workers participate in the labor market when the expected value of participating is higher. Then participation rates would also increase following an increase in the birth rate, as the value of workers' participation goes up when the firm-worker

ratio θ increases.

This model also sheds some light on why the unemployment rate is higher when the share of school children is higher. I offer two conjectures. First, new parents expend additional search effort to make sure that they find good jobs. By the time a large cohort of children is in school, the high average match quality will create a more rigid labor market, reducing job creation and raising unemployment. Alternatively, parents of school children are geographically less mobile, again creating a rigid labor market.

6 Testing the Fluid Labor Market Hypothesis

A first-order testable implication of the fluid labor market hypothesis is that there should be more voluntary quits in younger labor markets. Firms create more jobs to take advantage of mismatched workers, more of whom then quit their old jobs to take new ones in equilibrium. Indeed, the effect of the youth share of the labor market on quits may be much larger than the effect on the unemployment rate. In the numerical simulation in the previous section, the increase in the birth rate from 2% to 3% raised the quit rate from 1.2% to 1.9%.¹¹

An empirical test of this hypothesis would examine whether states with younger labor markets have more job-to-job movement. Unfortunately, reliable worker flow data is unavailable on a state level. An alternative is to focus on the flow of jobs. When a worker quits her job, one firm decreases its employment level, while another increases its employment. This should show up in the data as simultaneous job creation and destruction.

I test the theory with gross job creation and job destruction data in the manufacturing sector, using time series constructed from the Longitudinal Research Database by Davis, Haltiwanger, and Schuh (1996). This includes annual observations for 49 states from 1973–88.¹² Job creation in state i and year t measures

¹¹The effect on layoffs is ambiguous. With more firms in the market, a mismatched firm is less likely to contact a new potential employee; on the other hand, that potential employee is more likely to be unmatched or mismatched. In the numerical example, the layoff rate increased from 1.2% to 1.5%. In any case, a layoff (in the sense of the model) would not show up in the data set I use, because the same firm simultaneously creates and destroys a job.

the employment increase among expanding or newly created plants, expressed as a percentage of manufacturing employment in that state and year; job destruction measures the employment decrease (a positive number) among contracting or closing plants in state i and year t , again expressed as a percentage.

To test whether a younger state has more gross creation and destruction, I regress the log of job creation and job destruction on the log youth share and a full set of year and state dummies. The null hypothesis is that the elasticity of the creation and destruction variables should be equal to zero, while the theory predicts a positive value for *both* elasticities, and a larger value for the elasticity of creation than for destruction, reflecting an increase in net job creation. Table 6 shows the results from estimating the equation using OLS and IV. The point estimates are positive, and the estimated elasticity of job creation is indeed larger than that of job destruction; however, the estimated elasticity of job destruction is not statistically different from zero, nor are the estimated elasticities different from each other.

The theory also predicts that an increase in the youth share should have a larger effect on the creation and destruction rate than on total employment growth. A comparison of the point estimates here with the estimates for the manufacturing sector in Table 4 offers modest support for that prediction, although the difference between the estimates is again not statistically significant. I conclude that job creation and job destruction data do not contradict the fluid labor market hypothesis, and offer some modest support.

7 Reconciliation with Previous Studies

Previous studies have found that an increase in the youth share of the population raises the youth unemployment rate. For example, Korenman and Neumark (1997) conclude from a cross-country data set that a one percent increase in the youth share of the population raises the youth unemployment rate by about a third of a percentage point, although their estimate is not very significantly different than zero. The major methodological difference between that study and this one, is that they

¹²No data are available for Hawaii and Rhode Island.

include the prime age unemployment rate on the right hand side of the regression:

$$\log ur_{it}^{\text{youth}} = \alpha_i + \beta_t + \gamma \log \text{share}_{it} + \delta \log ur_{it}^{\text{prime}} + \varepsilon_{it} \quad (6)$$

The dependent variable is the youth (16–24 years old) unemployment rate. The prime age (25–54) unemployment rate is included on the right hand side,¹³ and the youth share is instrumented with lagged birth rates.

Unfortunately, it is not easy to interpret the estimate of the elasticity γ if changes in the youth share of the population cause changes in the prime age unemployment rate, as the evidence in this paper suggests. Indeed, there is a good reason to think that this methodology biases their estimate of γ upwards. Because there is a positive correlation between prime age and youth unemployment rates over the business cycle, one would expect to find a positive coefficient on δ , multiplying the prime age unemployment rate. Then if an increase in the youth share of the population reduces both the prime age and the youth unemployment rate, part or all of the effect on the youth unemployment rate will be captured by δ . The variable of interest γ may even become positive.

As confirmation of this reasoning, Table 7 shows estimates of equation (6) on the standard 51 state, 19 year data set, using both OLS and instrumental variables. If the prime age unemployment rate is excluded from the regression, one obtains estimates similar to those for 16–19 and 20–24 year olds in Table 2. But including this endogenous variable biases the estimated coefficient on the youth share towards zero, as predicted. Instrumenting the youth share with lagged birth rates yields a slightly positive coefficient estimate.

I have focused on Korenman and Neumark (1997) because their methodology is the most directly comparable with mine. However, many other authors have assumed that the youth share of the population does not affect the prime age unemployment rate. For example, Flaim (1979) interprets a positive correlation between the youth share of the population and the gap between the teenage and prime age

¹³I use these definitions of youth and prime age for consistency with Korenman and Neumark (1997). The results are not sensitive to this choice. Also, Korenman and Neumark include the prime age employment-population ratio on the right hand side. The results are not sensitive to including that variable.

unemployment rates in U.S. time series data as evidence that an increase in the youth share of the population raises the teenage unemployment rate. More recently, Shimer (1998) labels unemployment ‘demographic’ if it cannot be predicted from a linear regression on the prime age unemployment rate. Thus he presumes that the prime age unemployment rate is unaffected by demographic changes.

While this hypothesis certainly does not hold across states in the U.S., it does seem to hold in Korenman and Neumark’s (1997) cross-country data set, which has information on unemployment and participation rates, youth population shares, and lagged birth rates in 15 OECD countries¹⁴ for all or part 1970–94. To establish this, I estimate equation (1) on that data set. Panel A of Table 8 shows that although the point estimates from the basic OLS fixed effects regression of the unemployment rate on the youth share are negative, they are not significantly different from zero. Instrumenting the youth share with lagged birth rates reverses the sign of this elasticity. I conclude that there is no systematic relationship between the youth population share and the unemployment rate in that data set; and one can reject the elasticities estimated in U.S. data.

Panel B shows that the OLS-estimated impact of the youth share on the participation rate is ambiguous, switching signs for different age groups. However, the IV estimates are positive and similar to the estimates on U.S. data — and surprisingly, one can strongly reject exogeneity of the youth population share.

Why do cross-state data give such different results than cross-country in the unemployment rate regression? An interesting possibility is that there are effects missing from cross-state variation in the youth share of the population, which might be present in cross-country data sets. For example, in a closed economy neoclassical growth model, an increase in the population growth rate reduces the capital-labor ratio and thus the wage rate. With search frictions, this manifests itself as an increase in unemployment. However, with an open economy, capital would flow across regions so as to equalize the capital-labor ratio. Even if one state has a higher population growth rate than another, it would not have a lower capital-labor

¹⁴The countries are: Australia, Canada, Finland, France, Germany, Ireland, Italy, Japan, Netherlands, Norway, Portugal, Spain, Sweden, the UK, and the US. See their paper for details on the data set.

ratio or a higher unemployment rate. On the other hand, if all states have a higher population growth rate and international capital flows are imperfect, the capital-labor ratio would fall in the country, and the unemployment rate would increase. Using cross-state data, this would get pushed into a year fixed effect, but a cross-country time series would capture this effect of population growth.

8 Conclusion

This paper argues that an increase in the youth share of the population causes a sharp decline in the unemployment rate, particularly for older workers. It also leads to an increase in the participation rate, especially for younger workers. This is inconsistent with standard theories of unemployment, which either predict no relationship or the opposite relationship between these variables. However, it is consistent with a theory of the labor market in which mismatch of young workers is important, and firms prefer to locate in markets with a lot of mismatch, because it is easier to find good employees in thicker labor markets.

One can interpret the empirical results in this paper as a test for thick market externalities. The standard theory of equilibrium unemployment, as summarized in Pissarides (1990), assumes that the number of matches created in a period is a constant returns to scale function of the number of unemployed workers and the number of vacant firms. This yields many strong predictions. For example, the equilibrium of simple search and matching models is unique, and the economy rapidly converges towards a steady state. The standard model also predicts that an exogenous increase in the number of job searchers will have no effect on job creation and job destruction rates, although it will directly increase the aggregate unemployment rate (Shimer, 1998).

In contrast, Diamond (1982) allows the matching function to have increasing returns to scale. Multiple equilibria are then possible, and even with a unique equilibrium, the labor market may substantially amplify external shocks. For this reason, Hall (1989) declared that “economywide thick-market effects are one of the most promising ways to explain the business cycle.” Models with thick-market effects also predict that an increase in the number of job searchers will raise job

creation and reduce job destruction rates. The aggregate unemployment rate may even fall, as happens in the model developed in Section 5.

One can therefore test for thick market externalities by looking for exogenous variation in the number of job searchers. Lagged changes in birth rates are an ideal source of variation: First, there is substantial variation in birth rates over time and across regions. Second, birth rates are easily measureable, and good data is widely available. Third, lagged birth rates are unlikely to be affected by current labor market conditions, so there is hope of establishing a causal relationship. And finally, the nature of the shock is unambiguous, e.g. the entry of the new cohort should be anticipated. Exploiting this source of variation, this paper uncovers evidence that contradicts the standard model with constant returns to scale, but is completely consistent with the existence of thick market externalities.

References

- BLANCHARD, OLIVER, AND PETER DIAMOND (1989): “The Beveridge Curve,” *Brookings Papers on Economic Activity*, (1), 1–60.
- (1992): “The Flow Approach to Labor Markets,” *American Economic Association Papers and Proceedings*, 82(2), 354–359.
- BLANCHARD, OLIVIER, AND LAWRENCE KATZ (1992): “Regional Evolutions,” *Brookings Papers on Economic Activity*, (1), 1–61.
- BURDETT, KENNETH, AND DALE MORTENSEN (1998): “Equilibrium Wage Differentials and Employer Size,” *International Economic Review*, 39, 257–274.
- DAVIDSON, RUSSELL, AND JAMES MACKINNON (1993): *Estimation and Inference in Econometrics*. Oxford University Press, New York.
- DAVIS, STEVE, JOHN HALTIWANGER, AND SCOTT SCHUH (1996): *Job Creation and Destruction*. MIT Press, Cambridge, MA.
- DAVIS, STEVE, PRADASH LOUNGANI, AND RAMAMOHAN MAHIDHARA (1997): “Regional Labor Fluctuations: Oil Shocks, Military Spending and Other Driving Forces,” Northwestern/Chicago Joint Center For Poverty Research Working Paper 4.

- DIAMOND, PETER (1982): “Aggregate Demand Management in Search Equilibrium,” *Journal of Political Economy*, 90, 881–894.
- FLAIM, PAUL (1979): “The Effect of Demographic Change on the Nation’s Unemployment Rate,” *Monthly Labor Review*, 102, 13–23.
- (1990): “Population Changes, the Baby Boom and the Unemployment Rate,” *Monthly Labor Review*, 113, 3–10.
- GORDON, ROBERT (1982): “Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment,” in *Workers, Jobs and Inflation*, ed. by Martin Baily, pp. 89–152. Brookings Institute, Washington, D.C.
- HALL, ROBERT (1989): “The Beveridge Curve: Comment,” *Brookings Papers on Economic Activity*, (1), 61–64.
- KORENMAN, SANDERS, AND DAVID NEUMARK (1997): “Cohort Crowding and Youth Labor Markets: A Cross-National Analysis,” NBER Working Paper 6031.
- PERRY, GEORGE (1970): “Changing Labor Markets and Inflation,” *Brookings Papers on Economic Activity*, (3), 411–441.
- PISSARIDES, CHRISTOPHER (1990): *Equilibrium Unemployment Theory*. Basil Blackwell, Oxford.
- SHIMER, ROBERT (1998): “Why is the U.S. Unemployment Rate So Much Lower?,” in *NBER Macroeconomics Annual*, ed. by Ben Bernanke, and Julio Rotemberg, vol. 13. MIT Press, Cambridge, MA.

Dependent Variable	Youth Share	N
<u>A. Basic Regression</u>		
Unemployment Rate	-1.028 (0.158)	1377
Participation Rate	.052 (0.018)	1377
<u>B. AR Correction</u>		
Unemployment Rate	-.961 (0.261)	1290
Participation Rate	.083 (0.029)	1290

Table 1: OLS estimates of equation (1) using data from 51 states from 1970–96. All regressions include state and year fixed effects. Standard errors in parentheses.

Dependent Variable	Column 1: OLS			Column 2: IV		
	Youth Share	<i>N</i>		Youth Share	<i>N</i>	<i>p</i>
A. Unemployment Rate						
All Workers	-1.246 (0.251)	918		-1.796 (0.305)	892	0.032
	All	-1.268 (0.259)	916	-1.023 (0.308)	890	0.063
16–19	Men	-1.567 (0.400)	816	-1.098 (0.515)	790	0.077
	Women	-1.418 (0.374)	816	-0.601 (0.482)	790	0.002
	All	-1.633 (0.298)	918	-2.062 (0.350)	892	0.104
20–24	Men	-1.799 (0.352)	918	-2.154 (0.413)	892	0.300
	Women	-1.588 (0.306)	918	-1.872 (0.364)	892	0.345
	All	-1.526 (0.290)	918	-1.928 (0.347)	892	0.162
25–34	Men	-1.735 (0.362)	918	-2.152 (0.432)	892	0.265
	Women	-1.490 (0.280)	918	-1.689 (0.335)	892	0.417
	All	-1.832 (0.347)	918	-2.270 (0.419)	892	0.190
35–44	Men	-1.895 (0.398)	918	-2.495 (0.479)	892	0.078
	Women	-2.012 (0.374)	918	-2.082 (0.451)	892	0.967
	All	-2.455 (0.365)	918	-2.753 (0.440)	892	0.327
45–54	Men	-3.037 (0.476)	918	-3.399 (0.578)	890	0.351
	Women	-2.010 (0.379)	918	-2.268 (0.454)	892	0.428
	All	-2.472 (0.455)	918	-2.882 (0.552)	892	0.167
55–64	Men	-3.074 (0.548)	918	-3.621 (0.666)	892	0.168
	Women	-2.407 (0.589)	918	-2.533 (0.716)	890	0.685
	All	-2.440 (0.687)	893	-1.842 (0.844)	870	0.234
65+	Men	-2.984 (0.813)	843	-3.220 (0.965)	843	0.752
	Women	-2.152 (1.130)	751	-0.810 (1.426)	745	0.186

Table 2A

(continued)

Dependent Variable	Column 1: OLS		Column 2: IV			
	Youth Share	N	Youth Share	N	p	
<u>B. Participation Rate</u>						
All Workers	0.093 (0.029)	918	0.102 (0.035)	892	0.946	
16–19	All	0.365 (0.084)	916	0.451 (0.101)	890	0.169
	Men	0.414 (0.112)	816	0.558 (0.143)	790	0.135
	Women	0.427 (0.127)	816	0.643 (0.163)	790	0.032
20–24	All	0.185 (0.036)	918	0.233 (0.043)	892	0.067
	Men	0.172 (0.037)	918	0.193 (0.045)	892	0.383
	Women	0.211 (0.042)	918	0.282 (0.067)	892	0.039
25–34	All	0.060 (0.024)	918	0.071 (0.028)	892	0.763
	Men	0.026 (0.015)	918	0.027 (0.018)	892	0.947
	Women	0.105 (0.045)	918	0.151 (0.055)	892	0.220
35–44	All	0.055 (0.024)	918	0.065 (0.029)	892	0.811
	Men	0.040 (0.015)	918	0.029 (0.019)	892	0.243
	Women	0.078 (0.046)	918	0.125 (0.056)	892	0.221
45–54	All	0.048 (0.028)	918	0.055 (0.033)	892	0.721
	Men	0.056 (0.025)	918	0.046 (0.031)	892	0.816
	Women	0.062 (0.056)	918	0.097 (0.067)	892	0.495
55–64	All	0.164 (0.058)	918	0.163 (0.070)	892	0.885
	Men	0.169 (0.062)	918	0.178 (0.075)	892	0.861
	Women	0.150 (0.087)	918	0.131 (0.105)	892	0.662
65+	All	-0.209 (0.163)	918	-0.411 (0.195)	892	0.023
	Men	-0.143 (0.174)	918	-0.251 (0.205)	892	0.091
	Women	-0.303 (0.231)	918	-0.596 (0.278)	892	0.084

Table 2B

Table 2: OLS and IV estimates of equation (1) using 19 years and 51 states, for seven different age groups and both sexes. All regressions include state and year fixed effects and an AR(1) correction. Sample sizes N vary due to missing observations in 1994–96. Standard errors in parentheses. p is the p-value in a Davidson-MacKinnon (1993) exogeneity test.

Dependent Variable	Column I. Census Division/Year Fixed Effects	Column II. Aggregate Census Division	Column III. Large States
<u>A. Unemployment Rate</u>			
16–19	-1.065 (0.309)	-1.863 (0.680)	-1.704 (0.826)
20–24	-1.239 (0.322)	-2.004 (0.789)	-1.367 (0.975)
25–34	-1.130 (0.286)	-2.519 (0.827)	-1.436 (1.082)
35–44	-1.887 (0.382)	-2.358 (0.959)	-2.780 (1.162)
45–54	-2.163 (0.427)	-2.653 (1.023)	-3.063 (1.203)
55–64	-1.624 (0.532)	-3.386 (1.190)	-2.939 (1.234)
65+	-1.979 (0.937)	-2.967 (1.113)	-4.762 (1.198)
<u>B. Participation Rate</u>			
16–19	0.398 (0.114)	0.423 (0.204)	0.443 (0.257)
20–24	0.208 (0.048)	0.242 (0.084)	0.228 (0.105)
25–34	0.121 (0.030)	-0.048 (0.048)	-0.098 (0.057)
35–44	0.091 (0.032)	0.013 (0.054)	0.012 (0.061)
45–54	0.046 (0.037)	0.101 (0.056)	0.120 (0.078)
55–64	0.117 (0.081)	0.347 (0.122)	0.468 (0.182)
65+	-0.190 (0.223)	0.311 (0.322)	0.170 (0.406)

Table 3: OLS estimates of the elasticities of the unemployment and participation rates with respect to the youth share. Column I includes state and census division/year fixed effects. Column II is aggregated to the census division level and includes census division and year fixed effects. Column III is run only on the ten large states for which the unemployment rate is calculated directly from the CPS, and includes state and year fixed effects. All regressions have an AR(1) correction. Standard errors in parentheses.

Dependent Variable	Column I. OLS	Column II. IV	
	Youth Share	Youth Share	<i>p</i>
<u>A. Employment Levels</u>			
Total Employment	0.360 (0.052)	0.445 (0.071)	0.502
Construction	1.389 (0.189)	1.343 (0.265)	0.237
Manufacturing	0.531 (0.070)	0.947 (0.104)	0.001
Wholesale and Retail Trade	0.388 (0.050)	0.468 (0.072)	0.163
Services	0.319 (0.051)	0.457 (0.071)	0.707
Transportation & Public Utilities	0.252 (0.062)	0.183 (0.087)	0.015
Government	0.007 (0.039)	0.202 (0.061)	0.117
Finance, Insurance & Real Estate	0.240 (0.073)	0.059 (0.102)	0.001
Mining	0.126 (0.262)	-0.157 (0.390)	0.456
<u>B. Employment Rates</u>			
Total Employment	0.147 (0.039)	0.284 (0.053)	0.123
Construction	1.139 (0.175)	1.192 (0.223)	0.501
Manufacturing	0.344 (0.071)	0.810 (0.105)	0.000
Wholesale and Retail Trade	0.199 (0.040)	0.326 (0.056)	0.413
Services	0.091 (0.045)	0.298 (0.067)	0.021
Transportation & Public Utilities	0.058 (0.055)	0.049 (0.076)	0.416
Government	-0.160 (0.035)	0.078 (0.054)	0.000
Finance, Insurance & Real Estate	0.046 (0.065)	-0.076 (0.092)	0.023
Mining	-0.207 (0.249)	-0.402 (0.363)	0.642

Table 4: OLS and IV estimates of the elasticities of employment with respect to the youth share for the entire economy and in subsets of states. All regressions include state and year fixed effects and an AR(1) correction. Standard errors in parentheses. The source of employment data is official BLS time series for employment in 1-digit sectors. *p* is the *p*-value in a Davidson-MacKinnon (1993) exogeneity test.

Dependent Variable	share ⁵⁻¹⁵	share ¹⁶⁻²⁴	share ²⁵⁻³⁴
<u>A. Unemployment Rate</u>			
	1.443 (0.214)	—	—
	—	—	0.188 (0.230)
16-20	1.184 (0.236)	-0.669 (0.270)	—
	—	-1.813 (0.320)	-0.751 (0.268)
	1.072 (0.255)	-0.959 (0.363)	-0.322 (0.271)
	1.629 (0.273)	—	—
	—	—	0.704 (0.273)
20-24	1.190 (0.296)	-1.060 (0.323)	—
	—	-1.874 (0.386)	-0.326 (0.332)
	1.256 (0.319)	-0.889 (0.449)	0.189 (0.344)
	1.378 (0.271)	—	—
	—	—	0.699 (0.274)
25-34	0.931 (0.295)	-1.080 (0.318)	—
	—	-1.781 (0.385)	-0.340 (0.340)
	0.949 (0.316)	-1.032 (0.451)	0.053 (0.352)
	2.127 (0.185)	—	—
	—	—	0.804 (0.327)
35-44	1.317 (0.347)	-1.210 (0.375)	—
	—	-2.160 (0.459)	-0.439 (0.403)
	1.362 (0.372)	-1.095 (0.528)	0.129 (0.410)
	2.179 (0.326)	—	—
	—	—	1.096 (0.325)
45-54	1.516 (0.346)	-1.729 (0.389)	—
	—	-2.673 (0.466)	-0.290 (0.386)
	1.656 (0.372)	-1.368 (0.532)	0.393 (0.395)
	1.602 (0.420)	—	—
	—	—	1.461 (0.387)
55-64	0.789 (0.447)	-2.090 (0.503)	—
	—	-2.196 (0.565)	0.386 (0.467)
	1.092 (0.479)	-1.350 (0.671)	0.830 (0.500)
	2.660 (0.561)	—	—
	—	—	1.038 (0.547)
65+	2.198 (0.610)	-1.371 (0.735)	—
	—	-2.472 (0.828)	-0.046 (0.647)
	2.589 (0.663)	-0.462 (0.958)	1.010 (0.685)

Table 5A

(continued)

Dependent Variable	share ⁵⁻¹⁵	share ¹⁶⁻²⁴	share ²⁵⁻³⁴
<u>B. Participation Rate</u>			
	-0.348 (0.076)	—	—
	—	—	-0.322 (0.072)
16-20	-0.245 (0.084)	0.238 (0.094)	—
	—	0.213 (0.106)	-0.208 (0.091)
	-0.357 (0.087)	-0.076 (0.126)	-0.350 (0.095)
	-0.127 (0.034)	—	—
	—	—	-0.149 (0.031)
20-24	-0.063 (0.037)	0.154 (0.041)	—
	—	0.122 (0.046)	-0.087 (0.038)
	-0.109 (0.039)	0.037 (0.054)	-0.131 (0.041)
	-0.027 (0.022)	—	—
	—	—	-0.041 (0.021)
25-34	-0.002 (0.025)	0.059 (0.027)	—
	—	0.052 (0.031)	-0.011 (0.027)
	-0.007 (0.026)	0.047 (0.038)	-0.014 (0.029)
	-0.057 (0.022)	—	—
	—	—	-0.029 (0.021)
35-44	-0.043 (0.025)	0.033 (0.027)	—
	—	0.057 (0.031)	0.003 (0.027)
	-0.049 (0.027)	0.018 (0.038)	-0.017 (0.029)
	-0.057 (0.025)	—	—
	—	—	-0.045 (0.023)
45-54	-0.046 (0.028)	0.026 (0.031)	—
	—	0.025 (0.035)	-0.031 (0.030)
	-0.067 (0.030)	-0.027 (0.042)	-0.058 (0.032)
	-0.135 (0.053)	—	—
	—	—	-0.035 (0.050)
55-64	-0.085 (0.059)	0.122 (0.065)	—
	—	0.228 (0.074)	0.088 (0.063)
	-0.063 (0.063)	0.178 (0.089)	0.062 (0.068)
	-0.157 (0.150)	—	—
	—	—	0.460 (0.139)
65+	-0.309 (0.165)	-0.363 (0.180)	—
	—	0.212 (0.206)	0.575 (0.177)
	-0.130 (0.176)	0.108 (0.248)	0.521 (0.190)

Table 5B

Table 5: OLS regressions including additional measures of the youth share. All regressions include state and year fixed effects and an AR(1) correction. Standard errors in parentheses.

Dependent Variable	Column I. OLS	Column II. IV	
	Youth Share	Youth Share	<i>p</i>
Job Creation	0.775 (0.371)	1.506 (0.993)	0.492
Job Destruction	0.676 (0.416)	0.917 (1.009)	0.877

Table 6: OLS and IV estimates of the elasticity of job creation and destruction with respect to the youth share. Both regressions include state and year fixed effects and an AR(1) correction. Standard errors in parentheses. *p* is the p-value in a Davidson-MacKinnon (1993) exogeneity test.

Column I. OLS		Column II. IV		
Youth Share	Prime Unemploy.	Youth Share	Prime Unemploy.	p
-1.310 (0.269)	—	-1.354 (0.316)	—	0.868
-0.251 (0.171)	0.561 (0.027)	0.096 (0.211)	0.573 (0.028)	0.000

Table 7: Dependent Variable: Youth unemployment rate. OLS and IV estimates of the elasticity of the unemployment rate with respect to the youth share and the prime age unemployment rate. All regressions include state and year fixed effects and an AR(1) correction. Standard errors in parentheses. p is the p-value in a Davidson-MacKinnon (1993) exogeneity test.

Dependent Variable	Column I. OLS			Column II. IV		
	Youth Share		N	Youth Share	N	p
<u>A. Unemployment Rate</u>						
15–24	-0.173	(0.479)	319	0.443	(0.767)	317 0.351
25–54	-0.300	(0.546)	319	0.337	(0.930)	317 0.459
55–64	-0.173	(0.591)	314	0.013	(0.791)	312 0.392
<u>B. Participation Rate</u>						
15–24	0.130	(0.079)	319	0.463	(0.181)	317 0.025
25–54	-0.137	(0.028)	319	0.079	(0.071)	317 0.000
55–64	-0.108	(0.076)	314	0.146	(0.155)	317 0.066

Table 8: OLS and IV estimates of equation (1) using 25 years and 15 OECD countries for three different age groups. All regressions include state and year fixed effects, an AR(1) correction, and dummy variables to account for changes in some of the data series. Standard errors in parentheses. p is the p-value in a Davidson-MacKinnon (1993) exogeneity test.

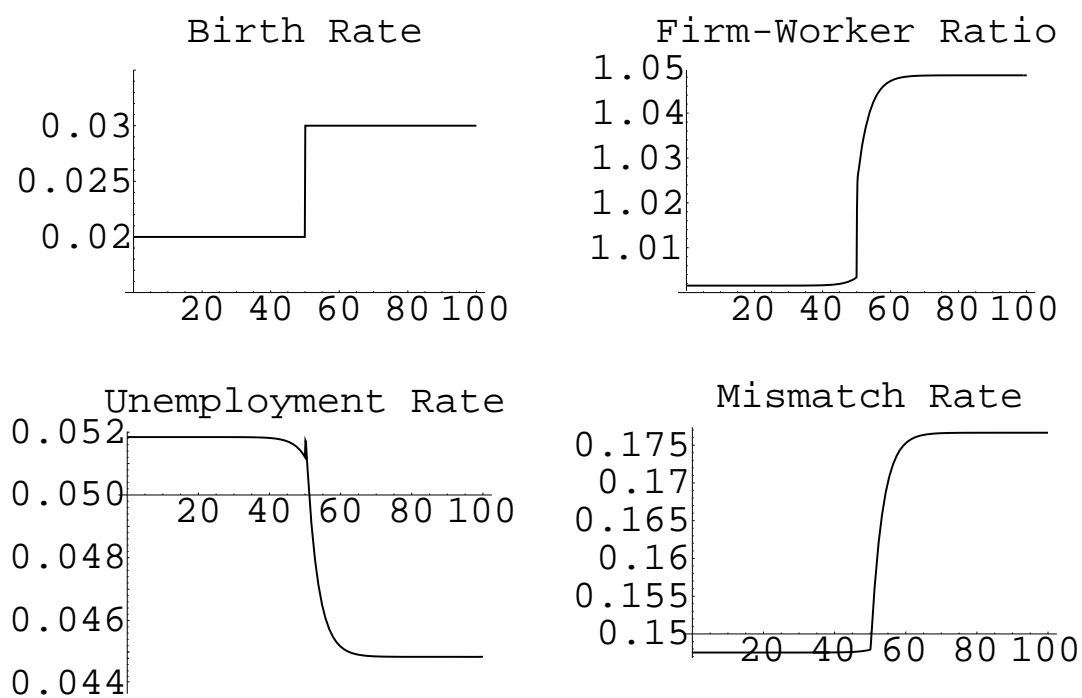


Figure 1: The dynamic response of the firm-worker ratio, unemployment rate, and mismatched rate to an anticipated permanent increase in the population growth rate from 0.02 to 0.03 in period 50.