

NBER WORKING PAPER SERIES

INTERPRETING RECENT QUASI-EXPERIMENTAL EVIDENCE ON THE EFFECTS
OF UNEMPLOYMENT BENEFIT EXTENSIONS

Marcus Hagedorn
Iouri Manovskii
Kurt Mitman

Working Paper 22280
<http://www.nber.org/papers/w22280>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2016

Support from the National Science Foundation Grant No. SES-1357903 is gratefully acknowledged. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Marcus Hagedorn, Iouri Manovskii, and Kurt Mitman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Interpreting Recent Quasi-Experimental Evidence on the Effects of Unemployment Benefit Extensions

Marcus Hagedorn, Iourii Manovskii, and Kurt Mitman

NBER Working Paper No. 22280

May 2016

JEL No. E24,J64,J65

ABSTRACT

We critically review recent methodological and empirical contributions aiming to provide a comprehensive assessment of the effects of unemployment benefit extensions on the labor market and attempt to reconcile their apparently disparate findings. We describe two key challenges facing these studies - the endogeneity of benefit durations to labor market conditions and isolating true effects of actual policies from agents' responses to expectations of future policy changes.

Marinescu (2015) employs a methodology that does not attempt to address these challenges. A more innovative approach in Coglianesi (2015) and Chodorow-Reich and Karabarbounis (2016) attempts to overcome these challenges by exploiting a sampling error in unemployment rates as an exogenous variation. Unfortunately, we find that this approach falls prey to the very problems it aims to overcome and it appears unlikely that the fundamental bias at the core of this approach can be overcome. We find more promising the approach based on unexpected policy changes as in the recent contributions by Johnston and Mas (2015) and Hagedorn, Manovskii and Mitman (2015). This approach by design addresses the problem of benefit endogeneity. It does not, however, fully address the effects of expectations and generally yields a lower bound on the actual effects of policies.

Marcus Hagedorn
Department of Economics,
University of Oslo,
Box 1095 Blindern,
0317 Oslo, Norway
marcus.hagedorn@econ.uio.no

Kurt Mitman
Institute for International Economic Studies
Stockholm University
106 91 Stockholm
SWEDEN
kurt.mitman@iies.su.se

Iourii Manovskii
Department of Economics
University of Pennsylvania
160 McNeil Building
3718 Locust Walk
Philadelphia, PA 19104
and NBER
manovski@econ.upenn.edu

1 Introduction

Unemployment insurance (UI) benefit extensions are one of the most prominent and actively used countercyclical stabilization policies. In the U.S., UI benefits have been extended in every recession since 1957, including an extension of unprecedented magnitude following the onset of the Great Recession - from a typical duration of 26 weeks up to 99 weeks. While previous extensions were largely discretionary, the President's 2016 fiscal year budget called for a new, permanent, federally-financed emergency unemployment insurance program which would automatically provide up to 52 weeks of additional benefits in recessions.

Of course, the wisdom of relying on this policy for macroeconomic stabilization depends on the assessment of its impact on the aggregate labor market variables, such as aggregate employment, unemployment, labor force, job vacancies, etc. The vast empirical micro literature had a considerably narrower focus of inferring the effect of benefit extensions on the duration of unemployment spells of UI claimants and generally found a large impact.¹ Moreover, the methodology in these studies was designed with the goal of inferring the effect of benefit extensions on the search effort of the unemployed. The equilibrium labor market theory implies, however, that unemployment benefit extensions also impact job creation decisions by employers.

To address these limitations, Hagedorn et al. (2013) proposed a semi-structural measurement methodology which allows to measure the aggregate effects of this policy that includes its effects on job creation. Their estimator is designed to overcome two key empirical challenges: First, job creation, as any other investment decision, depends not only on the current policy, but also on the expectations of future policies. Second, unemployment benefit extensions are endogenous with respect to economic conditions. Implementing this methodology in the data, they find that benefit extensions raise equilibrium wages and lead to a sharp contraction in vacancy creation, employment, and a rise in unemployment.

In an attempt to relax the theoretical restrictions underlying the more structural approach in Hagedorn et al. (2013), two strands have recently emerged in the literature that exploit quasi-experimental variation to infer the total effects of UI benefit extensions.

¹For example, Katz and Meyer (1990), whose findings became the widely accepted benchmark in the literature, conclude: "Our results indicate that a 1-week increase in potential benefit duration increases the average duration of the unemployment spells of UI recipients by about 0.16–0.20 week... An increase in potential benefit duration ... from 6 months to 2 years is predicted to generate a 13–16-week increase in unemployment duration." This roughly doubles the duration of unemployment spells in the authors' data. If the duration of all unemployment spells doubles, the overall unemployment rate doubles as well, e.g. rising from about 5 to 10 percent following the increase of benefit duration of this magnitude during the Great Recession. If only the spells of unemployed who actually collect benefits double, the implied increase in aggregate unemployment would be roughly half as large. While there is conflicting evidence in the literature on which of the two experiments is more appropriate, the direct effect of this policy is very sizable in either of them.

The first strand, which we refer to as “methodology based on mistakes” was pioneered by Coglianesi (2015) and refined by Chodorow-Reich and Karabarbounis (2016). These authors attempt to exploit the fact that unemployment benefit extensions in the US are determined at the state level as a function of the contemporaneously measured unemployment rate, which is subject to measurement error due to sampling variation. Coglianesi (2015) identifies the “mistakes” in setting benefit duration through a comparison with administrative UI claimant counts. Chodorow-Reich and Karabarbounis (2016), instead, exploit future revisions to state unemployment data for this purpose. Having identified “erroneous” changes in benefit duration, they proceed to infer the impact of these changes on unemployment or employment. In contrast to Hagedorn et al. (2013), they find small or even negative effects of benefit extensions on unemployment. Unfortunately, as we explain below, this happens because this methodology does not address either of the two challenges associated with measuring the effect of UI extension on forward looking job creation decisions. It does not address the role of expectations and, perhaps paradoxically, falls prey to the very benefit endogeneity problem that it attempts to overcome.

The second approach exploits unexpected cuts in benefit duration to infer their labor market implications. Two prominent contributions in this strand of the literature are Johnston and Mas (2015) and Hagedorn et al. (2015). The former paper exploits the unexpected cut in potential benefit duration by 16 weeks in Missouri in April 2011. The latter paper studies the effects of a nationwide elimination of all benefit extensions in December 2013. This methodology addresses the challenge of policy endogeneity but it is not designed to fully control for the effects of expectations. While the timing of the actual cut in benefits was unexpected, the benefit extensions were never designed to be permanent, so it was natural for economic agents to expect them to be cut at some point in the future. Thus the effect of the policy change was to bring forward the time of the cut relative to expectations. Economic theory then suggests that these policy experiments reveal a lower bound on the true effects of benefit extensions. Nevertheless, both experiments suggest even larger effects than those estimated by Hagedorn et al. (2013), thus further challenging the wisdom of relying on benefit extensions as a stabilization policy.

The remainder of the paper is organized as follows. In Section 2 we draw on Hagedorn et al. (2013) to highlight the key methodological challenges for inferring the economic consequences of UI benefit extensions with a particular emphasis on the less understood role of expectations. In light of these challenges, in Section 3 we review the methodology based on mistakes and in Section 4 we discuss the methodology based on unexpected cuts in benefits. Section 5 concludes.

2 Methodological Challenges: The Role of Expectations

A key lesson in Hagedorn et al. (2013) that sets it apart from the empirical literature is that to obtain unbiased and interpretable estimates of the effects of unemployment benefits on unemployment and vacancies, it is important to take into account that firms' job creation decisions depend on future policies. On the one hand, this finding does not come as a surprise as all investment decisions are affected by expectations of future productivity, demand, and economic policies. On the other hand, it is quite surprising that this channel was deemed unimportant a priori by the empirical literature.²

A simple example might be helpful in illustrating why taking expectations into account is necessary. Consider two otherwise identical states, one which passes a law extending benefits by 20 weeks for one year, and the other which extends benefits by 10 weeks permanently. The estimates in Hagedorn et al. (2013) imply that in the state with the temporary extension, unemployment would increase 0.5 percentage points, whereas in the state with the permanent extension unemployment would increase by 0.8 percentage points. The effect is higher in the state with the permanent change because firms expect that profits at all future dates will be lower because of the extension, whereas in the state with the temporary change, firms expect profits to be lower for one year only.³ A naive difference-in-differences analysis would erroneously suggest a significant negative impact of unemployment benefits on unemployment. Moreover, if employers anticipated these policy changes, say a year in advance, virtually all the adjustment of job creation (and of unemployment in a standard search model) would have occurred prior to the actual change in policy. Indeed, the basic optimality of firms' decisions implies no discrete jumps in vacancy posting at the time that expected policy changes are implemented. Thus, an observer may conclude that unemployment and vacancies are not related to benefits because they change dramatically when benefits do not and do not change when benefits change. It is only by controlling for the movements induced by the changes in expectations that the correct magnitude of the effect of UI policy can be identified. This ostensibly simple logic is remarkably misunderstood in the literature.⁴

²This is despite clear available evidence that economic agents react well in advance of expected UI policy changes. See, e.g., Carling et al. (2001).

³Profits decline in response to an increase in benefit generosity because wages rise for a given level of workers' productivity. The logic is standard: An improvement of workers' outside option leads to a wage increase to prevent shirking in the efficiency wage model and directly affects the wage bargain in the search model.

⁴For example, the identification strategy of Marinescu (2015) implicitly assumes that the effect of UI policies on vacancy creation happens contemporaneously and only contemporaneously and rules out any expectation effects. Furthermore, Marinescu (2015) focuses on just one benefit change per state to obtain a "before" and "after" period to apply a particular existing empirical methodology. Neither selecting the sample based on the endogenous outcome variable leading to no other benefit changes in the "before" or "after" windows nor




Figure 1: The Impact of Future Unemployment Benefit Durations on Current Unemployment.

The important role of expectations can also be seen directly in the data as illustrated in Figures 1 and 2. The key feature of the U.S. unemployment insurance system is that unemployment insurance policies are determined at the state level and apply to all locations within a state. Figure 1(a) plots the coefficients estimated in Hagedorn et al. (2013) from a regression of log state unemployment in quarter t on the log of contemporaneous benefit duration and changes in log benefit duration over the next eight quarters (the regression includes fixed state and time effects). We observe that current unemployment is significantly positively related to future changes in benefit duration. State-level evidence does not isolate the effect of expectations, however, because it also reflects the endogeneity of benefits: benefit duration tends to rise in response to past increases in unemployment at the state level. The same endogeneity problem implies that one cannot infer the effects of benefit extensions by simply relating benefit duration to unemployment in a panel of states.

To overcome the endogeneity problem, Hagedorn et al. (2013) employ the dominant methodology in the labor economics literature: they exploit a policy discontinuity at state borders and compare the evolution of unemployment in counties that border each other but belong to different states. This methodology was used, among others, by Holmes (1998) to identify the impact of right-to-work laws on location of manufacturing industry as well as by Card and Krueger (2000) and Dube et al. (2010) to identify the effect of minimum wage laws on earnings and employment of low-wage workers. It captured the imagination of labor economists because, as Dube et al. (2016) explain, “Measuring labor market outcomes from an immediately adjacent county provides a better control group, since firms and workers on either side are generally af-

implicitly assuming that all other benefit changes are irrelevant seems appealing. Obviously the fact that benefits change multiple times and agents change their behavior in anticipation of these changes requires a different methodology.

ected by the same idiosyncratic local trends and experience macroeconomic shocks at roughly the same time... The border discontinuity design also offers a way to address concerns about policy endogeneity. Minimum wage policies may react to shocks affecting the whole state, not just those affecting counties right at the border. Therefore, policy differences within cross-border pairs are unlikely to reflect endogeneity concerns that may severely bias studies using state-level variation.” Hagedorn et al. (2013) provide formal evidence that supports this assertion. Specifically, they verify that differences in state level productivities and demand (estimated using the Bartik methodology) do not help predict the difference in variables of interest between border counties. To ensure that the endogeneity is not induced by the county-level data construction, they also verify that the results are consistent across numerous variables and datasets: LAUS unemployment data, administrative UI claims data, administrative payroll counts from QCEW and QWI, as well as the universe of on-line job vacancies.⁵

This insight implies that the effects of expectations can be isolated using a similar regression but with all variables differenced between bordering counties (unemployment is now measured at the county level and the regression includes fixed effects for each border county pair). We observe that eliminating the endogeneity problem plaguing the state-level analysis indeed leads to a substantial reduction of estimated coefficients plotted in Panel 1(b) relative to Panel 1(a). Nevertheless, current unemployment at the county level continues to respond significantly to future state-level benefit changes.

Clearly, the presence of significant expectation effects implies an empirical analysis which fails to account for firms’ expectations would measure an uninterpretable mixture of the true policy effect and of the unknown effect of labor market participants’ expectations of future policies. To address this problem, Hagedorn et al. (2013) derive a quasi-difference estimator of the effect of unemployment insurance policies on variables such as vacancies and unemployment that controls for the effect of expectations.

Using the quasi-difference estimator, we can confirm that significant effects of future policies identified using border county comparisons in Figure 1(b) are indeed due to expectations. Specifically, Figure 2 plots estimates from the same regression with the dependent variable being the difference in quasi-differenced unemployment between border counties. Quasi-differencing

⁵Marinescu (2015) does not address the endogeneity problem in her work but criticizes the border-county approach by arguing that her results based on applying this methodology to the vacancy data from one website are different from the results in Hagedorn et al. (2013) who apply it to the data on the universe of on-line vacancies. Of course, this only indicates that the data she uses is non-representative. She also points out that benefit changes can induce workers to change the allocation of search effort between border counties. This was specifically incorporated in the analysis of Hagedorn et al. (2013) and found to induce a negligible downward bias on the effect of benefits on unemployment.




Figure 2: Difference of Quasi-Differenced Unemployment between Border Counties.

eliminates the effects of expectations and we observe that indeed, when expectations are controlled for, the difference in current unemployment between border counties is independent of future benefit durations.⁶

3 Methodology Based on Mistakes

The idea of this methodology is to infer the effects of benefit extensions by identifying the instances when the extensions occurred “by mistake.” The credit for developing this methodology is due to Coglianesi (2015). However, his proposed implementation for identifying mistakes, discussed below, is neither econometrically nor economically sound. This led Chodorow-Reich and Karabarbounis (2016) to propose a different implementation, which we discuss first.

3.1 Chodorow-Reich and Karabarbounis (2016)

The idea in Chodorow-Reich and Karabarbounis (2016) is very simple. Let $u_{s,t}$ be the unemployment rate in state s at time t , as measured in real time, i.e., at time t . The duration of benefits $T_{s,t}$ in state s depends mechanically (and discontinuously) on this measurement, that is when $u_{s,t}$ crosses a pre-determined threshold, extended benefits are triggered either on or off. The data on unemployment rates are revised over time as more data become available or measurement procedures and concepts are refined. Let $\tilde{u}_{s,t}$ be the measure of the unemployment rate in state s at time t as revised at some future date. Denote by $\tilde{T}_{s,t}$ the hypothetical duration

⁶Hagedorn et al. (2013) also use future benefits directly to compute the response of unemployment to changes in UI generosity. They find that both the measurement using direct evidence on future policy changes and the quasi-differenced estimator which uses current policy only deliver the same answer, implying first, that future expected policy changes affect current choices and second, that their estimator controls for those expectations very well.

of benefits which would have been implemented at date t had the revised data been available in real time (given the same mechanical rule – extended benefit triggers – mapping measured unemployment rate to benefit duration).

As the revised unemployment data differ from the real time unemployment data, the actual benefit duration $T_{s,t}$ differs from the hypothetical duration $\tilde{T}_{s,t}$ by an error $\eta_{s,t}$ due to the measurement error in the real time unemployment rate:

$$T_{s,t} = \tilde{T}_{s,t} + \eta_{s,t}. \quad (1)$$

The interpretation in Chodorow-Reich and Karabarbounis (2016) of $\eta_{s,t}$ as exogenous measurement error requires that $\tilde{T}_{s,t}$ and $\eta_{s,t}$ move independently of each other. An increase in $\tilde{T}_{s,t}$ triggered by an increase in $\tilde{u}_{s,t}$ then has to lead to an increase of equal size of $T_{s,t}$ induced by $u_{s,t}$ (otherwise $\tilde{T}_{s,t}$ and $\eta_{s,t}$ will co-move invalidating the assumption of η being measurement error). In terms of the unemployment rate, this co-movement of benefit means that whenever $\tilde{u}_{s,t}$ triggers a benefit change so does $u_{s,t}$. We now argue that the assumption of η being measurement error is demonstrably false and leads to a biased estimate. Denote therefore the difference between these two benefit series as

$$\hat{T}_{s,t} = T_{s,t} - \tilde{T}_{s,t}. \quad (2)$$

Chodorow-Reich and Karabarbounis (2016) interpret $\hat{T}_{s,t} = \eta_{s,t}$ as exogenous measurement error driven by the (incorrect) real-time data but not by the revised (true) unemployment rate. Such an exogenous variation in benefits can then be used to measure how a change in benefits affects unemployment. The type of regressions implemented in Chodorow-Reich and Karabarbounis (2016) regresses the revised unemployment rate on the measurement error in benefits while controlling for state and time fixed effects δ_s and δ_t ,⁷

$$\tilde{u}_{s,t} = \beta \hat{T}_{s,t} + \delta_s + \delta_t + \epsilon_{s,t}. \quad (3)$$

The key identification assumption to obtain an unbiased estimate is that the error term $\epsilon_{s,t}$ is uncorrelated with the regressor $\hat{T}_{s,t}$. The automatic increase in benefits in a state as a mechanical discontinuous function of state unemployment — a relationship the authors apparently overlook — immediately implies that this assumption is wrong. To see this consider a shock to the left hand side variable, revised unemployment $\tilde{u}_{s,t}$, which by assumption in Chodorow-Reich and

⁷Surprisingly, with exception of Hagedorn et al. (2013) and Hagedorn et al. (2015), this literature does not make use of the flexible interactive effects model of Bai (2009) which improves upon including just fixed effects as it controls for the heterogeneous impact of aggregate shocks across locations such as states or counties, provides a natural way to control for observed and unobserved spatial heterogeneity and allows for a very flexible model of local-level trends in variables.

Karabarbounis (2016) is not correlated with $\hat{T}_{s,t}$, that is it cannot induce a change in $\hat{T}_{s,t}$ at the same time. However, a shock that changes $\tilde{u}_{s,t}$ also changes (or can change) hypothetical benefits $\tilde{T}_{s,t}$. The question is then whether this induced change in $\tilde{T}_{s,t}$ leaves $\hat{T}_{s,t}$ unchanged. This is equivalent to the question whether changes in $\tilde{T}_{s,t}$ induce changes in $\eta_{s,t}$, which requires (as explained above) that the real time unemployment rate and the revised unemployment rate always cross the threshold at the same time. However, the measurement error in unemployment rates implies that this is not the case in the data.⁸

A Simple Example. A simple example illustrates the bias in this approach. Suppose there is only one unemployment threshold at 6%, so that benefit duration increases by 13 weeks from the regular 26 weeks to 39 weeks when unemployment crosses this threshold from below. Duration falls from 39 to 26 weeks if this threshold is crossed from above. At any point in time the economy can be described by one of four cases depending on whether revised and/or real-time unemployment are above or below the threshold:

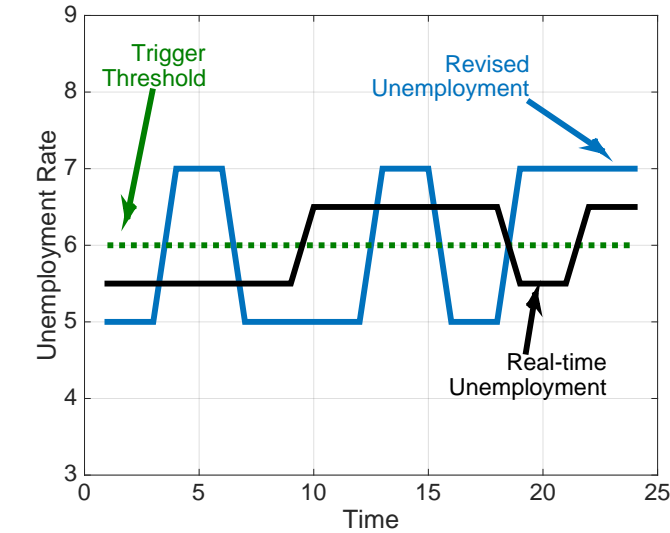
- Case 1. $u_t < 6\%$, $\tilde{u}_t < 6\%$: $T_t = 26$, $\tilde{T}_t = 26$, $\hat{T}_t = 0$.
- Case 2. $u_t > 6\%$, $\tilde{u}_t < 6\%$: $T_t = 39$, $\tilde{T}_t = 26$, $\hat{T}_t = 13$.
- Case 3. $u_t < 6\%$, $\tilde{u}_t > 6\%$: $T_t = 26$, $\tilde{T}_t = 39$, $\hat{T}_t = -13$.
- Case 4. $u_t > 6\%$, $\tilde{u}_t > 6\%$: $T_t = 39$, $\tilde{T}_t = 39$, $\hat{T}_t = 0$.

Now consider the shock ϵ from equation (3) which increases the revised unemployment rate \tilde{u} and triggers an increase in benefits \tilde{T} by 13 weeks. This shock moves the economy from Case 1 or 2 to either Case 3 or 4. Starting from Case 1, \hat{T} changes from 0 to either -13 or to 0. Starting from Case 2, \hat{T} changes from 13 to either -13 or to 0. So in all cases \hat{T} weakly decreases (strictly in 3 out of 4 scenarios) when \tilde{T} increases, establishing a negative co-movement. The case of a shock ϵ which decreases the revised unemployment rate \tilde{u} and triggers a decrease in benefits \tilde{T} is symmetric. This shock moves the economy from Case 3 or 4 to either Case 1 or 2. Starting from Case 3, \hat{T} changes from -13 to either 0 or to 13. Starting from Case 4, \hat{T} changes from 0 to either 0 or to 13. So in all cases \hat{T} increases (strictly in 3 out of 4 scenarios) when \tilde{T} decreases, once again establishing a negative co-movement.

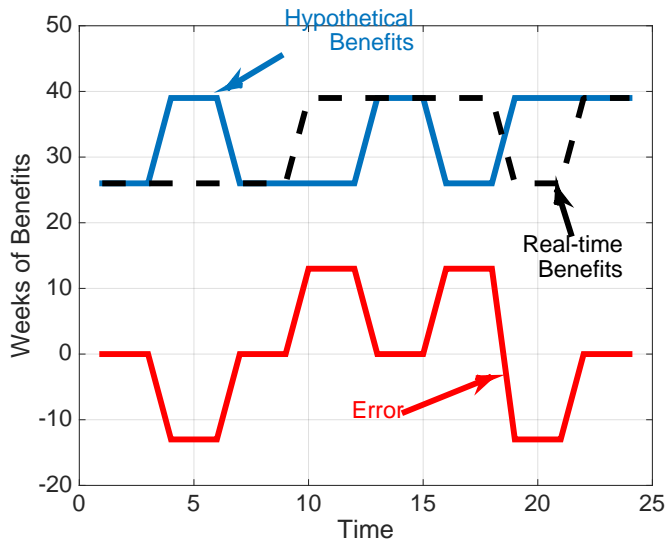
⁸The same arguments apply to regressions for lagged responses of unemployment

$$\tilde{u}_{s,t+k} = \beta^k \hat{T}_{s,t} + \delta_s + \delta_t + \epsilon_{s,t+k}, \quad (4)$$

since the shock to unemployment ϵ (and to most other macro variables) is highly persistent. An innovation ϵ_t which increases today's unemployment rate $\tilde{u}_{s,t}$ and triggers an increase $\hat{T}_{s,t}$ also increases $\epsilon_{s,t+k}$ and therefore $\tilde{u}_{s,t+k}$, implying a bias in the estimate β^k .



(a) Real-time and Revised Unemployment



(b) Real-time and Hypothetical Benefits and “Mistakes”

Figure 3: Mechanical Negative Correlation between the “Errors”, Hypothetical Benefit Duration, and Revised Unemployment Invalidating Empirical Methodology in Chodorow-Reich and Karabarounis (2016).

Figure 3 illustrates the mechanical bias that endogeneity of benefits at the state level induces on the measurement approach in Chodorow-Reich and Karabarounis (2016). Panel 3(a) plots a path of real time and revised unemployment that includes all possible transitions across the four Cases described above. Panel 3(b) plots the associated paths of real-time and hypothetical benefit durations and of the difference between real time and hypothetical benefit durations. The mechanical negative co-movement between the error and hypothetical benefits (and revised unemployment) is immediately evident. This, of course, invalidates the key identifying assumption in Chodorow-Reich and Karabarounis (2016) that innovations in revised unemployment

that trigger hypothetical benefit changes cannot lead to changes in the error \hat{T} appearing on the right hand side of the regression.

Evidence of Bias in the Data. The negative relationship between \hat{T} and \tilde{T} is also easy to see directly in the data. We use the same real time and revised measures of state unemployment rate during the Great Recession as Chodorow-Reich and Karabarbounis (2016). We regress the “measurement error” in benefit duration, $\hat{T}_{s,t}$, on the hypothetical benefit duration, $\tilde{T}_{s,t}$, based on the revised unemployment data:

$$\hat{T}_{s,t} = \beta^T \tilde{T}_{s,t} + \delta_s + \delta_t + \epsilon_{s,t}^T. \quad (5)$$

Violating the key identifying assumption in Chodorow-Reich and Karabarbounis (2016), $\hat{\beta}^T = -0.061$ is negative and strongly significant (s.e. 0.015) (focusing only on the Great Recession period of 2005–2013 we find an even stronger negative relationship $\hat{\beta}^T = -0.070$, s.e. 0.017). Thus, as expected, the two variables are strongly negatively correlated. Unsurprisingly, regressing the hypothetical benefit duration, $\tilde{T}_{s,t}$, on the revised unemployment rate, $\tilde{u}_{s,t}$,

$$\tilde{T}_{s,t} = \beta^u \tilde{u}_{s,t} + \delta_s + \delta_t + \epsilon_{s,t}^u. \quad (6)$$

yields a positive significant coefficient, $\hat{\beta}^u = 0.879$ (s.e. 0.105), reflecting partially that the benefit duration series was constructed using the revised unemployment rate series (focusing again on the Great Recession period we find an even stronger relationship $\hat{\beta}^u = 1.229$, s.e. 0.157). Taken together, these findings clearly imply a downward bias in the estimated β in equation (3). In case benefit extensions had a limited effect on unemployment, the estimated β would be significantly negative whereas a value of β equal to zero suggest a true positive effect of an increase in benefits on unemployment.

Errors vs. Innovations. Chodorow-Reich and Karabarbounis (2016) are not using $\hat{T}_{s,t}$ but instead use innovations to this series,

$$\nu_{s,t} = \hat{T}_{s,t} - \mathbb{E}_{t-1} \hat{T}_{s,t}. \quad (7)$$

Obviously, the same arguments apply and thus the same endogeneity problems arise with any innovation series as they do to the original series. When we regress the innovations $\nu_{s,t}$ on hypothetical benefits $\tilde{T}_{s,t}$,

$$\nu_{s,t} = \beta^I \tilde{T}_{s,t} + \delta_s + \delta_t + \epsilon_{s,t}^I, \quad (8)$$

we find a negative significant $\hat{\beta}^I = -0.014$ (s.e. 0.0046) (during the Great Recession period we




Figure 4: Errors \hat{T} and innovations in California and Nevada.

find $\beta^I = -0.014$, s.e. 0.00468).

Furthermore, constructing innovations instead of directly using the series of errors introduces an additional bias to the analysis, since the unemployment rate is used to generate predictions of future benefits (see Footnote 9 in Chodorow-Reich and Karabarbounis (2016)). More generally, an insurmountable inconsistency plagues the approach of using innovations. On the one hand, agents use the unemployment rate to predict future benefit levels, so omitting unemployment when constructing innovations will not yield unanticipated innovations to agents. When omitting unemployment, the econometrician would use regressors as innovations that are not innovations to agents, but an unknown convolution of current and past variables, which invalidates the approach. Thus, this requires using the unemployment rate when constructing innovations, but this again leads to an almost unavoidable bias. The econometrician has to make arbitrary choices about agents' information sets and their forecasting procedure, which are almost certainly not correctly reflecting how agents forecast, so that innovations depend on past variables such as the unemployment rate. As a result, a bias of unknown sign and size arises when the effect of an innovation on unemployment is estimated since innovations depend by construction on unemployment. In both cases — using unemployment as Chodorow-Reich and Karabarbounis (2016) apparently do or not using it — leads to the conclusion that additional biases are added when using innovations.

The conclusion that the construction of innovations has to be biased is confirmed through the experience of California and Nevada, as shown in Figure 4. The real-time and revised benefit duration fully agreed in these states so that at no point in time there was an error in setting benefits, i.e. $\hat{T}_t = 0 \forall t$. Remarkably, however, the innovations series computed by Chodorow-Reich and Karabarbounis (2016) for these states are almost never zero and move around quite inex-

plicably. It is unclear which behavior of unemployment or other macro variables in California and Nevada are causing movements in innovations but it is clear that these are erroneous movements not reflecting error in benefits. Clearly, one would expect that regressing unemployment on noise yields a coefficient of around zero.

3.1.1 A Placebo Test of the Chodorow-Reich and Karabarbounis (2016) Estimator

To illustrate the severity of the endogeneity bias plaguing the empirical strategy in Chodorow-Reich and Karabarbounis (2016) we apply the estimator to the data from a time period when there were no benefit extensions with an artificially created placebo measure of weeks of benefits available based on a hypothetical trigger of benefit extensions and using real-time and revised unemployment data.

Specifically, we consider data from 1996-2000 when no extended benefits were available in the US.⁹ In practice, a state triggers on a benefit extension in a given month if the three month average of the state seasonally adjusted unemployment rate exceeds a pre-determined threshold. Consequently, we use the real-time measure of unemployment provided by Chodorow-Reich and Karabarbounis (2016) and specify our placebo extension of 1.75 months (to correspond to the average non-zero benefit duration error in their sample) in any month when the preceding three month average of the state seasonally adjusted unemployment rate exceeded a pre-specified threshold, τ . Importantly, the 2015 state unemployment data revision by the BLS exploited by Chodorow-Reich and Karabarbounis (2016) affected not only the data from the Great Recession period, but also the data all the way back to the 1970s. Thus, we use the $\tilde{u}_{s,t}$ data from the 2015 revision to construct $\tilde{T}_{s,t}$ — the hypothetical duration of benefits which would have been implemented using the revised data given the same mechanically pre-specified thresholds τ . The two measures of unemployment and the implied benefit series allow us to construct the “measurement error” $\hat{T}_{s,t}$ as in Eq. (2).

We consider placebo trigger thresholds, τ , in 0.05% increments from 4% to 6% unemployment. This range was chosen to include thresholds that generate both higher and lower frequency and duration of triggered benefit extensions than observed during the Great Recession period, thus encompassing the plausible range of appropriate placebo experiments. We then regress the revised unemployment on the “measurement error” in benefits as in regression (3) for each τ . The estimates are all large and negative, and all but one of the individual coefficients are statistically significant. The coefficients corresponding to each threshold are plotted in figure

⁹Except for a brief extension of benefit duration in New Jersey studied by Card and Levine (2000). Removing that episode from the placebo sample has no impact on the findings.




Figure 5: Placebo Test of the Chodorow-Reich and Karabarbounis (2016) Methodology.

5. Of course, given that there were no actual unemployment benefit extensions, an estimator free of endogeneity bias must have recovered a zero effect of nonexistent benefit extensions of unemployment. Thus, the large negative impact of benefit extensions on unemployment found with the Chodorow-Reich and Karabarbounis (2016) estimator on the placebo sample, implies that the result is driven entirely by the endogeneity bias described above.¹⁰ To highlight the size of the bias, the red line in Figure 5 plots the minus one times the effect of an extension of this size estimated in Hagedorn et al. (2013). The average estimated negative bias in Chodorow-Reich and Karabarbounis (2016) is nearly four times as large as the estimated positive effect in Hagedorn et al. (2013) and is large for all values of τ . The size of the bias then suggests that the fact that Chodorow-Reich and Karabarbounis (2016) find no effect of benefit extensions on unemployment is consistent with significantly larger positive effects than those found in Hagedorn et al. (2013). The implied magnitude of the effect appears implausible, however.

3.1.2 The Role of Expectations

Up to this point in our discussion, we have given the benefit of the doubt to Chodorow-Reich and Karabarbounis (2016), who interpret the difference between $u_{s,t}$ and $\tilde{u}_{s,t}$ as measurement error. This is a questionable assumption because the data revision they rely on reflected not only better state unemployment data but a host of methodological changes, including, e.g. a

¹⁰Note that Hagedorn et al. (2013) use the same placebo experiment to verify the excellent performance of the estimator they proposed. Coglianese (2015) reaches a different conclusion because he does not correctly implement the estimator in Hagedorn et al. (2013) and because his placebo sample includes year 2001 despite the fact that unemployed in 2001 actually received extended benefits.

different seasonal adjustment procedure. If, however, the difference in the real-time and revised unemployment rate is not measurement error, then clearly the difference between the corresponding benefits cannot be measurement error either.

The measurement approach in Chodorow-Reich and Karabarbounis (2016) also does not take into account that expectations of the full sequence of future expected benefits matter for forward-looking decision makers. The strategy to identify innovations in the stochastic process for “policy errors” is not per se helpful in inferring agents’ expectations for at least two reasons. First, just considering unexpected current period innovations ignores changes in all future benefits. Obviously, a one-period change in benefits can only have very small effects since investment decisions by forward-looking agents respond mainly to persistent changes in future policies. Focusing just on innovations, therefore, ignores the quantitatively relevant part of policy and can even lead to wrong conclusions as we explained above. Second, inferring agents’ surprises from innovations in the stochastic process is problematic. This has been long realized in the literature which emphasized the distinction between the information set of the econometrician and the information set that the economic agents act on. For example, it is well understood that stochastic innovations in earnings identified by an econometrician need not correspond to surprise realizations to an agent but may well be anticipated by the agent long in advance. The key lesson from that literature is that one needs to impose considerably more structure to infer the information sets of forward looking individuals. We cannot imagine how this can be accomplished while imposing less structure than is required in Hagedorn et al. (2013).

3.2 Coglianesse (2015)

Coglianesse (2015) was the first to propose to identify the effects of benefit extensions based on what the author labels a “natural experiment ... whereby random sampling error in a national survey altered the duration of unemployment insurance in several states, resulting in random variation in the number of weeks of unemployment insurance available at the state level.”

However, what Coglianesse (2015) labels a sampling error is simply the difference between the insured unemployment rate (IUR) and the total unemployment rate (TUR). The IUR counts the number of continuing claimants in regular state UI programs (weeks 2 through 26 of unemployment). TUR, on the other hand, counts all unemployed. Since IUR refers to a small subset of all unemployed, it clearly differs from TUR.¹¹ Importantly, the difference is *endogenous*

¹¹The BLS provides a more detailed explanation on its website (http://www.bls.gov/cps/cps_htgm.htm): “While the UI claims data provide useful information, they are not used to measure total unemployment

to current and future economic conditions and UI policy. For example, the difference includes those who exhausted regular state benefits (and might be receiving extended benefits) and the share of such long-term unemployed is endogenous to both economic conditions and policy. It also depends on the decision of unemployed whether to apply for benefits, which clearly depends on the future economic conditions they expect to face. It also depends on the decision of whether to quit a job or leave schooling and enter unemployment ineligible for benefits. Thus, the difference is an aggregate of a myriad of endogenous choices and is certainly not a natural experiment.¹² The empirical strategy in Coglianesse (2015) that requires exogeneity of this difference is thus obviously fundamentally flawed.¹³ Presumably, this led Chodorow-Reich and Karabarbounis (2016) to propose a different implementation discussed above.

3.2.1 A Placebo Test of the Coglianesse (2015) Estimator

To illustrate the bias in Coglianesse (2015), we apply his methodology to the same placebo sample as in Section 3.1.1. Specifically, following Coglianesse (2015), we control for economic conditions by including 8 lags of IUR and instrument for UI duration with a dummy variables for whether a state was triggered on. The dependent variable is measured in quarter $t + 1$ “to avoid issues of simultaneous causality.” The regression includes state and time fixed effects and

because they exclude several important groups. To begin with, not all workers are covered by UI programs. For example, self-employed workers, unpaid family workers, workers in certain not-for-profit organizations, and several other small (primarily seasonal) worker categories are not covered.

In addition, the insured unemployed exclude the following:

1. Unemployed workers who have exhausted their benefits.
2. Unemployed workers who have not yet earned benefit rights (such as new entrants or reentrants to the labor force).
3. Disqualified workers whose unemployment is considered to have resulted from their own actions rather than from economic conditions; for example, a worker fired for misconduct on the job.
4. Otherwise eligible unemployed persons who do not file for benefits.

Because of these and other limitations, statistics on insured unemployment cannot be used as a measure of *total* unemployment in the United States. Indeed, over the past decade, only about one-third of the total unemployed, on average, received regular UI benefits.”

¹²Labeling the difference between TUR and IUR a “sampling error” is also misleading. First, the TUR is constructed by LAUS program as a partnership between states and the BLS and it specifically corrects for the sampling error in the CPS and the correlation structure of its errors that Coglianesse (2015) mentions. Second, TUR estimates and not IUR estimates are used to allocate billions of dollars by various federal programs (see <http://www.bls.gov/lau/lauadminuses.pdf> for a subset). Wouldn’t it be remarkable if policymakers have decided to base policies on a “sampling error” when counts of state UI claimants are readily available to them?

¹³Hagedorn et al. (2013) have already used the same administrative UI claims data that are used by Coglianesse (2015) to corroborate their findings based on total unemployment. They show how comparable statistics can be constructed from the two measures of unemployment and find that both measures of unemployment lead to the same conclusions.




Figure 6: Placebo Test of the Coglianesse (2015) Estimator.

standard errors are clustered at the state level.

$$\log(TUR_{s,t+1}) = \beta \log(UIDuration_{s,t}) + \sum_{k=1}^8 \gamma_k \log(IUR_{s,t-k}) + \theta_s + \delta_t + \epsilon_{s,t}. \quad (9)$$

The estimated coefficients for the full range of τ are plotted in Figure 6, revealing the systematic and highly significant bias across all placebo specifications. Despite being estimated on the placebo sample where no extensions actually took place, the average of the estimates (across values of τ) is $\hat{\beta} = 0.156$, implying large effect of placebo benefits on unemployment. These results clearly indicate that the methodology proposed in Coglianesse (2015) does not overcome the endogeneity problem. Unfortunately, this placebo experiment cannot fully describe the size and even the sign of the bias in the actual data where benefit extensions actually take place. This is because IUR is endogenous to current and expected benefit extensions, inducing an additional endogeneity bias not captured by this placebo analysis.

3.2.2 Final Comment on Coglianesse (2015)

While interpreting the clearly biased estimates of Coglianesse (2015) is probably a fool's errand, it seems worth pointing out that his preferred estimate of positive effects of benefit duration on employment growth is fully consistent with a Mortensen-Pissarides search and matching model. Indeed, we replicated this regression on the data generated from the model in Hagedorn et al. (2013) and found a coefficient of 0.103, somewhat larger than the empirical estimate in Coglianesse (2015). Of course, higher benefits in the model lead to a lower level of employment. At first glance this might appear like a contradiction but the logic is simple. Imagine that

benefits increase unexpectedly in period t and are expected to last for n periods. This leads to an upward revision in expected benefits in periods $t + 1$, $t + 2$, etc. Whenever the increase in period t benefits leads to an upward revision of expected future benefits, employment in period t falls relative to employment in period $t + 1$, employment in period $t + 1$ falls relative to employment in period $t + 2$ and so on until benefits return to their pre-period t level. Thus, higher benefits are associated with higher employment growth but this is due to benefits depressing current employment.

4 Methodology Based on Unexpected Cuts in Benefits

Another approach to assessing the effects of UI benefit extensions is based on studying the data surrounding unexpected changes in benefit duration, often occurring for political reasons.

4.1 The Natural Experiment in Missouri, Johnston and Mas (2015)

In a prominent recent quasi-experimental study in this methodological tradition, Johnston and Mas (2015) assessed the impact of the sudden and unanticipated cut in potential benefit duration by 16 weeks in Missouri in April 2011. The cut applied only to new claimants while those who claimed benefits prior to the reform were grandfathered into the old potential benefit duration schedule. Using individual-level administrative data on unemployment and employment, the authors find a very large positive effect of the cut in benefit duration on exit rate from unemployment (mainly into employment). This estimate is based on a comparison of a large number of individuals that differ with respect to potential benefit duration because of becoming unemployed shortly before or shortly after the reform. Thus, it plausibly identifies the effect on workers' search intensity because all these individuals arguably face the same labor market. Hagedorn et al. (2013) label the effect of UI benefit duration on individual search effort controlling for market-level conditions "the micro effect." This effect is well identified in Johnston and Mas (2015) although it is significantly larger than earlier estimates based on a similar methodology, e.g., Card and Levine (2000).¹⁴

Knowledge of the magnitude of the micro effect is, however, insufficient to assess the total effect of changes in benefit duration on the labor market. Chodorow-Reich and Karabarbounis (2016), among others, argue that evaluating a one time policy change in a single state provides effectively only one data point and one cannot draw econometrically sound conclusions from

¹⁴It is also larger than the estimates in the older literature reviewed in Krueger and Meyer (2002) and Nicholson and Needels (2004).




Figure 7: Dynamics of Aggregate Vacancies and Employment before and after 2014 Unemployment Benefit Duration Cut.

it that would apply to a nationwide cut in benefits. Despite this concern about their external validity, the patterns suggested by the Missouri experiment appear interesting. Johnston and Mas (2015) find that job creation must have increased following the reform just enough to absorb all the unemployed workers who increased their search effort following the reform. Hagedorn et al. (2013) consider direct evidence on the behavior of job vacancies in Missouri and document a sharp and discontinuous rise in the vacancy-unemployment ratio in Missouri at the time of the reform. Note that, at least on impact, the observed jump in vacancy creation cannot be in response to higher search effort of new claimants eligible for fewer weeks of benefits following the reform, simply because they represent a tiny fraction of all unemployed (only about a third of all unemployed were claiming benefits in Missouri at the time of the reform and it took a number of months for the claimants under the new rules to account for a meaningful share of all claimants). This is consistent with the presence of a significant “macro effect” introduced by Hagedorn et al. (2013) following the logic of the standard equilibrium search model where a weakening of the outside option of workers raises the hiring incentives for firms.

4.2 The 2014 Employment Miracle, Hagedorn et al. (2015)

In this section we discuss a quasi-experimental evaluation of a nationwide cut in benefits that overcomes the concern raised by Chodorow-Reich and Karabarbounis (2016) about the generality of findings based on a policy change in one state. Specifically, Hagedorn et al. (2015) exploit the unexpected failure by US Congress in December 2013 to reauthorize the unprecedented benefit extensions introduced during the Great Recession. All federal extensions were abruptly

cut to zero, implying the average decline of benefit duration across US states of about 25 weeks. Following the reform, the US labor market boomed. For example, Figure 7 illustrates a sharp increase in aggregate employment to population ratio and the number of job vacancies seeking to recruit a worker. While the evidence appears striking, it is nevertheless not convincing as it is still based on effectively one data point, albeit the one for the whole nation.¹⁵

To obtain more robust conclusions, Hagedorn et al. (2015) exploit the fact that, prior to being cut to zero in December 2013, federal benefit extensions ranged from 0 to 47 weeks across U.S. states and the fact that this policy change was exogenous to cross-sectional differences across U.S. states. The available benefit duration in a given state just prior to the reform, however, was endogenous to the economic conditions of the state. Thus, as in Johnston and Mas (2015), or indeed in any attempt to infer aggregate policy effects from variation induced by policy changes, the key challenge to a proper inference of the effects of benefits is to ensure that the effects are not confounded by pre-existing differences in employment or labor force trends. Hagedorn et al. (2015) use two distinct empirical methodologies and three different models of local-level trends for this purpose. They also consider employment data based on administrative job counts and on (primarily survey-based) counts of individuals with a job. Across all methodologies and data sources they consider, they find a large and statistically significant effect of the 2014 cut in benefits on employment. They provide a model that allows to aggregate local estimates to a nationwide impact and conclude that the cut in benefit duration accounted for about 50 to 80 percent of the aggregate employment growth in 2014. As Figure 8, summarizing their findings, indicates, over half of the aggregate employment growth was due to the increase in the labor force induced by the policy reform.¹⁶


4.3 Policy Reforms and Expectations

The advantage of using exogenous policy reforms to identify the effects of benefit extensions is that it partially breaks the mechanical link between unemployment and available benefit duration. This happens only partially because pre-reform benefit duration was indeed endogenous to economic conditions and it is important to measure the counterfactual employment trends that would have realized in the absence of the reform.

Moreover, exploiting this type of reforms cannot by itself fully control for the effects of expectations. While the reform itself is exogenous, what is relevant for the economic agents is

¹⁵Similarly striking evidence emerges from the labor market boom following the benefit duration cut (known as Hartz IV reform) in Germany. Krause and Uhlig (2012) review and structurally interpret the evidence.

¹⁶This is consistent with the theoretical prediction of the standard search model that an increase in job availability draws non-participants into the labor market. See Pissarides (2000) Ch. 7 for a textbook treatment.



Methodology and Specification

Figure 8: Implied aggregate effects of the benefit cut in December 2013 across methodologies and datasets used in Hagedorn et al. (2015).

the change in the full path of expected future benefit durations induced by the reform. Indeed, this experiment is likely to yield a lower bound on the true effect if job creators assigned, say, prior to the reform in December 2013, some probability to a future cut in benefit duration but were surprised by the specific timing of when the reform was implemented. Moreover, in case of evaluating a policy reform in an individual state, say, Missouri, there is a downward bias if, upon observing the policy change in that state, job creators in neighboring states assigned a higher probability to similar benefit duration cuts taking place in their states.

5 Conclusion

We reviewed recent methodological and empirical contributions aiming to provide a comprehensive assessment of the effects of unemployment benefit extensions on the labor market and attempted to reconcile their apparently disparate findings. The two key challenges facing the literature is to overcome the endogeneity problem due to the fact that unemployment benefit

durations set at the state level are a function of state’s labor market conditions. The second key challenge is driven by the fact that individuals’ decisions respond to expected changes in benefits before those changes actually take place. This makes it difficult to separate the effects of policies from the effects of expectations of future policy changes (perhaps induced by contemporaneous changes in policy).

Hagedorn et al. (2013) proposed an estimator designed to overcome these challenges. It draws on the classic approach in labor economics to the evaluation of (endogenous) state level policies based on the comparison of border counties belonging to different states. They also show that a very large class of modern equilibrium labor market models implies that the effects of expectations can be controlled for using an estimator based on the quasi-difference of the variables of interest. Following this approach, they find a significant positive effect of unemployment benefit extensions that took place during the Great Recession on overall unemployment. The quantitative magnitude of their findings is somewhat smaller than what can be expected from the earlier literature but this is likely due to methodological improvements.

The more recent literature has attempted to dispense with some of the complexities associated with the semi-structural approach in Hagedorn et al. (2013). One approach, pioneered by Coglianesi (2015) and refined in Chodorow-Reich and Karabarbounis (2016) attempts to overcome the endogeneity problem by identifying “mistakes” in setting policy, i.e., the instances when benefits were set in a way that was not justified by labor market conditions. Coglianesi (2015) identifies the “mistakes” as the difference between the insured unemployment rate and the total unemployment rate, but as we explained, the difference between these two variables cannot be interpreted as a mistake as the two variables and the difference between them measure endogenous choices affected by benefit policy. The approach in Chodorow-Reich and Karabarbounis (2016) measures the errors as the difference between the benefits set based on real time unemployment data and those that would have been set had the revised (in the future) unemployment data been available in real time. Unfortunately, their whole strategy for overcoming the mechanical correlation between real-time benefits and unemployment is based on ignoring the same mechanical correlation between revised benefits and unemployment. Their procedure flips the sign of the bias but does little to eliminate it. We show in a placebo experiment that the bias on the estimated effects of UI benefit extensions is large and negative. The assessment of the quantitative magnitude of the bias in the data when benefits actually change is even more complicated because this approach does not deal with the effects of expectations. Thus, while this approach is quite creative, its current implementations are fundamentally flawed at the very basic level. Unfortunately, we are doubtful that this strategy can be modified to overcome

the key identification challenges facing the literature.

A more promising approach to simplify the analysis is based on the assessment of unexpected policy changes. This quasi-experimental approach has a long tradition in economics and we argue that it can be helpful in overcoming some of the challenges facing this literature. For example, Johnston and Mas (2015) find large effects of an unexpected cut in benefit duration in Missouri. The external validity of the inference based on essentially one data point provided by this reform is perhaps questionable. However, this is the approach that will hopefully be applied by others to other policy reforms eventually enabling inference based on a range of such studies. A fruitful approach is to extend the scope of such studies and consider, as Hagedorn et al. (2015) a nationwide cut in benefit duration. Although such reforms are more rare, they allow to exploit the heterogeneity of impacts across locations to overcome the effects of sampling uncertainty.

The advantage of the studies based on this methodology is that they largely overcome the problem of policy endogeneity. They, however, are less able to isolate the effects of actual policy changes from the impact on expectations of future policies. As discussed above, these policy changes bring forward the timing of the policy changes that can be expected in the future. As all investment decisions are sensitive to expectations, the measured impact of these reforms is likely smaller than that of the policy change that not only changes the current benefits but also permanently changes the expected path of future benefits. Interestingly, even the likely lower bound on the true effects uncovered in these studies is quantitatively large.

Overall, we are encouraged by the recent efforts in the literature to rise to the challenge of measuring total effects of unemployment benefit extensions on the labor market.¹⁷ This is one of the most prominent and actively used countercyclical stabilization policies. Measuring its total impact is beyond the scope of earlier empirical micro work motivated by the narrow public finance questions of the impact of policy on the duration of insured spells. While obviously valuable, this literature left unanswered the key question of how benefit extensions affect the overall labor market. Our assessment of the existing findings points to the tentative conclusion that benefit extensions lead to significant increases in unemployment and to a reduction in employment and job vacancies. Given that these findings challenge the wisdom of relying on this policy instrument for cyclical stabilization, more research is clearly and urgently needed.

¹⁷An influential older literature with the same focus has concluded that UI generosity (with benefit duration being the key component) played a crucial role in explaining the cross-country differences in the levels and dynamics of unemployment between the US and many European economies. Scarpetta (1996) and Nickell et al. (2005) present original empirical analysis and review a large body of work based on cross-country regressions leading to this conclusion. An influential research program pioneered by Ljungqvist and Sargent (1998) aims to provide a structural interpretation to these findings.

Some approaches will ultimately not withstand the scrutiny but the issues are so important and the literature is so sparse, that the exploration of alternative identification strategies is clearly called for.

References

- BAI, J. (2009): “Panel Data Models with Interactive Fixed Effects,” *Econometrica*, 77, 1229–1279.
- CARD, D. AND A. B. KRUEGER (2000): “Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Reply,” *American Economic Review*, 90, 1397–1420.
- CARD, D. AND P. B. LEVINE (2000): “Extended Benefits and the Duration of UI Spells: Evidence from the New Jersey Extended Benefit Program,” *Journal of Public Economics*, 78, 107–138.
- CARLING, K., B. HOLMLUND, AND A. VEJSIU (2001): “Do Benefit Cuts Boost Job Finding? Swedish Evidence from the 1990s,” *The Economic Journal*, 111, 766–90.
- CHODOROW-REICH, G. AND L. KARABARBOUNIS (2016): “The Limited Macroeconomic Effects of Unemployment Benefit Extensions,” Working Paper 22163, National Bureau of Economic Research.
- COGLIANESE, J. (2015): “Do Unemployment Insurance Extensions Reduce Employment,” mimeo, Harvard University.
- DUBE, A., W. T. LESTER, AND M. REICH (2010): “Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties,” *Review of Economics and Statistics*, 92, 945 – 964.
- (2016): “Minimum Wage Shocks, Employment Flows and Labor Market Frictions,” *Journal of Labor Economics*, 34.
- HAGEDORN, M., F. KARAHAN, I. MANOVSKII, AND K. MITMAN (2013): “Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects,” National Bureau of Economic Research Working Paper 19499.
- HAGEDORN, M., I. MANOVSKII, AND K. MITMAN (2015): “The Impact of Unemployment Benefit Extensions on Employment: The 2014 Employment Miracle?” National Bureau of Economic Research Working Paper 20884.

- HOLMES, T. J. (1998): “The Effect of State Policies on the Location of Manufacturing: Evidence from State Borders,” *Journal of Political Economy*, 106, 667–705.
- JOHNSTON, A. C. AND A. MAS (2015): “Potential Unemployment Insurance Duration and Labor Supply: The Individual and Market-Level Response to a Benefit Cut,” mimeo, Princeton University.
- KATZ, L. F. AND B. D. MEYER (1990): “The Impact of the Potential Duration of Unemployment Benefits on the Duration of Unemployment,” *Journal of Public Economics*, 41, 45–72.
- KRAUSE, M. U. AND H. UHLIG (2012): “Transitions in the German Labor Market: Structure and Crisis,” *Journal of Monetary Economics*, 59, 64–79.
- KRUEGER, A. AND B. MEYER (2002): “Labor Supply Effects of Social Insurance,” *Handbook of Public Economics*, 4, 2327–2392.
- LJUNGQVIST, L. AND T. J. SARGENT (1998): “The European Unemployment Dilemma,” *Journal of Political Economy*, 106, 514–550.
- MARINESCU, I. E. (2015): “The General Equilibrium Impacts of Unemployment Insurance: Evidence from a Large Online Job Board,” mimeo, University of Chicago.
- NICHOLSON, W. AND K. NEEDELS (2004): “Extended Unemployment Benefits: A Review of the Literature,” Working paper, Amherst College.
- NICKELL, S., L. NUNZIATA, AND W. OCHEL (2005): “Unemployment in the OECD since the 1960s. What Do We Know?” *The Economic Journal*, 115, 1–27.
- PISSARIDES, C. (2000): *Equilibrium Unemployment Theory*, Cambridge, MA: MIT Press.
- SCARPETTA, S. (1996): “Assessing the Role of Labour Market Policies and Institutional Settings on Unemployment: A Cross County Study,” *OECD Economic Studies*, 26, 43–98.