

SPEECH



DATE: 29/05/2018
SPEAKER: Deputy Governor Henry Ohlsson
VENUE: Örebro University

SVERIGES RIKSBANK
SE-103 37 Stockholm
(Brunkebergstorg 11)

Tel +46 8 787 00 00
Fax +46 8 21 05 31
registratorn@riksbank.se
www.riksbank.se

Data, correlations and economic policy challenges*

At the beginning of the Riksbank's monetary policy reports, there is a page describing the monetary policy strategy. The basis of the strategy is that the Riksbank shall strive towards an annual change of 2 per cent in the consumer price index with a fixed interest rate (CPIF). Given that there is confidence in the inflation target, monetary policy shall also support good growth in output and employment. The Riksbank therefore conducts what is generally referred to as flexible inflation targeting.

This is the basic framework for monetary policy. Within this framework, the Riksbank must cope with the fact that the economy and the way it works change over time. My intention here today is to discuss some of the challenges facing monetary policy and economic policy in general as a result of changes that have occurred over the last decade. In brief, it is a question of surprisingly subdued wage increases, the apparent difficulty in reducing unemployment further via expansionary policy alone, and the apparent deterioration in matching on the labour market.

I intend to discuss this with the help of a simple survey, in which I look at how five key variables covary. The time period I'm studying is 2007–2018 and I'm using monthly data. The five key variables are:

- inflation, measured as the CPIF, the target variable for monetary policy
- wage increases, according to short-term wage statistics from the National Institute of Economic Research (NIER)
- unemployment, according to the Labour Force Surveys (LFS)
- job openings, according to statistics from the Swedish Public Employment Service (Arbetsförmedlingen) on remaining job openings

* I would especially like to thank Mikael Apel and Emelie Theobald for their help with this speech. Meredith Beechey, Charlotta Edler, Jesper Hansson, Ann-Leena Mikiver, Åsa Olli Segendorf, Cecilia Roos-Isaksson and Ulf Söderström have all submitted valuable comments.

- the krona exchange rate, according to the trade-weighted KIX index.

In order to focus on the trends in these five variables, I remove seasonal variations and general noise in the series by using moving averages over twelve months throughout. We will follow the time series for these variables two at a time. By doing so, we can see the covariation between inflation and unemployment, for example.

Looking at data is a good place to start

It is important to understand that covariation between two variables does not necessarily mean that there is a causal relationship, a causality, between them. Instead, the covariation may, for example, depend on a third variable that is not shown. A classic example is the strong correlation between sleeping with your shoes on and waking up with a headache. But this presumably does not imply that the shoes cause the headache. There is instead an underlying variable explaining why a person both sleeps with their shoes on *and* wakes up with a headache.

We should also be aware that when plotting two variables against each other, it is very rare to expect to find a clear-cut correlation of the type we see in textbook illustrations. The reason is, of course, that the variables are not just linked to each other but are also constantly influenced by a number of other factors. For example, inflation can be affected by changes in energy prices that do not have any significant impact on unemployment. Unemployment can, in turn, be affected by reforms on the labour market, for example, which do not have any significant effect on inflation, at least not immediately. In other words, we cannot expect to see an exact relationship but often have to be content with just barely being able to detect one.

On the other hand, the absence of a correlation that is visible to the eye does not mean that there isn't one. It may still be there and be detectable only if we control for variables other than just the two we are studying.

To gain a better understanding of what happens behind the scenes, therefore, we often have to delve deeper into the data, consider more variables and use more sophisticated analytical methods than “eyeball econometrics”. But studying exactly what the data looks like is *always* a good starting-point.

So let's now look at the figures!

The Phillips curve – central but controversial

Allow me to begin with the correlation that is normally called “the Phillips curve”. The name comes from the New Zealand economist, William Phillips, who, in an article in 1958, showed that there was a negative correlation between unemployment and wage increases in British data for the period 1861–1957. High unemployment tended therefore to be associated with low wage increases and vice versa, as in Figure 1. When many people compete for jobs, demand for higher wage increases falls while a shortage of labour leads to greater wage increases.

Figure 1. The Phillips curve

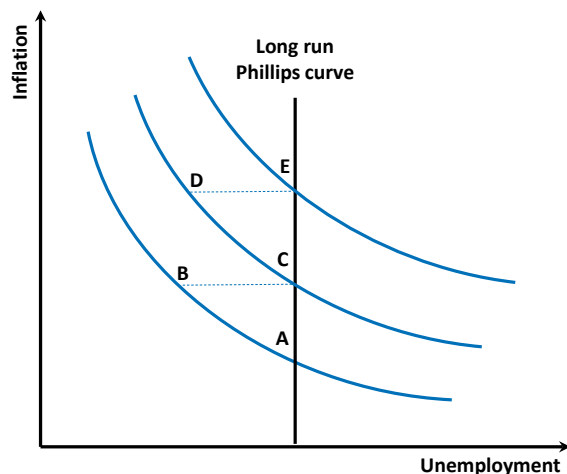


Subsequently, the Phillips curve concept was broadened to accommodate slightly different, though similar, correlations. In 1960, two American economists, Paul Samuelson and Robert Solow, both awarded the Prize in Economic Sciences in Memory of Alfred Nobel, reformulated the curve in terms of unemployment and inflation rather than unemployment and wage increases. As there is a close relationship between unemployment and output, the Phillips curve was also often formulated in terms of the correlation between inflation and output – or, as it is commonly expressed nowadays, resource utilisation. But it is basically the same thing we want to identify, namely a correlation between economic activity and the rate of increase in prices and wages.

For a time, the Phillips curve was interpreted as economic policy-makers being able to choose a combination of inflation and unemployment. This was thought to be possible if the aim was a more permanent reduction in unemployment and there was a willingness to accept higher inflation in return.

But about ten years after the launch of the Phillips curve, Milton Friedman and Edmund Phelps, Nobel laureates in economics themselves, each lay the foundations for the so-called “expectations-augmented Phillips curve”. In brief, this implies that there is a short-term negative correlation between inflation and unemployment, but no long-term relationship. In other words, it is not possible to choose a certain combination of inflation and unemployment and then expect it to last. The reasoning behind this can be illustrated using Figure 2.

Figure 2. The expectations-augmented Phillips curve



The short-term correlation basically depends on the fact that higher inflation is often unexpected inflation. An expansionary policy can lead to lower unemployment and higher inflation; we move from A to B. But gradually, agents in the economy realise that inflation has risen and demand higher wage increases as compensation. In other words, the higher inflation is incorporated in the expectations and becomes the norm, and the higher wage increases cause unemployment to rise. The short-term correlation shifts upwards and we move from B to C. We are then back to the same unemployment as when we started, but with higher inflation. If the same thing were to be repeated, we would move in a similar way from C to D to E. Once again, unemployment is the same but inflation is higher. This unemployment is normally referred to as “equilibrium unemployment”.

A consequence is thus that an expansionary monetary policy cannot permanently create high employment and low unemployment. The long-run Phillips curve is therefore assumed to be vertical, as shown in the figure.

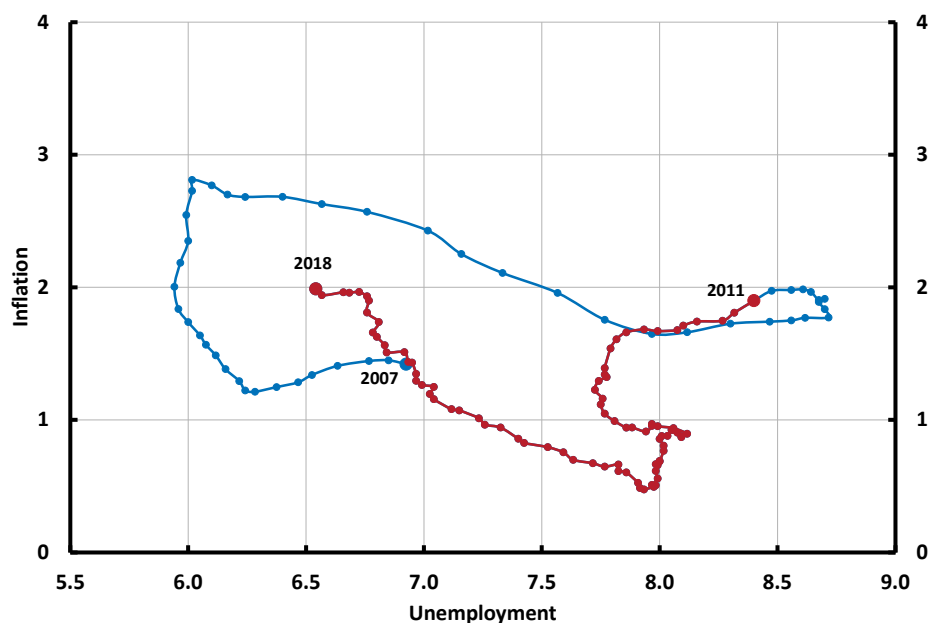
It’s good to keep this in mind, but I won’t discuss the long-run Phillips curve any more here. Instead, I will concentrate on the short-run Phillips curve, as it is an important piece of the puzzle for economic policy-makers, not least central banks. It is the link between demand in the economy, which central banks are assumed to be able to influence with their policy rates, and inflation.

Negatively sloped Phillips curve with inflation – but flat with wage increases

So what does the correlation look like in Swedish data? As I mentioned in my introduction, I will look at data for the last decade. To make the figures clearer, I have marked the period 2007–2010 in blue and the period from 2011 in red.

Figure 3. Inflation and unemployment

Annual percentage change and as a share of the labour force



Note. Moving averages over 12 months.

Sources: Statistics Sweden and the Riksbank

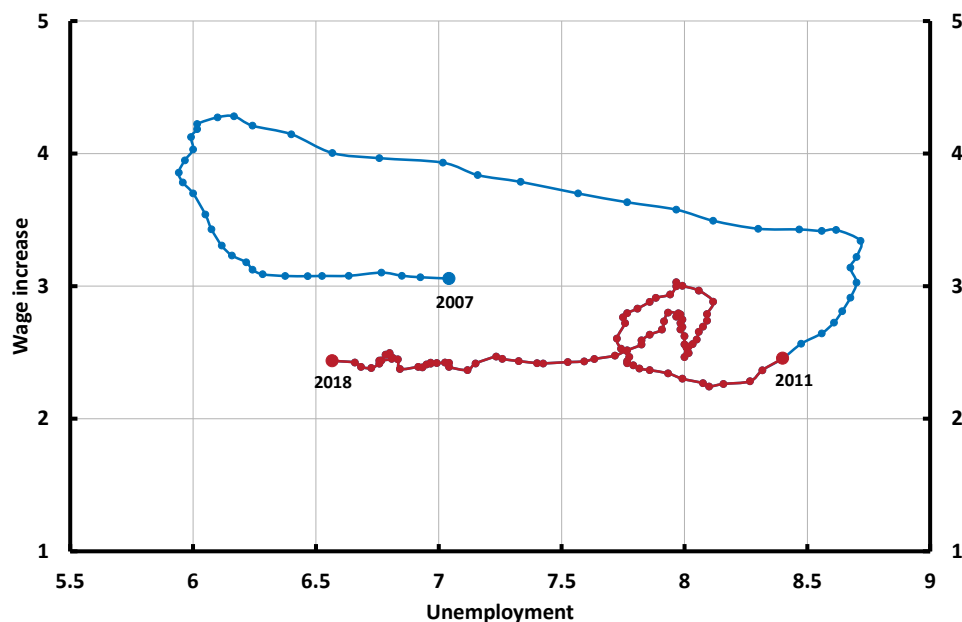
As we see in Figure 3, there is a reasonably clear negative correlation during the whole period. The following additional observations can be made if we look a little closer:

- (i) During the period 2015–2017, inflation rose while unemployment fell.
- (ii) Over the last six months, unemployment has fallen while inflation has remained on a level consistent with the inflation target.
- (iii) For a given level of unemployment, inflation has tended to be lower in recent years than it was during the period 2007–2010.

As I noted, the original Phillips curve was formulated in terms of wage growth and not inflation. So what does such a curve look like for Sweden over the last decade?

Seen over the entire period, the correlation also seems to be negative in Figure 4. But if we look at developments since 2011 in particular, there does not appear to be a correlation – The Phillips correlation seems to be more or less horizontal. Despite good economic developments especially in recent years, and many believe we are in the midst of an economic boom, wage growth has not increased. Why is this?

Figure 4. Wage increase and unemployment
Annual percentage change and as a share of the labour force



Note. Moving averages over 12 months.

Sources: Statistics Sweden, the Swedish National Mediation Office and the Riksbank

First of all, we should remember here that we are, after all, talking about a rather short period of time. In other words, it is still slightly too early to draw the conclusion that some form of structural change has taken place on the labour market that will affect developments in the future. But obviously, it is nevertheless interesting to reflect on the causes of the development.

It is also worth pointing out the fact that the unusually low wage growth, given the activity in the economy, is not a specifically Swedish phenomenon. The situation is much the same in most other industrialised countries and the underlying reasons are currently the subject of rather intensive discussion.¹

A more flexible labour supply a possible explanation

In my view, an important overarching explanation is that the labour supply has probably become more flexible and elastic than previously. In other words, the supply curve on the labour market is flatter. This implies that greater resource utilisation does not necessarily lead to higher wage increases. When demand in the economy rises, it can be satisfied by labour being drafted into the production process, without much of a pressure to increase wages.

As I see it, there are primarily three probable causes behind the more flexible and elastic labour supply. The first is globalisation and increasing economic integration. Globalisation, and hence rising trade, increases labour supply in an indirect way. By this, I mean that products previously manufactured inside the

¹ See for instance Leduc and Wilson (2017).

country by domestic labour can now often be manufactured abroad, often more cheaply, and then be imported. This implies that global labour supply affects domestic wage increases, especially regarding the less educated workforce.

Economic integration within the EU affects labour supply in a direct way as a result of the free movement of labour within the Union. Many utilise the opportunities provided by free movement – some by moving permanently, others by working temporarily in other countries within the Union.

The other cause is technological development. The electronic revolution has made it possible to organise work in a different way than previously. For example, it has made teleworking much more possible. This occurs not only inside Sweden but also in the form of teleworking from other countries in organisations that are located in Sweden. Teleworking makes it possible to use the labour force more flexibly over the business cycle, so that the work can be part-time when economic activity is low and full-time when activity is high.

The third cause has to do with changes of institutions in a broad sense. A concrete Swedish example is that more and more collective agreements contain agreements on working time “banks”. According to these agreements, the number of hours worked for an employee can be redistributed over the business cycle. Another example is the emergence of temporary employment agencies in Sweden. From once upon a time being illegal, the industry is currently about 80,000 strong. In 2017, 250,000 were employed by the temporary employment industry in the areas of temporary staffing, career transition and direct recruitment.

In Sweden, we have also seen greater labour supply among older people. According to the LFS for the first quarter of 2018, almost four per cent of those employed were between 65 and 74 years old. One can assume that their labour supply is very flexible. If we add all these explanations together, it is perhaps not so strange that the correlation between resource utilisation and nominal wage increases seems to have weakened in recent years.

Other explanations also feasible

But there are also other feasible explanations for the observation in Figure 4. One is that what is normally referred to as “structural unemployment”, the unemployment level in a situation with normal activity in the economy, has fallen.² An overestimation of structural unemployment implies that there is more spare capacity in the economy than originally thought and it is thereby possible to stimulate demand more before wage growth kicks in.

The fact that structural unemployment cannot be observed and also varies over time obviously makes it difficult to verify the hypothesis. Purely theoretically, the observation that wage increases are unexpectedly low is well in line with an overestimation of structural unemployment. Personally, I am a little sceptical

² For example, Bell and Blanchflower (2018) think this is the case in the United Kingdom.

about this hypothesis as far as Sweden is concerned. I will come back to this later on.

Another possible explanation for the low wage increases has to do with how inflation has developed after the financial crisis. For about six years, between 2011 and 2017, CPIF inflation was below and often well below target. It would not be particularly surprising if this has had a certain impact on what economic agents generally believe about inflation.

It is true that inflation expectations, as they are measured in surveys, indicate that economic agents are now once again expecting inflation to remain on target in the long term. It also appears that employee organisations are increasingly seeing the inflation target as the basis for wage negotiations. At the same time, however, employers are claiming that the inflation target is no longer an obvious starting-point. According to the Swedish Association of Industrial Employers, it is instead competitiveness and wage outcomes in the euro area that constitute their benchmark.³ The implications of this are not obvious. Ultimately, the outcome depends on each respective party's negotiating strength.

But since wage increases in Germany, for example, have tended to be quite modest and since the industrial sector's wage agreements in turn constitute a benchmark for other collective agreements in Sweden, it is conceivable that this could have contributed to the flat Phillips curve since 2011 in Figure 4. In the long term, however, this should not be a major problem as the euro area has an inflation target of "close to 2 per cent", i.e. very close to our own target. This should imply that wage increases will also be higher in the euro area going forward. But tensions may arise periodically, of course.

All in all, therefore, there are numerous feasible explanations for why wage increases have been relatively low despite good economic development. But so far we have no definitive answer to why it is like this nor to how permanent the phenomenon is.

Research on the Phillips curve ongoing

As I mentioned, the discussion about the Phillips curve is international to a high degree. Quite lively debate has been taking place in various countries recently centred on intensive estimation of Phillips curves, using slightly different specifications and different methods.⁴ This is what I mean when I say that we need more advanced methods than studying pairwise plots, in order to try to come to grips with what the correlation actual looks like.

In this context, it can be interesting to mention that a few fresh contributions to the debate originate from the School of Business here at Örebro University, where Sune Karlsson and Pär Österholm have estimated Phillips curves for the US and Sweden respectively in terms of inflation and unemployment. For Sweden during the period 1995–2017, the Phillips curve has, perhaps not so surprisingly, been found not to have been stable over time. It is also found that the Phillips curve,

³ See, for example, Kinnwall and Rune (2018).

⁴ See for example Blanchard (2016), Cunliffe (2017) and Murphy (2018).

specified in this way, has not become flatter in recent years. The Phillips curve has also been found to have been unstable for the US. It appears to have been flatter in 2005–2013 but thereafter to have had a clearer negative slope.

It should be understood that economic research helps to improve our knowledge in that findings from many different quarters are thrashed out and hopefully eventually all point in a particular direction. Findings have so far varied a great deal. Off and on, the Phillips curve has been said to be dead, still alive but flatter than before or about the same as usual. My interpretation is that “the jury is still out” but that the most common finding in the studies performed seems to be that the correlation between the economic situation and inflation or wage increases has weakened, but is still there.

Correlation between job openings and wage increases

The number of people unemployed is a measure indicating an unsatisfied labour supply, i.e. the number of people who would like to have a job but have yet to obtain one. The unsatisfied labour demand, i.e. job openings that remain unfilled, can also be measured in a similar way. The unsatisfied labour demand, or labour shortage as it is normally referred to, is normally measured in terms of job vacancies.

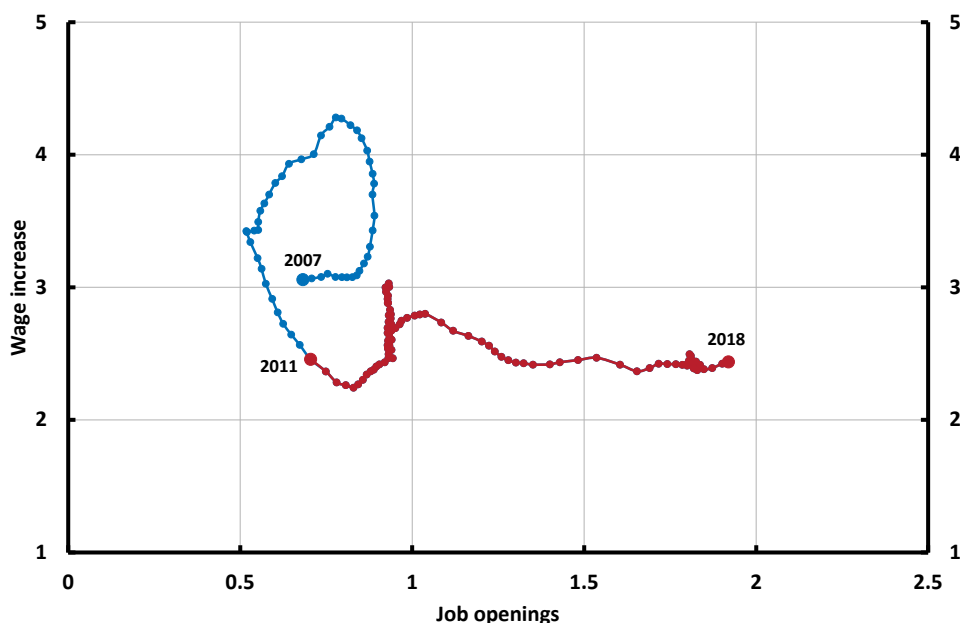
In the statistics, the discrepancy is made between job vacancies and job openings. The fact that an employer has a job opening means that the employer has started the external recruitment of an employee but has yet to employ someone. This is a wider concept than vacancy, which is defined as a “vacant, unstaffed job that can be taken up immediately”. The difference is that a job opening can be staffed, for example by a supply worker performing the work temporarily. But it can also be unstaffed without it necessarily being possible for someone to fill it immediately. One reason can be that the employer has planned well in advance and starts recruitment in good time before the needs become acute.

As I noted in my introduction, the fact that I am pointing out these rather subtle and technical discrepancies is only because I have chosen to highlight various correlations over the past decade by using monthly data. In Swedish statistics, there are monthly data on job openings, while job vacancies are only reported quarterly. Even though it would be desirable in many ways to use vacancies, I therefore have to use job openings as an approximation. If we look at the data, it does not make much difference in this context, but it can still be good to know that I am simplifying a little here as job openings and job vacancies are not really the same thing.

If we now view job openings as a measure of labour demand, we could see the relationship between job openings and wage increases as a kind of inverse Phillips curve. In the same way as the relationship between wage increases and unemployment can be expected to be negative, the correlation between wage increases and job openings should be positive.

Figure 5. Wage increase and job openings

Annual percentage change and as a share of the labour force



Note. Moving averages over 12 months.

Sources: Swedish National Mediation Office, Swedish Public Employment Service and the Riksbank

Figure 5 indicates that this is not really the case during this period. Seen over the entire period, the correlation, if anything, is negative. If we look again at developments since 2011, the curve is almost horizontal, once again reflecting that wage increases have been low and stable.

We have now looked at the correlations between wage increases and unemployment and between wage increases and the number of job openings. Let us complete the cycle and also study the third possible combination, that is, the correlation between the number of job openings and unemployment.

The Beveridge curve shows how well the labour market is functioning

The correlation between job openings, or job vacancies to be exact, and unemployment is another well-known curve within economics. It is called the Beveridge curve and like the Phillips curve was developed in 1958. In contrast to the Phillips curve, however, William Beveridge was not the originator of the curve, but a British economist who gave his name to the curve as a result of his commitment to labour issues such as unemployment and skills matching.

The Beveridge curve is a way of measuring how well the labour market is functioning. It is reasonable to assume that when demand in the economy is high, the number of job openings rises while unemployment falls, and vice versa. The correlation should therefore be negative. But it is above all the position on the Beveridge curve that is interesting. If there is a large number of job openings at the same time as unemployment is high, it suggests that those who are

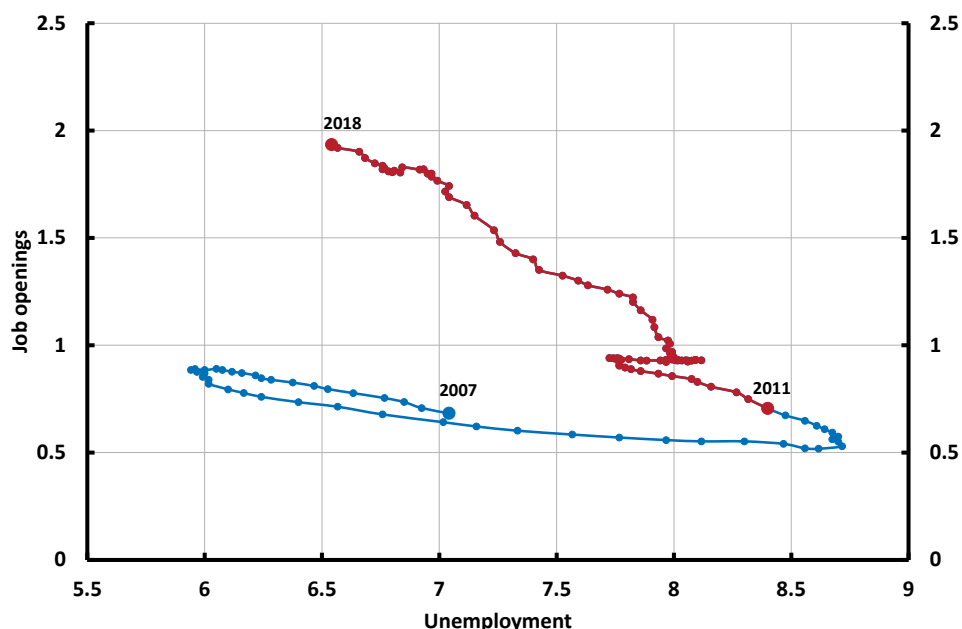
unemployed cannot for some reason fill the openings. In other words, there are problems matching job-seekers to the job openings on the labour market, i.e. the labour market is not functioning particularly well.

What then does the Beveridge curve look like in Swedish data over the last decade?

As we see in Figure 6, the correlation is negative, exactly as expected. But we also see that the correlation from 2011 is higher than previously, i.e. for each given level of unemployment, there is a greater number of job openings. In other words, it seems as though matching on the Swedish labour market is now worse than it was before.

Figure 6. Job openings and unemployment

Annual percentage change and as a share of the labour force



Note. Moving averages over 12 months.

Sources: Swedish Public Employment Service, Statistics Sweden and the Riksbank

What overall conclusions can we therefore draw from these simple correlations I have presented? Well, if we look merely at the data, lower unemployment has not led to higher wage increases. It has, on the other hand, been associated with a rise in the number of job openings. It is not altogether easy to interpret this development. One interpretation could, however, be that companies after the financial crisis have, to a greater extent, been looking to recruit workers with different types of cutting-edge expertise. This is unfortunately a type of labour that is often difficult to find. This is also the impression I get when I have travelled around the country and talked to different companies. If higher wages are offered, it is for this type of labour, but this nevertheless does not appear to have prevented a shortage of labour from emerging in this segment of the labour market.

On the increasingly dominant “traditional” labour market, however, labour shortage does not seem to have been such a major problem. There, a flexible labour supply has helped to restrain wage increases, in the way I described earlier.

Why has inflation risen?

The Phillips correlation measured in terms of inflation has been negatively sloped since 2011, while in terms of wage increases, it has been completely flat. How then has inflation risen despite wage increases not having done so?

Inflation rose steadily between 2014 and 2017, measured in terms of the CPIF, from around 0.5 per cent to a level consistent with the inflation target of 2 per cent. It is of course difficult to identify the exact impact of individual factors on inflation, but it is nevertheless possible to highlight some things that have played a part. One factor that affects the CPIF in rather a direct and straightforward way is changes in the international market prices for energy. Particularly in 2017, higher energy prices helped fuel CPIF inflation. Earlier in the period, the contribution from energy prices was instead negative.

Another factor, which is more difficult to capture and had an inhibitory effect for a long time, is international economic activity. In Sweden, resource utilisation has been rising for quite a long time, which should have made it easier for companies to raise their prices. At the same time, however, a weak and uncertain economic situation abroad has contributed to lower external price pressures. Recently, however, increasingly strong international economic activity has pushed inflation upwards in many countries. More inflation impulses from abroad should also have further improved the conditions for an inflation rate in line with the target.

A factor that has probably also contributed to higher inflation is the rise in both short- and long-term inflation expectations, as they are measured in surveys, since 2015. Since the end of 2016, more long-term inflation expectations have been close to the inflation target of 2 per cent. Actual inflation obviously affects expectations, but the converse is also true. When confidence in the inflation target is strong, it becomes easier to reach the target, as price and wage formation are adapted to these expectations.

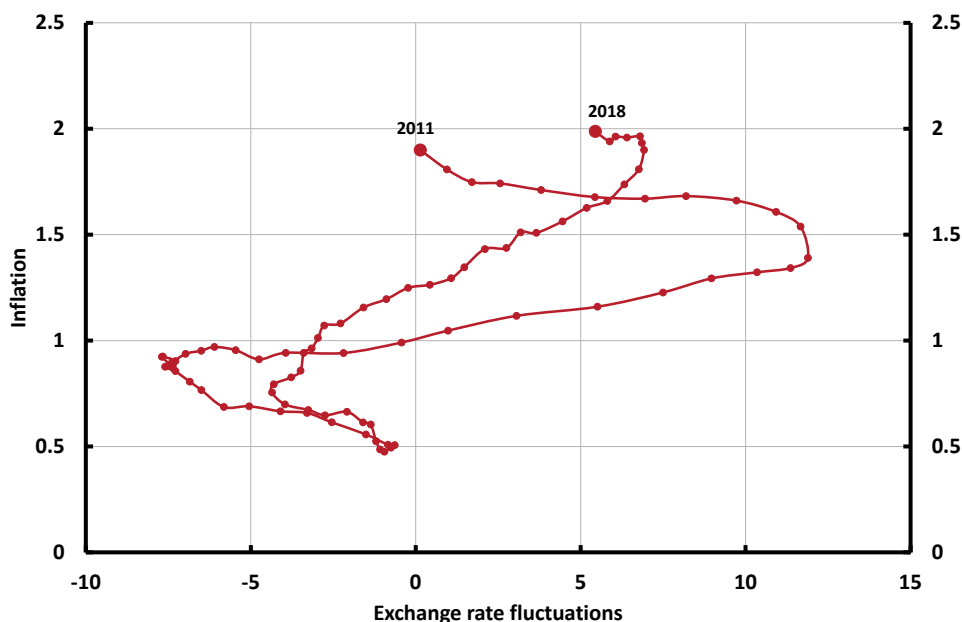
The development of the exchange rate has also contributed to the upturn in inflation. In a simple theoretical framework, the *change* in price level, i.e. inflation, is affected by fluctuations in the exchange rate.⁵ But such a correlation is difficult to find. If we compare exchange rate fluctuations and inflation during the same period, the correlation is zero. However, the exchange rate influences inflation with a time lag. The correlation between inflation and exchange rate fluctuations twelve months previously is 0.19 in my data material. If we lag the exchange rate fluctuations a further twelve months, the correlation increases to 0.43. The maximum correlation (0.46) occurs when the time lag is 27 months. It now becomes clear that there is a positive correlation between exchange rate

⁵ In models where prices are rigid, the inflation rate is instead affected by the level of the exchange rate, or by the real exchange rate's deviation from its long-term value.

fluctuations and inflation, see Figure 7. When the krona weakens, i.e. when the KIX-index increases, inflation also rises.

Figure 7. Inflation and exchange rate fluctuations, 27-month time lag

Annual percentage change



Note. Moving averages over 12 months.

Sources: Statistics Sweden and the Riksbank

Of course, the correlation between the exchange rate and inflation is neither clear-cut nor straightforward. It depends, for example, on what has caused the exchange rate to fluctuate and how long it is expected to last.⁶ But the fact that there is a correlation, insofar as a weakening of the exchange rate tends to result in rising inflation and vice versa, is obvious.

Finally, it should be emphasised that monetary policy has of course been important in this development. From the ongoing debate, the impression can sometimes be that monetary policy does not play such a major role and that economic activity had been just as strong and inflation had been on target regardless of the policy conducted by the Riksbank. This is of course not the case.

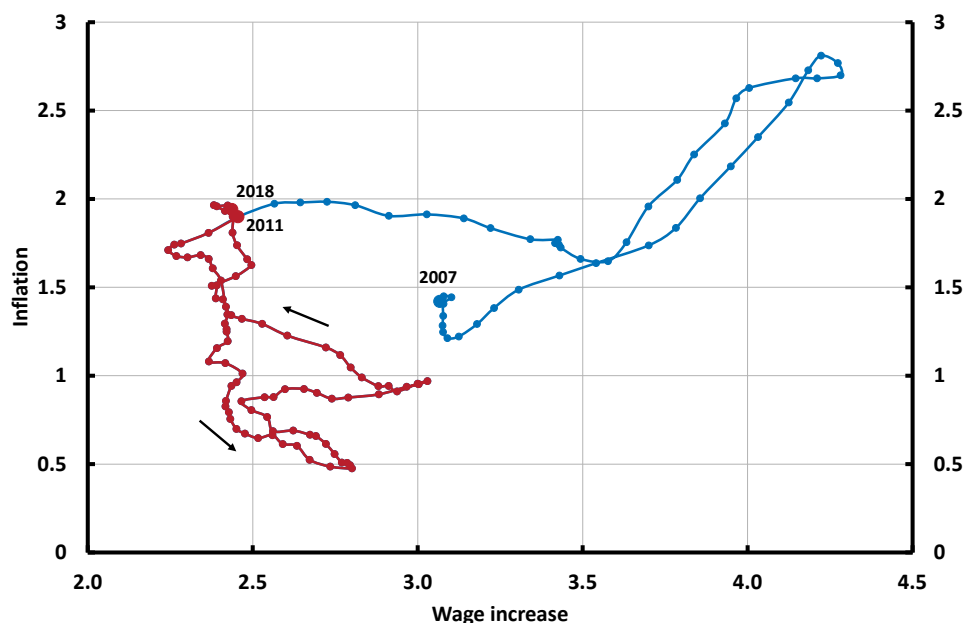
No causal relationship between wages and prices

What then does the fact that inflation has risen but not wages actually signify? To begin with, and for the sake of completeness, we can also plot inflation against wage increases (see Figure 8). As we see, there is a tendency towards a positive correlation, at least periodically.

⁶ See for instance the articles “The impact of the exchange rate on inflation” in the December 2016 Monetary Policy Report and “The exchange rate and inflation” in the April 2018 Monetary Policy Report.

Figure 8. Inflation and wage increase

Annual percentage change



Note. Moving averages over 12 months.

Sources: National Mediation Office, Statistics Sweden and the Riksbank

But what about the causal relationships? There sometimes appears to be a tendency to see wage increases as the source of all inflation. The perceived chain of events is that an increase in inflation is always preceded by a tighter labour market pushing up wage increases and companies then passing on the increased costs in the form of higher prices. Can we therefore expect modest wage increases to push down inflation going forward?

No, such an unshakeable causal relationship between wage increases and inflation does not exist. It can just as easily be the case that stronger demand can allow companies to increase their prices. This increases their profits and margins, which in turn provides scope for them to raise wages. The interplay between prices, wages and activity in the economy – and of course a number of other variables, not least productivity – is complex.

There are quite a few studies that have examined the links between wage increases and future inflation and whether wage increases tend to precede or lag behind inflation. Not unexpectedly, the findings vary somewhat, but a fairly common finding seems to be that wage increases are not particularly informative when it comes to forecasting inflation.⁷ The fact that wage increases have been modest for a while need not therefore signify that inflation will fall in the period ahead. Personally, I believe that both inflation and long-term inflation expectations will remain close to 2 per cent going forward, even if monetary policy is made slightly less expansionary.

⁷ See for instance Peneva and Rudd (2017).

With inflation on target – new challenges

So allow me to summarise. I think this survey presents three important economic policy challenges. An important starting-point is that inflation is on target. It has now been around the target of 2 per cent in outcomes and expectations for some time. Target attainment is good. Occasional monthly variations by the odd tenth of a percentage point around the 2 per cent mark must not determine monetary policy, as I see it. And given that inflation is on target, there is scope for considering other important macroeconomic variables.

The first challenge is that for several years there has been no correlation between unemployment and vacancies on the one hand and wage increases on the other. Unemployment has fallen but wage increases have remained steady at around 2 per cent as an annual rate. Inflation has indeed climbed back to target without any help from wage increases, which shows that inflation does not necessarily have to be wage-driven. But there is still reason to ask oneself how permanent the subdued wage increases are.

The second challenge has to do with unemployment. The key question is: Has unemployment bottomed out? Or can it fall further with an expansionary monetary policy? My assessment is that it will not be possible to bring unemployment down much further with monetary policy alone. In my view, the increasing demand for labour that we are currently witnessing is more likely to lead to increased labour supply than to lower unemployment. This is also what Labour Force Survey flow statistics show. Further reduction in unemployment requires measures from other policy areas.

The third and final challenge is the fact that skills matching on the labour market has deteriorated. This implies that output and employment are not as high as they could have been without the matching problems. Monetary policy can scarcely improve matching on the Swedish labour market. Here, too, measures in other policy areas are required. It is a question of increasing productivity among unemployed persons and reducing the cost of employing them. Labour market policy has a crucial role to play here.

References

- Bell, D.N.F. and D.G. Blanchflower (2018), "The Lack of Wage Growth and the Falling NAIRU", NBER Working Paper 24502.
- Blanchard, O. (2016), "The Phillips Curve: Back to the '60s?", *American Economic Review: Papers & Proceedings* 106(5), 31–34.
- Cunliffe, J. (2017), "The Phillips curve: lower, flatter or in hiding?", speech at the Oxford Economics Society, 14 November, Bank of England.
- Dow, J.C.R. and L.A. Dicks-Mireaux (1958), "The Excess Demand for Labour. A Study of Conditions in Great Britain, 1946-56", *Oxford Economic Papers* 10, No. 1, 1-33.
- Forder, J. (2014), *Macroeconomics and the Phillips Curve Myth*, Oxford University Press.
- Friedman, M. (1968), "The Role of Monetary Policy", *American Economic Review* 58(1), pp. 1–17.
- Karlsson, S. and P. Österholm (2018), "Is the US Phillips Curve Stable? Evidence from Bayesian VARs", Working Paper 2018:5, School of Business, Örebro University.
- Karlsson, S. and P. Österholm (2018), "A Note on the Stability of the Swedish Phillips Curve", Working Paper 2018:6, School of Business, Örebro University.
- Kinnwall, M. and A. Rune (2018), "Teknikföretagen: Riksbanken skapar bubblor" [Association of Swedish Engineering Industries: The Riksbank is creating bubbles], Debate article in Swedish in *Dagens Industri*, 29 March.
- Leduc, S. and J.D. Wilson (2017), "Has the Wage Phillips Curve Gone Dormant?", *Economic Letter* 2017-30, Federal Reserve Bank of San Francisco.
- Murphy, A. (2018), "The Death of the Phillips Curve?", Working Paper 1801, Federal Reserve Bank of Dallas.
- Peneva, E.V. and J.B. Rudd (2017), "The Passthrough of Labor Costs to Price Inflation", *Journal of Money, Credit and Banking* 49(8), pp. 1777-1802.
- Phelps, E.S. (1967), "Phillips Curves, Expectations of Inflation, and Optimal Employment over Time", *Economica* 34(3), pp. 254–281.
- Phillips, A.W. (1958), "The Relation between Unemployment and the Rate of Change of Money Wages in the United Kingdom", *Economica* 25(4), pp. 283–289.
- Samuelson, P. and R. Solow (1960), "Analytical Aspects of Anti-inflation Policy", *American Economic Review* 50(2), pp. 177–194.