

Why Does Unemployment Differ Persistently Across Metro Areas?

By Jordan Rappaport

Unemployment rates differ widely across U.S. metropolitan areas. In 2007, they ranged from 3.1 percent or less among the 25 lowest unemployment metropolitan areas to 6.6 percent or more among the 25 highest unemployment metropolitan areas. Such large differences in metro unemployment rates have held continuously since at least 1990. Moreover, those metro areas that had a relatively high unemployment rate in one year tended to have a high unemployment rate 10 years and even 20 years later.

How can such large differences in unemployment rates persist over such long periods? Why don't households move from high long-term unemployment metros to low long-term unemployment metros in order to improve their chances of finding a job? Why don't firms move from low long-term unemployment metros to high long-term unemployment metros in order to more easily hire workers?

While such moves might seem sensible, they might not actually improve household welfare or firm profitability. A first possibility is that the skills of workers in high unemployment metro areas poorly match the hiring needs of firms elsewhere. Conversely, the needs of firms in

Jordan Rappaport is a senior economist at the Federal Reserve Bank of Kansas City. Michael Connolly, a research associate at the bank, helped prepare the article. This article is on the bank's website at www.KansasCityFed.org.

low unemployment metros may not match the skills of workers elsewhere. A second possibility is that some metros may have intrinsic characteristics that make households and firms unwilling to move. A third possibility is that high moving costs dampen the mobility of households and firms and, therefore, allow metro unemployment rates to diverge over long periods.

Empirical analysis finds support for all three explanations. Metro workforce characteristics are able to account for the largest share of the variation in metro unemployment rates measured over complete business cycles. Characteristics more intrinsic to metro areas themselves account for much of this variation as well, though not as much as workforce characteristics. And moving costs are estimated to be high enough that some households will be unwilling to move away from high-unemployment metros.

Section I describes in more detail the dispersion and persistence of metro unemployment rates. Section II describes the generally weak effect of metro employment growth on metro unemployment. Sections III and IV describe the separate and combined correlations between the metro workforce and intrinsic characteristics and metro unemployment. Section V describes estimates of the size of moving costs. The concluding section discusses some public policy implications of the results.

I. THE DISPERSION AND PERSISTENCE OF UNEMPLOYMENT ACROSS METRO AREAS

To set the stage for explaining the persistent, large unemployment differences across metros, this section describes these differences in more detail. Throughout the analysis, metro unemployment rates are based on the monthly Bureau of Labor Statistics survey and definition. A person is unemployed if he does not have a paid or unpaid job, was available for work, and sought work over the previous four weeks. The unemployment rate is the number of unemployed workers divided by the sum of employed plus unemployed workers. The analysis that follows focuses on average unemployment rates over the two most recent national business cycles, with approximate peaks in 1990, 2000, and 2007.¹

The high dispersion of unemployment rates

The dispersion of unemployment rates across metropolitan areas is typically quite large. Even though it was significantly dampened by a

business cycle peak, the dispersion among metro unemployment rates in 2007 was substantial. Sixteen of 308 metros had unemployment rates more than 1.5 percentage points *below* the national rate; 31 had unemployment rates between 1.5 percentage points and 3.5 percentage points *above* the national rate; and 11 had unemployment rates more than 3.5 percentage points above the national rate (Chart 1).

The dispersion of metro unemployment rates remained high from at least 1990 through 2011. During this period, the difference between the 95th-percentile and 5th-percentile metro unemployment rates was always at least 4.5 percentage points (Chart 2). In some years it was more than twice as high as this.

The persistence of relative metro unemployment rates over time

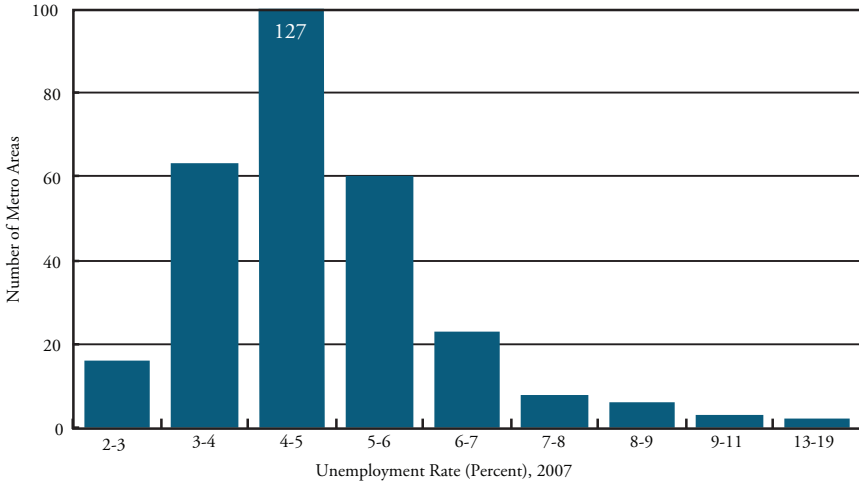
In addition to being highly dispersed, metro-area unemployment rates are highly persistent over time. A metro that had relatively high unemployment in an initial year was likely to have high unemployment many years later.

The high persistence of long-run unemployment is immediately evident from a scatter of metro unemployment rates averaged over the years 2000 to 2007 against metro unemployment rates averaged over the years 1990 to 1999 (Chart 3). The metro areas in the bottom left of the chart had *low* unemployment rates in both time periods. They are disproportionately located in upper Midwest states such as the Dakotas, Wisconsin, Minnesota, Iowa, and Nebraska. The metro areas in the top right of the chart had *high* unemployment rates in both time periods. They are disproportionately located in central California and the border region of Texas.

The upward-sloping line in Chart 3 represents the predicted value of 2000-to-2007 average metro unemployment rates based on 1990-to-1999 average metro unemployment rates. The predicted values are based on a simple regression, which determines the best linear fit of the later-period rates by the earlier-period rates. As is evident in the chart, the actual and predicted 2000-to-2007 unemployment rates are relatively close. The variation in 1990-to-1999 unemployment can account for 76 percent of the variation in 2000-to-2007 unemployment as measured by the regression R-square statistic.²

Chart 1

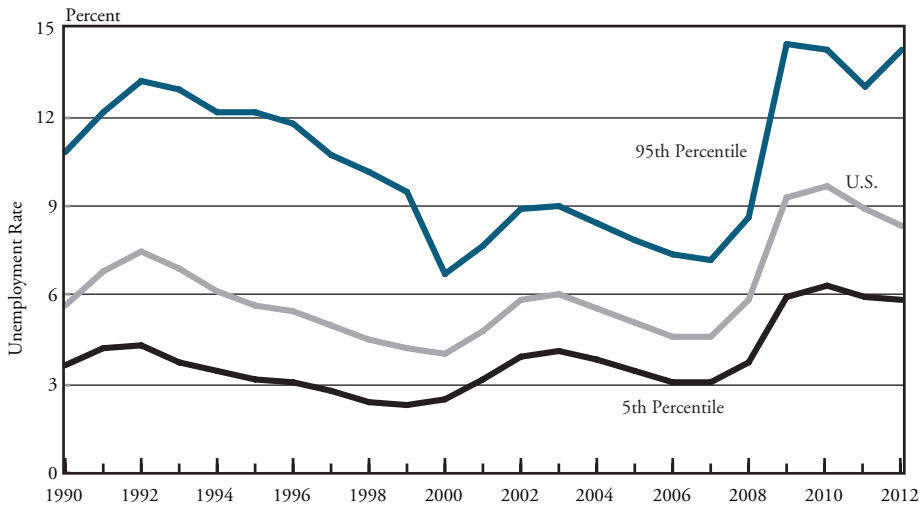
UNEMPLOYMENT RATE FREQUENCIES ACROSS U.S. METRO AREAS IN 2007



Sources: Bureau of Labor Statistics and author's calculations

Chart 2

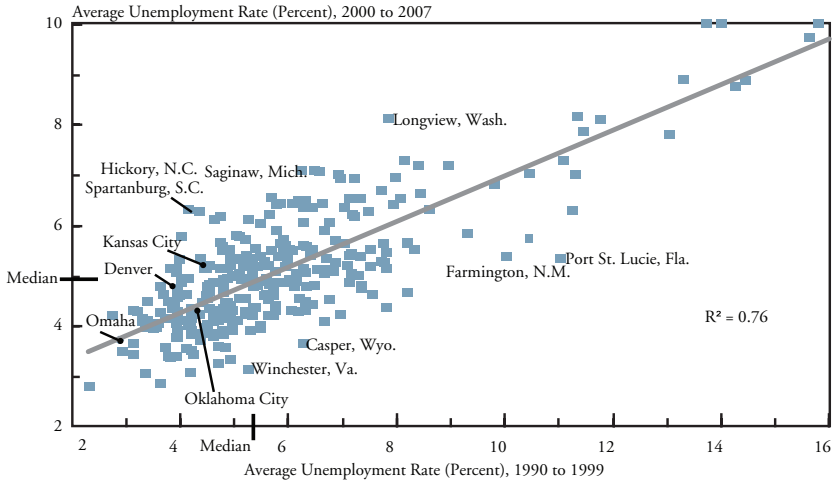
5TH AND 95TH PERCENTILES OF METRO UNEMPLOYMENT OVER TIME



Sources: Bureau of Labor Statistics and author's calculations

Chart 3

METRO AVERAGE UNEMPLOYMENT 2000 TO 2007 VERSUS METRO AVERAGE UNEMPLOYMENT 1990 TO 1999



Note: Four metros have unemployment rates outside the chart boundaries and are not shown.
Source: Bureau of Labor Statistics and author's calculations

Metro areas *below* the regression line had actual 2000-to-2007 unemployment lower than predicted. For example, Port St. Lucie, Fla.; Farmington, N.M. Casper, Wyo.; and Winchester, Va.; each had a 2000-to-2007 average unemployment rate that was at least 1.6 percentage points below its predicted value. Metro areas *above* the regression line had 2000-to-2007 unemployment higher than predicted. For example, Saginaw, Mich.; Spartanburg, S.C.; Hickory, N.C.; and Longview, Wash.; each had a 2000-to-2007 average unemployment rate that was at least 1.8 percentage points above its predicted value.³

One intuitive explanation for the persistent large differences in unemployment across metros is that low unemployment rates are caused by persistently fast employment growth and high unemployment rates are caused by persistently slow employment growth. However, the next section documents that metro unemployment rates are largely uncorrelated with metro employment growth. As a consequence, metro characteristics are likely to affect long-term metro unemployment for reasons other than their effect on employment growth.

II. METRO UNEMPLOYMENT AND JOB GROWTH

For the United States as a whole, unemployment and job growth are strongly negatively correlated. As national job growth accelerates national unemployment declines. Intuitively, a similar negative correlation would seem likely to characterize metro unemployment and job growth. What this intuition misses is the large flow of workers across metropolitan areas.⁴ As a result of worker migration, along with changes in labor force participation, long-term metro unemployment rates are largely uncorrelated with long-term metro job growth.

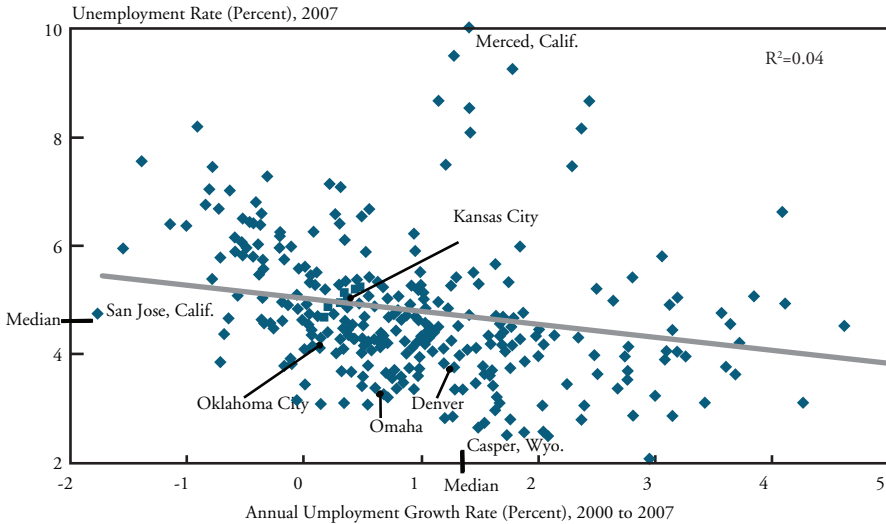
The weak correlation between metro unemployment and metro job growth

The ability of job growth over some multiyear interval to predict the unemployment rate in the final year of the interval has historically been close to zero. For example, metro employment growth from 1990 to 2000 can account for less than 1 percent of the variation in metro unemployment rates in 2000. Similarly, metro employment growth from 2000 to 2007 can account for only 4 percent of the variation in metro unemployment rates in 2007.⁵

The near-zero correlation of metro unemployment and metro employment growth is clearly visible in a scatter diagram of the two variables (Chart 4). For any specific 2000-to-2007 employment growth rate (horizontal axis), the range of 2007 unemployment rates is quite wide. For example, among metros that had employment growth between 1 percent and 2 percent, ending-period unemployment ranged from 2.5 percent (Casper, Wyo.) to 10.0 percent (Merced, Calif.). Similarly, for any specific ending-period unemployment rate, the range of average annual employment growth rates is quite wide. For example, San Jose, Calif., and Bend, Ore., both had 2007 unemployment rates just under 5 percent. But their 2000-to-2007 employment growth rates differed by almost 6 percentage points.

An important exception to the lack of correlation between the unemployment rate and employment growth concerns metro areas whose employment declined from 2000 to 2007. Over this period (but not for 1990 to 2000), metros that experienced declining employment were characterized by a negative, statistically significant correlation between employment growth and unemployment.⁶ For each average

Chart 4

METRO UNEMPLOYMENT IN 2007 VERSUS
EMPLOYMENT GROWTH 2000 TO 2007

Note: Three metros have unemployment rates outside the chart boundaries and are not shown.

Source: Bureau of Labor Statistics and author's calculations

annual percentage point by which employment declined from 2000 to 2007, predicted 2007 unemployment rose by approximately 1 percentage point. Variations in employment growth rates for these metros with declining employment account for 36 percent of the variation in ending-period unemployment rates.⁷

Decomposing metro job growth

The low correlation between unemployment and employment growth can be explained in part by decomposing each metro's change in employment into three components. The first component is the change in the labor force due to net migration. The second component is the change in the labor force due to the change in adult residents' labor force participation. The third component is the change in the number of unemployed workers. As a simple accounting identity, any net increase in metro employment must equal the net inflow of workers plus the net entry into the labor force by existing residents plus the net decrease in the number of unemployed workers. For present purposes, it is helpful to normalize each of the components to represent the number of persons per 100-person increase in metro employment.⁸

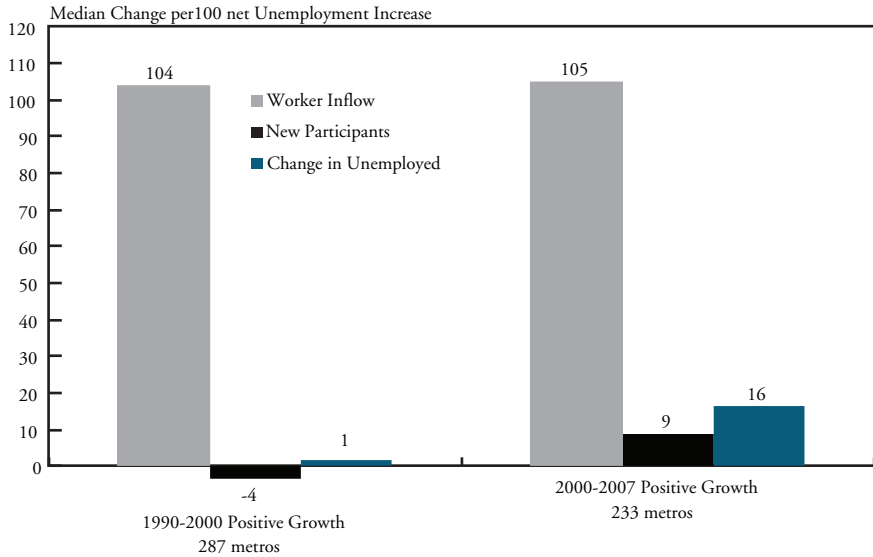
On average, increases in metro employment were overwhelmingly accounted for by worker in-migration. Among the 287 metros where employment increased from 1990 to 2000, the median net inflow of workers per 100 jobs added was 104 (Chart 5).⁹ Among the 233 metros where employment increased from 2000 to 2007, the median net inflow of workers per 100 jobs added was 105. Thus for both of these time periods, the worker inflow exceeded the employment gain. Essentially, nearly all new jobs were filled by workers migrating into the metro area rather than by workers who already lived there. For 1990 to 2000, the excess of the inflow of workers over the number of new jobs was largely offset by a modest decrease in labor force participation. But for 2000 to 2007, labor force participation actually increased moderately as well. Hence for this later time period unemployment rose. Specifically, a median 16-person increase in the number of unemployed workers accompanied each 100 jobs created.¹⁰

In sharp contrast, labor flows played a much smaller role in adjusting to decreases in metro employment. Equivalent to the decomposition above, any decrease in metro employment must equal the net number of workers who depart the metro plus the number of remaining residents who exit the labor force plus the net increase in unemployed workers. Among metro areas that experienced net employment decreases from 1990 to 2000, the median labor outflow accounted for only 26 workers per 100 jobs lost (Chart 6). In other words, most workers who lost their jobs remained in the metro area. Rather than a worker outflow, metros' adjustment to job losses primarily took the form of a fall in labor force participation. The median decrease in labor force participation accounted for 67 workers per 100 jobs lost. As a result, the median implied increase in unemployed workers amounted to just two workers per 100 jobs lost.

The labor outflow response to net job losses was almost completely absent from 2000 to 2007. During this period, net decreases in employment were matched about equally by large decreases in labor force participation and by a large increase in the number of unemployed workers.

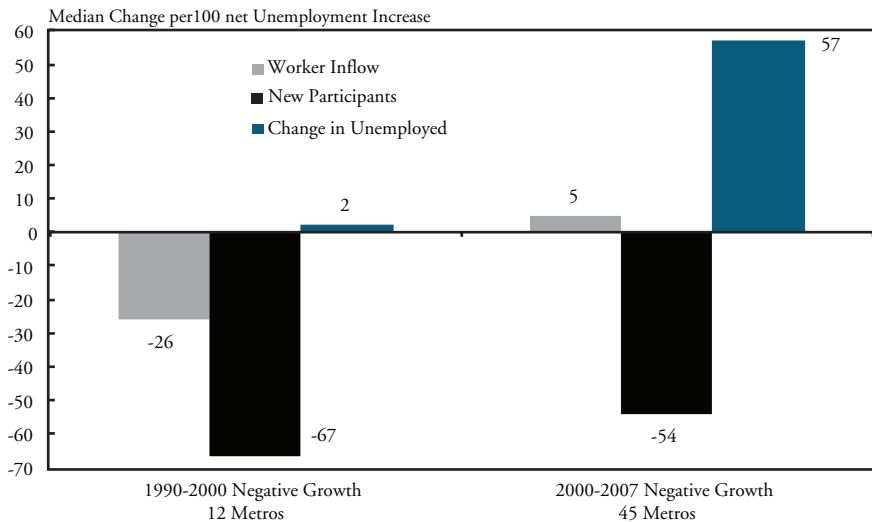
The decompositions of metro employment growth into the median migration, participation, and unemployment components together describe a sobering dynamic. Metro areas that experienced employment increases typically saw no decrease in unemployment due to the inflow of workers from elsewhere. But metro areas that experienced

Chart 5
DECOMPOSITION OF NET POSITIVE EMPLOYMENT CHANGES



Note: Numbers may not sum to 100 due to rounding because these are median values.
 Sources: Census Bureau and author's calculations

Chart 6
DECOMPOSITION OF NET NEGATIVE EMPLOYMENT CHANGES



Note: Numbers may not sum to 100 because these are median values.
 Sources: Census Bureau and author's calculations

employment declines, at least from 2000 to 2007, experienced a significant increase in their unemployment. A key reason is that laid-off workers did not move to other metro areas in search of better job prospects. Hence metros with rising unemployment and declining employment over an extended period may not see unemployment fall if and when employment starts increasing again.

A description of differences in employment growth, labor force participation, and other aggregate metro outcomes for low-, medium-, and high-unemployment metros is included as Appendix Table A. A description of differences in workforce and intrinsic characteristics across metro areas is included as Appendix Table B.

III. WORKFORCE SKILLS

A possible explanation for large and persistent differences in metro unemployment is that they reflect differences in the skills of metro workers. This hypothesis can be tested, at least in part, by calculating the extent to which workforce characteristics can predict long-run metro unemployment rates. If workforce characteristics cannot predict unemployment well, they are unlikely to be an important determinant of metro unemployment. If they can predict unemployment well, which is indeed the case, one interpretation is that differences in workforce characteristics across metro areas are *causing* differences in unemployment rates across metro areas. Alternatively, differences in workforce characteristics across metro areas may reflect differences in non-workforce characteristics that are the true cause of unemployment differences. A problem with this alternative interpretation is that it fails to explain why workers do not move from high-unemployment metros to low-unemployment metros. Under either interpretation, the tight correlation between metro unemployment and workforce characteristics implies that any successful strategy to bring down a metro's unemployment rate will almost certainly involve an upgrading of its workers' skills.

Table 1 reports results from regressions of average annual metro unemployment rates from 1990 to 2000 and from 2000 to 2007 on five sets of workforce characteristics. Each set of characteristics captures different skills that directly or indirectly may make a worker a good match for some jobs and a poor match for others. The first set is the share of metro employment in each of 16 industries (in 1990) or 19 industries (in 2000). The second set is the share of metro employment in each

Table 1

REGRESSIONS OF MULTIYEAR AVERAGE UNEMPLOYMENT ON WORKFORCE CHARACTERISTICS

Characteristic		(1) Average Unemployment Rate, 1990-2000	(2) Average Unemployment Rate, 2000-2007
Industries	Worker Share Variables	16	19
	R-squared alone	0.52	0.47
	Marginal R-squared	0.03	0.03
Occupations	Worker Share Variables	11	13
	R-squared alone	0.43	0.56
	Marginal R-squared	0.02	0.08
Education	Attainment Level Variables	5	7
	R-squared alone	0.38	0.40
	Marginal R-squared	0.01	0.02
Age	Age Bracket Variables	8	8
	R-squared alone	0.28	0.24
	Marginal R-squared	0.02	0.04
"English Language"	Proficiency Level Variables	2	2
	R-squared alone	0.55	0.30
	Marginal R-squared	0.06	0.02
"FULL REGRESSION"	Combined Variables	42	49
	R-squared	0.78	0.70
	Metros	316	308
"FULL REGRESSION" Split Samples	R-squared, init U \leq median	0.55	0.57
	R-squared, init U \geq median	0.84	0.75

Sources: Bureau of Labor Statistics, Census Bureau and author's calculations

of 11 occupations or 13 occupations. The third set is the share of the adult population with highest educational attainment at one of five levels or seven levels. The fourth set is the share of the population in each of eight age groups ranging from 18 to 24 through 85 and older. The last set is made up of two measures of English language proficiency.¹²

The combined sets of workforce characteristics are able to closely predict metro unemployment rates. The 42 workforce characteristics measured in 1990 account for 78 percent of the variation in average 1990-to-2000 unemployment as measured by the R-square statistic. The 48 workforce characteristics measured in 2000 account for 70 percent of the variation in average 2000-to-2007 unemployment. These R-squares are high, even after taking account of the large number of workforce characteristics used to explain unemployment.

The tightness of the fit between annual unemployment and the workforce characteristics is readily apparent in a scatter plot of actual average 2000-to-2007 unemployment rates against average 2000-to-2007 unemployment rates as predicted by the 48 workforce characteristics in 2000 (Chart 7). Of the 308 metros in the underlying regression, more than 80 percent have actual unemployment rates within 1 percentage point of their predicted value. Only seven metros have actual unemployment rates more than 2 percentage points away from their predicted value.¹³ Additional workforce characteristics not included in the underlying regression may be able to improve the fit of these outlier metros as well as the remaining metros with smaller differences between actual and predicted values.

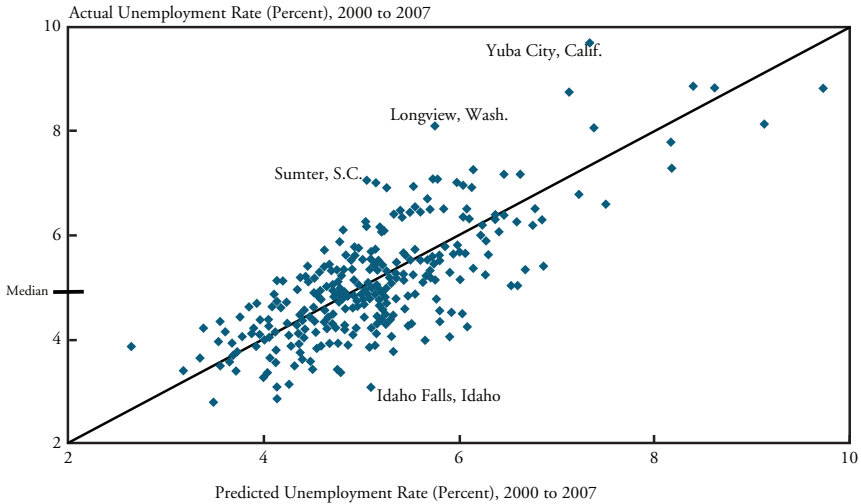
A second measure of the workforce characteristics' ability to predict metro unemployment is the ability of each of the five sets of characteristics to predict unemployment on its own. The predictive ability is measured by the R-squares from separate regressions of long-run unemployment on each set of workforce characteristics. For average unemployment from 1990 to 2000, these R-squares range from 0.28 (age) to 0.52 (industry). For average unemployment from 2000 to 2007, they range from 0.24 (age) to 0.56 (occupation). These moderate-to-tight fits establish that each of the workforce characteristics groups on its own could feasibly cause a large share of the variation in long-run unemployment.

A third measure of the workforce characteristics' ability to predict metro unemployment is the improvement in predictive power that comes from adding each of the characteristic sets to the other four combined. This measure is simply the R-square from a regression of unemployment on all five sets of workforce characteristics minus the R-square from a similar regression that excludes one set of workforce characteristics. The resulting marginal R-square values range from 0.01 to 0.03 for the 1990-to-2000 unemployment regressions and from 0.02 to 0.08 for the 2000-to-2007 unemployment regressions (Table 1, third row of each grouping). They establish that each characteristic set is able to fit a portion of the variation in unemployment that none of the other characteristic sets can. However, most of the variation in unemployment can be fit by at least two of the workforce characteristics sets.¹⁴

The fit of unemployment by the workforce characteristics is tighter for metros with high unemployment rates (bottom of Table 1). For

Chart 7

ACTUAL 2000-TO-2007 UNEMPLOYMENT VERSUS 2000-TO-2007 UNEMPLOYMENT PREDICTED FROM A REGRESSION ON WORKFORCE CHARACTERISTICS (49 VARIABLES IN TOTAL)



Note: Seven metros have unemployment rates outside the chart boundaries and so are not shown.
Sources: Bureau of Labor Statistics, and author's calculations

metros with 1990 unemployment above its median value, the workforce characteristics can account for 84 percent of the variation in average 1990-to-2000 unemployment. For metros with 2000 unemployment above its median value, the workforce characteristics can account for 75 percent of the variation in average 2000-to-2007 unemployment. These values are respectively 29 percentage points and 18 percentage points above the comparable R-squares for metros with 1990 unemployment below its median value. One interpretation of this asymmetry is that high metro unemployment follows from a fairly specific combination of characteristic values. Low metro unemployment, in contrast, may be consistent with a broader array of characteristic values.¹⁵

An important concern tempering interpretations of the correlation between long-run unemployment and workforce characteristics is that most of the characteristics are endogenous. A high share of manufacturing workers in the Midwest and of agriculture workers in central California is hardly coincidental. In part, these concentrations are the result of the industrial development of the United States over the 19th

and 20th centuries. Hence, the correlation of metro unemployment with workforce characteristics may actually be picking up a causal link from characteristics of metros themselves, rather than of their workers, to unemployment. For example, what matters may be not the share of manufacturing *workers* in a metro but rather the share of manufacturing *jobs*. The two differ because workers can decide on their own whether to move elsewhere in search of employment. Relocating manufacturing jobs, in contrast, is a decision of firms.¹⁶

Such concern about endogeneity is clearly valid. Even so, the evidence supporting the workforce characteristic hypothesis remains strong. The reason is that non-workforce metro characteristics generally cannot account for why unemployed workers choose to remain in high-unemployment metros. For the workforce variables to *not* play a significant causal role requires that substantial moving costs keep workers from moving from high-unemployment metros. As will be discussed later, the evidence for this is mixed.

Consistent with the view that workforce characteristics are likely to be a main determinant of differing metro unemployment rates, national unemployment rates vary considerably across each of the worker-characteristic categories. For example, 2000-to-2007 average national unemployment by educational attainment ranged from 2.4 percent for workers with a bachelor's degree or higher to 7.6 percent for workers who lacked a high school diploma. Similarly, 2000-to-2007 average national unemployment by industry ranged from 2.4 percent for public-sector workers to 9.0 percent for agriculture workers. A worker's characteristics are generally not changed by moving. Hence, workers with characteristics that have high national unemployment rates may not be able to significantly improve their work prospects by moving. (Appendix Table C enumerates unemployment rates for selected workforce characteristics.)

IV. INTRINSIC METRO CHARACTERISTICS

As suggested in the previous section, long-run metro unemployment rates may follow from characteristics intrinsic to the metro area itself rather than from the characteristics of its workers. But for intrinsic characteristics to drive metro unemployment, they must be accompanied by some sort of obstacle to job mobility that prevents workers from easily moving between metros and job types

Intrinsic metro characteristics include a large range of attributes and public policies that distinguish one metro area from another. Such characteristics might include exogenous attributes such as weather, coastal proximity, and adjacency to mineral and energy deposits. And they might include endogenous characteristics valued by households or firms such as low taxes, high-quality government services, wide-ranging entertainment and cultural opportunities, and easy access to land and air transportation.

For assessing the contribution of intrinsic characteristics on long-run unemployment, the analysis focuses solely on exogenous characteristics. One reason is that exogenous characteristics, by definition, cannot be caused by unemployment or by any variable excluded from a regression. Another reason is that measuring endogenous characteristics typically requires considerable subjectivity (Rappaport).

Table 2 reports results from regressing average metro unemployment rates from 1990 to 2000 and from 2000 to 2007 on three sets of intrinsic metro-area characteristics that households and firms are likely to consider in choosing where to locate. The first set of characteristics describes metro weather including summer and winter temperature, annual days below 32 degrees Fahrenheit and above 90 degrees Fahrenheit, annual rainfall, and annual rainy days. The second set is made up of characteristics that describe metros' proximity to ocean and Great Lakes coasts, major and navigable rivers, ocean seaports, and the length of any coastline. Several of the ocean characteristics are entered separately depending on whether the characteristic is along the North Atlantic, South Atlantic, Gulf, or Pacific coasts. The third set of variables measures metro hilliness based on how much altitude varies within each metro. Many of these characteristics are included in the regression both linearly and quadratically (squared).

Regressing long-run unemployment on the combined three sets of intrinsic metro characteristics yields moderately tight fits. The intrinsic characteristics can account for 55 percent of the variation of 1990-to-2000 unemployment and 53 percent of the variation of 2000-to-2007 unemployment. While high, this explanatory power is nevertheless about 20 percentage points below the explanatory power of the work-force characteristics.

Table 2

REGRESSIONS OF MULTIYEAR AVERAGE UNEMPLOYMENT ON METRO INTRINSIC CHARACTERISTICS

Characteristic		(1) Avg Unmpl Rate 1990-2000	(2) Avg Unmpl Rate 2000-2007
Weather	Variables	14	14
	R-squared alone	0.40	0.36
	Marginal R-squared	0.32	0.31
Coastal Proximity	Measures	15	15
	R-squared alone	0.08	0.07
	Marginal R-squared	0.06	0.06
Hilliness	Measures	4	4
	R-squared alone	0.16	0.16
	Marginal R-squared	0.07	0.08
FULL REGRESSION	Characteristics	33	33
	R-squared	0.55	0.53
	Metros	316	308
FULL REGRESSION, Split Samples	Rsq, init U < mdn*	0.42	0.40
	Rsq, init U > mdn*	0.76	0.70

Sources: Bureau of Labor Statistics, The Climate Source Inc. and author's calculations

Among the three sets of intrinsic characteristics, the weather characteristics have the highest explanatory power. Alone, they account for 40 percent of 1990-to-2000 unemployment and 36 percent of 2000-to-2007 unemployment. Added to the coastal proximity and hilliness characteristics, the weather characteristics increase explanatory power by more than 30 percentage points in each of the two time periods.¹⁷

The coastal proximity and hilliness characteristics also help explain variations in metro unemployment. Explanatory power is about 8 percent for the coastal proximity characteristics on their own and 16 percent for the hilliness characteristics on their own. Adding the coastal variables to the hilliness and weather variables boosts explanatory power by 6 percentage points. Adding the hilliness variables to the coastal proximity and weather variables boosts explanatory power by about 8 percentage points.

As with the workforce characteristics, the intrinsic metro characteristics account for a higher share of the variation in unemployment among high-unemployment metros than they do among

low-unemployment metros (bottom of Table 2). This asymmetry again suggests that high long-run metro unemployment may follow from a relatively specific set of intrinsic metro characteristics. In contrast, low long-run metro unemployment may be associated with a broader array of intrinsic metro characteristics.

A comparison of the workforce and intrinsic-characteristics regressions along with combined regressions that include both types of characteristics suggests that long-run unemployment is more closely tied to the workforce characteristics (Table 3). One reason is that workforce characteristics account for considerably more of the variation in unemployment than the intrinsic metro characteristics. In particular, for average 1990-to-2000 and 2000-to-2007 unemployment, they account for 22 percentage points and 19 percentage points more of the variation. A second reason is that a considerable share of the variation in unemployment can be accounted for by the workforce characteristics but not the intrinsic characteristics (marginal R-squares for the two time periods are 0.29 and 0.27). Only a more modest share of the variation in unemployment can be accounted for by the intrinsic characteristics but not the workforce characteristics (marginal R-squares of 0.07 and 0.09). Together, the high separate and marginal R-squares of the workforce characteristics make it unlikely that they are only picking up causal relationships from intrinsic variables excluded from the regression.

An important concern when comparing the workforce and intrinsic characteristics is the likelihood that important characteristics of each type are excluded from the regressions. Excluding a variable from a regression, either workforce or intrinsic, that is an important determinant of metro unemployment can bias estimates. An additional concern with the intrinsic characteristics is that the inclusion of additional attributes might allow these characteristics to match the explanatory power of the workforce characteristics.

Many endogenous metro characteristics also doubtlessly contribute to firm and household location decisions and therefore help determine metro unemployment. Some of the many likely candidates include government tax rates, services, and regulations; civic amenities, restaurants, and sports teams; and traffic, pollution, and crime. Interpreting correlations between unemployment and these endogenous characteristics is difficult. Such correlations may be caused by unemployment. Or they may be caused by some characteristics excluded from the analysis.

Table 3

REGRESSIONS OF MULTIYEAR AVERAGE UNEMPLOYMENT
ON BOTH WORKFORCE AND INTRINSIC
CHARACTERISTICS

Characteristic Type		(1)	(2)
		Avg Unmpl Rate 2000-2007	Avg Unmpl Rate 2000-2007
Workforce	variables	42	49
	R-squared alone	0.78	0.70
	Marginal R-squared	0.29	0.27
Fixed	variables	33	33
	R-squared alone	0.55	0.53
	Marginal R-squared	0.07	0.09
FULL REGRESSION	Combined variables	75	82
	R-squared	0.84	0.79
	Metros	316	308
FULL REGRESSION, Split Samples	R-squared, init U \leq median	0.77	0.75
	R-squared, init U $>$ median	0.93	0.88

Sources: Bureau of Labor Statistics, The Climate Source Inc. and author's calculations

V. MOVING COSTS

The costs associated with moving to a new metro area can clearly be substantial for some households and therefore dampen labor mobility. Hence it is likely that moving costs contribute to the persistence of unemployment differences across metros.

The typical cost of moving is unclear. Renting a large truck for a do-it-yourself move between metro areas can cost less than \$1,000. For a professional move between metros of a four-person household, the cost is probably closer to \$12,000 (Worldwide ERC). Adding in the costs of home-finding trips, temporary living, travel and lodging during the move, and miscellaneous other expenses can easily double this. Homeowners face additional expenses including preparing their existing house for sale, the selling commissions on it, and mortgage origination costs if they purchase a new house. For the median-priced house in 2000, these homeowner-specific moving costs would be about \$16,000 (2011 dollars). Thus a homeowner may face total monetary moving costs of \$40,000 or more. For comparison, median household income in 2010 was \$50,000.

A broader interpretation of moving costs suggests that they may be several times higher. Moving to a new metro area may require leaving behind family, friends, professional networks, schools, doctors, and more. Putting a price on these separations is difficult. Of particular importance is that such costs are likely to vary considerably across households. However, a theoretical estimate of average moving costs can be calculated as the cost required to limit migration across metro areas to the observed rate.¹⁸ Different applications of this methodology produce estimates of moving costs ranging from negligible to moderate (\$20,000) to extremely high (\$300,000) (Gallin; Bayer and Juessen; Kennan and Walker).

One hypothesis that can be directly tested is whether the additional moving costs associated with homeownership cause metros that have higher homeownership rates to have higher unemployment rates. The empirical evidence strongly rejects this hypothesis. For both the 1990-to-2000 and 2000-to-2007 periods, metro average unemployment was typically lower where starting-year homeownership was higher rather than the reverse (Table 4, columns 1 and 2). For both of these multiyear periods, the share of unemployment variation accounted for by homeownership rates was less than 4 percent. Similarly for the years 2007 to 2011, which span the recent collapse in housing prices, metro unemployment was essentially uncorrelated with homeownership rates in 2007 (Table 4, column 3). Consistent with these results, researchers have also documented that recent migration patterns for homeowners and for renters have been fairly similar (Malloy, Smith, and Wozniak).¹⁹

More generally, observed high mobility suggests that moving costs are likely not to be a large friction to a significant share of American households. In the mid-2000s, prior to the most recent recession, approximately 3.3 percent of Americans migrated from one metro area to another each year. Also in the mid-2000s, almost one-third of Americans were living in a state different from the one in which they were born (Malloy, Smith, and Wozniak). The large worker inflows that accompanied metro employment gains discussed in Section II serve as further evidence of high household mobility.

High household mobility does not, however, rule out the likelihood that American households face a range of moving costs from low to high. For households with low moving cost, small improvements

Table 4

CORRELATIONS OF MULTIYEAR METRO
UNEMPLOYMENT AND HOMEOWNERSHIP

Dependent Variable	(1)	(2)	(3)
	Average U 1990-2000	Average U 2000-2007	Average U 2007-2011
Ownership Share (initial year)	<i>-0.097</i> <i>(0.045)</i>	<i>-0.052</i> <i>(0.029)</i>	-0.087 0.054
Metros	316	308	303
R-Squared	0.04	0.03	0.03

Note: Bold and italic type signify a coefficient that respectively differs from 0 at the 0.05 or 0.10 levels.
Sources: Bureau of Labor Statistics, Census Bureau, and author's calculations

in employment prospects may be sufficient to cause a move. But for households with high moving costs, even relatively large improvements in employment prospects may not suffice to induce a move.

Differing moving costs can account for the asymmetric response of worker flows and unemployment to employment changes. When metro employment declines, laid-off workers who face low costs may move elsewhere, dampening the rise in unemployment. But laid-off workers who face high costs may not move, thereby contributing to an increase in metro unemployment. Conversely, when metro employment increases, low-moving-cost workers from throughout the nation may flood into the metro, keeping unemployment from falling.²⁰

Moving costs should have a large impact on metro unemployment differences when workforce differences are small but intrinsic differences are large. In this case, low moving costs allow most workers to choose a metro area with advantageous intrinsic characteristics. As a result, unemployment rates will tend to equalize across metros. With high moving costs, in contrast, unemployment rates will not equalize.

On the other hand, when workforce differences are large and intrinsic differences are small, moving costs may have a relatively small impact on unemployment rates. A first reason is that workers with some skill sets may face low job opportunities regardless of where they locate. Hence they may choose not to move, even if they face very low moving costs. A second reason is that whatever a worker's skills, metros with similar intrinsic characteristics are likely to offer similar employment prospects. As a result, the benefit from moving from one metro to another is likely to be small as well. Only workers with very low

moving costs would likely choose to move. For both of these reasons, the number of households that choose to move will be small regardless of whether moving costs are moderate or high.

VI. SUMMARY AND CONCLUSIONS

Unemployment rates vary widely and persistently across U.S. metropolitan areas. Possible reasons for this variation fall into three categories. First, worker skill sets vary considerably across metro areas. Persistently high unemployment in some metropolitan areas reflects a poor match between its workers' skill sets and the hiring needs of firms throughout the country. Second, the intrinsic characteristics of metropolitan areas cause unemployment rates to differ. Some characteristics may cause workers to prefer to live where they face repeated spells of unemployment. Other characteristics may cause firms to locate in metros where they face stiff competition to hire and retain good workers. Third, moving costs may be high for many households and firms, preventing the migration of workers and jobs needed to equalize metro unemployment rates.

Evidence supports each of these hypotheses. Metro workforce and intrinsic characteristics measured in 1990 and 2000 are each able to account for a large share of the variation in average unemployment rates from 1990 to 2000 and from 2000 to 2007. The correlation with the workforce characteristics is stronger than the correlation with the intrinsic characteristics. A significant share of the variation in unemployment can be accounted for by the workforce characteristics but not the intrinsic ones. But the relative importance of the workforce variables is likely to be overstated because of their endogeneity. Estimates of moving costs suggest that they are probably high for some households. Hence moving costs will likely slow the dissipation of unemployment rate differences across metro areas, especially for metros that experience employment declines.

The public policy implications of these various explanations are wide-ranging. To the extent that workforce characteristics are the key determinant of the persistent large differences in metro unemployment rates, policies that help workers upgrade their skills are the obvious remedy. An important question concerns whether local governments have sufficient incentive to provide such training. Specifically, getting

local voters to support funding programs that help workers find a job in another metro is likely to be difficult.

To the extent that *exogenous* intrinsic metro characteristics help determine unemployment differences, many *endogenous* intrinsic metro characteristics are likely to do so as well. For endogenous characteristics that can feasibly be changed, public and private efforts to do so may be desirable. But, again, getting local voter support to fund changes that create jobs that will be filled by workers from other metros is likely to be difficult

Finally, if moving costs are the key culprit, the federal government may have a role in making moving more affordable. For example, the federal government might offer some sort of tax incentive to laid-off workers or workers collecting unemployment insurance if they move from their current metro area to accept work elsewhere. Doing so would obviously benefit unemployed workers. But it may also serve the national interest. Maximizing U.S. wealth creation depends in part on achieving the best match between workers and jobs. Enhancing the geographic mobility of workers could help to achieve this.

*Appendix Table A***SELECTED SUMMARY STATISTICS FOR LOW-, MIDDLE-, AND HIGH-UNEMPLOYMENT METROS**

Variable	Low U	Med U	High U	(High-Low)
Unemployment, avg. annual rate, 1990-2000	3.7	5.3	9.7	6.0
Unemployment, avg. annual rate, 2000-2007	3.8	4.9	7.4	3.7
Unemployment, annual rate, 2011	7.1	8.9	12.0	4.9
Population (1990)	434,900	684,200	306,500	-128,400
Population (2000)	504,600	894,400	345,700	-158,900
Population (2007)	558,400	998,700	369,000	-189,400
Population (2011)	585,700	996,400	377,400	-208,300
employment-to-population ratio (pct., 1990)	49.2	45.1	41.2	-8.0
employment-to-population ratio (pct., 2000)	48.4	46.1	41.9	-6.4
employment-to-population ratio (pct., 2007)	49.1	46.8	42.8	-6.3
Labor force participation rate (pct., 1990)	68.0	63.9	60.0	-7.9
Labor force participation rate (pct., 2000)	66.2	63.5	60.3	-5.8
Labor force participation rate (pct., 2007)	66.5	64.1	60.6	-5.8
Employment grwth, avg. annual rate, (pct., 1990 to 2000)	2.5	1.9	1.7	-0.8
Employment grwth, avg. annual rate, (pct., 2000 to 2007)	1.5	0.7	0.4	-1.0
Employment grwth, avg. annual rate, (pct., 2007 to 2010)	-1.7	-1.9	-2.1	-0.5
Labor force grwth, avg. annual rate, (pct., 1990 to 2000)	1.8	1.3	1.2	-0.5
Labor force grwth, avg. annual rate, (pct., 2000 to 2007)	1.6	1.0	0.8	-0.9
Labor force grwth, avg. annual rate, (pct., 2007 to 2011)	0.0	0.2	0.8	0.8
Household real income, median (1990)	51,300	48,700	45,200	-6,100
Household real income, median (2000)	55,400	52,900	48,500	-6,900
Rent, median (real, monthly, 1990)	690	690	690	0
Rent, median (real, monthly, 2000)	720	710	670	-50
Rent, median (real, monthly, 2007)	800	780	730	-70
Homeowner-estimated house real value, median (1990)	120,100	126,400	129,000	8,900
Homeowner-estimated house real value, median (2000)	144,800	143,800	132,500	-12,300
Homeowner-estimated house real value, median (2007)	205,100	200,200	195,300	-9,800
Homeownership rate (pct., 1990)	67.9	66.8	66.1	-1.8
Homeownership rate (pct., 2000)	69.7	68.6	68.5	-1.2
Homeownership rate (pct., 2007)	70.4	68.8	68.1	-2.3
growth home prices, avg annual nom rate (pct., 1990 to 2000)	4.2	3.6	3.2	-1.1
growth home prices, avg annual nom rate (pct., 2000 to 2007)	6.8	6.1	6.8	0.1
growth home prices, avg annual nom rate (pct., 2007 to 2011)	-2.8	-3.0	-5.5	-2.7

Notes: Low unemployment metros are those with unemployment rates in bottom 20 percent. Medium are metros with in the middle 60 percent. High are metros in the top 20 percent. For 1990 and 1990 statistics, unemployment is measured as the average from 1990 to 2000. For remaining statistics, it is measured as the average from 2000 to 2007. All monetary amounts are real (normalized to 2011).

Sources: Bureau of Labor Statistics, Federal Housing Finance Agency, Census Bureau and author's calculations.

*Appendix Table B***SELECTED CHARACTERISTICS IN 2000 FOR LOW-, MIDDLE-, AND HIGH-UNEMPLOYMENT METROS**

WORKFORCE CHARACTERISTICS	Low Unmplymnt	Medium Unmplymnt	High Unmplymnt	(High -Low)
Selected Industry Shares				
Manufacturing	13.1	15.8	16.4	3.2
Agriculture, forestry, fishing and hunting	1.6	1.3	3.6	2.0
Professional, scientific, and technical services	5.1	4.4	3.4	-1.6
Finance and insurance	4.9	4.3	3.2	-1.7
Selected Occupation Shares				
Production occupations	8.1	9.7	10.7	2.6
Farming, fishing, and forestry occupations	0.7	0.6	2.5	1.9
Transportation and material moving occupations	6.3	6.5	7.4	1.1
Sales and related occupations	11.8	11.5	10.9	-0.8
Computer and mathematical occupations	2.1	1.8	1.0	-1.1
Business and financial operations occupations	4.1	3.7	3.0	-1.1
Office and administrative support occupations	15.5	15.2	14.3	-1.2
Management occupations (non-farm)	8.1	7.7	6.7	-1.5
Educational Attainment				
Bachelors degree	33.2	29.1	22.6	-10.6
Masters degree or higher	17.7	15.9	11.6	-6.1
Age Characteristics				
Population, all, pct aged 0 to 17	25.5	25.2	27.3	1.8
Pop 18 and older, pct aged 35 to 44	21.9	21.4	21.3	-0.5
English Language Proficiency				
Do not speak English very well	4.2	5.0	9.7	5.5
Do not speak English well	2.2	2.6	5.9	3.7

Table B continued

INTRINSIC CHARACTERISTICS				
Weather				
Maximum daily January temperature (°F)	41.5	43.9	44.5	3.0
Annual Days Temperature Falls below 32°F	114.5	99.5	93.1	-21.4
Maximum daily July heat index (°F)	97.0	97.5	95.9	-1.2
Maximum daily July relative humidity	66.3	66.3	61.3	-5.0
Annual Days Temperature Rises above 90°F	39.2	41.4	46.5	7.3
Annual rainy days	94.4	99.6	94.6	0.1
Annual rainfall	37.9	40.3	35.5	-2.5
Coastal Proximity				
Within 80 km of an ocean coast	29.5	25.3	21.3	-8.2
Within 80 km of a medium or large ocean seaport (percent)	18.0	15.1	9.8	-8.2
Within 80 km of a Great Lakes' coast	1.6	12.4	16.4	14.8
Within 40 km of a navigable river	19.7	19.9	11.5	-8.2
Within 40 km of a major river	49.2	39.8	41.0	-8.2
Hilliness				
within metro altitude range per unit area	380	444	889	509
within metro altitude standard deviation per unit land area	72	84	176	103

Notes: Low, medium, and high unemployment metros are those with average 2000-to-2007 unemployment rates in bottom 20 percent, middle 60 percent, and top 20 percent. Weather values are based on 30 year averages, 1960 to 1990. Reported values are the mean over the metro areas in each of these groups.

Sources: Bureau of Labor Statistics, Census Bureau, The Climate Source Inc. and author's calculations

*Appendix Table C***NATIONAL UNEMPLOYMENT RATES FOR WORKERS
WITH SELECTED CHARACTERISTICS**

WORKFORCE CHARACTERISTICS	Unemployment Rate (2000-2007)
INDUSTRY:	
Agriculture	9.0
Leisure and hospitality	7.8
Construction	7.7
Professional and business services	6.4
Wholesale and retail trade	5.3
Manufacturing	5.1
Information	5.0
Other services	4.7
Mining, quarrying, and oil and gas extraction	4.4
Transportation and utilities	4.3
Education and health services	3.1
Financial activities	3.1
Government	2.4
OCCUPATION:	
Farming and fishing and forestry occupations	10.8
Construction and extraction occupations	7.8
Natural resources construction and maintenance occupations	6.7
Transportation and material moving occupations	6.7
Production occupations	6.5
Service occupations	6.2
Sales and related occupations	5.1
Office and administrative support occupations	4.5
Installation maintenance and repair occupations	3.9
Professional and related occupations	2.5
Management business and financial operations	2.3
EDUCATIONAL ATTAINMENT:	
Less than high school diploma	7.6
High school graduate no college	4.6
Some college or associate degree	3.8
Bachelors degree and higher	2.4

Table C Continued

AGE:	
Pop, pct aged 16 to 17	18.0
Pop, pct aged 18 to 19	14.3
Pop, pct aged 20 to 24	8.7
Pop, pct aged 25 to 34	5.0
Pop, pct aged 35 to 44	3.9
Pop, pct aged 45 to 54	3.4
Pop, pct aged 55 and higher	3.3

Sources: Bureau of Labor Statistics and author's calculations

ENDNOTES

¹According to the NBER Business Cycle Dating Committee, these peaks occurred in July 1990, March 2001, and December 2007.

²An R-square statistic measures the share of variation in one variable accounted for by variation in one or more other variables.

³The prediction of 2000-to-2007 unemployment by 1990-to-1999 unemployment is not the same as a forecast. A prediction is based on knowing the distribution of both earlier- and later-period unemployment rates. A forecast is based on knowing the distribution of only earlier-period unemployment. Nor does the accurate prediction of later-period unemployment imply that it was *caused* by earlier-period unemployment. While this might be the case, it may also be true that some other highly persistent factor is determining unemployment rates in both periods.

⁴“Workers”, as used herein, denotes labor market participants regardless of whether they are employed or unemployed.

⁵Based on the regression of the unemployment rate in 2007 on employment growth from 2000 to 2007, there is a small negative, but statistically significant, correlation between the unemployment rate and employment growth. Its magnitude implies that each percentage point by which annualized employment growth was above average was associated with an ending-period unemployment rate that was 0.2 percentage point below average. No such correlation was found in the regression of unemployment in 2000 on employment growth from 1990 to 2000.

⁶From 1990 to 2000, there are too few metros with which to estimate the correlation between the unemployment rate and employment *declines*. But, for the 60 metros with average annual employment growth of less than 1 percent, the unemployment rate and employment growth were almost completely uncorrelated.

⁷The negative correlation between unemployment in 2007 and employment growth from 2000 to 2007 is evident in Chart 4, beginning at the top-left in the vicinity of Flint, Mich., and sloping downward and to the right ending in the vicinity of Oklahoma City. An alternative specification finds a negative, statistically-significant correlation between employment growth and unemployment for metros with increasing employment from 2000 to 2007 (but not from 1990 to 2000). Among these metros, faster growth from 2000 to 2007 was associated with larger unemployment rate decreases over the same period. For the 225 metros for which employment increased over this period, employment growth can account for 17 percent of the variation in the change in the unemployment rate. However, the magnitude of the correlation is very small: faster employment growth of 1 percentage point annually over the seven years was associated with a decrease in unemployment of 0.06 percentage point. A decrease in unemployment by a significant amount was associated with implausibly fast employment growth

⁸Let E_{t1} and P_{t1} denote a metro's employment and adult population at time $t1$. Let l_{t2} and u_{t2} denote a metro's labor force participation rate and number of unemployed persons at $t2$. It is straightforward to show that $E_{t2} - E_{t1} = l_{t2} * (P_{t2} - P_{t1}) + P_{t1} * (l_{t2} - l_{t1}) - (u_{t2} - u_{t1})$.

⁹Metros included in the 1990-to-2000 and 2000-to-2007 decompositions were additionally required to have employment changes of at least 1 percent in absolute value. Retaining metros with very small absolute-value employment changes causes component values to blow up when normalizing.

¹⁰Component values are the medians across metros. As a result, they do not sum to 100. For any specific metro, they sum exactly to 100.

¹¹All variables are constructed from the 1990 and 2000 decennial censuses. Examples of industry classifications include construction, retail trade, and food services. Examples of occupation classifications include construction trades workers, sales occupations, and food preparation. Examples of education attainment levels include high school graduate, bachelor's degree, and professional degree (e.g., law, medicine). The language proficiency characteristics are the share of the working-age population that didn't speak English very well and the share that didn't speak it well.

¹²Metros with actual 2000-to-2007 unemployment more than 2 percentage points above predicted values were Sumter, S.C.; Longview, Wash.; Yuba City, Calif.; El Centro, Calif.; and Yuma, Ariz (El Centro and Yuma's unemployment rates place them outside the boundaries of Chart 7; hence they are not visible.) Metros with actual 2000-to-2007 unemployment more than 2 percentage points below predicted value were Salinas, Calif. (not visible in Chart 7), and Idaho Falls, Idaho.

¹³The actual causal contribution by a characteristic set, should one exist, may be considerably above its marginal R-square value. If one of the sets of characteristics was indeed determining unemployment, it could take credit for those portions of unemployment variation that can also be accounted for by a non-causal characteristic set. Separately, the various regressions discussed in the main text all yield estimates of the partial correlations of unemployment with the various workforce characteristics. But with so many variables in the regression, these correlations are difficult to interpret and so are not reported.

¹⁴The ability of the workforce characteristics to fit average unemployment is approximately the same for metros with fast employment growth and those with slow employment growth.

¹⁵In addition, manufacturing workers include those who are unemployed in addition to those with jobs. However, the Census Bureau industry classifications used herein are based on only employed workers.

¹⁶Among the weather variables, the average annual number of days that the temperature falls below 32 degrees Fahrenheit and the average annual number of days it rises above 90 degrees Fahrenheit have especially high explanatory power. Entered linearly and quadratically, the resulting four variables can account for 28

percent of the variation in 1990-to-2000 unemployment and 20 percent of the variation in 2000-to-2007 unemployment. Compared to the remaining intrinsic characteristics (including the other weather characteristics) these four variables improve the fits to unemployment by 9 percentage points and 14 percentage points. The associated magnitudes in the 2000-to-2007 combined intrinsic characteristic regression are large. The difference in expected unemployment from having the 80th percentile number of cold (32 degrees Fahrenheit or less) days rather than the 20th percentile number of cold days (149 cold days versus 47 cold days) is -4.6 percentage points. The difference in expected unemployment from having the 80th percentile number of hot (90 degrees Fahrenheit or more) days rather than the 20th percentile number of hot days (77 hot days versus 11 hot days) is 3.7 percentage points.

¹⁷After taking account of individual workers' characteristics and differences in housing costs between metros, any remaining difference in wages is assumed to be possible only because workers are not fully mobile. Comparing the lifetime value of the higher wages that could be achieved by moving, researchers can estimate the magnitude of the moving cost that would prevent workers from taking advantage of this gain.

¹⁸In contrast to the lack of correlation between unemployment and homeownership, changes in unemployment are modestly negatively correlated with changes in house prices from 2000 to 2007 and strongly negatively correlated with changes in house prices from 2007 to 2011. For the latter years, a 1 percent decline in house prices was associated with an approximate 0.1 percent increase in metro employment ($R^2 = 0.50$). This high explanatory power arises mostly from those metros where house price decreases were largest. The endogeneity of both of these contemporaneously-determined metro outcomes makes interpreting the correlation difficult. Partly addressing this endogeneity concern, regressing the 2007-to-2011 change in unemployment on the 2007-to-2011 change in house prices as predicted by the 2000-to-2007 change in house prices yields an equally strong negative correlation.

¹⁹It is possible that worker moves can equalize unemployment rates across metros even if a large share of workers faces high moving costs. What is required is that workers must consider not just whether they themselves are unemployed but also what are their chances of becoming unemployed in the intermediate future. In other words, if a low-moving-cost employed worker in a high-unemployment metro area thinks that there is relatively high chance that he will be laid off within a year or two, he may choose to take a job elsewhere. This would then allow for the employment of a worker who had previously been laid off.

REFERENCES

- Bayer, Christian, and Falko Juessen. 2008. "On the Dynamics of Interstate Migration: Migration Costs and Self-Selection." Institute for the Study of Labor Discussion Paper No. 3330, February.
- Gallin, Joshua H. 2004. "Net Migration and State Labor Market Dynamics." *Journal of Labor Economics* 22, pp. 1-23
- Kennan, John, and James R. Walker. 2011. "The Effect of Expected Income on Individual Migration Decisions." *Econometrica* 79, 1 (January), pp. 211-251.
- Malloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2011. "Internal Migration in the United States." *Journal of Economic Perspectives*, 25, 3 (Summer), pp. 173-196.
- Rappaport, Jordan. 2008. "Consumption Amenities and City Population Density." *Regional Science and Urban Economics*, 38, 6 (November), pp. 533-552.
- Worldwide ERC. 2011. *U.S. Transfer Volume & Cost Survey*. Downloaded April 9, 2012, from <http://www.worldwideerc.org/Resources/Research/Pages/reports-and-surveys.aspx>.

