THE IMPACT OF YOUNG WORKERS ON THE AGGREGATE LABOR MARKET*

ROBERT SHIMER

An increase in the share of youth in the working age population of one state or region relative to the rest of the United States causes a sharp reduction in that state's relative unemployment rate and a modest increase in its labor force participation rate. This is inconsistent with many theories of the labor market, but can be easily explained by a model of frictional unemployment with on-the-job search. The theory makes strong predictions regarding the behavior of wages which are shown to be consistent with the data. The paper also reconciles its findings with an existing body of apparently contradictory empirical evidence.

I. INTRODUCTION

The baby boom has profoundly altered the demographic structure of the U. S. population during the past 50 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 verified by Flaim [1990] and Shimer [1998].

The effects of the baby boom on unemployment can be grouped into two categories. First, since the aggregate unemployment rate is a weighted average of the unemployment rates of workers of different ages, 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 84 basis point (0.84 percentage point) increase in the aggregate unemployment rate from 1954 to 1978, and an 81 basis point decline from 1978 to 1998.

* I have received helpful comments from Daron Acemoglu, Olivier Blanchard, Andrew Caplin, Henry Farber, Christopher Foote, Edward Glaeser, Bo Honore, Lawrence Katz, Alan Krueger, Christina Paxson, Harvey Rosen, Robert Topel, two anonymous referees, and many seminar participants. Thanks to Susan Gorel 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 and SES-0079345 for financial support.

The Quarterly Journal of Economics, August 2001

Second, changes in demographics may have "indirect" effects on age-specific unemployment rates. 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 imperfect substitutes, an increase in the youth labor supply may have a differential impact on young and prime age workers. Shimer [1998] estimates that the indirect effects of the baby boom were about as large as the direct effect, so that in total the baby boom caused a 180 basis point increase in the aggregate unemployment rate before 1978, and a subsequent 145 basis point decline.

Unfortunately, while the direct effect of this supply shock can be precisely estimated, most estimates 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 rate to the baby boom. Although the timing of this increase coincides with the entry 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 were responsible for this correlation, a possibility that cannot be precluded using time series evidence from a single baby boom. Second, if young and prime age workers are complements in production, the baby boom may have simply *reduced* the prime age unemployment rate. If that is the case, the indirect effects of the baby boom at least partially offset the direct effects.

Recent work by Korenman and Neumark [2000] addresses the first issue by using time series data on unemployment rates for fifteen countries. They look at the relative unemployment rate of young and prime age 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 caused by an increase in the birthrate two decades before raises the youth unemployment rate relative to the prime age unemployment rate, with an elasticity of approximately 0.5. The cross-country data set 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 rate.

This paper's primary innovation is to focus on data from within the United States. I use annual state unemployment rates 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 both 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 -1.5; and that the effect on the prime age unemployment rate is even larger in magnitude. A 1 percent increase in the youth share of the population reduces the prime age unemployment rate by more than 2 percent. At the same time, labor market participation rates increase. One possible explanation is that young workers migrate to states with low unemployment rates; however, controlling for the endogeneity of the age structure of the population by instrumenting with lagged birthrates does not alter the results. Not only are the signs of these estimates contrary to the existing literature, 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.

The second task of this paper is to develop a theory that can explain these findings. 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. One possible explanation is that a younger population is also a positive aggregate demand shock, as would be the case if people in younger states want to consume more. As evidence against this possible explanation, I show that there is more growth in manufacturing employment than in any other sector. Since the demand for manufactured goods is not determined at the state level, explanations based on aggregate demand are unlikely to be successful.

This leads me to look to the supply side for an explanation. I begin with a plausible extension to a standard [Mortensen and Pissarides 1994; Pissarides 2000] search model, on-the-job search. There is no inherent difference between young and prime age workers; however, because it takes time to find a good match, most young workers are mismatched in their current employment and prepared to take another job. The second key assumption is that hiring is easier when a larger proportion of the labor force is willing to accept the job, a possibility that Diamond [1982] labeled a *trading externality*. This implies that firms will find creating jobs in younger states to be more profitable, boosting job creation and reducing the unemployment rate of both young and prime age workers.

The model has a number of testable implications. First, average wages should be lower in states with more young workers. I confirm that a 1 percent increase in the youth share of the working age population reduces the average hourly wage in a state by one- or two-tenths of a percent. Although this is consistent with the model, it can also be explained by the simple observation that younger workers earn lower wages. A more interesting prediction is that age-conditional wages should be higher, particularly some time after the increase in the youth population share. This distinguishes the model from another supply-side hypothesis, that employment rises in younger states because of an increase in labor supply, for example if parents work harder to finance their children's education. I find strong evidence in support of my model: a 1 percent increase in the youth population share raises the wage of many groups of workers by 1 or 2 percent, but with a ten-year lag. Finally, the model predicts that labor markets with more young workers should have more turnover. Since reliable worker flow data do not exist at the state level, I instead use job flow data for the manufacturing sector [Davis, Haltiwanger, and Schuh 1996]. Consistent with the hypothesis, I find that a 1 percent increase in the youth share of the population raises the job creation rate by 1.5 percent and the job destruction rate by 0.8 percent.

The final task is to try 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. The results presented here sharply contrast with Korenman and Neumark's [2000] conclusion from their review of the literature that "there seems to be evidence of an adverse effect of cohort size on youth unemployment, employment, and wages across a number of countries." One difference between this paper and the existing literature, is that most previous studies presume that the youth share of the population does not affect the prime age unemployment rate, an assumption that my analysis rejects.

I show that this biases these studies toward finding evidence that large cohorts suffer high unemployment rates. But when I apply my methodology to Korenman and Neumark's fifteen-country data set, I find no significant correlation between unemployment rates and the youth share of the population, which supports neither the cohort crowding hypothesis nor mine. I argue that the inconsistency of state- and national-level findings remains a bit of a puzzle, but may be explained by the relative importance of competing effects at different levels of aggregation.

Section II describes the data used in the main empirical analysis, whose results are presented in Section III. I develop a model in Section IV 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 V offers tests of this theory using wage data and job creation and destruction data from manufacturing. Section VI reconciles these results with existing evidence on the relationship between the youth share of the population and labor market outcomes. Section VII concludes by exploring the broader implications of the findings.

II. Data

The main empirical analysis draws on cross-state differences in birthrates within the United States, and the consequent impact on the youth share of the population and on unemployment and participation rates. The basic source of unemployment and participation 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.[1] This yields an official series for the state rates, which is publicly available from the BLS web site ⟨http://stats.bls.gov/⟩. In addition, the BLS constructs (but does not make generally available) data on state unemployment and participation rates for

1. ⟨http://stats.bls.gov/laumthd.htm⟩ describes the BLS's local area unemployment methodology. This article points out that the BLS smoothes the monthly estimates using a "signal-plus-noise" approach, although this should not have much of an effect on annual data. I have also constructed unemployment and participation rates directly from the March CPS to verify that my results are not driven by the BLS methodology.

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.[2] Although the BLS produces analogous numbers, the 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 in my main analysis. I have also verified that my results are robust to using the BLS estimates.

The third piece of data is birthrates for each state from 1954 to 1980 (16 to 24 years before my sample period). These come from various years of the Statistical Abstract of the United States, and are measured in births per thousand residents. Whenever possible, I use birthrates corrected for undercounting, rather than the official birth census. Birthrate data are unavailable for Alaska and Hawaii before 1960. Since this affects a third of the observations for these two states, those before 1984, I drop them from the empirical analysis, although I have verified that this is unimportant for my results. This leaves me with a panel of 49 states, including the District of Columbia.

III. Empirical Analysis

A. Empirical Model

The main empirical model looks at how the unemployment rate in state i and year t depends on the youth share of the working age population, share_{it}, defined as the number of 16 to 24 year olds divided by the number of 16 to 64 year olds:

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

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

2. These numbers are not simply an interpolation between decennial censuses. Instead, they include estimates of migration patterns for different age cohorts using projections from available data on the migration of school-age children. The Census methodology is described in Byerly and Deardorff [1995] and similar volumes from previous years, and summarized at ⟨http://www.census.gov/population/methods/stage98.txt⟩.

hypothesis is that the elasticity γ is zero. I later look for endogenous variation in the youth share by instrumenting it with lagged birthrates. I also estimate similar regressions with the employment-population ratio (hereinafter employment ratio) or the labor market participation rate on the right-hand side.[3]

With a nineteen-year sample, the year dummies are needed to dispose of any macroeconomic shocks, such as monetary policy, consumer confidence, and oil prices, that may create a spurious correlation between the youth population share and the unemployment rate. At a national level, the unemployment rate averaged 7.5 percent during the first half of the sample period, and fell to 6.1 percent during the second half. The age structure of the population shows similar temporal variation, with the share of youths in the working age population declining from 26 percent at the beginning of the sample to 19 percent at the end, ensuring a negative correlation between the two variables. The ability to include year dummies is the main advantage of using state data.

Similarly, unemployment and demographics show considerable cross-sectional variation. The youth share of the working age population averaged 27 percent in Mississippi and 22 percent in Connecticut during the sample period, while the state unemployment rates averaged 8 percent and 6 percent, respectively. It is not obvious whether there is a causal relationship, and if so, in which direction the causation goes. To sidestep this issue, I also include state dummies in my regressions.

Including the state and year dummies comes at a cost: they soak up 78 percent of the variation in the youth population share. Thus, although the difference between the ninetieth percentile log youth population share and the tenth percentile is about 0.37, the dummies reduce the interdecile range to 0.07. Similarly, the interdecile range for the birthrate after accounting for state and year fixed effects is 0.09. Since the remaining variation disproportionately reflects measurement error, there will be some attenuation bias in my estimates.

Before moving on to the formal empirical analysis, Figure I summarizes the data that drive the results. There are nine graphs, each depicting the time series behavior of the unemploy-

3. Let U denote the number of unemployed workers, E denote the number of employed workers, and P denote the population. Then the unemployment rate is $ur = U/(U + E)$; the participation rate is $pr = (U + E)/P$; and the employment ratio is $E/P = pr(1 - ur)$.

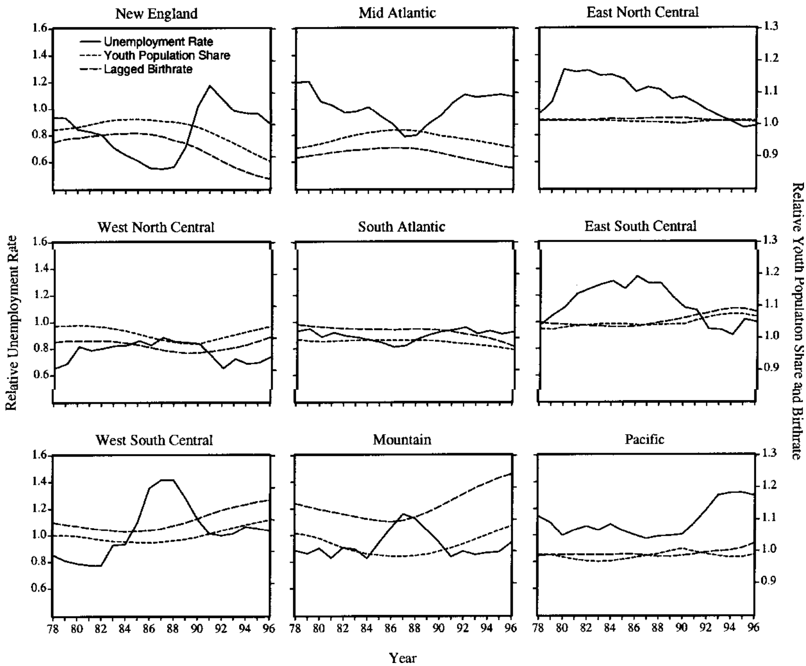

FIGURE I
Unemployment Rate, Youth Population Share, and Lagged Birthrates, Relative to the Contemporaneous National Average

ment rate, youth population share, and lagged birthrates in a Census division from 1978 to 1996.[4] The annual figures are expressed as ratios of the contemporaneous national rates, which eliminates cyclical fluctuations from the graphs, analogous to the year dummies. All three time series are persistent, and the youth population share and birthrates change only gradually. Moreover, no time series is monotonic in any division.

More to the point, lagged birthrates are highly correlated with the contemporaneous youth population share, giving it

4. The U. S. Census Bureau divides the country into nine divisions: New England (ME, NH, VT, MA, RI, CT), Mid Atlantic (NY, NJ, PA), East North Central (OH, IN, IL, MI, WI), West North Central (MN, IA, MO, ND, SD, NE, KS) South Atlantic (DE, MD, DC, VA, NC, SC, GA, FL), East South Central (KY, TN, AL, MS), West South Central (AR, LA, OK, TE), Mountain (MT, ID, WY, CO, NM, AZ, UT, NV), and Pacific (WA, OR, CA, AK, HI). These divisions are relatively homogeneous, and therefore provide a useful level of aggregation. The divisions are further aggregated into four regions; however, these are much less homogeneous.

TABLE I
THE EFFECT OF YOUNG WORKERS ON UNEMPLOYMENT, PARTICIPATION, AND EMPLOYMENT

Dependent variable	Column I Youth Share	Nobs	Column II 1970–1978	1979–1987	1988–1996	Nobs
A. Basic regression						
Unemployment rate	−1.221	1293	−2.012	−0.942	−1.057	1293
	(0.160)		(0.228)	(0.256)	(0.163)	
Participation rate	0.032	1293	−0.062	−0.041	0.058	1293
	(0.018)		(0.026)	(0.029)	(0.018)	
Employment ratio	0.105	1293	0.062	0.014	0.118	1293
	(0.023)		(0.033)	(0.037)	(0.024)	
B. AR correction						
Unemployment rate	−1.219	1244	−2.709	−1.065	−1.129	1244
	(0.264)		(0.511)	(0.497)	(0.264)	
Participation rate	0.086	1244	−0.004	−0.016	0.106	1244
	(0.029)		(0.056)	(0.055)	(0.029)	
Employment ratio	0.162	1244	0.182	0.055	0.173	1244
	(0.037)		(0.076)	(0.073)	(0.038)	

OLS estimates of equation (1) using data from 49 states from 1970 to 1996. All regressions include state and year fixed effects. Panel B corrects for AR(1) residuals. Column II allows for breaks in the elasticities in 1978 and 1987. Standard errors are in parentheses. Unemployment, employment, and participation data are constructed by the BLS from the CPS and information in the UI system. The youth population share is the number of 16 to 24 year olds divided by the number of 16 to 64 year olds in a state, according to U. S. Census Bureau estimates.

power as an instrumental variable. On the other hand, these time series tend to move in the opposite direction of the unemployment rate. Two divisions, New England and the Mid Atlantic, show a hump-shaped youth share and birthrate, with a peak in the mid-1980s; and both divisions realized their lowest unemployment rate at that time. The West North Central, West South Central, and Mountain divisions experienced troughs in the youth population share and lagged birthrates in the mid-eighties, and a simultaneous peak in their unemployment rates. In the East South Central and Pacific divisions, there was little change in the youth population share until the 1990s, at which time the unemployment rate moved in the opposite direction. Finally, in the South Atlantic and East North Central divisions, there is little variation in either birthrates or the youth population share.

B. OLS Results

Column I, panel A of Table I shows the results from estimating equation (1) using aggregate unemployment rate data from

1970 to 1996. The estimated elasticity of the unemployment rate with respect to the youth share of the population is approximately -1.2, significantly different from zero at any standard confidence level. The elasticities of the participation rate and employment ratio have the opposite sign and are much smaller, although they are still statistically significantly different from zero. A 10 percent increase in the youth share of the working age population is correlated with a 0.3 percent increase in the participation rate and a 1 percent increase in the employment rate in that state. In each case, the state and year fixed effects are highly significant, justifying the panel data analysis.

The estimates in panel A suffer from autocorrelated residuals, because an increase in the unemployment rate only gradually disappears as the laid-off workers find new jobs. A regression of the OLS residual $\hat{\epsilon}_{it}$ on itself lagged one period, $\hat{\epsilon}_{it-1}$, yields a coefficient of 0.73 in the unemployment rate equation, 0.71 in the participation rate equation, and 0.75 in the employment ratio equation. There is no evidence of higher ordered autocorrelation, so panel B quasi-differences equation (1) to correct for first-order autocorrelation in the residuals. The estimated elasticity of the unemployment rate is unchanged, while the elasticities of participation and employment increase. Although the standard errors are somewhat larger, this does not alter the significance of the results. The results throughout the remainder of the paper correct for first-order autocorrelation in the residuals.

Column II of Table I allows the elasticity γ to vary over three time intervals, 1970–1978, 1979–1987, and 1988–1996, in order to show that the estimates are reasonably stable. For example, the unemployment rate elasticity is stable from 1978–1996, my primary sample period, and is more negative during the first nine years of the sample. The elasticity of the participation rate, on the other hand, is insignificant in the first two-thirds of the sample. The estimate in column I is driven entirely by what happened after 1988. And the elasticity of the employment ratio is positive throughout the sample, but only significant in the first and last thirds.

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

To address these concerns, I reestimate equation (1) separately for seven different age groups and both sexes. In each case, the right-hand-side variable of interest is the youth share of the working age population. The left-hand-side variable is the unemployment rate for the appropriate age and sex. For the data availability reasons described in Section II, I use the sample period 1978–1996. Column (1) in Table IIa shows the results, including a regression of the aggregate unemployment rate on the youth share using this shorter sample period. They are qualitatively similar for 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. Quantitatively, an increase in the youth share of the population has twice as large an effect on the unemployment rate of older men as on younger workers and women. The elasticity of the teenage unemployment rate is −1.5, and the elasticity of the unemployment rate for men over 45 is in excess of −3. As a result, the relative unemployment rate of young workers, i.e., the ratio of youth to prime-age unemployment rates,[5] rises in response to an increase in the youth share of the population. This is consistent with the cohort crowding hypothesis and the empirical evidence in support of it, although of course the absolute decline in the youth unemployment rate is not.

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, and a somewhat larger response of women than of men (Table IIb). This is consistent with a greater labor supply elasticity for these groups. Also, the decline in unemployment and increase in participation suggest that the employment ratio should rise. I have confirmed this, but do not report the results.

5. The youth unemployment rate is typically several times higher than the prime age unemployment rate. As a result, the *difference* between the youth and prime age unemployment rate falls in response to an increase in the youth share of the population. Throughout this paper I define relative unemployment rates using ratios, in keeping with the most common usage in the related literature.

TABLE IIa
THE EFFECT OF YOUNG WORKERS ON UNEMPLOYMENT, BY AGE AND SEX

Age/sex cohort		(1) OLS	(2) IV	Nobs	p
All workers		−1.466	−1.807	882	0.045
		(0.255)	(0.307)		
16–19	Men	−1.661	−1.012	784	0.067
		(0.407)	(0.512)		
	Women	−1.489	−0.597	784	0.001
		(0.379)	(0.486)		
20–24	Men	−2.018	−2.180	882	0.452
		(0.347)	(0.419)		
	Women	−1.752	−1.878	882	0.510
		(0.304)	(0.359)		
25–34	Men	−1.963	−2.173	882	0.351
		(0.360)	(0.424)		
	Women	−1.593	−1.698	882	0.553
		(0.285)	(0.336)		
35–44	Men	−2.109	−2.511	882	0.110
		(0.404)	(0.476)		
	Women	−2.126	−2.101	882	0.913
		(0.382)	(0.450)		
45–54	Men	−3.215	−3.425	880	0.498
		(0.480)	(0.571)		
	Women	−2.102	−2.275	882	0.474
		(0.382)	(0.451)		
55–64	Men	−3.346	−3.994	882	0.095
		(0.611)	(0.725)		
	Women	−2.367	−2.429	882	0.879
		(0.652)	(0.771)		
65+	Men	−3.379	−2.833	882	0.527
		(1.344)	(1.599)		
	Women	−3.081	−2.718	872	0.784
		(1.925)	(2.334)		

OLS and IV estimates of equation (1) using data from 49 states from 1978 to 1996, for seven different age groups and both sexes. All regressions include state and year fixed effects and an AR(1) correction. Standard errors are in parentheses. p is the p-value from a Davidson and MacKinnon [1993] exogeneity test. Unemployment data are constructed by the BLS from the CPS and information in the UI system. The youth population share is the number of 16 to 24 year olds divided by the number of 16 to 64 year olds in a state, according to U. S. Census Bureau estimates. The instrumental variable is birthrates per capita lagged 16–24 years, from various years of the Statistical Abstract of the United States. The following observations are missing: 16–19 year olds in 1995 and 1996; 45–54 year old men in Utah in 1994; and women 65 and over in Delaware, Idaho, Mississippi, Tennessee, and Utah in 1994.

C. IV Results

One possible explanation for these results is that young workers move to states with low unemployment rates. Note that the story is not as simple as saying that state i always has low

TABLE IIb
THE EFFECT OF YOUNG WORKERS ON PARTICIPATION, BY AGE AND SEX

Age/sex cohort		(1) OLS	(2) IV	Nobs	p
All workers		0.104	0.102	882	0.938
		(0.029)	(0.035)		
16–19	Men	0.433	0.565	784	0.123
		(0.116)	(0.145)		
	Women	0.453	0.647	784	0.041
		(0.129)	(0.163)		
20–24	Men	0.185	0.197	882	0.603
		(0.037)	(0.044)		
	Women	0.226	0.285	882	0.087
		(0.056)	(0.067)		
25–34	Men	0.026	0.026	882	0.979
		(0.015)	(0.018)		
	Women	0.117	0.150	882	0.247
		(0.046)	(0.055)		
35–44	Men	0.040	0.028	882	0.248
		(0.016)	(0.019)		
	Women	0.095	0.126	882	0.297
		(0.048)	(0.056)		
45–54	Men	0.048	0.044	882	0.838
		(0.026)	(0.030)		
	Women	0.068	0.097	882	0.411
		(0.057)	(0.067)		
55–64	Men	0.175	0.179	882	0.900
		(0.063)	(0.075)		
	Women	0.145	0.129	882	0.770
		(0.090)	(0.106)		
65+	Men	−0.059	−0.250	882	0.073
		(0.172)	(0.203)		
	Women	−0.319	−0.596	882	0.047
		(0.232)	(0.274)		

OLS and IV estimates of equation (1) using data from 49 states from 1978 to 1996, for seven different age groups and both sexes. All regressions include state and year fixed effects and an AR(1) correction. Standard errors are in parentheses. p is the p-value from a Davidson and MacKinnon [1993] exogeneity test. Unemployment data are constructed by the BLS from the CPS and information in the UI system. The youth population share is the number of 16 to 24 year olds divided by the number of 16 to 64 year olds in a state, according to U. S. Census Bureau estimates. The instrumental variable is birthrates per capita lagged 16 to 24 years, from various years of the Statistical Abstract of the United States. The observations for 16–19 year olds are missing for 1995 and 1996.

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.

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 birthrate in that state 16 to 24 years before. More precisely, birth_{it} in year t = 1978 is the sum of the number of births per person in state i from 1954 to 1962. Regressing log share_{it} on log birth_{it} and state and year fixed effects from 1978 to 1996 yields an elasticity estimate of 0.689 with a standard error of 0.016. Lagged birthrates explain 68 percent of the residual variation in the youth share after accounting for the fixed effects. Together, lagged birthrates and the fixed effects explain 98 percent of the variation in the youth share. The instrument is an excellent predictor of future youth shares. This result is puzzling in its own right, since Blanchard and Katz [1992] have shown that U. S. workers are quite mobile in response to economic incentives. One possible explanation is that the Census Bureau uses birthrates to estimate the age structure of the population between decennial censi. According to Byerly and Deardorff [1995], this is not the case; and even if it were, this would not affect the validity of the birthrate instrument, just the interpretation of OLS estimates.

Column 2 of Table IIa and IIb then reports the results from estimating the basic regression using IV. If young workers move to low unemployment states, the estimated elasticities should be smaller in magnitude or may even switch signs. On the other hand, if the youth share is mismeasured, the magnitude of the estimates may grow. In fact, none of the results are significantly changed, and the magnitudes are generally larger, consistent with measurement error. A 1 percent increase in the youth share of the population caused by an increase in the birthrate 16 to 24 years before reduces the unemployment rate of teenagers by up to 1 percent and of prime age workers by 2 or 3 percent, and by more for men than for women. It raises the participation rate of teenage workers by about half a percent, and slightly raises the participation rate of prime age workers. Because the standard errors are somewhat larger using IV, a few of the elasticity estimates, e.g., those for some teens and older workers in Table IIa, are not significantly different from zero.

Since lagged birthrates and state and year fixed effects predict 98 percent 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 inefficient. If the youth share of the population is exogenous and correctly measured, 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 birthrates 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 and correctly measured.

The final column in Table II shows the p-values from this test, which are just the p-values for t-tests of whether the coefficient on the predicted youth share is different from zero. One can only reject "exogeneity" at the 10 percent confidence level in 8 of 30 cases, and at the 5 percent level in 3 cases. This suggests that OLS is likely to be a consistent and efficient estimator. Nevertheless, I report the IV estimates in the remainder of the paper, since they are certainly consistent and they are precise enough for the purposes of this paper.

D. Robustness

This section describes a number of robustness checks on the results in Table II.[6] To reduce clutter, I simply report the results for three age categories, 16–24, 25–54, and 55+ , and do not break them down by sex. Column 1 of Table III replicates the baseline IV analysis for these three groups.

The first concern is that birthrates might not be a valid instrument because family planning anticipates future economic conditions. In its most plausible form, this is an argument about economic convergence. Suppose that there are two regions, North and South. Initially, the North is more developed than the South. Employment in the North is predominantly in sectors of the

6. I have verified the robustness of the results on a number of other dimensions. I ran the regression in levels rather than logs, weighted the regressions by state populations, ran quantile regressions to control for outliers, and performed a variety of subsample estimates. I have also run the basic regression using the BLS measure of the youth share of the population and a measure of state unemployment rates that I constructed directly from the March CPS.

TABLE III
ROBUSTNESS CHECKS

Age cohort	Column 1	Nobs	Column 2	Nobs	Column 3	Nobs	Column 4	Nobs
				UNEMPLOYMENT RATE				
16–24	−1.500	881	−2.890	881	−2.064	162	−1.004	195
	(0.311)		(0.384)		(0.740)		(0.368)	
25–54	−2.346	880	−3.421	880	−2.979	162	−2.028	195
	(0.356)		(0.474)		(0.936)		(0.444)	
55+	−2.562	872	−4.219	872	−2.621	156	−2.267	191
	(0.521)		(0.650)		(1.306)		(0.626)	
				PARTICIPATION RATE				
16–24	0.281	881	0.364	881	0.275	162	0.241	195
	(0.057)		(0.065)		(0.130)		(0.068)	
25–54	0.068	880	0.084	880	0.008	162	0.040	195
	(0.024)		(0.021)		(0.048)		(0.034)	
55+	0.200	872	0.307	872	0.365	156	0.143	191
	(0.099)		(0.119)		(0.173)		(0.124)	
	State and year fixed effects, AR(1) correction.		State and year fixed effects, state time trends, AR(1) correction.		Aggregated to the Census division level. Division and year fixed effects, AR(1) correction.		Five year averages. State and period fixed effects.	

Age cohort	Column 5 Short run	Column 5 Long run	Nobs	Column 6 1978–1987	Column 6 1988–1996	Nobs	p
			UNEMPLOYMENT RATE				
16–24	−0.835	−1.962	881	−1.142	−1.471	881	0.598
	(0.172)	(0.385)		(0.755)	(0.318)		
25–54	−1.030	−2.931	880	−2.677	−2.365	880	0.651
	(0.187)	(0.472)		(0.819)	(0.361)		
55+	−1.909	−2.808	872	−2.212	−2.536	872	0.750
	(0.403)	(0.566)		(1.231)	(0.529)		
			PARTICIPATION RATE				
16–24	0.158	0.346	881	0.435	0.294	881	0.239
	(0.034)	(0.072)		(0.144)	(0.060)		
25–54	0.026	0.079	880	0.064	0.068	880	0.921
	(0.012)	(0.036)		(0.057)	(0.025)		
55+	0.111	0.256	872	0.012	0.187	872	0.367
	(0.054)	(0.126)		(0.232)	(0.100)		
	Lagged dependent variable. State and year fixed effects.			Allows for a break in the elasticity in 1987. State and year fixed effects, AR(1) correction.			

IV estimates of equation (1) using data from 49 states from 1978 to 1996. Standard errors are in parentheses. In Column 5, "Short run" indicates the immediate effect of a change in the youth population share, and "Long run" indicates the long-run effect of a permanent change. In Column 6, p indicates the p-value of a test for parameter stability. Unemployment data are constructed by the BLS from the CPS and information in the UI system. The youth population share is the number of 16 to 24 year olds divided by the number of 16 to 64 year olds in a state, according to U. S. Census Bureau estimates. The instrumental variable is birthrates per capita lagged 16 to 24 years, from various years' of the Statistical Abstract of the United States. The following observations are missing from the BLS unemployment and participation data: 16–19 year olds in D.C. in 1996; 45–54 year olds in Utah in 1994; and 65 and over in Delaware, Idaho, Mississippi, Tennessee, and Utah in 1994.

economy that do not experience much unemployment, and, as is common in more developed economies, fertility rates are quite low. During the post-World War II period, however, the South gradually caught up with the North. Then during my sample period, I will observe a relative decline in the youth population share and in the unemployment rate in the South, although both are caused by a third factor. If this story is correct, the introduction of separate linear time trends for each state should substantially weaken the results. Instead, column 2 of Table III shows that the estimated elasticities are actually significantly larger than in the benchmark, eliminating this concern.

Next I turn to the possibility that neighboring states are not independent observations, as would be the case if the results are driven by regional fluctuations and regional birthrate shocks. For example, suppose that the recession in New England in the late 1980s happened to be correlated with a decline in the youth share of the population. Since there are six states in that Census division, this single event may have a substantial effect on the estimates. To see whether this drives the results, aggregate the state data up to the Census division level and rerun the basic regression. Column 3 of Table III shows that the point estimates are mostly unchanged, and remain statistically significant, except for the prime age participation rate. Another specification, including a separate year dummy for each census division in the basic regression in addition to the full set of state dummies, yields similar results (unreported).

A third issue is the dynamic relationship between the youth population share and economic conditions. Figure I shows that both of these variables change only slowly, a problem that I have addressed by correcting for first-order autocorrelation in the residuals. Tests indicate that higher order autocorrelation is not a problem, and in any case the results are not very sensitive to the specification. For example, aggregate the annual data to obtain five-year averages. This yields four observations for each state, 1978–1982, 1983–1987, 1988–1992, and 1993–1996. Column 4 of Table III shows the coefficients from the basic regression (1) run on this short panel (without any autoregressive correction). They are slightly smaller than the benchmark, but statistically unchanged and significantly different from zero.

I have also tried different specifications of the annual data. In Column 5, I put a single lag of the dependent variable on the right-hand side of the regression. The "Short-run" coefficient estimates are much smaller in magnitude. However, these represent the immediate effect of a change in the youth population share. To the extent that changes in the youth population share are permanent, the numbers in the “Long run” column indicate the eventual impact on the unemployment and participation rate. Once again, these are close to the benchmark estimates. Other unreported tests, such as running the regression in first differences, yield similar results.

Verifying the stability of the elasticity estimates over different time periods serves as a further robustness check. In column 6 of Table III, I allow the elasticity γ to take on different values in the first and second half of the sample. The resulting estimates are similar to those in the full sample, and, as indicated by the p-value, I cannot reject the null hypothesis that the elasticity is unchanged.

Finally, I verify that the main result holds in a completely different survey. The BLS uses the Current Employment Statistics (CES), an establishment survey, to construct an alternative measure of labor force activity. Although the CES does not contain information on unemployment and participation rates, as in the CPS household survey, it can be used to calculate total employment levels in the 49 states from 1970 to 1997. I then convert this into employment ratios using the Census measure of the 16–64 population. When I regress the log of this on the log Census youth population share and state and year fixed effects using the birthrate instrument, I obtain an elasticity of a quarter (Table IV, first row), somewhat larger than the estimated elasticity of the employment ratio in the CPS (Table I, last row).

Using the CES, it is also possible to calculate state employment levels in the eight one-digit industries. Let the state industry employment-population ratio be the level of employment in that industry divided by the Census population 16–64 years old. I regress the log of this on the contemporaneous log youth share and fixed effects, and show the results in the remaining rows of Table IV. The effect of the young workers is strongest on construction employment, with a slightly weaker impact on manufacturing, wholesale and retail trade, and services. However, since construction employed less than a quarter as many people as each of these three sectors during this 28-year period, the implied impact on the level of employment is largest in manufacturing, then in trade, services, and construction. This ordering is interesting, because one would expect that if the results were

TABLE IV
THE EFFECT OF YOUNG WORKERS ON THE EMPLOYMENT RATIO BY SECTOR

Industry	Youth share	Nobs
Total employment	0.251	1323
	(0.048)	
Construction	1.203	1311
	(0.224)	
Manufacturing	0.794	1323
	(0.105)	
Wholesale and retail trade	0.327	1323
	(0.057)	
Services	0.304	1229
	(0.068)	
Government	0.077	1307
	(0.054)	
Transportation and public utilities	0.044	1318
	(0.077)	
Finance, insurance, and real estate	−0.077	1317
	(0.093)	
Mining	−0.449	1234
	(0.348)	

IV estimates of the elasticity of the employment ratio for the entire economy and in eight one-digit industries using data from 49 states from 1970–1997. All regressions include state and year fixed effects and an AR(1) correction. Standard errors are in parentheses. The dependent variable is the (sectoral) employment ratio. Sectoral employment is computed by the BLS from the CES. For state population I use the number of 16–64 year olds according to Census estimates. The youth population share is the number of 16 to 24 year olds divided by the number of 16 to 64 year olds in a state, according to U. S. Census Bureau estimates. The instrumental variable is birthrates per capita lagged 16 to 24 years, from various years of the Statistical Abstract of the United States.

driven by an increase in aggregate demand, e.g., for new offices and housing, resources would be shifted from "tradables" to "non-tradables," leading to a relative decline in manufacturing employment. Instead, it seems that the explanation for this phenomenon must lie on the supply side.

E. Summary

Cross-state evidence suggests that a predictable 10 percent increase in the youth share of the population in a state reduces the unemployment rate by about 20 percent, with the strongest effect on prime age unemployment and a somewhat weaker effect on young workers. It also raises the participation rate of young workers by 3 percent, with a smaller and less significant effect on prime age workers. This implies a 2 or 3 percent increase in the employment ratio, with employment growth in a wide variety of sectors, notably manufacturing.

These elasticities not only have an unexpected sign, but they are very large. Consider a regression of the log youth share of the population, share_{it}, on state and year dummies. The residual has a standard error of 0.034, 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 6 percent and raise the participation rate of young workers by about 1 percent.

The implied impact of the baby boom on the aggregate unemployment rate is enormous. The youth population share bottomed out at 18 percent in 1953, rose to 27 percent by 1976, and has since fallen back to 19 percent. Roughly speaking, this change should have first halved and then doubled the prime age unemployment rate, in sharp contrast to the findings of Shimer [1998]. There are a number of reasons to be cautious with this calculation, however, 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. The level of aggregation may be important for the results; or international demographic shocks may primarily affect relative prices, while state shocks affect quantities, since prices are determined in the world market. Section VI explores this issue further.

IV. Theoretical Hypothesis

Why does an anticipated increase in the youth population share reduce the unemployment rate and raise the labor market participation rate? The standard model of equilibrium unemployment [Mortensen and Pissarides 1994; Pissarides 2000] implies that younger workers, and hence younger states, will have higher unemployment rates. This is due to the direct effect of young workers—they are unemployed more often. Indirect effects, i.e., changes in age-specific unemployment rates, are absent from the standard model, because the number of vacancies increases proportionately with unemployment, leaving both job creation and job destruction rates unchanged, see Shimer [1998] for details.

On-the-job search, a natural modification to the standard model, changes this conclusion. It becomes time-consuming for firms to locate the truly mobile workers when anybody is potentially interested in the job, a trading externality [Diamond 1982]. Since older workers are typically less mobile, this means hiring is easier in a younger economy. The presence of young workers generates a more fluid labor market, which stimulates vacancy creation. The aggregate unemployment rate may fall, but the many employed workers who continue searching for better jobs justify the creation of the new vacancies.

A. Model

There are two types of agents, workers and firms. Let $L(t) \equiv L_{t_0} e^{\int_{t_0}^{t} n(s)ds}$ denote the exogenous measure of the labor force at time t as a function of the initial population L_{t_0} and the subsequent growth rate $n(t) > 0$.[7] Let $\theta(t)L(t)$ denote the endogenous measure of firms, each of which can hire at most one worker. Agents are risk neutral, infinitely lived, and discount the future at rate $r > 0$. Each agent, worker or firm, may be in one of three states: unmatched, mismatched, or well matched, with exogenous flow payoffs $0 < x_1 < x_2$ in each of the respective states. New firms enter the market by paying a one-time fixed cost c, implying a perfectly elastic supply of jobs.[8]

Let $1 - \alpha_1(t) - \alpha_2(t)$, $\alpha_1(t)$, and $\alpha_2(t)$ denote the endogenous fraction of workers who unmatched, mismatched, and well matched at time t. Simple accounting shows that at time t, a fraction of firms $(\theta(t) - \alpha_1(t) - \alpha_2(t))/\theta(t)$ are unmatched. These fractions are limited by search frictions. 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 $\mu(\theta) \equiv \theta\eta(\theta)$.[9] 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

7. I assume that the entire time path of $n(t)$ is known. This reflects the empirical finding that the size of new cohorts can be forecast twenty years in advance using lagged birthrates. The results are not substantially affected if instead $n(t)$ is known only a few years in advance.

8. The elastic supply of jobs is driven by the empirical analysis, which looks at changes in the birthrate in one state relative to the national average. This source of variation may affect the flow of capital across states, but it should not change the "cost of capital" c.

9. All unmatched workers search for jobs, and are considered unemployed. Introducing an exogenous worker-specific value of not participating in the labor market would allow the model to address the labor market participation decision. More workers would participate in the labor market when the value of participating is higher, as I show will be the case in economies with high birthrates.

and over time. The 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.[10] A mismatch is accepted only if both agents are unmatched.[11] This means that it is easier to match when few agents are in good matches, a trading externality. Whenever a matched agent accepts a new partner, her old partner becomes unmatched. 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.

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 short-term employment. In the Burdett-Mortensen model, the marginal product of labor is constant, so firms never fire workers. One can show that without short-term employment, this framework cannot explain a decline in the aggregate unemployment rate in response to an increase in the labor force growth rate.

B. Equilibrium

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 a pair of initial conditions $\alpha_i(t_0) = \alpha_{it_0}$, $i \in \{1,2\}$. And second, the

10. Well-matched workers never leave their match, but they reduce the efficiency of the search process since searchers meet them. This assumption may seem unappealing, but it is important to the model as it is written down here. However, in a more general model with many match qualities, the results will still go through if workers can reduce the time and effort devoted to search as they climb the quality ladder.

11. This model may have another "perverse" equilibrium in which everyone leaves good matches for mismatches, making good matches unstable. This will be an equilibrium if the cost of instability, becoming unmatched, is high relative to the productivity difference across matches, $x_2 - x_1$ [Burdett, Imai, and Wright 2000; Webb 2000]. Also, the contemplated matching patterns are efficient if $x_2 \geq 2x_1$. Otherwise, output might be higher if mismatched agents did not match with other mismatched agents.

expected value of creating a new unmatched firm is always equal to the startup cost c.

It is easy to show that the state variables evolve according to a pair of differential equations:

$$\alpha_2(t) = \mu(\theta(t))(1 - \alpha_2(t))\frac{\theta(t) - \alpha_2(t)}{\theta(t)}p - n(t)\alpha_2(t) \tag{2}$$

(3)

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

The fraction of matches that are good increases with the confluence of four independent events: a worker meets a firm (rate μ), the worker is not in a good match (fraction $1 - \alpha_2$), the firm is not in a good match (fraction $1 - \alpha_2/\theta$), and they realize a good match (fraction p). It decreases only because of the entry of new unmatched workers. The fraction of matches that are bad takes a similar form, but it decreases due to new entrants (rate n), workers quitting for good matches (rate $\mu(\theta)(1 - \alpha_2/\theta)p$), and firms quitting for good matches (rate $\eta(\theta)(1 - \alpha_2)p$).

Next I check what drives the value of creating a new unmatched firm $W_0(t)$ to zero. To do this, I must calculate the value of mismatched and well-matched firms, $W_1(t)$ and $W_2(t)$, respectively. These are related by three Bellman equations: $W_2(t) \equiv x_2/2$ (since this is an absorbing state);

$$\begin{aligned}rW_1(t) - \dot{W}_1(t) &= x_1 + \eta(\theta(t))(\theta(t) - \alpha_2(t))p(W_0(t) - W_1(t)) \\ &\quad + \eta(\theta(t))(1 - \alpha_2(t))p(W_2(t) - W_1(t));\end{aligned} \tag{4}$$

and

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

These equations say that the flow value of a firm in a particular state is its flow payoff while it remains in that state, x_1 while mismatched and zero while unmatched, plus terms that are the product of the flow probability of switching states times the resulting capital gain or loss.

A simple algorithm delivers a numerical approximation to the equilibrium. Conjecture a path for the firm-worker ratio θ, 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 guess toward 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.

Note that the size of an economy, measured by the initial population L_{t_0}, does not enter these equations, and so the model does not feature "scale effects" or agglomeration externalities, as were explored by Ciccone and Hall [1996] or Coles and Smith [1996]. Instead the efficiency of the matching process depends on the proportion of agents who are actively searching, as in Diamond [1982], a distinct and potentially complementary source of increasing returns to scale in the labor market. I focus on it because it matches the empirical evidence presented in Section III.

C. Simulation

Consider the qualitative effect of an increase in the birthrate n. Holding the firm-worker ratio θ fixed, the direct effect of this is to increase the unemployment and mismatch rates, since young workers take a while to find good jobs. This would raise the value of unmatched firms, inducing entry and increasing θ, until the zero profit condition is restored. The increase in θ may be enough to push the aggregate unemployment rate below its preshock level; in that case, the new job creation is sustained by a large mismatched population. Simulations of the model using a wide variety of parameters show that this is in fact the "likely" scenario.[12]

I provide details from one of these simulations to illustrate how the economy responds to a shock. Think of a period as

12. If both the unmatched and mismatched rates fall below their preshock levels, the value of an unmatched firm would fall, which cannot be an equilibrium. So at least one of these rates must increase; and equations (2) and (3) imply that if the birthrate, firm-worker ratio, and unmatched rates all rise, the mismatch rate must rise as well. Putting this logic together, an increase in the birthrate can induce a decline in the unmatched rate, but can *never* induce a decline in the mismatch rate. It is possible that both the unmatched and mismatched rates are higher in a high birthrate economy, as would be the case if firms' meeting rate η is sufficiently responsive to changes in the firm-worker ratio; but simulations indicate that this is unlikely to occur.

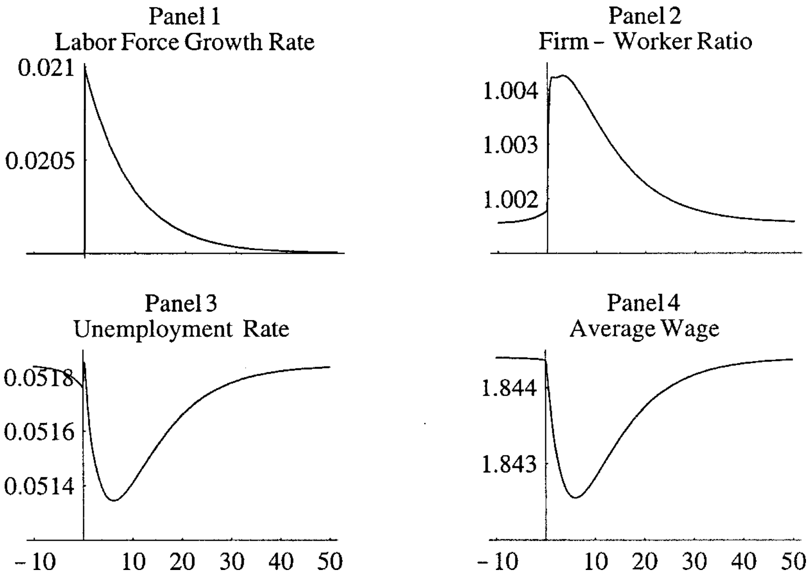

FIGURE II

Response to a Labor Force Growth Rate Shock

The dynamic response of the labor force growth rate, firm-worker ratio, unemployment rate, and average wage to an anticipated increase in the labor force growth rate from 0.020 to 0.021 in year 0.

representing a year. 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 its steady state; the firm-worker ratio is $\theta = 1$, about 5.2 percent of workers are unemployed, 14.8 percent are mismatched, and the remaining 80.1 percent have found good matches.

Now suppose that there is an anticipated 5 percent increase in the labor force growth rate in year 0 which gradually disappears. For $t \geq 0$, $n(t) = 0.02 + 0.001e^{-0.110t}$ (Figure II, panel 1).[13] In response, the firm-worker ratio jumps up at date zero (panel 2), causing the aggregate unemployment rate to fall (panel

13. Regress birthrates for the 49 states from 1944 to 1991 on state and year fixed effects. The residuals follow an AR(1) process $\eta_{it} = 0.896\eta_{it-1} + \nu_{it}$. Translating this difference equation into continuous time yields this process.

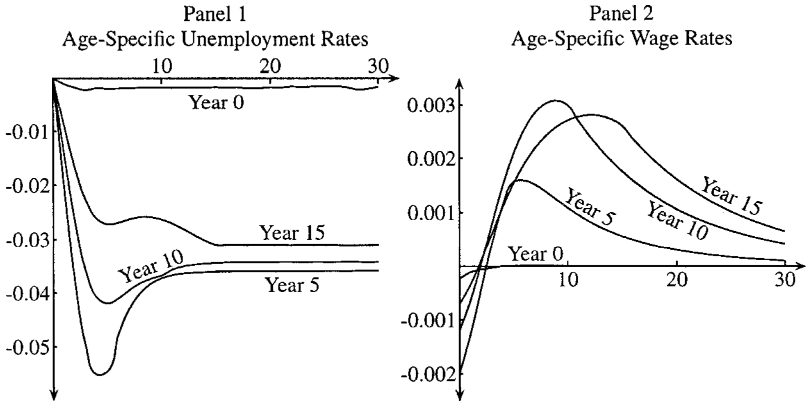

FIGURE III
The Change in the Unemployment Rate and Average Wage as a Function of Age

The log change in the unemployment rate (Panel 1) and the average wage (Panel 2) compared the steady state, as a function of the worker's age in years 0, 5, 10, and 15. The shock is an anticipated increase in the labor force growth rate from 0.020 to 0.021 in year 0.

3). This is driven by a sharp decrease in "age"-specific unemployment rates. Panel 1 of Figure III shows the difference between the log unemployment rate of an age $s \in [0,30]$ worker in year $t = \{0,5,10,15\}$, compared with the steady state log unemployment rate of age s workers. Five years after the shock, most workers enjoy a 4 percent reduction in their unemployment rate. This gradually fades away, although workers who were around at the time of the shock, e.g., those over fifteen years old in year 15, have persistently less unemployment. This matches the central empirical finding uncovered in the previous section.

The model also makes predictions about the behavior of wages. The additional entry is driven by an increase in mismatch and a decrease in good matches. Since mismatched workers earn less than well-matched ones, it follows that the average wage rate among employed workers, $(x_1\alpha_1(t) + x_2\alpha_2(t))/(\alpha_1(t) + \alpha_2(t))$, must fall when labor force growth is faster (Figure II, panel 4). But this prediction does not distinguish the model from a number of alternatives. Young workers earn less on average than older workers, and so the birthrate shock directly reduces wages.

More interesting is the model's prediction regarding age-specific wage rates (Figure III, panel 2). Five years after the shock, the wage of many workers has increased, but not those of the newest entrants. Instead, the increase in the firm-worker ratio θ implies that many more firms are unmatched, hence willing to create a mismatch, which actually reduces the average wage of these workers. Over time, however, age-conditional wages start to increase, peaking ten to fifteen years after the shock. Wages stay high in later years, particularly for those workers who entered immediately following the shock. This prediction is stark, and appears to be unique to models that predict that an increase in labor supply will create more than its own demand. Subsection V.A finds support for this prediction in U. S. data.

The model also predicts that in young states, more workers should move from one job to another. This is a first-order prediction, since it reflects the fact that firms create jobs in young states because they can attract employed workers. Subsection V.B finds some support for this prediction using job flow data. However, the power of the test is weak, since many other models predict that young workers are more likely to move from job-to-job, for example, Neal [1999].

V. Testing the Hypothesis

A. Wages

According to the model, an increase in the labor force growth rate should cause a modest decline in the average wage, but should gradually push up age-conditional wages. I test this using wage data for workers sixteen and over from the 1978–1998 March CPS. Define an individual's hourly wage to be her annual earnings divided by the product of the number of weeks worked and the usual number of hours per week.[14] Throw away individuals for whom either the numerator or denominator is zero. Also trim the top and bottom 1 percent of the sample in each year. Finally, construct the average log hourly wage in each state and year using the CPS sampling weights.

The first row of Table V shows the results from an instrumental variables regression of a state's average log wage on the

14. To be precise, since the 1994 CPS redesign, the numerator is a recode for the previous year's total wage and salary earnings (P243). The denominator is the product of the answer to two questions: "During [the previous year] in how many weeks did . . . work even for a few hours. Include paid vacation and sick leave as work?" [P171]; and "In the weeks that . . . worked, how many hours did he/she usually work per week?" [P181]. I use the answers to similar questions throughout the sample period. Note that since the questions concern the previous year, the sample period covered is 1977–1997.

TABLE V
THE EFFECT OF YOUNG WORKERS ON WAGES

Wage measure	Youth share	Nobs
Average log hourly wage	−0.184 (0.067)	980
Average log weekly earnings	−0.155 (0.078)	980
Log average annual pay	−0.102 (0.028)	832

IV estimates of the elasticity of log wage rates with respect to the youth population share for 49 states. All regressions include state and year fixed effects and an AR(1) correction. Standard errors are in parentheses. The first two rows include data from 1977–1997, and the dependent variable was computed by the author from the 1978–1998 March CPS. The last row includes data from 1981–1998, and the dependent variable comes from the Covered Employment and Wages (ES-202) program. New Jersey data are unavailable for 1998. The youth population share is the number of 16 to 24 year olds divided by the number of 16 to 64 year olds in a state, according to U. S. Census Bureau estimates. The instrumental variable is birthrates per capita lagged 16 to 24 years, from various years of the Statistical Abstract of the United States.

contemporaneous log youth population share and a full set of state and year dummies. The estimated elasticity is negative and fairly small in magnitude. The next row shows that a similar result obtains using average log weekly earnings. Calculating the average wages in levels before taking logs also yields similar (unreported) results. The final row of Table V uses a completely different dependent variable, the logarithm of Average Annual Pay in state i and year t. These data come from the Covered Employment and Wages or ES-202 program, an establishment survey that includes all employed workers covered by unemployment insurance, about 96 percent of total wage and salary income in 1998. It is the ratio of total pay in a particular state and year divided by the average employment level.[15] State data from 1981 to 1996 are available from historical issues of the Statistical Abstract of the United States, and 1997 and 1998 data are available online at the BLS web site ⟨http://stats.bls.gov/⟩. Despite the very different source of data, the estimated elasticity is statistically unchanged and significantly negative.

The more striking prediction of the model is that age-contingent wages should rise only slowly after an increase in the popu-

15. Annual Pay generally includes bonuses, tips, retirement contributions, stock options, and the cash value of meals and lodging. This is aggregated to the state level, and then divided by the average number of employees in that state on a particular week of each month, including workers on paid vacation and part-time workers. Therefore, Average Annual Pay is actually a measure of average weekly earnings. For details, see ⟨http://stats.bls.gov/news.release/annpay.tn.htm⟩.

lation growth rate, peaking perhaps a decade later than the trough in age-contingent unemployment rates (Figure III). To test this, I use the methodology described above to construct state average log hourly wages for seven different age groups and both sexes from the March CPS. I then regress this on the contemporaneous log youth population share and on the log youth population share in that state a decade earlier. Since state demographic data are unavailable between decennial censuses in the 1960s, I can only construct the latter variable beginning in 1980, and so restrict my sample to 1980–1997. Finally, I instrument each youth population share variable using the sum of appropriately lagged birthrates.

Table VI shows that a large youth population share causes a modest contemporaneous increase in wages for most of the fourteen groups, and has a much larger effect ten years later. For example, a 20-year-old male who is part of a cohort that is 10 percent larger than normal can expect 5.3 percent higher wages now and 11.6 percent higher wages by the time he is 30. A man who is 30 when the big cohort enters expects 2.7 percent higher wages now and 3.9 percent higher wages when he is 40 (although neither of these estimates is significantly different from zero). These results are qualitatively consistent with the model,[16] although some of the quantitative predictions of the model do not appear in this data set. For example, the model predicts that wages will decrease for the youngest workers when birthrates increase, while the data show that virtually all groups enjoy higher wages. Note, however, that teenagers do get smaller wage increases than 20–24 year olds, which is again qualitatively consistent with the model. Also, the model predicts that workers

16. One might be concerned that these results are spurious, particularly since many of the point estimates are not significantly different from zero. While I cannot rule out the possibility that the relationships reported in Table VI are driven by a third factor, there are at least two reasons to believe that this is unlikely. First, I ran the same regression with age-specific unemployment rates as the dependent variable, and found that the lagged youth population share has little explanatory power, as the model predicts. This suggests that there is no general relationship between labor market conditions and lagged demographic variables. Second, I regressed the average log hourly wage on the contemporaneous youth share of the population and the contemporaneous share of 25–34 year olds in the state. In each of the fourteen cases, I found that the estimated elasticity of the 25–34 year old share in this regression was smaller than the elasticity of the 16–24 year old share from ten years earlier reported here. This is what one would expect if the share of 25–34 year olds today is simply a noisy measure of the share of 16–24 year olds a decade ago and it is the lagged youth share that causes high wages.

TABLE VI
THE EFFECT OF YOUNG WORKERS ON WAGES, BY AGE AND SEX

Age cohort	Youth share	Lagged youth share	Nobs
	MEN		
16–19	0.511	1.242	833
	(0.270)	(0.662)	
20–24	0.529	1.938	833
	(0.281)	(0.708)	
25–34	0.266	1.160	833
	(0.195)	(0.494)	
35–44	0.151	0.386	833
	(0.175)	(0.438)	
45–54	0.089	0.551	833
	(0.199)	(0.493)	
55–64	0.296	0.997	833
	(0.280)	(0.696)	
65+	0.119	2.038	833
	(0.570)	(1.398)	
	WOMEN		
16–19	0.252	0.029	833
	(0.246)	(0.603)	
20–24	0.710	1.513	833
	(0.222)	(0.553)	
25–34	0.541	1.793	833
	(0.206)	(0.524)	
35–44	0.191	1.052	833
	(0.185)	(0.463)	
45–54	0.543	1.946	833
	(0.224)	(0.561)	
55–64	0.242	1.155	833
	(0.365)	(0.657)	
65+	−0.112	0.360	833
	(0.495)	(1.211)	

IV estimates of the elasticity of average log hourly wages with respect to current and lagged youth population share for 49 states from 1980–1997. All regressions include state and year fixed effects and an AR(1) correction. Standard errors are in parentheses. The dependent variables were computed by the author from the March CPS. The youth population share is the number of 16 to 24 year olds divided by the number of 16 to 64 year olds according to Census estimates. The lagged youth population share is the same ratio taken ten years earlier. The instruments are birthrates per capita lagged 16 to 24 and 26 to 34 years, from various years of the Statistical Abstract of the United States.

who enter the labor force as members of large cohorts will later earn higher wages than surrounding cohorts, which is not borne out by the data. In particular, women's wage gains are spread throughout the age distribution. This is sensible if women exit

TABLE VII
THE EFFECT OF YOUNG WORKERS ON JOB CREATION AND DESTRUCTION

Dependent variable	OLS	IV	Nobs
Job creation	1.006	1.454	720
	(0.404)	(0.767)	
Job destruction	0.932	0.817	720
	(0.459)	(0.875)	

OLS and IV estimates of the elasticity of job creation and destruction with respect to the youth share for 48 states from 1973–1988. Both regressions include state and year fixed effects and an AR(1) correction. Standard errors are in parentheses. The dependent variables were computed from the Longitudinal Research Database, and are described in Davis, Haltiwanger, and Schuh [1996]. The youth population share is the number of 16 to 24 year olds divided by the number of 16 to 64 year olds in a state, according to U. S. Census Bureau estimates. The instrumental variable is birthrates per capita lagged 16–24 years, from various years of the Statistical Abstract of the United States.

and reenter the labor market, a possibility ignored by the theory, which is better at modeling and predicting men's wages.

B. Job Flows

I test the prediction that younger states should have more workers moving from one job to another using job creation and destruction data for manufacturing.[17] The data were constructed by Davis, Haltiwanger, and Schuh [1996] from the Longitudinal Research Database, and include annual observations from 1973–1988 for 48 states (Rhode Island data are unavailable). Job creation in state i and year t measures 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. I regress the log of job creation or job destruction on the log youth share and a full set of year and state dummies. The null is that the elasticity of the creation and destruction variables should be equal to zero, while the model predicts a positive value for *both* elasticities. Table VII shows that the OLS and IV estimates are positive, although the IV estimate of the elasticity of job destruction is not statistically different from zero. The model 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

17. Ideally, I would like to use worker flow data, but they are unavailable at the state level.

estimates here with the estimates for the manufacturing sector in Table IV offers modest support for that prediction, although the difference between the estimates is again not statistically significant. I conclude that job creation and destruction data do not contradict the model, but rather offer some support.

VI. Reconciliation with Previous Studies

A number of previous studies have found that an increase in the youth share of the population raises the youth unemployment rate. A common feature of all these studies is that they explicitly or implicitly assume that demographic changes do 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 unemployment rates in U. S. time series data as evidence that an increase in the youth share raises the teenage unemployment rate.[18] More recently, Shimer [1998] attributes to demographics any change in the aggregate unemployment rate that cannot be predicted from a linear regression on the prime age unemployment rate.

Korenman and Neumark [2000] make the same assumption in their cross-country panel analysis. They regress the youth (16–24) unemployment rate on the youth share of the population instrumented by lagged birthrates and on the prime age (25–54) unemployment rate:

$$\log ur_{it}^{\text{youth}} = \alpha_i + \beta_t + \gamma \log \text{share}_{it} + \delta \log ur_{it}^{\text{prime}} + \epsilon_{it}. \tag{6}$$

This last variable is intended to control for macroeconomic fluctuations. Unfortunately, if the youth share of the population changes the prime age unemployment rate, the estimate of the effect of the youth share on the youth unemployment rate γ will be biased up. The positive cyclical correlation between prime age and youth unemployment rates ensures a positive coefficient δ multiplying the prime age unemployment rate. Then if an in-

18. This relationship still appears in monthly U. S. data. Regress the difference between the youth (16–24) and prime age (25–54) unemployment rates on the youth share of the working age (16–64) population using a single U. S. time-series from 1948–1998. The estimated coefficient on the youth share is 0.308 (standard error 0.017), meaning that a ten percentage point increase in the youth share of the population is correlated with a three percentage point increase in the difference between the youth and prime age unemployment rates. The correlation that Flaim [1979] first uncovered is not in doubt, merely the interpretation of it.

crease in the youth share of the population reduces both the prime age and the youth unemployment rate, part or all of the effect will be captured by δ. The estimated value of γ may even become positive. As confirmation of this reasoning, estimate equation (6) using instrumental variables with the 49 state, nineteen-year (1978–1996) panel that I have exploited throughout this paper. The elasticity of the youth unemployment rate with respect to the prime age unemployment rate is precisely estimated to be $\delta = 0.57$ (standard error 0.03), while the estimated effect of the youth population share is a biased $\gamma = 0.01$ (0.21).

This suggests using my methodology on other studies' data sets to see whether the conclusions change. Korenman and Neumark [2000] have constructed a fifteen-OECD country, 25-year (1970–1994) data set, to my knowledge the only other panel containing the necessary data. I therefore estimated equation (1) on this data set using both OLS and IV. Table VIII shows that the coefficient estimates often have the opposite sign of my estimates using cross-state data (Tables I and II). There is possibly some evidence that an increase in the youth share of the population raises the labor market participation rate and the employment ratio. However, the indirect effects are too weak to counteract the direct effects, and so the aggregate labor market participation rate and employment ratios are statistically unchanged. And there is no evidence that an increase in the youth share of the population reduces the unemployment rate.

There are a couple of possible interpretations for these findings. First, in many OECD countries, employment ratios may better reflect labor market outcomes than do unemployment rates, because changes in labor market policies and institutions move the ephemeral line between workers who are unemployed and those who are not in the labor force [Juhn, Murphy, and Topel 1991]. If this is correct, then we should focus on the employment ratio elasticities in Table VIII, panel C. The large standard errors, a consequence of the importance of country-specific macroeconomic shocks and the relatively small sample of fifteen countries, force us to conclude that we cannot reject the null hypothesis that the cross-country and cross-state results are the same. In other words, the cross-country data set does not contradict the earlier findings using U. S. data, that an increase in the youth share of the population caused by a birthrate shock raises the employment ratio.

An alternative interpretation is to accept that the unemploy-

TABLE VIII
THE EFFECT OF YOUNG WORKERS ON UNEMPLOYMENT AND PARTICIPATION, OECD DATA

Dependent variable	OLS	IV	Nobs
	A. UNEMPLOYMENT RATE		
All	0.813	1.418	287
	(0.410)	(0.664)	
15–24	0.774	1.552	287
	(0.423)	(0.650)	
25–54	0.683	1.249	287
	(0.466)	(0.717)	
55–64	−0.360	−0.200	280
	(0.625)	(0.809)	
	B. PARTICIPATION RATE		
All	−0.175	−0.012	287
	(0.028)	(0.054)	
15–24	−0.042	0.624	287
	(0.085)	(0.295)	
25–54	−0.161	0.136	287
	(0.031)	(0.084)	
55–64	−0.086	0.080	287
	(0.081)	(0.167)	
	C. EMPLOYMENT RATIO		
All	−0.274	−0.053	287
	(0.049)	(0.136)	
15–24	−0.089	0.410	287
	(0.129)	(0.372)	
25–34	−0.219	0.086	287
	(0.045)	(0.126)	
55–64	−0.101	0.139	287
	(0.094)	(0.229)	

OLS and IV estimates of equation (1) using fifteen OECD countries from 1970–1994. All regressions include country and year fixed effects, an AR(2) correction, and dummy variables to account for changes in some of the data series. Standard errors are in parentheses. The data were constructed and provided to me by Sanders Korenman and David Neumark, except as indicated. The countries are Australia, Canada, Finland, France, Germany, Ireland, Italy, Japan, the Netherlands, Norway, Portugal, Spain, Sweden, the United Kingdom, and the United States. In most countries, data are available for all workers and for three age groups: 15–24, 25–54, and 55–64. In Norway, Spain, Sweden, the United Kingdom, and the United States, the youngest age group is 16–24. In Italy the youngest age group is 14–24, and the middle one is 25–59. Unemployment, participation, and employment rates were calculated from OECD statistics, except for the aggregate numbers for AU, CA, FR, GE, IT, JA, SW, UK, and US, which I calculated from the Bureau of Labor Statistics "Foreign Labor Statistics." The youth population share is the ratio of the size of the youngest group to all three groups combined. Birthrate data for Europe come from *International Historical Statistics: Europe, 1850–1988;* for Australia and Asia from *International Historical Statistics: Africa, Asia and Oceania;* and for the United States from *Vital Statistics of the United States, 1991.* The following observations are missing: Germany and Italy, 1994; Ireland, 1970, 1972–1974, 1976, 1978, 1980, 1982, and 1994; Japan and Spain, 1970–1971; the Netherlands, 1970; Norway, 1970–1977; Portugal, 1970–1973 and 1994; and the United Kingdom, 1970–1972, 1974, and 1978–1979. For more details see Korenman and Neumark [2000].

ment rate is informative, but to recognize that additional factors may be important in generating the cross-country data. 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 its marginal product.[19] In terms of the theory developed in Section IV, even if one state has a higher population growth rate than another, capital mobility ensures that the cost of job creation c will be the same across states. This was my justification for assuming that a change in the population growth does not affect c. But if one country has a high population growth rate and international capital flows are imperfect, the cost of job creation will be driven up, reducing the capital-labor ratio in that country, and therefore increasing the unemployment rate.

An appealing feature of this explanation for the difference between the cross-country and cross-state results, is that it also reconciles this paper with the bulk of the cohort crowding literature (discussed in the introduction) which relies on time series evidence. In post-World War II time series, the baby boom provides the bulk of the variation in the independent variable, a measure of the youth share of the population. But the baby boom was a phenomenon throughout the developed world. For example, Korenman and Neumark's [2000] data show a decline in the youth share of the population from 1970 to 1994 in each of the fifteen OECD countries that they study. Thus, even if capital is perfectly mobile across international boundaries, the global capital-labor ratio should have been low at the peak of the baby boom, pushing up the cost of job creation c, and therefore raising the unemployment rate.[20] If this reconciliation of the literature is valid—at this point, it is only a conjecture—then the existing

19. Borjas, Freeman, and Katz [1996] point out an analogous issue in their analysis of the labor market impact of immigration into the United States. Regressing state labor market outcomes on the share of immigrants and time dummies yields an underestimate of the impact of immigration if some of the impact is spread around the country through the migration of capital or native-born workers.

20. In Shimer [1998] I show that the prediction of the textbook [Pissarides 2000] search model is that an increase in the youth share of the population should not affect age-specific unemployment rates. I reconcile this with evidence of a relative increase in the youth unemployment rate during the 1970s, by arguing that labor demand may not be perfectly elastic; i.e., that the cost of job creation may have increased during that decade [page 56]. This is the same argument that I am making here.

literature may be correct in concluding that the baby boom raised the aggregate unemployment rate; yet this paper is correct in concluding that regions with relatively young populations will enjoy relatively low unemployment rates.

This paper is also related to a large literature that examines the impact of international migration on labor market outcomes. Friedberg and Hunt [1995] and Borjas, Freeman, and Katz [1996] review a number of "area-studies" that regress wage or employment rates on measures of immigration. Other papers control for the endogeneity of migration decisions using natural experiments, including the migration of Russian Jews to Israel in the early 1990s [Friedberg 1997], the repatriation of French Algerians in 1962 [Hunt 1992], and the flight of Cubans to Miami in 1980 [Card 1990]. Most of these papers focus on the effect of immigration on native-born workers, a notable contrast to the "baby boom" literature's presumption that the size of the youth cohort does not affect prime age workers. Although some papers conclude that immigration reduces wages and employment rates for native-born workers, especially less skilled ones [Altonji and Card 1991], many papers find no discernible effect.

It might seem that the evidence and theory developed in this paper would suggest that immigration should raise native-born wages and employment rates. But there is a crucial difference between the two types of labor supply shocks: an increase in the youth share of the population can be forecast twenty years ahead, while the immigration literature has focused on unexpected changes, e.g., the breakup of the Soviet Union, Algerian independence, and more generally, the portion of immigration that cannot be predicted by area fixed effects. Introducing time-to-build or adjustment costs for capital into the model in Section IV can reconcile these results. Suppose that the measure of jobs may only increase slowly following an unanticipated increase in the return to capital W_0. Then in the short run, an increase in immigration will reduce the firm-worker ratio θ. This will make it harder for native workers to find and retain jobs, raising their unemployment rate and lowering their wages. These indirect effects then reinforce the direct effects of immigrants on labor market outcomes: high unemployment and low wages. However, once capital has a chance to adjust, the mechanism described in this paper will prevail.

As some confirmation of this reasoning, consider what happens if one includes area fixed effects in an estimate of the impact of immigration. When Borjas, Freeman, and Katz [1996] regress the change in wages for native-born workers on the change in the percentage of foreign-born workers, they obtain a positive correlation; but this is reversed when fixed effects are included (see their Table 2). They argue that this discrepancy reflects immigrants' tendency to move to permanently high wage areas. According to this logic, the fixed effects soak up permanent and exogenous differences in labor market conditions. But there is an alternative interpretation: fixed effects absorb the anticipated component of immigration. Since only anticipated immigration induces job creation, the resulting estimates describe the negative relationship between unanticipated immigration and native-born workers' wages. They omit an important channel by which immigration raises native wages, the theory described in this paper. Turning Borjas, Freeman, and Katz on its head, I conjecture that an important reason why there are significant permanent differences in labor market conditions across regions, is that some regions regularly receive more immigration.

VII. Conclusion

This paper shows that a relative increase in the youth share of the population in one state causes an immediate decline in the unemployment rate and an increase in the labor market participation rate, and later causes an increase in wages for women and younger men. This is inconsistent with standard theories of unemployment, which predict either 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 these labor markets. It is also inconsistent with a substantial body of empirical evidence on the impact of cohort size on labor market outcomes. The reason for the different results remains a puzzle, although it may be due to the use of time fixed effects in a panel of states where there is excellent cross-sectional capital mobility. The fixed effects absorb a source of variation that was important to existing studies, changes in the cost of capital that are correlated with the youth share of the population. Whether this conjecture is correct remains an open question.

One can interpret the empirical results in this paper as a test for search externalities. The standard theory of equilibrium unemployment, as summarized in Pissarides [2000], 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 toward 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 due to trading externalities. 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 trading externalities "are one of the most promising ways to explain the business cycle." Models with trading externalities 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 IV.

One can therefore test for search externalities by looking for exogenous variation in the number of job searchers. Lagged changes in birthrates are an ideal source of variation. First, there is substantial variation in birthrates over time and across regions. Second, birthrates are easily measurable, and good data are widely available. Third, lagged birthrates 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 can be explained by the existence of trading externalities in search markets.

PRINCETON UNIVERSITY AND NATIONAL BUREAU OF ECONOMIC RESEARCH

REFERENCES

Altonji, Joseph, and David Card, "The Effects of Immigration on the Labor Market Outcomes of Less-Skilled Natives," in John Abowd and Richard Freeman, eds., *Immigration, Trade and the Labor Market* (Chicago: University of Chicago Press, 1991), pp. 201–234.

Blanchard, Olivier, and Lawrence Katz, "Regional Evolutions," *Brookings Papers on Economic Activity* (1992), 1–61.

Borjas, George, Richard Freeman, and Lawrence Katz, "Searching for the Effect of Immigration on the Labor Market." *American Economic Review,* LXXXVI (1996), 246–251.

Burdett, Kenneth, Ryoichi Imai, and Randall Wright, "Unstable Relationships," University of Pennsylvania Working Paper, 2000.

Burdett, Kenneth, and Dale Mortensen, "Equilibrium Wage Differentials and Employer Size," *International Economic Review,* XXXIX (1998), 257–274.

Byerly, Edwin, and Kevin Deardorff, *National and State Population Estimates: 1990 to 1994.* U. S. Bureau of the Census, Current Population Reports, P25–1127 (Washington, DC: U. S. Government Printing Office, 1995).

Card, David, "The Impact of the Mariel Boatlift on the Miami Labor Market," *Industrial and Labor Relations Review,* XLIII (1990), 245–247.

Ciccone, Antonio, and Robert Hall, "Productivity and the Density of Economic Activity," *American Economic Review,* LXXXVI (1996), 54–70.

Coles, Melvyn, and Eric Smith, "Cross-Section Estimation of the Matching Function: Evidence from England and Wales," *Economica,* LXIII (1996), 589–597.

Davidson, Russell, and James MacKinnon, *Estimation and Inference in Econometrics* (New York: Oxford University Press, 1993).

Davis, Steve, John Haltiwanger, and Scott Schuh, *Job Creation and Destruction* (Cambridge, MA: MIT Press, 1996).

Diamond, Peter, "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy,* XC (1982), 881–894.

Flaim, Paul, "The Effect of Demographic Change on the Nation's Unemployment Rate," *Monthly Labor Review,* CII (1979), 13–23.

——, "Population Changes, the Baby Boom and the Unemployment Rate," *Monthly Labor Review,* CXIII (1990), 3–10.

Friedberg, Rachel, "The Impact of Mass Migration on the Israeli Labor Market," Brown University Working Paper, 1997, *Quarterly Journal of Economics,* forthcoming.

Friedberg, Rachel, and Jennifer Hunt, "The Impact of Immigrants on Host Country Wages, Employment and Growth," *Journal of Economic Perspectives,* IX (1995), 23–44.

Gordon, Robert, "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment," In Martin Baily, ed., *Workers, Jobs and Inflation* (Washington, DC: Brookings Institution, 1982), pp. 89–152.

Hall, Robert, "The Beveridge Curve: Comment," *Brookings Papers on Economic Activity* (1989), 61–64.

Hunt, Jennifer, "The Impact of the 1962 Repatriates from Algeria on the French Labor Market," *Industrial and Labor Relations Review,* XLV (1992), 556–572.

Juhn, Chinhui, Kevin Murphy, and Robert H. Topel, "Why Has the Natural Rate of Unemployment Increased over Time?" *Brookings Papers on Economic Activity* (1991), 75–126.

Korenman, Sanders, and David Neumark, "Cohort Crowding and Youth Labor Markets: A Cross-National Analysis," in David Blanchflower, and Richard Freeman, eds., *Youth Employment and Joblessness in Advanced Countries* (Chicago: University of Chicago Press, 2000), pp. 57–105.

Mortensen, Dale, and Christopher Pissarides, "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies,* LXI (1994), 397–415.

Neal, Derek, "The Complexity of Job Mobility among Young Men," *Journal of Labor Economics,* XVII (1999), 237–261.

Perry, George, "Changing Labor Markets and Inflation," *Brookings Papers on Economic Activity* (1970), 411–441.

Pissarides, Christopher, *Equilibrium Unemployment Theory* (Cambridge, MA: MIT Press, 2000).

Shimer, Robert, "Why Is the U. S. Unemployment Rate So Much Lower?" in Ben Bernanke and Julio Rotemberg, eds., *NBER Macroeconomics Annual,* XIII (Cambridge, MA: MIT Press, 1998), 11–61.

Webb, Tracy, "Marriage and Divorce," University of Bristol Working Paper, 2000.