DISCUSSION PAPER SERIES

DP16535

The Elusive State of Full Employment

Antonio Fatás

MONETARY ECONOMICS AND FLUCTUATIONS

CEPR

ISSN 0265-8003

The Elusive State of Full Employment

Antonio Fatás

Discussion Paper DP16535
Published 11 September 2021
Submitted 09 September 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

The Elusive State of Full Employment

Abstract

During the expansion phase of the business cycle unemployment declines towards levels that are usually associated to the natural rate of unemployment. Unfortunately, those low levels of unemployment are short lived. In the last 60 years, the US has not witnessed a long period of stable and low unemployment. We argue that this result is consistent with an asymmetric view of the business cycles (as in Friedman's plucking model) where many expansions are unfinished business, they end too early. The speed of labor market recovery is too slow and before the economy reaches full employment a recession starts, either because of imbalances or bad luck. Our analysis has implications for how we think about labor market slack and potential output. From a policy perspective it calls for a focus on faster recoveries while avoiding the type of imbalances that can be the seed of the next crisis.

JEL Classification: E32, J0

Keywords: Potential Output, business cycles, Natural rate of unemployment, recoveries

Antonio Fatás - antonio.fatas@insead.edu
INSEAD and CEPR

The Elusive State of Full Employment[*]

Antonio Fatas
INSEAD

Abstract: *During the expansion phase of the business cycle unemployment declines towards levels that are usually associated to the natural rate of unemployment. Unfortunately, those low levels of unemployment are short lived. In the last 60 years, the US has not witnessed a long period of stable and low unemployment. We argue that this result is consistent with an asymmetric view of the business cycles (as in Friedman's plucking model) where many expansions are unfinished business, they end too early. The speed of labor market recovery is too slow and before the economy reaches full employment a recession starts, either because of imbalances or bad luck. Our analysis has implications for how we think about labor market slack and potential output. From a policy perspective it calls for a focus on faster recoveries while avoiding the type of imbalances that can be the seed of the next crisis.*

[*] A previous version of this paper was circulated under the title "Is full employment sustainable?". I would like to thank Olivier Blanchard, John Fernald, Dimitrije Ruzic, Sanjay Singh and Alan Taylor for discussions during the development of the project and Wee Yang Ng for providing support as a research assistant.

1. Introduction

The most recent recession in the US started in February 2020 and was associated to a global pandemic caused by COVID-19, an event that came at a time when the US was enjoying very low levels of unemployment. Despite its depth, the 2020 recession became the shortest recession according to the NBER Business Cycle Dating Committee.

This recession was preceded by the longest expansion on record, starting in June 2009 and lasting for 128 months. During the expansion, unemployment came down from a peak of 10% in October 2009 to 3.5% in February 2020. At multiple times during this trajectory, and as unemployment declined to historically low levels, there was an ongoing debate on whether the US had already reached full employment. This debate is partly a result of the lack of consensus on how to define and measure equilibrium or natural levels of the unemployment rate or potential output. In the 2009-2020 expansion, and as time passed, a combination of further reductions in unemployment and increases in the labor force participation, without inflationary pressures, led to upward revisions of estimates of the natural rate of unemployment or what was considered a sustainable level of the employment to population ratio. The low level of unemployment achieved in the last years of the expansion delivered gains to segments of the labor market where progress was thought to be almost impossible, bringing some marginal workers from out of the labor force back into employment (Cajner, Coglianese, and Montes (2021)). This led to a conversation among policy makers and academics about the benefits of running a high-pressure economy.[1]

What is interesting is that despite the extraordinary length of the expansion, the potential gains from this high-pressure labor market turned out to be short lived. The expansion ended right at the time when unemployment rates had reached historically low levels. While a crisis caused by a global pandemic represents a unique event, the pattern of unemployment we have just described is also present in all previous US cycles. In this paper we show that US unemployment exhibits a V shape around all recessions. The right side of the V is known, recessions trigger a sudden increase in unemployment.[2] But the left side of the V is less known or understood. Unemployment is typically declining in the quarters prior to the beginning of a recession. In other words, the US does not display long periods of stable low unemployment rate.

There are several hypotheses that can explain the shape of unemployment around recessions. First, it might be that as we reach full employment financial or real imbalances are built and they lead to a recession.[3] The second hypothesis is that we have never seen an expansion that is long enough to witness a long period of full employment. For this to happen, we need both a slow speed of reduction of unemployment and the natural

---

[1] See Yellen (2016) of Aaronson et al. (2019). The notion of a high-pressure economy is not new, see Okun (1973).

[2] In fact, the moment we observe unemployment climbing up, it is a sign that a recession started (Sahm (2019)).

[3] There is a large academic literature on imbalances preceding recessions or financial crisis. See Sufi and Taylor (2021) for a recent review of this literature.

occurrence of frequently-enough recessionary shocks. Their combination makes the possibility of sustained periods of low unemployment highly unlikely.

We find empirical evidence that supports both hypotheses. Imbalances increase as unemployment decreases. But even when they do not, shocks happen (like a global pandemic) ending an expansion that had just brought the labor market to full employment. The slow speed of labor market recovery means that even a 10-year expansion cannot deliver a sustained period of full employment. In fact, our evidence supports the view that most recoveries are not finished by the time a next recession starts. If this interpretation is correct, from a policy point of view the focus should be on finding ways to accelerate the speed of labor market recoveries during an expansion to provide for longer periods of low unemployment. But that needs to be done in a way that avoids the buildup of imbalances via macroprudential policies.

This paper is related to several strands of the academic literature. A set of papers study the question of whether expansions die of old age. There is limited evidence that as time passes the end of an expansion is more likely; expansions do not seem to die of old age (Rudebusch (2016)). While this is true, the evidence of this paper suggests that as unemployment becomes low and reaches levels around or below what we consider the natural rate of unemployment, a recession tends to happen. In other words, if we measure time as the decreasing path of unemployment during an expansion, it seems that after a certain point the chances of a recession increase very fast.

The fact that unemployment displays such a decreasing trend during an expansion lends support to the idea that the business cycle is asymmetric, recessions can be seen as downward deviations from potential and the expansion phase is the time it takes for the economy to return to potential (Friedman (1993), Dupraz, Nakamura, and Steinsson (2019), Jordà, Schularick, and Taylor (2020), or Fatas and Mihov (2013)). What is new in this paper is that the economy seems to spend very little time close to potential, the episodes of low unemployment are short-lived.

Finally, the results of this paper also cast doubt on the empirical relevance of the concepts of full employment or the natural rate of unemployment. It seems that only if we are lucky and an expansion lasts for many years, as in the expansion starting in 2009, the economy finds itself at low enough rates of unemployment. This is linked to a literature that has shown the slow speed at which unemployment decreases during an expansion (Hall and Kudlyak (2021),Hall and Kudlyak (2020), Cole and Rogerson (1999), Amior and Manning (2018)). If this interpretation is correct, our estimates of the natural rate of unemployment are influenced by the length of expansions. As an example, if the global pandemic had happened in 2017 when unemployment was around 4.5%, it is very likely that we would be thinking of unemployment rates as low as 3.5% as unachievable.

Section 2 of the paper presents some stylized facts on the behavior of the labor market over the business cycle. Section 3 produces support for our main hypothesis, that low unemployment rates are not sustainable. Section 4 discusses the speed of labor market recoveries. Section 5 concludes and discusses policy implications.

2. Cyclicality of the Labor Market

2.1. Measuring slack in the labor market

During recessions and as economy activity declines, employment falls and unemployment increases fast. Unemployment typically reaches a peak shortly after the recession ends and as the expansion phase starts, we observe a declining unemployment rate towards a low level. The low level that the economy reaches by the end of the expansion is often associated to the notion of full employment of the natural rate of unemployment.[4] Figure 1 displays monthly unemployment for the US where the grey columns represent the months defined as recessions by the NBER.[5]

Figure 1. Unemployment rate

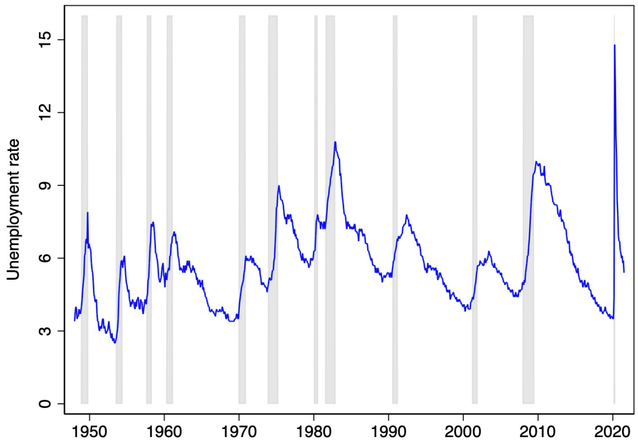

While the unemployment rate is the most common measure of the slack in the labor market, it can be inaccurate at times. Statistics on unemployment can, in some cases, underestimate the true amount of slack because of the presence of discouraged workers waiting for the labor market to get strong enough to re-enter the labor market. In these instances, the employment to population ratio can be more meaningful.

Figure 2 displays employment to population ratio for those over 16 years old as well as those between 25-54 years old.[6]

---

[4] In this paper we use terms such as full employment, the natural rate of unemployment or potential output to loosely describe situations where economic activity is at the level that employs the maximum amount of resources and that is sustainable. In the academic literature there are many definitions of such concepts (see, for example, Blanchard (2018)), but the nuances behind these definitions are not the focus of this paper. Empirically, there is great uncertainty about the level of any of these definitions. For example, estimates of the natural rate of unemployment were revised downwards as unemployment fell during 2009-2020 US expansion.

[5] Source for the unemployment rate: Unemployment Rate, Index 2012=100, Percent, Seasonally Adjusted. https://fred.stlouisfed.org/series/UNRATE. Monthly 1948M1-2021M7.

[6] Sources: Employment-Population Ratio 25-54 Yrs. Percent, Seasonally Adjusted. Monthly 1948M1-2021M7. https://fred.stlouisfed.org/series/LNS12300060. Employment-Population Ratio 16 and over. Percent, Seasonally Adjusted. Monthly 1948M1-2021M7. https://fred.stlouisfed.org/series/EMRATIO

Figure 2. Employment-to-population ratio

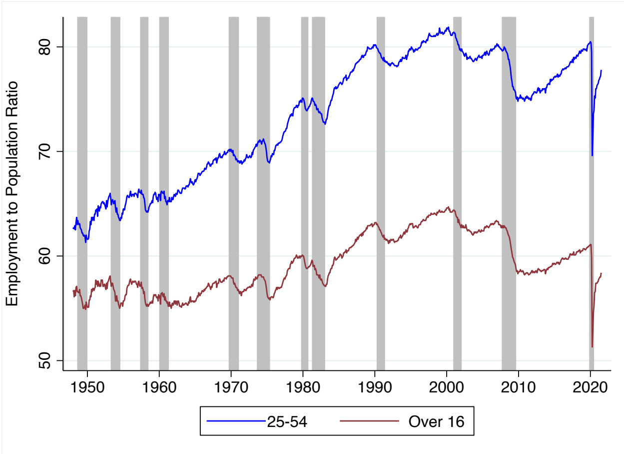

Employment to population ratios also come with their own measurement and interpretation problems. There is a strong upward trend in both series in the earlier years due to increases in female labor force participation. In addition, participation rates can be influenced by aging. After 1990 the 25-54 group seems to reach a plateau while the 16 and over displays a decreasing trend, possibly because of the effect of aging on participation rates. These trends make the analysis of the cyclical state of the labor market challenging.

Trends associated to aging of the population are commonly addressed by focusing on prime-age workers, 25 to 54 years old, for whom retirement decisions are less likely to be relevant. We will follow that strategy and focus on the employment to population ratio for this group.

But even for that group there is still a visible trend, in particular in the early decades. In order to remove the influence of this early trend we detrend the series by assuming a linear trend from January 1948 to July 1990 that corresponds with the peak of an NBER expansion, and it matches the time when the employment-to-population ratio starts to flatten out. Once we detrend the employment-to-population ratio for the 25-54 category we obtain the series plotted in Figure 3.

While our detrending makes it easier to see the return towards a level consistent with full employment, it does not deliver a cyclical shape that is as clean as in the case of the unemployment rate. Interestingly we can see larger swings in the employment to population ratio in recent cycles.

For the remaining of the paper we will refer to the series depicted in Figure 3 (detrended employment to population ratio for 25 to 54 years old) as the **employment rate**.

Comparing Figure 1 and Figure 3 we can see that the two measures of slack give us, at times, a different perspective on the labor market. As an example, in the 2007 peak the unemployment rate had fully recovered and was as low as in the previous peak (in 2001) but the employment rate was still far away from its previous peak. We will not make a call on which of these two series is a better indicator of the slack in the labor market and we will simply present results for both, providing two competing measures.

Figure 3. Detrended employment-to-population 25 to 54 years old

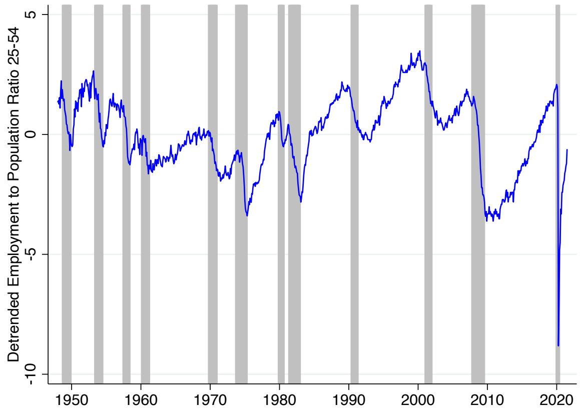

2.2. Dating the cycle

The analysis of this paper makes use of the framework that characterizes business cycles as the movement of the economy between two alternating phases of the cycle: expansions and recessions. This is the long-standing methodology used by the NBER, which produces dates based on the definition of turning points of economic activity, broadly defined. GDP is used as the "single best measure of aggregate economic activity" but other variables, such as payroll employment, are also considered "carefully".[7]

Because our focus is on the labor market a relevant question is how close the NBER dates are to turning points in either the unemployment rate or the employment rate. In their analysis of the cyclical properties of unemployment, Dupraz, Nakamura, and Steinsson (2019) calculate an alternative dating of the business cycle that matches better the peaks and troughs of the unemployment rate series. Their methodology looks for minima and maxima of the unemployment rate while avoiding small movements around these values. Their analysis reveals that while the two set of dates are quite close, there are some minor differences in the timing of those turning points.

We replicate and update the results of Dupraz, Nakamura, and Steinsson (2019) and also extend them by applying the methodology to the employment rate.[8] Figure 4 summarizes our results where DNS refers to the Dupraz, Nakamura, and Steinsson (2019) algorithm. Overall, the DNS methodology produces dates that are broadly consistent to the NBER dates with the exception of merging the 1980 and 1981 recessions into one. In addition, recessions tend to be longer as the labor market seems to turn before the NBER recession begins and does not pick up until a few months after the expansion starts.

The last recession, even if unusual, is picked up similarly by the three methodologies. The NBER declared February 2020 as the end of the expansion and April 2020 as the end of the recession, making it the shortest recession on record. Interestingly, the DNS methodology produces similar dates. When applied to the unemployment rate suggests that the expansion

[7] https://www.nber.org/business-cycle-dating-procedure-frequently-asked-questions
[8] We use a threshold value of 1% for the employment rate. Variations around this number have very little impact on the dating of the cycle.

ended in October 2019 and the subsequent recession finished by May 2020. When applied to the employment rate the recession almost matches perfectly the NBER call, as it starts in February 2020 and ends three months later.

Figure 4. Comparison of business cycle dates

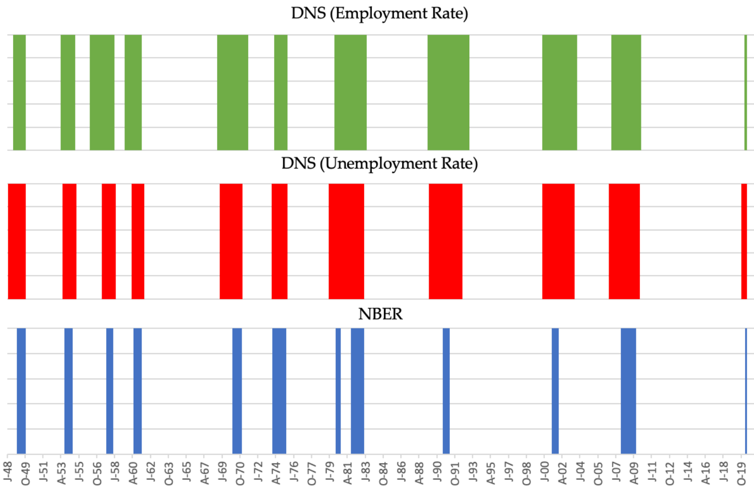

In our analysis we will be making comparisons of historical patterns across each of these cycles. As we do so, we need to acknowledge the limited number of observations that are available (between 10 and 11 cycles). In addition, the latest recession is an outlier when it comes to several dimensions. Because of this, in some instances we will highlight how the results are affected when we remove it from our analysis.

We start with some basic description of the length of the expansion and recession phases as defined by each of the three criteria (Table 1). We present results for both the full sample as well as the sub sample after the 1982 recession (to illustrate the effects of the Great Moderation). We include the expansion that was in place in January 1948 in our analysis, but we exclude the current ongoing expansion as it is not over.

Expansions have gotten longer since 1982 regardless of the methodology being used. Recessions have gotten shorter only when we use the NBER cycles (and this is partly the result of how short the last recession was). However, when using the DNS methodology, recessions are getting longer after 1982, even when including the influence of the very short recent recession. This increase is more marked when looking at employment rates given how long it took for this variable to increase during the expansions following the 1990 and 2007 recessions.

Table 1. Length of expansions and recessions (in months)

	NBER		DNS (Unemployment)		DNS (Employment Rate)	
	Full Sample	Post-1982M12	Full Sample	Post-1982M12	Full Sample	Post-1982M12
Recessions	10.3	9	25.1	30	27.3	33.2
Expansions	66.6	103.3	59.1	82.3	56.1	79.5

2.3. Asymmetries

The movements of labor market variables during the two phases of the business cycle are clearly not symmetric. The fast speed at which unemployment increases during recessions contrasts with the slow speed at which it returns to normal during expansions. This asymmetric nature of unemployment is clearly visible in Figure 1 and has often been discussed in the academic literature (Neftçi (1984), Dupraz, Nakamura, and Steinsson (2019), Jordà, Schularick, and Taylor (2020)).

A simple way to describe the asymmetry is to calculate the average speed of change in either unemployment rates or employment rates during expansions and recessions. Results for each of the dating criteria are presented in Table 2 and Table 3.

The destruction of employment during recessions happens at a speed that is much faster than the creation during expansions. The asymmetry is stronger when using the NBER dates because recessions are estimated to be shorter. Using the DNS algorithm, recessions include months where the destruction of jobs is less intense. There is no significant difference between using the unemployment or the employment rate.

Table 2. Absolute monthly average change in unemployment rate

	NBER		DNS	
	Full Sample	Post-1982M12	Full Sample	Post-1982M12
Recessions	0.339	0.508	0.176	0.185
Expansions	0.041	0.035	0.063	0.055

Table 3. Absolute monthly average change in employment rate

	NBER		DNS	
	Full Sample	Post-1982M12	Full Sample	Post-1982M12
Recessions	0.262	0.475	0.140	0.173
Expansions	0.030	0.026	0.056	0.051

The results of Table 2 and Table 3 are very much linked to the difference in duration of recessions and expansions. From Figure 1 and Figure 3 we know that the labor market tends to return to a similar level of activity towards the end of each expansion. Given the slower speed of job creation during an expansion compared to the speed of job destruction during a recession, it must be that expansions are longer than recessions. Using this logic, we can see that longer expansions after 1982 are a consequence of a decrease in the speed of labor market recoveries. It now takes longer to get to the same level of employment.

We can also explore further the fact that the depth of a recession does not seem to have an effect on the level of unemployment or employment that is achieved at the end of the expansion that follows. This means that if we measure the amplitude of a phase as the total change in unemployment or the employment rate, there must be a correlation between the amplitude of expansions and the depth of the previous recession. Recessions that damage the labor market by a larger amount must be followed by expansions that improve labor market outcomes by a larger amount as well.

A regression confirms this hypothesis (Table 4). When we regress amplitude of expansions against depth of previous recessions, we get a coefficient of about 1 with a high R-squared.[9] A coefficient of 1 suggests that expansions bring back the labor market to a similar state to the one existing before the previous recessions started. This is true regardless of the method used for dating recessions or the use of unemployment or employment rates.

Figure 5 displays the scatterplot of the data used in the first column of Table 4. The positive correlation is clearly present in the data.

Table 4. Explaining amplitude of expansions

	Unemployment		Employment Rate	
	NBER	DNS	NBER	DNS
Depth Previous Recession	1.124**	1.125***	1.406**	1.045***
	(0.353)	(0.307)	(0.459)	(0.195)
Constant	-0.320	-0.477	-0.750	-0.156
	(1.027)	(1.337)	(1.130)	(1.804)
Observations	11	10	11	10
R-squared	0.488	0.592	0.490	0.441

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Figure 5. Amplitude of expansions vs depth of previous recession

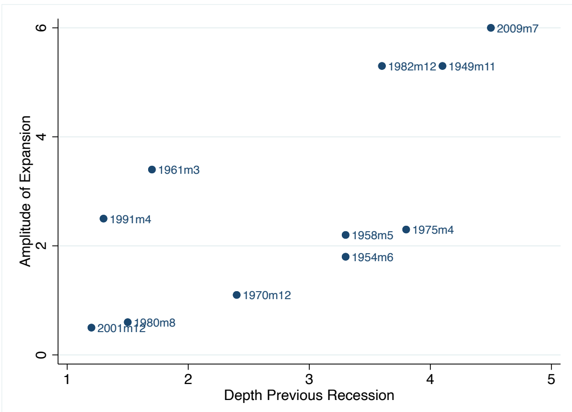

Interestingly, this result is not true the other way around, the depth of recessions is unrelated to the amplitude of the previous recovery. With a caveat: introducing the last recession in this analysis can create a spurious correlation.

[9] Amplitude of expansions is defined as the absolute change in the unemployment rate or the employment rate from the first to the last month of the expansion. Similarly, the depth of a recession is defined as the absolute change of either variable during the recession months.

In Table 5 we display the results of regressing the depth of a recession as a function of the amplitude of the expansion that preceded it. Columns (1) to (4) include the coefficients using the full sample. The coefficient tends to be positive and in the case of column (4) close to 1, although insignificant.

This patterns seems to contradict the result of Dupraz, Nakamura, and Steinsson (2019) that shows no correlation between the depth of a recession and the amplitude of the previous expansion. The reason for the divergence is that we are including one additional observation, that of the last recession. The last recession is the deepest on record and follows and expansion with the largest amplitude. Columns (1)† and (4)† reproduce the results of columns (1) and (4) after removing the last recession. The coefficient of column (1) has now turned even negative for the regression using unemployment rates and the NBER recessions dates, consistent with Dupraz, Nakamura, and Steinsson (2019). And it is close to zero for the regression using employment rates (column (4)†).

Table 5. Explaining depth of recessions

	Unemployment			Employment Rate		
	(1)	*(1)†*	(2)	(3)	(4)	*(4)†*
	NBER	*NBER†*	DNS	NBER	DNS	*DNS†*
Amplitude Previous Expansion	0.483	*-0.346*	0.556	0.556	0.920	*0.0194*
	(0.669)	*(0.222)*	(0.726)	(0.825)	(0.711)	*(0.279)*
Constant	2.085	*3.526****	2.346	1.643	1.060	*3.066***
	(1.405)	*(0.587)*	(2.226)	(1.097)	(1.766)	*(1.074)*
Observations	11	*10*	10	11	10	*9*
R-squared	0.109	*0.252*	0.126	0.100	0.299	*0.001*

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

†Excludes the last recession

Figure 6. Depth of recessions vs amplitude of previous expansion

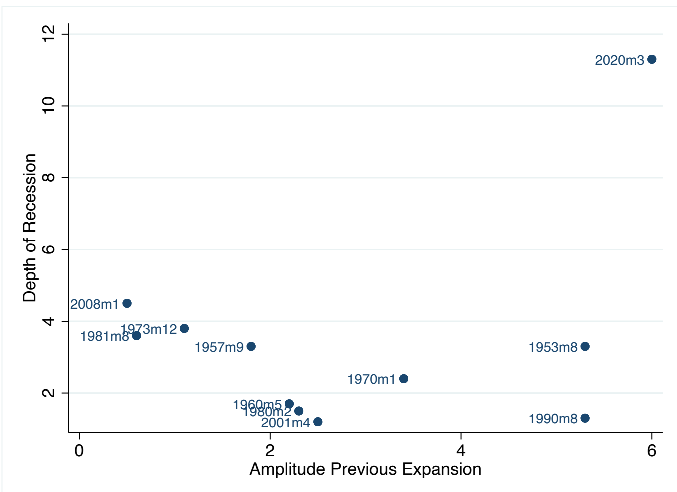

To illustrate how much of an outlier the last recession is, we plot in Figure 6 the data from the regression included in column (1). As it is clear from the figure, the correlation is mostly negative but adding the recession that started in March 2020 forces the regression line to show a positive slope.

For the purpose of this analysis on asymmetries (Table 4 and Table 5), we will treat the last recession as an outlier. We conclude that, from a labor market perspective, recessions are indeed independent of what happened in the previous expansions while expansions do depend on the depth of the previous recession. These results provide strong support for the Friedman plucking model of the business cycle (Friedman (1993), Dupraz, Nakamura, and Steinsson (2019)). In that model, recessions are sudden events that quickly deteriorate the economy and the labor market. The economy resembles a string that has been stretched as it moves away from its normal state. Later, as the expansion starts, we see a path of increasing employment and a return back to full employment.

3. The Short-lived Full Employment

3.1. The V-shape around recessions

The description of the labor market in the previous section suggests an asymmetric behavior that matches well the predictions of the plucking model of economic fluctuations. But the plucking model says very little about what comes after the economy reaches potential or full employment. In some sense we are simply waiting for the next recession.

But when we look at the behavior of the US labor market we see that these periods of full employment do not last long. Very little time passes before a new shock takes the economy away from those levels and back towards higher unemployment. This is a stylized fact that is not part of the plucking model and that the previous literature has not paid enough attention to.

To highlight these dynamics let's start with the last expansion that ended in February 2020. Unemployment declined after the recession that started in 2007. It did so for 128 months before it bounced back up, as a result of the 2020 recession (Figure 7).

Figure 7. Unemployment and employment rates during the last NBER Expansion

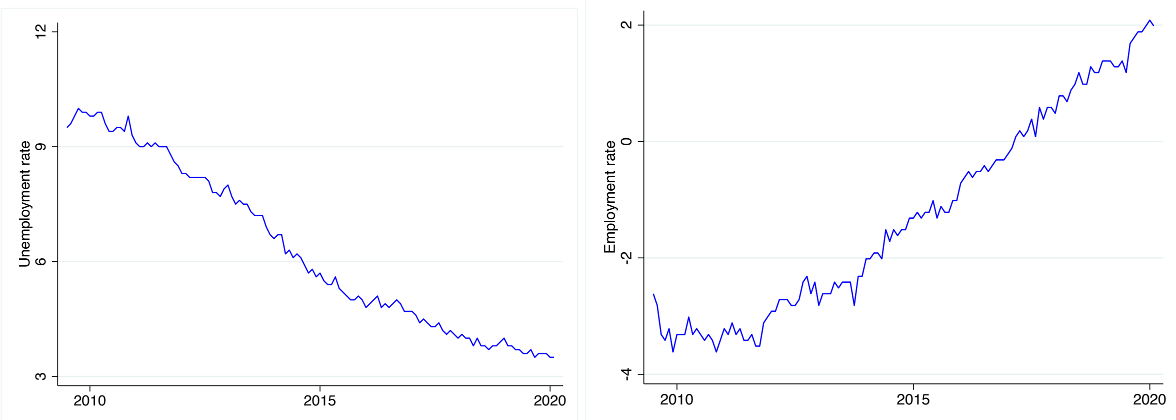

What is interesting in Figure 7, is that the unemployment rate was on a downward trajectory for practically every month during the expansion. And the employment rate was on an upward trajectory during the same period. Only in the most recent months we find some sense of stability in the case of the unemployment rate. One would expect that as we reach full employment both of these variables reach a plateau, and the economy remains in such a state for a while. It is even more remarkable that we do not observe such a plateau during the longest expansion on record. How long an expansion has to be so that we observe a long period of full employment?

There is nothing unique about this expansion, the same pattern applies to every previous US expansion. Figure 8 displays unemployment rates around the peak of each of the previous five expansions (where we use NBER dates for recessions). For easier comparison we calculate unemployment as deviations from the level achieved the month before the recession started (i.e. zero unemployment rate is associated to that month). We plot 5 years before the recession started and 10 months after the recession. Some series are shorter because either the previous expansion was shorter than 5 years or the recession lasted less than 10 months.

Figure 8. Unemployment Rate Around Recessions (US)

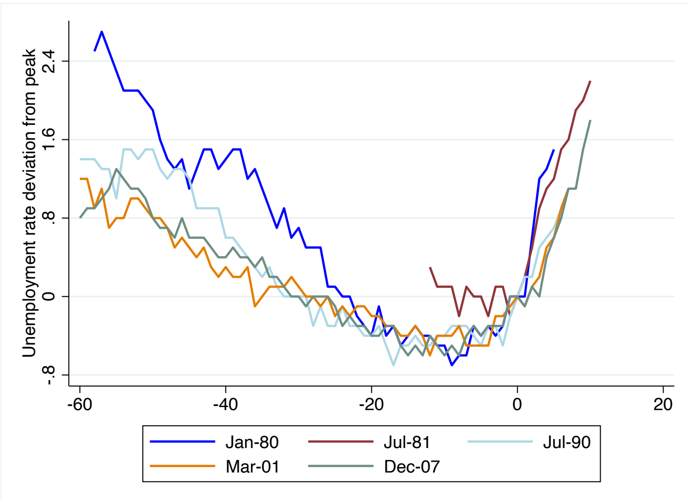

All cycles display a V-shape evolution for unemployment. Unemployment declines steadily in each of the expansions and it reaches its lowest point around 12-18 months before the recession starts. In most cases, unemployment was already increasing in the months preceding the recession. This relates to our statement earlier about how NBER turning points do not coincide with the turning points in unemployment rates. Unemployment rates turning points (as captured by our DNS methodology) are a few months before the NBER dates. But beyond the exact date of the turning point, what is important is the absence of a single long episode of stable low unemployment (or full employment). It seems as if reaching a low level of unemployment is *always* followed by a recession.

The V-shape that unemployment follows during each of the previous cycles is partly related to the "Sahm rule" to identify recessions and trigger automatic stabilizers in the real time.

Sahm (2019) proposes a rule based on the observation that once unemployment starts increasing, a recession follows. The rule uses a 3-month average of monthly unemployment rates and then calculated the difference between the current level of that average and the minimum over the previous 12 months. If this difference is larger than 0.5% then the US economy is heading into an imminent recession.[10] Sahm rule is based on the fact that do not see episodes of increasing unemployment that do not qualify as a recession. In other words, we do not see short-lived episodes of increasing unemployment that get quickly reversed, we do not observe "mini-recessions".

But our description of the V-shape around recessions goes beyond the facts behind this rule. In principle, unemployment could follow a wide U shape before and around recessions but, in this sample, we never see this shape.

One way to illustrate the significance of our results is to show that they are not universal. To see examples of different behavior we need to look at data from other countries. For example, Australia has sustained a low unemployment rate for decades until the 2020 recession. After a recession in the early 1990s, unemployment increased and then started a decline through a path similar to any US expansions. By the year 2000 unemployment reached a low level that has remained mostly flat for many years (Figure 9).[11] In other words, the unemployment rate in that cycle does not display V-shape dynamics but looks more like an open L-shape, waiting for the next recession. Similarly, the first recession in 1971 is preceded by a period of stable and low unemployment. Of course, some of the other cycles in Australia do resemble the US pattern.

Figure 9. Unemployment Rate (Australia)

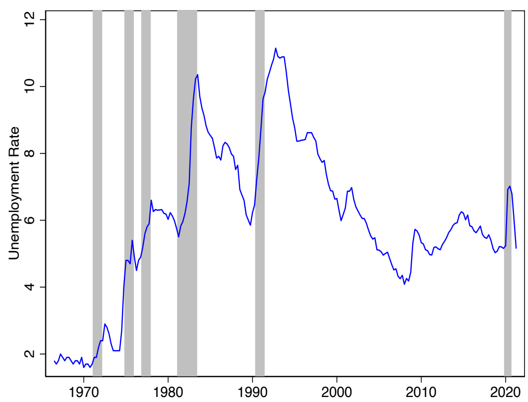

[10] The Sahm rule using real-time data can be found at https://fred.stlouisfed.org/graph/?g=qhzx The indicator crosses the red 0.5% threshold before each of the previous recessions. The only times the indicator gives a potential false positive is in Nov 1976 when the index reached exactly 0.5 but a recession had not started or was not about to start for another 3.5 years.

[11] We plot quarterly data on unemployment rate. Shaded areas correspond to recessions according to the dates provided by an updated version of Claessens, Kose, and Terrones (2012) provided by Ayhan Kose. We have added a recession starting in 2020Q2 and lasting for 2 quarters. Australia unemployment rates from the OECD short-term labor market statistics (https://stats.oecd.org). Sample 1966Q3-2021Q2.

Another international example comes from the evolution of unemployment right before the 2007Q4 recession. We compare in Figure 10 the evolution of US, UK and Ireland labor markets.[12] In all cases we calculate the deviation of the unemployment rate relative to 2007Q3, the last quarter of that expansion. In the case of Ireland and the UK we see four years of stable unemployment. In the US we see the familiar declining path.

These international examples are simply to illustrate the fact that the V shape that characterizes all the US cycles in the last five decades is not universal and that other shapes are possible.

Figure 10. Unemployment rate before 2007Q4 recession

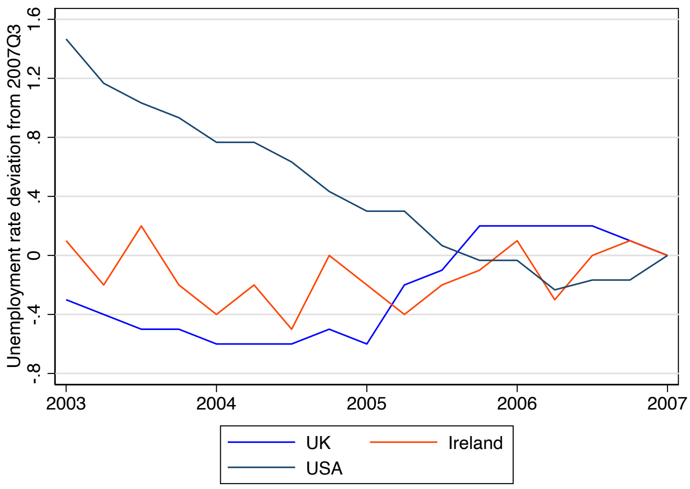

3.2. Does low unemployment cause recessions?

The V-shape that unemployment displays during recessions suggests that as unemployment decreases there is an increase in the probability of a recession. But how strong is this effect quantitatively? Can low unemployment be seen as a *strong* predictor of recessions?

We start with a simple OLS regression to see if the current level of unemployment can help predict future changes in US unemployment rates. Table 6, first column, shows the coefficients of running a regression of changes in unemployment rate over the next 36 months (three years) on the initial level of unemployment:

$$(U_{t+36} - U_t) = \alpha + \beta U_t + \varepsilon_t$$

The coefficient is negative and significant with a relatively high R-squared. We also replicate the same analysis using the detrended employment rate in the second column and the results are consistent. In both cases we see that periods where the labor market is strong tend to precede periods where the labor market deteriorates, and the other way around. The direction of this correlation should not be a surprise as unemployment rates are mean reverting, we do expect high levels of unemployment to precede decreases in

[12] Ireland and UK data on unemployment rates from the OECD short-term labor market statistics (https://stats.oecd.org).

unemployment while lower levels are likely to be observed before a crisis. But quantitatively the coefficient and the R squared are large. To establish a point of comparison, we look at other countries. In Table 7 we run the same regression as in Table 6 but now for Australia. We maintain the same horizon of three years. While we see also a negative coefficient, it is half the size and the R-squared is also smaller.

Table 6. Low unemployment (or high employment) rate as a predictor of increases in unemployment (or employment) rate (US)

US	$U_{t+36} - U_t$	$E_{t+36} - E_t$
U_t	-0.739***	-0.616***
	(0.025)	(0.029)
Constant	4.293***	-0.042
	(0.172)	(0.051)
Observations	847	847
R-squared	0.348	0.282

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 7. Low unemployment as a predictor of increases in unemployment (Australia)

Australia	$U_{t+12} - U_t$
U_t	-0.348***
	(0.0342)
Constant	2.301***
	(0.231)
Observations	208
R-squared	0.235

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

But in these regressions we are mixing positive and negative, small and large changes in the unemployment rate that make the interpretation of the results problematic. What we want to do is to focus on recessionary episodes. We do that in the next section.

3.3. A quantitative assessment of unemployment risk

In this section we focus on the cyclical risk of low levels of unemployment. When we talk about cyclical risk, we typically want to focus on recessionary episodes in which unemployment typically increases by a large amount at a fast pace.

Estimating the determinants of the tail risk of unemployment changes is related to the concept of value at risk used by financial institutions: focusing on the worst potential outcomes over a specific time window. The same analysis has been applied to business cycles by focusing on the tail risk of unemployment or GDP growth and we refer to this concept as GDP at risk (Cecchetti and Levin (2008), Kiley (2018), Adrian, Boyarchenko, and Giannone (2019)).

Empirically, we can isolate the effects on a particular part of the distribution by using quantile regressions. Unlike OLS that finds the best fit for the average of the changes in unemployment, quantile regression weights errors differently by putting larger weight on the errors near the quantile of the distribution that is of interest to the researcher. See Koenker and Hallock (2001) for details on the methodology and Kiley (2018) for a similar analysis of unemployment dynamics for the US.

We are interested in the tail risk of sharp unemployment increases, which are associated with recessions, and we will capture that by the 90th percentile coefficient. The question is whether low unemployment rates are followed by large increases in unemployment. We take as a starting point the results of Table 6 and we run a similar specification using quantile regressions. We switched to quarterly data because of additional variables we will include later in our analysis, but we maintain the three year horizon (also used in Kiley (2018)). In Table 8 we show, for the US, the results for three quantiles: the bottom 10%, the median (q50) and the 90th percentile of the distribution.[13] Not only the three coefficients are negative and significant but, more importantly, their size increases as we move from small changes in unemployment to large changes. This means that low unemployment rates are strong predictors of the tail risk of large increases in unemployment.

Table 8. Quantile regression unemployment (US)

	Dependent Variable: $U_{t+12} - U_t$		
Quantiles	q10	q50	q90
U_t	-0.509***	-0.744***	-1.029***
	(0.0349)	(0.0677)	(0.148)
Constant	1.319***	4.023***	8.402***
	(0.197)	(0.462)	(1.175)
Observations	282	282	282

Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Interestingly, the same phenomenon is not present in other countries. Table 9 shows the same analysis for Australia and the coefficient does not increase in size as we move towards the 90th percentile of the distribution. In fact, the coefficient becomes smaller and it is not significant for the 90th percentile. This suggests that, compared to the US, low unemployment rates are weak predictors of recessionary episodes in Australia (beyond the mean reversion effect present for all quantiles).[14]

[13] Estimation is done using simultaneous-quantile regression with standard errors calculated via bootstrapping (using 200 repetitions).

[14] If we apply the same analysis to other countries for which quarterly unemployment rates is available, we can see a combination of both of these patterns. Some resemble the US pattern where the coefficient on unemployment rates increases as we move from the 10th to the 90th percentile (e.g. Greece or Sweden). While in other cases the coefficients are of similar size across the different parts of the distribution (e.g. Germany or the UK). Results are available upon request.

Table 9. Quantile regression unemployment (Australia)

	Dependent Variable: $U_{t+12} - U_t$		
Quantiles	q10	q50	q90
U_t	-0.276***	-0.390***	-0.220
	(0.0202)	(0.0627)	(0.179)
Constant	0.397***	2.255***	3.693***
	(0.149)	(0.466)	(0.773)
Observations	208	208	208

Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

An alternative way to look at the predictive power of low unemployment rates for recessions is simply to include it in a probit model. We follow Cooper, Fuhrer, and Olivei (2020) and run a probit model to assess the power of unemployment rates as a predictor of recessions.

$$\Pr(NBER_{t+4} = 1) = \Phi(\beta_0 + \beta_1 U_t)$$

We use current US quarterly unemployment rates as a predictor of whether the economy is in a recession a year later (Table 10). Unsurprisingly, the results of the probit model confirm our earlier conclusion, that quarterly unemployment is a good predictor of recessions in the US.

Table 10. Predicting recessions

	$\Pr(NBER_{t+4} = 1)$
U_t	-0.323***
	(0.0795)
Constant	0.652
	(0.411)
Observations	294

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

3.4. Unemployment risk and imbalances

The US pattern of unemployment before recessions suggests that low levels of unemployment are a strong predictor of sudden increases in unemployment, those associated to recessions. But why is low unemployment unsustainable? What leads to a recession? The empirical literature on "Growth at Risk" analyzes factors leading to a recession and it emphasizes two set of variables: those associated to macroeconomic imbalances (such as inflation) and those associated to financial imbalances. For example, Kiley (2018) studies unemployment risk and uses credit growth, bond spreads and inflation as explanatory variables. Adrian, Boyarchenko, and Giannone (2019) study GDP growth risk and use the National Financial Conditions Index as their indicator of financial imbalances. There is an even larger literature that has looked at predictors of financial crisis (the literature has been recently surveyed by Sufi and Taylor (2021)). This literature has also emphasized the role of credit growth or asset price growth as a determinant of crises (Schularick and Taylor (2012)).

It is very likely that many of these risk factors are correlated with unemployment. As unemployment decreases, inflationary pressures might build. As the expansion lengthens, credit growth or asset price growth could become unsustainable. We make use of these additional variables to see if they can account for some or all of the observed correlation between low levels of unemployment and future recessions. We first introduce them in our quantile regression as additional risk factors. Once we control for these risks, do we still have low levels of unemployment predicting sudden and large increases in unemployment?

We include two variables in our analysis. We compute inflation as the (log) change in the PCE index over the last 4 quarters.[15] We then calculate the (log) change of the ratio of total private credit as a percentage of GDP over the last 8 quarters.[16] The reason to include two years for the growth of credit is that the results of Schularick and Taylor (2012) suggest the importance of the second (annual) lag of credit variables. We include each of these variables separately in our quantile regression in Table 11. Columns (1) to (3) present the results of including the previous 4 quarter level of inflation. The coefficient on inflation is positive (as expected) but does not show a particular pattern across the three quantiles we present. But interestingly, the coefficients on the unemployment level, while negative and still significant, are flatter than before. The coefficient does not increase as much from the q10 to the q90 quantiles. Inflation seems to be taken away some of the predictive power of low unemployment.

Table 11. Unemployment Quantile Regression with Controls (US).

	Dependent Variable: $U_{t+12} - U_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
Quantiles	q10	q50	q90	q10	q50	q90
Unemployment	-0.648***	-0.826***	-0.962***	-0.498***	-0.667***	-0.539***
	(0.0829)	(0.0363)	(0.165)	(0.0744)	(0.0549)	(0.158)
Inflation	28.38***	44.82***	31.26***			
	(5.519)	(5.427)	(9.666)			
Credit				0.237	5.775***	29.46***
				(2.427)	(2.127)	(5.358)
Constant	1.530***	3.154***	6.662***	1.273***	3.455***	4.393***
	(0.425)	(0.260)	(1.398)	(0.431)	(0.397)	(1.143)
Observations	278	278	278	274	274	274

Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

[15] Source: Personal Consumption Expenditures: Chain-type Price Index, Index 2012=100, Quarterly, Seasonally Adjusted. https://fred.stlouisfed.org/series/PCECTPI

[16] Source: Total Credit to Private Non-Financial Sector, Adjusted for Breaks, for United States, Percentage of GDP, Quarterly, Not Seasonally Adjusted: https://fred.stlouisfed.org/series/QUSPAM770A

This result is even stronger when we introduce the growth of credit over the last 8 quarters. Results are displayed in columns (4) to (6). Not only the coefficients on credit are significant, but they increase as we move from q10 to q90, confirming that fast credit growth is a strong risk factor for recessions.

But what is more interesting is that the introduction of the credit variable changes the pattern of the coefficients on unemployment levels. Unemployment still displays a negative coefficient, highlighting the reversion to the mean that was apparent in the simple OLS regression, but now the coefficients for the different quantiles are similar. In fact, the q90 coefficient is now smaller than the q50 and similar to the q10 one. Low levels of unemployment do not seem to be good predictors of the tail risk associated to recessions (beyond the fact that there is reversion to the mean at any point in time).[17]

Finally, Table 12 presents the results of including the two variables at the same time. The results are very similar to those of the previous table. The coefficient on the q90 unemployment quantile is now smaller than any of the previous two.

Table 12. Unemployment Quantile Regression with Controls (US). (cont.)

	Dependent Variable: $U_{t+12} - U_t$		
	(1)	(2)	(3)
Quantiles	q10	q50	q90
Unemployment	-0.659***	-0.792***	-0.477***
	(0.0807)	(0.0300)	(0.0936)
Inflation	25.31***	42.35***	29.12***
	(5.095)	(5.476)	(6.015)
Credit	-3.090	4.150***	29.64***
	(2.380)	(1.359)	(5.123)
Constant	1.770***	3.046***	3.066***
	(0.464)	(0.155)	(0.705)
Observations	274	274	274

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

We now check whether these insights are replicated when we include these variables in a probit model to predict recessions. The results are presented in Table 13. Both variables come with the right sign, inflation and fast credit growth both help predict recessions. We also include an additional variable in these regressions, one that has been used multiple times in predictive models of recessions: the slope of the yield curve. We measure it as the difference between a 10-year bond and the federal funds rate.[18]

[17] We have also reproduced these results including the National Financial Conditions Index as in Adrian, Boyarchenko, and Giannone (2019). The NFCI is positive and also significant, but it does not change the coefficients on the unemployment rate as much as credit growth does.

[18] Source: 10-Year Treasury Constant Maturity Minus Federal Funds Rate, Percent, Quarterly, Not Seasonally Adjusted (T10YFFM). https://fred.stlouisfed.org/series/T10YFFM

Table 13. Predicting recessions with controls

	$\Pr(NBER_{t+4} = 1)$			
U_t	-0.437***	-0.290***	-0.229**	-0.297**
	(0.0800)	(0.0792)	(0.0956)	(0.120)
Inflation	19.90***			7.341
	(4.517)			(9.366)
Credit		0.943		8.925**
		(2.111)		(3.945)
Yield Curve			-0.459***	-0.400***
			(0.0797)	(0.146)
Constant	0.531	0.430	0.293	-0.00773
	(0.417)	(0.430)	(0.502)	(0.530)
Observations	290	284	268	266

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

In all the regressions presented, unemployment remains a significant predictor although as controls are introduced the size of the coefficient decreases. Overall, these results suggest that during episodes of low unemployment, we also tend to observe higher inflation or faster credit growth. Once we account for these variables, the predictive power of low unemployment diminishes, although it does not completely disappear in all the specifications. These results are consistent with the literature that studies factors that precede recessions. If credit growth is indeed important to understand crises, it is likely that we observe faster credit growth in the later years of an expansion phase when unemployment is low.

4. The Slow March Towards Full Employment

4.1. Imbalances or external forces as drivers of recessions

Our results so far have shown that as unemployment decreases, there is an increasing risk that a recession will happen. The quantile and probit regressions have demonstrated that some of this risk is related to the existence of imbalances such as inflation or fast credit growth, imbalances that grow as the expansion phase becomes mature. But imbalances are not the only source of recessions, external forces also matter. We do not need to go too far to prove this point: a pandemic in 2020 precipitated a global recession.[19] In addition, the academic literature on business cycles heavily relies on the notion of external, unexpected, shocks as a source of cycles.

If we ignore the risk caused by imbalances, and assume that the only cause of recessions are random shocks, how likely is that an economy never displays long periods of stable and low unemployment? Mechanically the answer depends on how frequent those shocks are but also on the speed at which the labor market returns back to normal. Even if shocks are rare, a slow healing labor market could fail to deliver significant periods of full employment.

[19] Of course, one can still debate when a US recession would have happened in the absence of the pandemic.

Analyzing the speed of labor market recoveries is also important even if imbalances are the ultimate cause of recessions. Then we see a race between the buildup of imbalances and the recovery in the labor market. What we want is a fast recovery that does not lead to the imbalances that could put growth at risk

This is the focus of this section: what is the normal speed of recovery in the labor market? Has it changed over time? Is it faster following some particular recessions?

We start by calculating some basic statistics on the average speed of labor market improvements during each of the past expansions. We use both the change in the unemployment rate and the change in the employment rate. And, as before, we use both the NBER dates as well as the dates calculated using the DNS methodology. We ignore the last expansion because it is ongoing.

The results are presented in Figure 11 and Figure 12. The results are consistent across the four charts although there is some interesting variation depending on the dating criteria we use for expansions. A pattern that is consistent across the four charts is that the early expansions saw faster improvements in labor markets. Even if we focus on the last five, the last three expansions have consistently slower speeds than the previous two.

Figure 11. Speed of labor market improvement (NBER expansions)

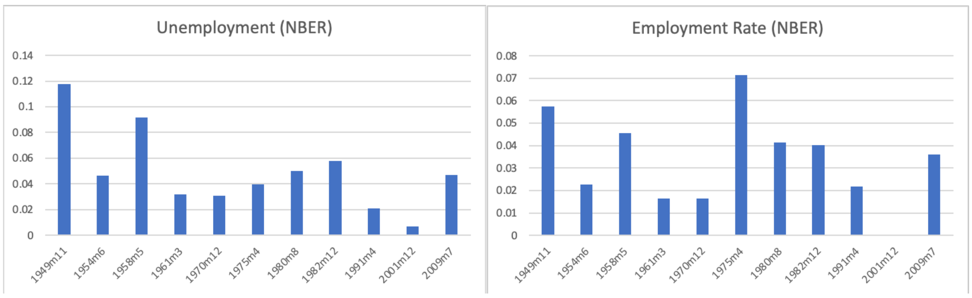

Figure 12. Speed of labor market improvement (DNS expansions)

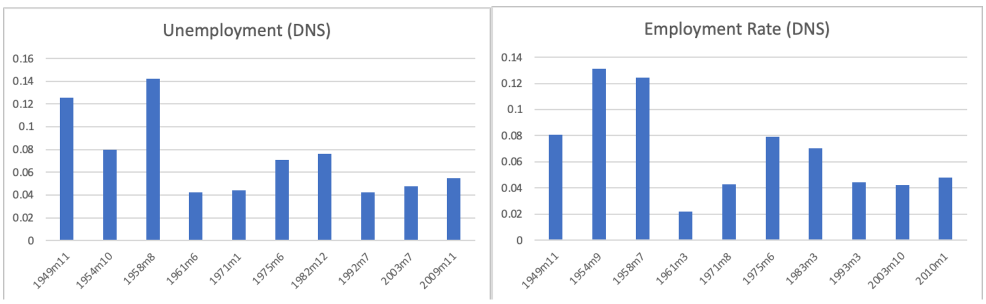

But what is behind these differences? Could it be that the speed is faster in expansions that start with a higher unemployment rate? And what is the pattern of this speed as the expansion gets older? Do we see the speed slowing down?

Hall and Kudlyak (2021) and Hall and Kudlyak (2020) have looked in detail at this question using the unemployment rate and have reached two important conclusions. First, the speed of recoveries is very consistent across different expansions, despite the differences in those cycles. Second, the speed of recovery is too linear during the expansion months. In Figure 13, left panel, we reproduce one of the figures in their paper using the unemployment rate and the DNS dates for cycles. We also extend their analysis by looking at the employment rate in the right panel. In both cases we see similar speeds of recovery across expansions. And the speed seems to be fairly constant over time. This is at odds with standard theory that would suggest that at the beginning of the expansion we should see faster improvements than later in the expansion.

Figure 13. Speed of labor market recovery during expansions (DNS dates).
Unemployment rate (left) Employment rate (right)

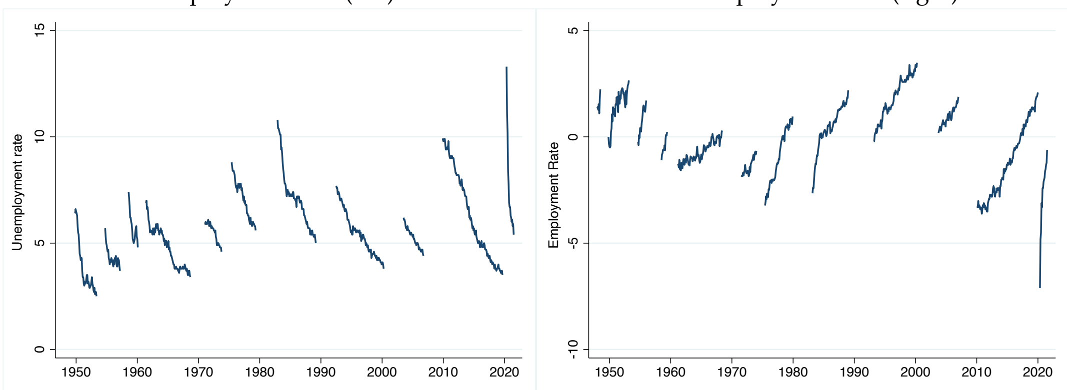

In the next section we extend the analysis of Hall and Kudlyak (2021). We add the employment rate into the analysis, and we check whether the speed of labor market improvement during expansions depends on the initial state of the labor market (which is an outcome of the depth of the previous recession). We also quantify how the speed of recovery changes as time passes and the expansion becomes more mature and compare it to the predictions of a standard model of the labor market.

4.2. Speed of recovery and depth of recessions.

We first address the question of whether the speed of recovery is faster after particularly deep recessions. We characterize these recessions using two different metrics. Either the level of unemployment or employment rate when the recession ends or by the change in unemployment or employment rate during recession phase. Both the level and the overall change are likely to be correlated but the second indicator potentially allows for changes in the natural rate of unemployment over time.

Table 14 shows the result when we use the initial level of unemployment or the employment rate to characterize the state of the labor market when the expansion starts. We then regress the average speed of a recovery, measured by the absolute average decrease in unemployment rate or the average increase in the employment rate during the recovery that follows.

The results show that there is no obvious correlation. Except when we use detrended employment rates and we look at the NBER definition of rates. But even there, the size of the coefficient is small. A recession that ends with an extra 1% lower employment rate, is likely to lead to an increase in the speed of the recovery of about 0.11 percentage points per year. Remember that the typical speed of recovery is around 0.4 per year.

Table 15 presents a similar exercise. We compare the speed of a recovery to how large the deterioration of the labor market was in the previous recession. This comparison might be more appropriate if the natural rate of unemployment or the steady state employment rate have changed over time. Once again most of the coefficients are not significant. The only case where we find a significant coefficient is when we look at unemployment and we use the NBER dates for business cycles, first column.

Table 14. Speed of recovery and initial conditions

	Unemployment		Employment Rate	
	NBER	DNS	NBER	DNS
Initial Conditions	0.00706*	0.00205	-0.00962**	0.00332
	(0.00377)	(0.00390)	(0.00325)	(0.00738)
Constant	-0.00414	0.0565	0.0229***	0.0738***
	(0.0294)	(0.0368)	(0.00560)	(0.0212)
Observations	11	10	11	10
R-squared	0.132	0.009	0.396	0.015

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 15. Speed of recovery and depth of previous recession

	Unemployment		Employment Rate	
	NBER	DNS	NBER	DNS
Depth Previous Recession	0.0163*	0.0121	-0.0111	-0.00324
	(0.00721)	(0.00836)	(0.00746)	(0.0113)
Constant	0.00363	0.0275	0.0115	0.0585
	(0.0163)	(0.0279)	(0.0143)	(0.0381)
Observations	11	10	11	10
R-squared	0.390	0.157	0.179	0.008

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

However, this significant correlation is driven by the earliest two cycles in our dataset, the expansions that started in 1949M11 and 1958M5 (Figure 14). For the rest of the cycles, the speed of the expansion is remarkably stable regardless of the depth of the previous recession, a result highlighted in Hall and Kudlyak (2021).

The fact that deeper recessions do not affect the speed of the recovery that follows might seems to be in contradiction with our earlier result stating that the amplitude of unemployment changes during an expansion was correlated with the depth of the recession (a result that supported the plucking model of the business cycle). There is no contradiction. Our new result says, using Friedman (1993) plucking analogy, that as the string is stretched it bounces back towards its natural state but the speed at which it does is unrelated to the distance at which the string was stretched. This means that it just takes longer to climb back to full employment after a deeper recession. Or, in other words, deeper recessions would need to be followed by longer expansions if the labor market was to return to its full employment levels. The expansion that followed the 2008 crisis might be a perfect example of that.

Figure 14. Speed of Expansions and Previous Recession (Unemployment)

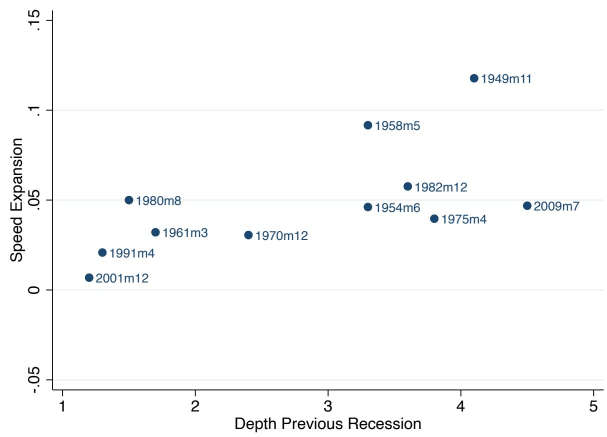

Another way to present the same intuition is to think about the labor market conditions that prevail at the end of a given expansion. During the last part of an expansion we see unemployment rate at the lowest level and this level is sometimes associated to the natural rate of unemployment (similar logic applies to the employment rate). But if the speed of recovery is too slow and expansions end too early, we might not see the true level of unemployment or employment that is consistent with full employment. Under this hypothesis, the unemployment rate at the end of an expansion would be a function of how bad the previous recession was and how long the expansion was.

Table 16. Explaining end of expansion labor market

	Unemployment		Employment Rate	
	NBER	DNS	NBER	DNS
Initial Conditions	0.411**	0.434***	0.0183**	0.550*
	(0.176)	(0.121)	(0.00606)	(0.256)
Duration	-0.0206*	-0.0219***	0.239	0.0234*
	(0.00901)	(0.00516)	(0.143)	(0.0119)
Constant	3.049*	2.135*	-0.0414	0.927*
	(1.365)	(0.973)	(0.408)	(0.449)
Observations	11	10	11	10
R-squared	0.427	0.513	0.510	0.413

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The data is consistent with our hypothesis and the results are displayed in Table 16. We regress the unemployment rate or the employment rate at the end of an expansion on the value it had at the beginning of the expansion and the duration of months of the expansion. The logic is that if we start with a more deteriorated labor market it will take longer to reach a particular level of low unemployment. Both variables have the right sign and are almost always significant at least at 10%. When we start with a high unemployment rate, everything else equal, we will end the expansion with also a higher unemployment rate. And only expansions that are long enough, manage to heal a broken labor market. This is fully consistent with our earlier result showing that the speed of recovery is not a function of the initial state of the labor market.

Notice that the R squared is as around 50% for the four specifications. We can explain more than 50% of the state of the labor market at the end of an expansion by using the initial level as well as its duration. This is a strong sign that some expansions are unfinished business and that estimating the equilibrium state of the labor market by assuming that the end of expansions represent their natural states is likely to bias our analysis towards a pessimistic view of what outcomes are feasible in the labor market.

4.3. Recovery speed during the expansion phase.

We have seen that the speed of recovery during expansions is not strongly correlated with how deep the previous recession was. We now look inside of each of the expansions and see whether the speed varies over time as the labor market normalizes and we move further away from the end of the recession. Hall and Kudlyak (2021) have emphasized the linear nature of the labor market recovery during expansions. But is the speed truly constant during the expansion phase? Does the same result apply when we look at employment rates?

We analyze how the annual change in the state of the labor market, measured by changes in the employment rate or the unemployment rate, depends on two variables that capture the age of an expansion: the number of months the expansion has lasted and the lagged levels of unemployment or employment rate.

Table 17 presents the results for the unemployment rate. Overall, we observe that as time passes the speed of recovery decreases. The speed (measured per month) is reduced by about 0.006684 every 10 months, compared to the average speed of recovery of about 0.04. Once we introduce the lagged unemployment rate, this mechanical effect disappears, consistent with the idea that it is the unemployment rate that matters for the job finding rate, not simply the age of the expansion.

When we do the same analysis for the employment rate (Table 18), the results are similar. Overall, the speed of recovery slows down as time passes and once we control for the lagged employment rate the age of the expansion does not matter anymore. If we use the coefficient from column (3), the change in the speed of recovery every 12 months would be around 0.0038, compared to an average of 0.063, a smaller relative magnitude relative to the results using the unemployment rate. One way to think about the difference in results is that as time passes, the employment rate does not decrease as fast because we see an inflow back into the labor force.

Table 17. Speed of recovery during expansion (Unemployment Rate)

	Full Sample				Excludes Last Recovery			
	NBER		DNS		NBER		DNS	
Months	0.000986***	-9.86e-05	0.00137***	0.000159	0.000557***	0.000192	0.000811***	0.000429
	(0.000254)	(0.000312)	(0.000326)	(0.000338)	(0.000190)	(0.000238)	(0.000236)	(0.000275)
U_{t-1}		-0.0355***		-0.0380***		-0.0124**		-0.0127**
		(0.0122)		(0.0129)		(0.00518)		(0.00536)
Constant	-0.0973***	0.155**	-0.128***	0.138*	-0.0672***	0.0202	-0.0942***	-0.00661
	(0.0171)	(0.0783)	(0.0190)	(0.0811)	(0.0125)	(0.0380)	(0.0134)	(0.0385)
Observations	748	748	606	606	733	733	591	591
R-squared	0.023	0.077	0.032	0.098	0.011	0.021	0.018	0.029

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 18. Speed of recovery during expansion (Employment Rate)

	Full Sample				Excludes Last Recovery			
	NBER		DNS		NBER		DNS	
Months	-0.000487*	0.000453	-0.000847**	-8.48e-05	-0.000103	0.000244	-0.000337	-0.000192
	(0.000275)	(0.000313)	(0.000361)	(0.000342)	(0.000220)	(0.000278)	(0.000279)	(0.000325)
$ERate_{t-1}$		-0.0358***		-0.0305**		-0.0142**		-0.00646
		(0.0123)		(0.0134)		(0.00582)		(0.00579)
Constant	0.0622***	0.0204	0.0996***	0.0659***	0.0351**	0.0203	0.0679***	0.0618***
	(0.0179)	(0.0165)	(0.0201)	(0.0169)	(0.0138)	(0.0162)	(0.0147)	(0.0168)
Observations	748	748	576	576	733	733	561	561
R-squared	0.005	0.052	0.011	0.051	0.000	0.010	0.002	0.005

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

These two tables show that the speed of labor market recoveries during recessions is almost constant but there is a small deceleration over time. The labor market does improve faster when unemployment rates are higher. But how does this effect compare to the prediction of a standard labor market model? We base on our analysis on the simplest version of the Diamond, Morten and Pissarides (DMP) model of search and matching in the labor market (see Hall and Kudlyak (2021) or Mercan, Schoefer, and Sedláček (2021) for a more detailed analysis of the predictions of the model).

In such a model, the job finding rate depends on the tightness of the labor market defined as the ratio of vacancies (v) to unemployment (u)

$$\theta = \frac{v}{u}$$

Assuming a Cobb-Douglas matching function

$$m(u, v) = \mu u^{(1-\gamma)} v^{\gamma}$$

Where μ is the matching efficiency and γ is the matching elasticity, the job finding rate takes the form

$$f(\theta) = \mu\theta^{1-\gamma}$$

In a standard market without inefficiencies and where wages are set to split the surplus of the bargaining so that productivity and wages move together, the tightness of the labor market will be constant and equal to steady state (θ_{ss}), same as the job finding rate ($f(\theta_{ss})$).

Under that assumption, the dynamics of unemployment can be summarized by

$$u_t = \left(1 - f(\theta_{ss})\right)u_{t-1} + \delta_t(1 - u_t)$$

Where δ_t is the separation rate at time t.

If we assume a one-time shock to the separation rate so that it is higher for just one period and then returns to steady state level (δ_{ss}}, the expression above for any period after the shock can be rewritten as

$$u_t - u_{t-1} = \delta_{ss} - (f(\theta_{ss}) + \delta_{ss})u_{t-1}$$

Using the calibration of Mercan, Schoefer, and Sedláček (2021) for the US, $f(\theta_{ss}) \approx 0.143$ and $\delta_{ss} \approx 0.011$, we get a coefficient on lagged unemployment equal to 0.154. This number is about 10 times larger than the estimates of Table 17 for the sample without the last recession and about 5 times larger than the estimates for the full sample.

What our estimates show is that while the speed of recovery is larger in the earlier months of an expansion, the magnitude is much smaller than what is predicted by the standard DMP model. The DMP model predicts a strong recovery in the earlier months that decelerates fast as the labor market improves. In other words, job market healing is too linear and too slow in the early days of the expansion.

In order to match the data, we need to introduce frictions that make the tightness of the labor market cyclical, which will make the expression above nonlinear. For higher levels of unemployment the degree of tightness will be lower leading to a lower job finding rate and therefore a slower recovery in the first months of the expansion.

Hall and Kudlyak (2021) provide an extensive discussion of alternative models that could explain the countercyclicality of the tightness of the labor market. One of the potential avenues is to think about congestion in the labor market. Mercan, Schoefer, and Sedláček (2021) follow that approach and develop a model of congestion in the labor market that generate similar dynamics to what we observe in the data. Congestion means that there is a limit on the speed at which companies can absorb workers (or an increasing cost for a higher hiring rate). This is not only consistent with our results but also with other empirical studies. For example Mercan and Schoefer (2020) show how shocks to local employment among a selected group of firms does not lead to a corresponding hiring by the firms not being affected.

If labor markets display congestion, we could observe differential speeds of recoveries across different groups of workers. There is evidence that when unemployment is high employers seem to demand greater skills for the same type of jobs, at the cost of lower-skills workers (Modestino, Shoag, and Ballance (2020)). Within this framework Cajner, Coglianese, and Montes (2021) or Mason, Konczal, and Melodia (2021) analyze the slow recovery of employment or unemployment rates after recessions and suggest that differences in the speed at which different segment of the labor market recover during an expansion can explain such a slow speed. They reach that conclusion because indicators of structural unemployment are procyclical. For example, differences in unemployment rates among different groups in the labor market are the widest right after the recession and they narrow over time as the recovery gains momentum (Aaronson et al. (2019)).

Our results as well as this related literature points towards a labor market that returns slowly towards equilibrium. Where higher levels of slack do not generate a much higher rate of reduction in unemployment and the overall speed of healing is too slow compared to a frictionless model. Under this scenario, deep recessions require longer expansions to return to full employment and supports the main result of this paper regarding the absence of long periods of stable and low unemployment.

5. Conclusions and Policy Implications

Our analysis of US business cycles has uncovered a fact that has been overlooked by the previous literature: low levels of unemployment are short lived. This result raises questions about the notion of a well-defined natural or equilibrium state for the labor market, at least to the extent that such a state is anything but "natural" from the statistical point of view. Our empirical results are consistent with an asymmetric view of the business cycle where recessions are downward deviations from potential output, what is known as the plucking model of the business cycle. Recessions can then be seen as large or small disasters interrupting the natural growth path of activity (Jordà, Schularick, and Taylor (2020)). But we add a twist to these models, the US economy spends very little time close to potential output. When an expansion is mature, as measured by a low unemployment rate or a high employment rate, the probability of a recession increases fast. The outcome is an economy that cannot enjoy the benefits of a high-pressure labor market for long periods of time.

We have explored two complementary hypotheses to explain this phenomenon. First, it is possible that as the expansion matures, the economy builds imbalances that represent the seed of a recession. Our analysis using quantile regressions or predictive models of recessions produces support for this hypothesis. But this hypothesis cannot be enough to explain all cycles. Some recessions are driven by external shocks, such as the 2020 recession driven by a global pandemic. And what is interesting about this recession is that while it followed the longer expansion on record it was unable to deliver a long period of full employment. This raises a second hypothesis that we explore in our paper: the speed of labor market improvements during an expansion is too slow to take the economy back to full employment in a reasonable period of time. This is consistent with models of the labor markets where frictions or congestion effects limit the speed at which workers can be absorbed back into employment, generating a quasi-constant speed during the expansion phase.

The slow speed of labor market improvement has implications for both of our hypotheses. Recessions end in two ways, either because they get old and create a large enough imbalance or because an external shock happens. In either case, the slow speed of improvement in the labor market reduces the likelihood of seeing periods of low unemployment. It can also lead to overestimation of the lowest rate of unemployment that are feasible or consistent with a stable inflation rate. For example, if a global pandemic had happened in 2017, we would have seen a recession three years earlier and our empirical estimates of the natural rate of unemployment would be higher as a result.

Under this interpretation, periods of higher unemployment, as in the 1970s, should not be seen as periods where structural unemployment increased but, instead, as years with more frequent shocks and shorter expansions that never delivered sustained periods of low unemployment. We provide evidence that this is the case by showing that more than 50% of the variation in unemployment rates at the end of an expansion can be explained by the state of the labor market at the beginning of the expansion and the length of the expansion phase.

From a policy point of view there are two fundamental lessons. Anything that policy can be done to accelerate the recovery of the labor market, in particular in the early years, will result in longer periods of full employment where we can benefit from a high-pressure economy. Second, macroprudential policy should ensure that the faster growth does not lead to the types of imbalances that can cut an expansion short. None of these lessons are completely new to policy markets, but the results of this paper make explicit the cost of inaction. Expansions that end too early are unfinished recoveries that keep the economy away from its potential. This strengthens the argument for policies that focus on ensuring quicker recoveries after every recession.

The exact form that those policies should take depends on the reasons why the economy is slow to recover. This paper has not explicitly explored this issue. If the issue is one of congestion in the labor market, policies should be focused on removing those barriers or creating the necessary financial incentives to speed up the rate of hiring. If the slow speed is related to firms being reluctant to hire due to uncertainty about the level of demand, more aggressive demand policies early in the recovery phase might be the right tool.

6. References

Aaronson, Stephanie R, Mary C Daly, William L Wascher, and David W Wilcox, 2019, Okun revisited: Who benefits most from a strong economy?, *Brookings Papers on Economic Activity* 2019, 333–404.

Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone, 2019, Vulnerable Growth, *American Economic Review* 109, 1263–89.

Amior, Michael, and Alan Manning, 2018, The persistence of local joblessness, *American Economic Review* 108, 1942–70.

Blanchard, Olivier, 2018, Should we reject the natural rate hypothesis?, *Journal of Economic Perspectives* 32, 97–120.

Cajner, Tomaz, John M Coglianese, and Joshua Montes, 2021, The Long-Lived Cyclicality of the Labor Force Participation Rate, *FEDS Working Paper*.

Cecchetti, Stephen G., and Andrew Levin, 2008, Measuring the macroeconomic risks posed by asset price booms, *Asset prices and monetary policy* (University of Chicago Press).

Claessens, Stijn, M Ayhan Kose, and Marco E Terrones, 2012, How do business and financial cycles interact?, *Journal of International economics* 87, 178–190.

Cole, Harold L, and Richard Rogerson, 1999, Can the Mortensen-Pissarides Matching Model Match the Business-Cycle Facts?, *International Economic Review* 40, 933–959.

Cooper, Daniel H, Jeffrey C Fuhrer, and Giovanni P Olivei, 2020, Predicting Recessions Using the Yield Curve: The Role of the Stance of Monetary Policy, *Federal Reserve Bank of Boston*.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson, 2019, A plucking model of business cycles, *NBER Working Paper* 26351, 28.

Fatas, Antonio, and Ilian Mihov, 2013, Recoveries, *CEPR Discussion Papers*.

Friedman, M., 1993, The plucking model of business fluctuations revisited, *Economic Inquiry* 31, 171–177.

Hall, Robert E, and Marianna Kudlyak, 2020, The Inexorable Recoveries of US Unemployment, National Bureau of Economic Research.

Hall, Robert E, and Marianna Kudlyak, 2021, Why Has the US Economy Recovered So Consistently from Every Recession in the Past 70 Years?, *NBER Macroeconomics Annual*.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor, 2020, Disasters Everywhere: The Costs of Business Cycles Reconsidered, National Bureau of Economic Research.

Kiley, Michael T., 2018, Unemployment Risk. Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System (US).

Koenker, Roger, and Kevin F. Hallock, 2001, Quantile Regression, *Journal of Economic Perspectives* 15, 143–156.

Mason, J.W., Mike Konczal, and Lauren Melodia, 2021, Reimagining Full Employment: 28 Million More Jobs and a More Equal Economy, Roosevelt Institute.

Mercan, Yusuf, and Benjamin Schoefer, 2020, Jobs and matches: Quits, replacement hiring, and vacancy chains, *American Economic Review: Insights* 2, 101–24.

Mercan, Yusuf, Benjamin Schoefer, and Petr Sedláček, 2021, A Congestion Theory of Unemployment Fluctuations, National Bureau of Economic Research.

Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance, 2020, Upskilling: Do employers demand greater skill when workers are plentiful?, *Review of Economics and Statistics* 102, 793–805.

Neftçi, Salih N., 1984, Are Economic Time Series Asymmetric over the Business Cycle?, *Journal of Political Economy* 92, 307–328.

Okun, Arthur M., 1973, Upward mobility in a high-pressure economy, *Brookings Papers on Economic Activity* 1973, 207–261.

Rudebusch, Glenn D., 2016, Will the Economic Recovery Die of Old Age? Federal Reserve Bank of San Francisco Economic Letters.

Sahm, Claudia, 2019, Direct Stimulus Payments to Individuals, *Recession Ready: Fiscal Policies to Stabilize the American Economy*.

Schularick, Moritz, and Alan M. Taylor, 2012, Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008, *American Economic Review* 102, 1029–61.

Sufi, Amir, and Alan M Taylor, 2021, Financial Crises: A Survey, *NBER Working Paper*.

Yellen, Janet L., 2016, Macroeconomic research after the crisis, *Board of Governors of the Federal Reserve System*.