

Some Observations on Hagedorn, Karahan, Manovskii, and Mitman, “Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects”

Robert E. Hall
Hoover Institution and Department of Economics
Stanford University
National Bureau of Economic Research
rehall@stanford.edu; stanford.edu/~rehall

November 6, 2013

This paper has attracted a huge amount of attention, much of it skeptical. I think it is an imaginative and potentially important contribution, but needs a lot of work to convince a fair-minded skeptic (like me). I’m not under the illusion that it could ever convince the diehard skeptics.

References here are to NBER working paper 19499, October 2013.

1 The Big Picture

Hagadorn and co-authors (HKMM) invite the interpretation that most of the large increase in unemployment since 2007 can be blamed on the rising generosity of unemployment insurance benefits, leaving little to blame on the financial crisis. But the paper does not make that claim. And the claim is incorrect. The right way to explain the claims of the the paper is that a policy of extending UI benefits when unemployment rises can amplify the effects of a small initial impulse that raises unemployment. That impulse, according to the paper, causes benefits to improve, which in turn results in a shift in bargaining power toward workers, lowering job-creation incentives and raising unemployment further.

Figure 1 shows the amplification in a graph. The steeper blue line describes the policy of extending the duration of benefits in response to high unemployment. The paper discusses and documents the combination of standing policies based on unemployment triggers, gener-

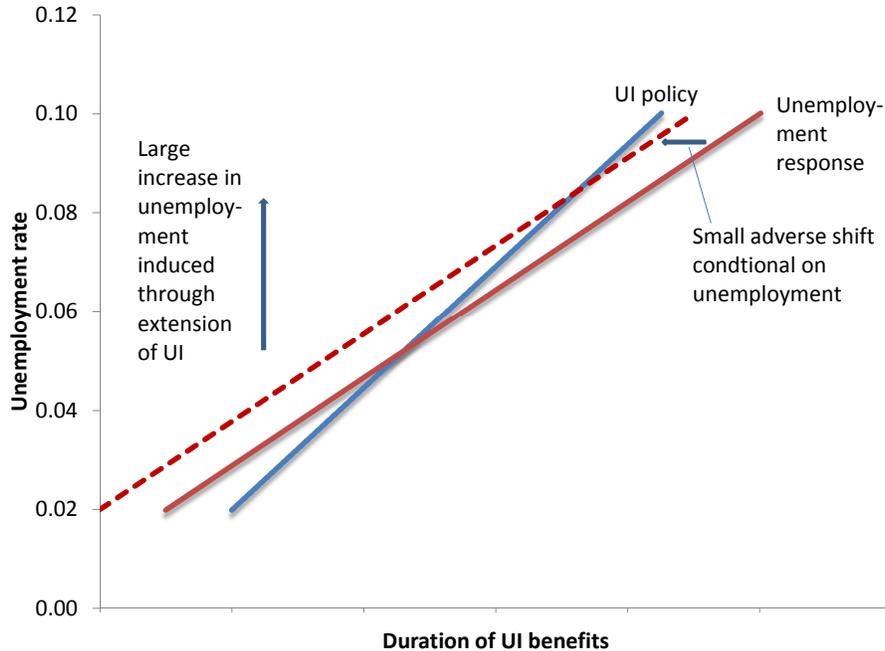


Figure 1: The Amplification Process According to HKMM

ally state by state, and further extensions adopted in the aftermath of the crisis. The flatter solid red line describes the mechanism that the paper focuses on—the higher equilibrium unemployment that occurs when UI benefits become more generous. Its position reflects pre-crisis conditions. The dashed red line shows a small adverse shift thanks to the crisis. Unemployment is higher at any given level of UI benefits because of the harm associated with the crisis. Equilibrium unemployment rises by much more than the shift in the red line, thanks to the positive feedback effect from the upward slopes of both lines.

Thus the right way to explain the paper is that the crisis caused a burst of unemployment, but unemployment grew mainly because of the policy response to the crisis. It would not be appropriate to say that the paper explains a large part of the rise in unemployment, leaving only a small part to be attributed to the crisis.

The paper does not pursue this point, though it can be read into the first paragraph of the introduction and into the conclusion.

The third paragraph of the paper introduces a crisp distinction between two sources of unemployment increase. The first—search intensity—operates at the personal level. A large body of existing research, most recently by Farber and Valletta, shows that UI benefit extensions have only small effects at the personal level, adjusting for labor-market conditions. Farber and Valletta study the diminution of job-seeking success and exit rates

from the labor force, in personal-level equations that control for local unemployment and employment growth. The second, labor-market tightness, has escaped the attention of prior empirical researchers, though it has surfaced occasionally in discussions of models in the Diamond-Mortensen-Pissarides tradition, which provides a coherent framework for thinking about tightness.

HKMM do not pursue the crisp distinction in this paper. Rather, they focus on measuring the entire effect on unemployment and then merely mention that because they find much larger effects associated with UI benefit extensions than have earlier authors such as Farber and Valletta, they must be finding large market-tightness effects.

2 How Big is the Effect of UI Extensions on Unemployment?

HKMM focus on a particular metric for the effects they measure of benefit duration on unemployment. It is the coefficient on log duration with the quasi-difference in log unemployment as the left-hand variable. Column (1) of Table 1 gives their basic result—the coefficient is 0.06. The discussion on page 14 shows that the capitalized effect over 16 quarters is 0.54, implying that the unemployment rate would rise from 5 percent to 8.6 percent on impact.

Table 5 uses the same metric for other labor-market indicators. One is vacancies, V , and the other is tightness, $\theta = V/U$. The effect for unemployment is the elasticity for vacancies, -0.0631 , less the elasticity for tightness, 0.1067 , which equals 0.0436 . The authors' explanation of the discrepancy, in a private communication, is that (1) the estimation for vacancies uses a smaller sample of county pairs than does the estimation for unemployment and (2) the specification for aggregate effects differs among the equations.

3 The Macro Effect

The discussion of Table 5 calculates the implications for unemployment of the measured effect of UI extensions on tightness alone, in a setup where UI has no effect on search effort (the parameter μ is taken as a constant; it corresponds to the search intensity s way back in equation (1)). The calculations also ignores any effect of UI on the unemployment entry rate (often mistakenly called the separation rate). Nonetheless, the calculated unemployment effect, with the matching elasticity taken to be the reasonable value of 0.5, is 0.05, right at

the average of the two direct estimates I just discussed. The paper treats this finding as a simple validation of the estimation and neglects the opportunity to observe that the small implied shift in μ lines up with earlier research such as Farber and Valletta.

It is possible to extend the calculation in the current version of the paper mentioned at the beginning of this section. This would require keeping the assumption that UI benefits have no effect on the inflow to unemployment. Then the value of γ that is consistent with no change in μ , that is, no effect on search effort, is 0.40 and the value of the elasticity of μ with respect to benefit duration based on $\gamma = 0.5$ is -0.0105 . These results confirm the paper's general conclusion and the results of earlier research that the personal search-effort effect of UI benefits is small.

4 Data

Many of the key results in the paper rest on the measurement of unemployment in the Local Area Unemployment Statistics database of the Bureau of Labor Statistics. A basic data source for LAUS is the number of workers drawing unemployment benefits as reflected in state records. Though there are some conceptual issues in translating claimants into unemployed people, at least the data are genuine. But the BLS's website warns as follows for data on unemployment of those not drawing benefits (www.bls.gov/lau/laumthd.htm):

The second category, "new entrants and reentrants into the labor force," cannot be estimated directly from UI statistics, because unemployment for these persons is not immediately preceded by the period of employment required to receive UI benefits. In addition, there is no uniform source of new entrants and reentrants data for States available at the LMA level; the only existing source available is from the CPS at the State level. Separate estimates for new entrants and for reentrants are derived from econometric models based on current and historical state entrants data from the CPS. These model estimates are then allocated to all Labor Market Areas (LMAs) based on the age population distribution of each LMA. For new entrants, the areas proportion of 16-19 years population group to the State total of 16-19 years old population is used, and for reentrants, the handbook areas proportion of 20 years and older population to the State total of 20 years and older population is used.

The authors have informed me that they hope to receive clarification of the BLS’s imputation procedure. They believe that the procedure does not rely on state-level data, but rather uses reliable county-level demographic data from the census of population together with reliable data on UI claims at the county level.

To investigate this issue, I downloaded the county pairs from the *Review of Economics and Statistics* replication data website (<http://thedata.harvard.edu/>) for the paper by Dube and Reich and the 2007 LAUS data from <http://www.bls.gov/lau/#cntyaa>. That year preceded any important changes in UI benefits, so the data should reveal something about the relative importance of the statewide average and the conditions in an adjacent county in another state in determining a border county’s unemployment. I estimated a regression equation with the unemployment rate of the second listed county as the left-hand variable and the average unemployment rate in that county’s state and the unemployment rate in the adjacent county (first listed county) as right-hand variables. The estimated coefficient on the state average is 0.908 with a standard error of 0.039 and the coefficient on the adjacent out-of-state county is 0.350 with a standard error of 0.023. The large role of the in-state average is consistent with the idea that the construction of the county estimates leans heavily on state aggregates. Even if the estimates do not have this character, the finding that the adjacent out-of-state county has little predictive power for its in-state counterpart points away from the basic identification strategy, which assumes that the adjacent county fully controls for the idiosyncratic unemployment disturbance in the county. The regression does not compel either of these conclusions, however.

HKMM investigated this issue by running a similar regression for employment—a variable free from suspicion related to imputation from state totals—and finding quite similar results, which points away from the concern about imputation but does not affect the concern about the relatively low predictive power of the adjacent county.

HKMM believe that state-level determinants of employment—mainly tax policy—were sufficiently important in 2007 to account for the superior predictive power of the state aggregate over the adjacent out-of-state county. I believe the topic deserves further investigation.

5 Identification

The central issue of identification is that any random upward shift in state-wide unemployment triggers more generous UI benefits. To the extent that the in-state county in a pair

shares in that upward shift while the out-of-state adjacent county does not, HKMM’s estimation strategy will overstate the causal role of UI benefits. The disturbance in equation (14), $\nu_{p,t}$, will be positively correlated with the right-hand variable, $\Delta b_{p,t}$. The logic of the zero-correlation assumption is that the left-hand variable, the difference between unemployment in the paired adjacent counties, removes the idiosyncratic element because the two counties share the same determinants of unemployment and thus have the same idiosyncratic elements.

HKMM’s discussion of identification, on pages 15 through 18, defends this zero-correlation assumption explicitly. Column (2) of Table 1 gives regression results testing the assumption by adding a variable to the estimating equation. The variable is the difference in log-productivity between the two states in the pair. The logic is that, if there is feedback from unemployment to UI duration, and if productivity affects unemployment, either through the positive feedback mechanism involving UI duration or in some other way, then the regression can distinguish between the two channels. If there is feedback, it will stop at the state border. If some other channel, it should affect the cross-border pairs equally. Thus with positive feedback through UI, the coefficient of a regression of the unemployment differences between county pairs on the difference in state productivity should have a negative coefficient, while if that channel is not an issue, the differencing of adjacent counties will result in a zero coefficient.

A leading example of the county pairs in the study comprises Cook County, Illinois and Lake County, Indiana. Cook County (Chicago) accounts for 39 percent of employment in Illinois, so any shock in Cook County is material for triggering UI benefits in the state. The remaining question is whether such a shock would affect Lake County by the same number of percentage points of unemployment.

One way to get at this issue is to study the distance between the population centers of the paired counties. To find the distances, I downloaded the coordinates of the centers from the website for the 2010 Census of Population and applied the standard formula from spherical trigonometry to calculate the great circle distance separating the pair. The population center of Lake County is 30 miles from the center of Cook County. This seems consistent with partial but not full offset to shocks originating in Chicago.

The pairs include a large number with long distances separating the population centers. The longest—279 miles—is between Washoe County, Nevada (Reno) and tiny Harney

County, Oregon. Clark County, Nevada (Las Vegas) accounts for two pairs, both distant: Mohave County, AZ (90 miles) and Inyo County CA (187 miles and the 12,000 foot White Mountain range). All told, 180 of the 1172 pairs exceed 50 miles and 593 are 30 or more miles apart.

Cook County is hardly the only county that accounts for a large fraction of its state's employment, as Table 1 shows. Footnote 17 of the paper notes the issue: "We can expect to see some impact on the estimate as there might be at least some correlation between the measured productivities of the county and of the state it belongs to since the number of counties in a state may be too small for the Law of Large Numbers to apply." Certainly the law of large numbers does not apply to the District of Columbia, which has only one county. In many counties, a county-level unemployment shock obviously feeds into the state unemployment rate and thus into its UI duration.

I think it would be a good idea to examine the effects on the results of removing the county pairs that are more than, say, 20 miles apart, and, separately, removing the pairs containing a county that is more than, say, 15 percent of its state's population. In both cases, the diminution in the number of observations would be moderate. The paper does contain estimates for a subset of the counties, in column 6 of Table 1. It would be useful to study further whether it is plausible that the adjacent counties in that subset control fully for county-level shocks to unemployment.

6 Recommendations

6.1 Positive feedback

Something like Figure 1 in this document would be a good way to introduce the topic and to sidestep the misunderstanding that the paper claims that the crisis was not the cause of high unemployment.

6.2 Framework

Equation (1) promises a nice framework for understanding this paper and its predecessors. But the current version of the paper does not pursue the distinction. Instead, it just finds a larger effect than the search effort literature has found and attributes the residual, in quite a casual way, to the GE effect. The paper contains results that bear on the distinction directly,

<i>County</i>	<i>Instances</i>	<i>Percent of state</i>
District of Columbia	4	100
Clark County, NV	3	71
New Castle County, DE	6	62
Providence County, RI	4	57
Cook County, IL	1	39
Hillsborough County, NH	3	31
Douglas County, NE	1	27
Fairfield County, CT	3	25
Chittenden County, VT	2	25
Hartford County, CT	1	24
Middlesex County, MA	1	24
Cass County, ND	2	24
Rockingham County, NH	1	23
Minnehaha County, SD	2	23
Sussex County, DE	4	21
Johnson County, KS	2	20
Multnomah County, OR	2	20
Jefferson County, KY	3	18
St. Louis County, MO	2	17
Montgomery County, MD	3	17
Kent County, RI	2	17
Kent County, DE	3	17
Washoe County, NV	7	17
York County, ME	2	16
Prince George's County, MD	2	15

Table 1: Counties in the Pairs that Account for at Least 15 Percent of Their State's Employment

so it would be easy to make equation (1) be a continuing organizing principle. The paper would be a lot more persuasive with this change.

6.3 Vacancies

The results in Table 5 are not integrated into the paper. As I point out, they imply a different effect on unemployment from the one in Table 1, which needs to be sorted out. Much more important, because the vacancy result does not depend on the potentially questionable unemployment data, it gives badly needed support to the general conclusion of the paper. Of course, that requires much more discussion of how vacancies are measured, a topic absent from the paper now.

6.4 Calculation of the implied GE and search effort effects

The discussion in the current version of Table 5 fails to note that it assumes no effect of UI benefits on the entry rate to unemployment, which is hard to swallow.

6.5 Data

I look forward to the next version of the paper, which I understand will explain the details of the LAUS data and eliminate the suspicion that state-level data have a role in constructing the county-level estimates of unemployment.

6.6 Identification test

The explanation of the identification test is quite opaque. I certainly did not follow it correctly when I first read it. Normally we don't think of fundamental identifying assumptions as susceptible to testing. Rather we offer arguments to support the assumptions. The fundamental assumption in this paper is that a county across the state border but sharing that border with a given county fully controls for the shock to unemployment in that country. Partial control is not enough. My study of the actual county pairs raises some questions about the inclusion of some of them. I think that, at a minimum, further investigation of the effects of exclusion of some of the questionable pairs would be appropriate.