Macro Theme

Read my mind, instantaneously: AI central bank watching

Central banks influence markets and economies not only by what they do, but also by what they say. Analysts have long tried to parse written central bank communications for signals of future policy directions. We show how artificial intelligence (AI) techniques can quickly, reliably and objectively convert text into data and provide unprecedented 'virtual assistance' to financial market participants.

What central bankers care about

We are able to tell a very intuitive story about the monetary policy strategy at the Fed, the ECB and the Riksbank by training an algorithm to distribute the text of the meeting minutes into topics. This is a granular and real-time complement to 'traditional' Taylor estimates of central bank reaction functions.

The Hawk-o-Meters: Cut, raise, or stay still

Our 'Thermostat' and 'Timing' Hawk-o-Meters indicate the most probable next policy action in terms of size and direction during the next two meetings based on the words used and the topics discussed in the minutes.

Uncertainty and policymaking

Increased uncertainty makes the appropriate policy action harder to decide and prone to more frequent changes. We build an 'uncertainty index' that offers an important cross-check on eg the Hawk-o-Meter presented in this report.

AI central bank watching at your fingertips, in minutes

All analyses will be available to clients within a couple of minutes after the publication of the relevant text sources. We will publish through the Twitter account @NordeaMacroLab, which can also be accessed through Bloomberg.

Nordea Markets - Market Research
Kristin Magnusson Bernard
Global Head of Macro Research
+46 1570026
kristin.magnusson.bernard@nordea.com

Anders Svendsen, Chief Analyst
+45 55471527
anders.svendsen@nordea.com

Mats Hyden, Chief Analyst
+46 8 4079104
mats.hyden@nordea.com

Inge Klaver, Analyst
+46 720848268
inge.klaver@nordea.com

Sofia Fröjd, Associate Analyst

Daniel Bäumler, Associate Analyst

RIKSBANK'S MINUTES – TOPIC DISTRIBUTION

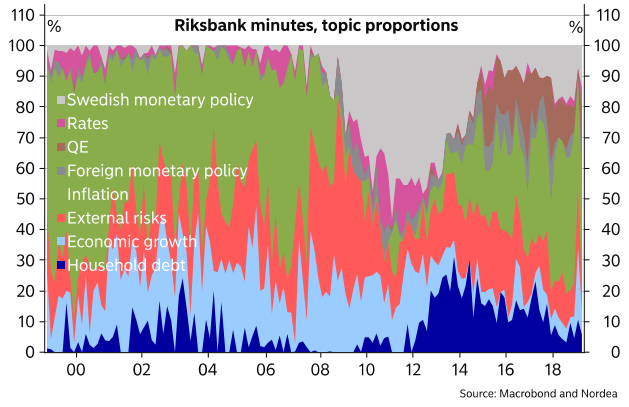

FED HAWK-O-METER

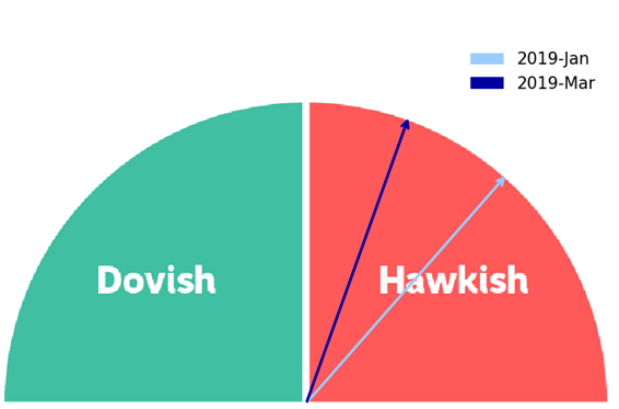

Source: Nordea and Macrobond

Central bank talk: Not empty words

Central bank communication has emerged as a policy tool in its own right, with the content of the information provided now going far beyond macroeconomic variables. In this report, we discuss how machine learning can be used to: 1) visualise the most discussed topics at the Federal Reserve, the Riksbank and the ECB; 2) predict the most likely upcoming policy rate changes, and, 3) measure the degree of uncertainty involved in making policy decisions. And all within a minute or two of central bank minutes being published!

The sky is seemingly the limit in terms of the importance of central bank communication

Not very long ago, central banks did not announce their interest rate changes. They were inferred from market reactions instead. Since the Federal Reserve started publishing rate decisions in 1994, the importance of central bank communication as a policy tool has intensified. Firstly, the general trend towards increased independence from central banks around the world made it much more important for banks to communicate with, and be accountable to, the general public. Secondly, with liberalised financial markets, the effectiveness of monetary policy became much more contingent on its ability to steer expectations of future policy action. Last but not least, over the past five years, we have witnessed the rise of "forward guidance", ie verbal commitment to keeping interest rates at a certain low level either for an unspecified time frame or until a specific condition is met.

"If I turn out to be particularly clear, you've probably misunderstood what I've said."

Alan Greenspan (1998)

"Within our mandate, the ECB is ready to do whatever it takes to preserve the euro. And believe me, it will be enough."

Mario Draghi (2012)

Huge benefits to getting the message right, but not always easy

Long gone are the days when Fed Chairman, Alan Greenspan, made an art out of making mysterious statements. Today's central bankers put considerable effort into communicating as clearly as possible. However, this does not mean that financial market participants always find it easy to get across what they are trying to say. The volume of central bank communication (eg meeting minutes, speeches and newspaper interviews) has increased significantly, as have the potential gains for investors who figure out the message faster and more accurately than competitors.

Machine learning and topic modelling

This is what we do: LDA model is trained in policy minutes and press conferences

We apply a machine learning algorithm called Latent Dirichlet Allocation (LDA) to both Fed and Riksbank minutes from policy meetings. This has already been done a couple of times for Fed minutes, but this is only the second attempt at such for Riksbank minutes to our knowledge. The results for the Riksbank are an extension of our co-author Sofia's MSc thesis: Measuring the information content of Riksbank meeting minutes. The ECB does not publish detailed minutes of its policy meetings and we instead use the introductory statements from its press conferences and the "Monetary accounts".

The algorithm merely recognises the number of latent topics in a text, which we have specified, and cannot decipher the meaning of the individual words. It does not even know what the correct answer to its endeavour would be, which makes it an unsupervised learning algorithm. It is then trained to allocate words to topics and topics to the paragraphs of the minutes under a probability distribution called the Dirichlet distribution.

This may all seem a little complicated and our technical appendix may make things even more perplexing. However, the main thing is that this model allows us to perform four analyses that will be presented in the remainder of this research theme.

Meeting topics can be used to tell a story about the history of monetary policy at the Fed, the Riksbank and the ECB

Firstly, we are able to tell a very intuitive story about the history of monetary policy at the Fed, the ECB and the Riksbank by looking at the topics that have been discussed at their meetings. We do this by labelling the latent topics through analysing the words that the model has assigned to each topic and by visualising the topic distribution over time.

Our Hawk-o-Meter is a contextual word count that helps us predict upcoming rate changes

Secondly, we use the latent topics to build a Hawk-o-Meter – a measure of the hawkishness/dovishness of the text in the minutes – to predict upcoming rate changes. This works kind of like a contextual word count and is based on the idea that some words could have a different meaning depending on the context in which they are used.

Sentiment score from lists of hawkish/dovish words

Thirdly, we build a sentiment score from lists of dovish and hawkish words that we have manually chosen. The sentiment score is likely to be more intuitive than the Hawk-o-Meter given that it is based on words that we have chosen.

Uncertainty can be measured by the changes in words from minutes to minutes

Fourthly, we try to measure uncertainty using lexical innovations, based on a research paper by the Bank of Canada (If Text Could Speak, What Would It Say?). This basically measures how much the words change from minutes to minutes as a means of identifying the degree of uncertainty involved in making policy decisions. To our knowledge, this is the first time that such an analysis has been performed for the Fed, the Riksbank and the ECB.

Twitter Bot and quant reports provide near-instantaneous analysis

Lastly, if you want to follow the output from this project, we have set up a Twitter Bot to publish the results of the model immediately after the text from new minutes or press conferences is processed, ie within a minute or two. You can also follow @NordeaMacroLab via Twitter or your Bloomberg terminal. Furthermore, you can subscribe to a slightly more extensive quantitative report, which will also be available within minutes, through our Nordea e-Markets research portal. We will of course also use relevant outputs from these models in our regular research notes.

What central bankers talk about

The issues on top of central bankers' minds naturally vary by central bank and over time. Complementing traditional analyses of reaction functions with AI techniques should allow for much more granular and real-time detection of the changing priorities of central bankers. Here is a history of the Fed, the ECB and the Riksbank through the eyes of an algorithm.

"Topic clustering": A powerful tool for uncovering hidden structure in texts

Different central banks have different mandates, ie they care about inflation or labour market outcomes to different extents. However, different decision makers at the same central banks also vary in terms of how they interpret the mandate of their institutions. Finally, the time-varying state of the world naturally impacts what are considered to be the key issues of the day. Traditional ways of gauging central bankers' decision-making, such as estimation of Taylor rules (Central Bank Themes: Return to your rules, please) are not always well suited for detecting such shifts. So-called "topic clustering" (which is a certain type of natural language processing, see Technical Appendix for details) can add important information in real time about what is actually on top of central bankers' minds.

A brief history of the Fed – from the eyes of an algorithm

Modelling Fed minutes from 1993

In machine learning, a model is trained on a set of data. In this case, the full text of FOMC minutes from 1993 to today. In unsupervised learning algorithms, there is no correct answer. We do not know exactly what the algorithm should be seeking; therefore, we cannot tell whether it is right or wrong. It will have to learn by itself. The only assumption that we make is about the number of topics in the minutes – eight, in this case. The algorithm models the topics without any comprehension of the meaning of the words. By looking at the words and the paragraphs in each topic in the resulting model, we chose the eight labels in the chart below.

TOPIC DISTRIBUTION – FED'S MINUTES (INTERACTIVE CHART)

Distribution of topics

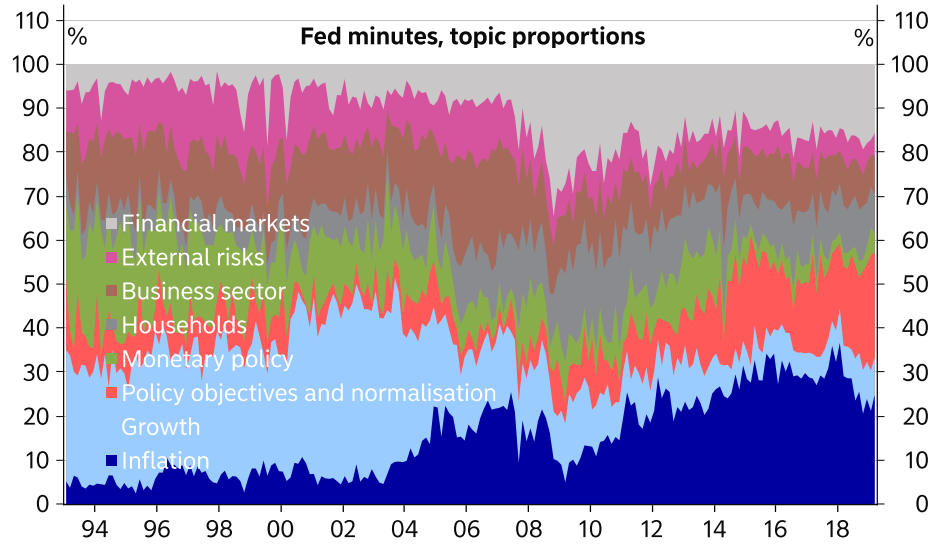

Source: Macrobond and Nordea

The topic distribution – or the percentage of each meeting's minutes devoted to each topic – tells a story about the Fed. The extent to which the different topics have been discussed varies, in some cases substantially, over the last 26 years. The changes in topics relate first and foremost to changed economic situations in the US. However, changes in monetary policy methods and changes in how policy is communicated are also visible. Different FOMC members and chairpersons over time may also impact which topics are in focus during different meetings.

Growth has become less important

Economic growth-related paragraphs, which also include discussions about, for instance, productivity or the economy as a whole, have historically been the most prevalent part of the Fed's meetings. This topic occupied about one fifth of every meeting's minutes on average. From the 1990s until the mid-2000s, growth-related passages were by far the most prominent until other topics, such as financial markets

or inflation, started to receive relatively more attention after Greenspan retired as Fed chairman. "Growth talk" peaked in the early 2000s after the dot-com bubble burst and the US entered a recession. The topic has lost in prominence since then. During the last two to three years, less than 10% (on average) of FOMC meeting minutes were used to discuss growth. Interestingly, business sector-related topics have lost about half of their weight over time since the 1990s. As one conclusion, it appears that the Fed talks much less about the real economy than it did in the past.

Focus has shifted to inflation and households

Instead, the Fed's focus has turned towards inflation – an upward trending topic – especially since the global financial crisis. The increased prevalence of inflation-related discussions in the 2010s is notable and illustrates the FOMC's concern about inflation being too low for too long. Moreover, the Fed has only had a formal inflation target since 2012.

Households, including the labour market and the housing market, is another topic that has gained prominence in recent years; this could relate to the historically low unemployment rate and the lack of wage growth. As a topic, households also received a lot more attention during the mid-2000s, the spring of the US mortgage crisis. At its peak immediately after the financial crisis, close to one fifth of every meeting was devoted to households.

Finally, monetary policy discussions can be divided into words that are closer to longer-run policy objectives and balance-sheet normalisation, and to words that are closer to standard monetary policy, including policy rates and QE. Especially during Jerome Powell's period as Fed Chair, the minutes increasingly contain discussions about the Fed's objectives, and this is unlikely to change given that the Fed has started a process to review its policy framework.

Passages relating to financial markets, banks and credit spreads suddenly gained prominence during FOMC meetings once the financial crisis hit. The fact that these discussions took more than twice as much space during the 2010s as they did before the financial crisis clearly illustrates how the Fed has become more aware and more conscious of financial market developments.

CHANGE IN TOPIC PROPORTIONS – MARCH VS JANUARY 2019

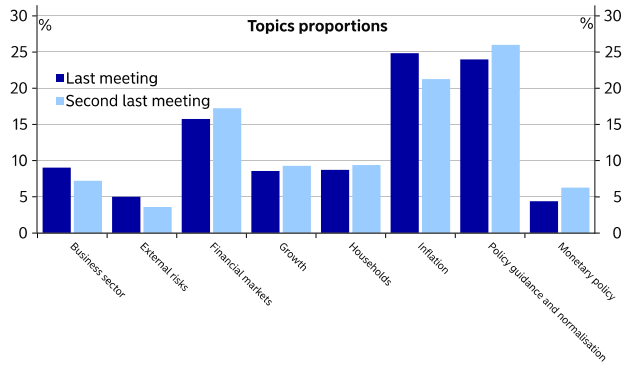

Source: Macrobond and Nordea

CHANGE IN TOPIC PROPORTIONS OVER TIME

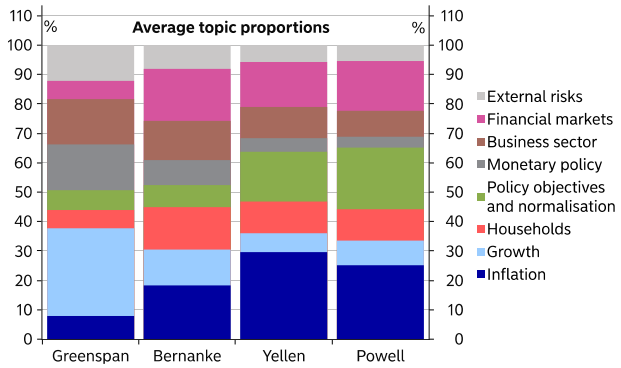

Source: Macrobond and Nordea

Reading the Riksbank

Increase in Swedish monetary policy

When performing the same exercise for the Riksbank, several time periods stand out in the resulting chart (see below). One notable period is the enormous increase in discussions about "Swedish monetary policy", which started around 2007 and grew substantially until 2009, after which it started to decline in prominence. This topic covers, among other things, the rate decision, some elements of the economic outlook and unemployment. In late 2007, we note that the minutes were changed to identify each board member by name, which may have increased the prevalence of the topic in the minutes to some extent. Moreover, we note that during the period when the topic had the largest proportion in the minutes, Lars E.O. Svensson was on the Executive Board (2007-13). He advocated a more expansionary monetary policy amid the gradual rate hikes that occurred from mid-2010 to the end of 2011. Looking at the sections that the algorithm identified as having a high probability of belonging to this topic, we find that many of the sections contained discussions brought forward by Svensson.

External risks increased as a topic around the financial crisis

As expected, the topic of "external risks" showed an increase around the financial crisis of 2008. During this time, there was little discussion about inflation. The topic of "QE" (quantitative easing) first appeared around 2014. This was around the time when the Riksbank started purchasing government bonds to make monetary policy more expansive in addition to lowering the policy rate. "Household debt" as a topic also increased in prevalence in recent years. Owing to low rates, households have started to borrow money in greater amounts, thereby leaving household finances quite sensitive to increasing rates. Increasing household debt has been a worry of the Riksbank's for some time.

Inflation has re-emerged as the dominant topic in recent years

The topic of "inflation" dominated the minutes until 2007, together with "growth". Before 2007, the policy document that was regularly published by the Riksbank was called "Inflation Report" and not "Monetary Policy Report", as it has been called since 2007. It contained fewer forecast variables, and it could be that the policy discussion was less about forecast details and more focused on a qualitative discussion of the classical trade-off between inflation and the business cycle. The topic model appears to confirm this. Since 2013, when the responsibility for macroprudentials was put in the hands of the Swedish FSA, the "inflation" topic has gained in importance once again. The battle against low inflation expectations added to the topic's importance in recent years.

TOPIC DISTRIBUTION – RIKSBANK'S MINUTES (INTERACTIVE CHART)

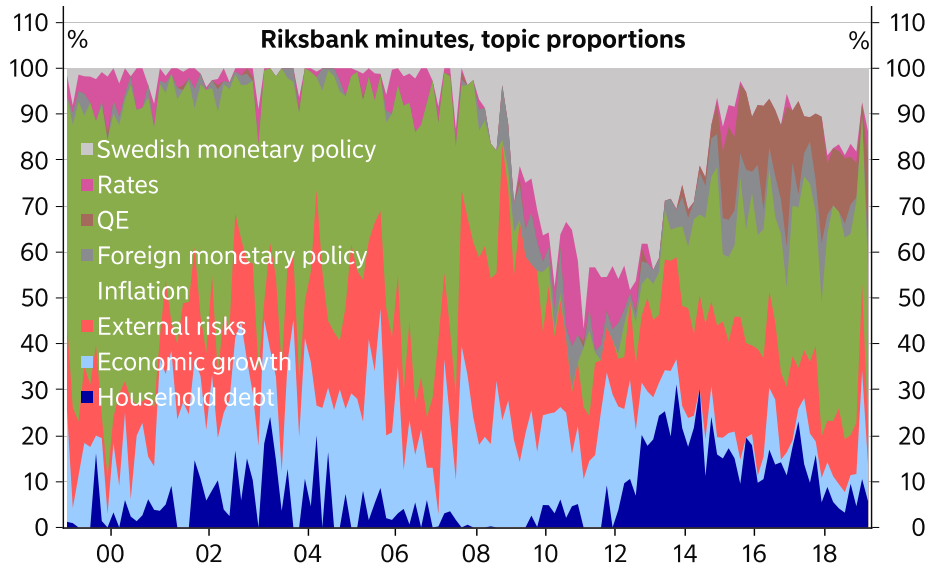

Source: Macrobond and Nordea

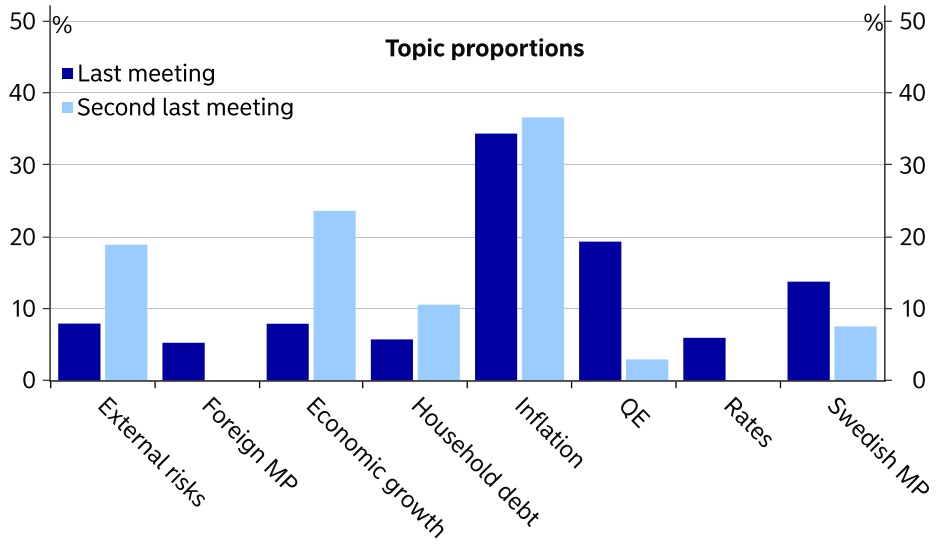

Source: Macrobond and Nordea

Mixed messages from the ECB

We look at two text sources from the ECB

As the ECB started publishing accounts of its monetary policy meetings only in 2015, we have included the introductory statements in our topic analysis for the ECB.

The introductory statements provide a longer-term perspective on ECB communication

If we first look at the longer-term perspective, the topic distribution of the introductory statements (the part of the press conference before the Q&A) again shows that topics are dynamic and change over time, depending on economic events and changes in the central bank's policy course. Also, the topics appear to be strongly dependent on the person delivering the statement, being the president of the ECB. During the outset of the ECB under President Duisenberg, there was still a strong emphasis on monetary analysis, in line with the two-pillar approach. When Trichet took over in 2003, however, this topic received much less attention, and the focus was on explaining the policy assessment in the statement. In 2011, when Draghi took the reins, this marked the end of that, alongside a gradual increase in attention to the second pillar, economic analysis.

Some recurring topics in the statements, and some special attention points

The call for structural reforms is a relatively consistent item on the agenda of the press conference, and once every quarter the new staff macroeconomic projections are presented. The low-inflation and low-rate environment has led to more discussion about inflation in the past years, and monetary analysis has regained ground since the start of QE in 2015.

TOPIC DISTRIBUTION - ECB (INTERACTIVE CHART)

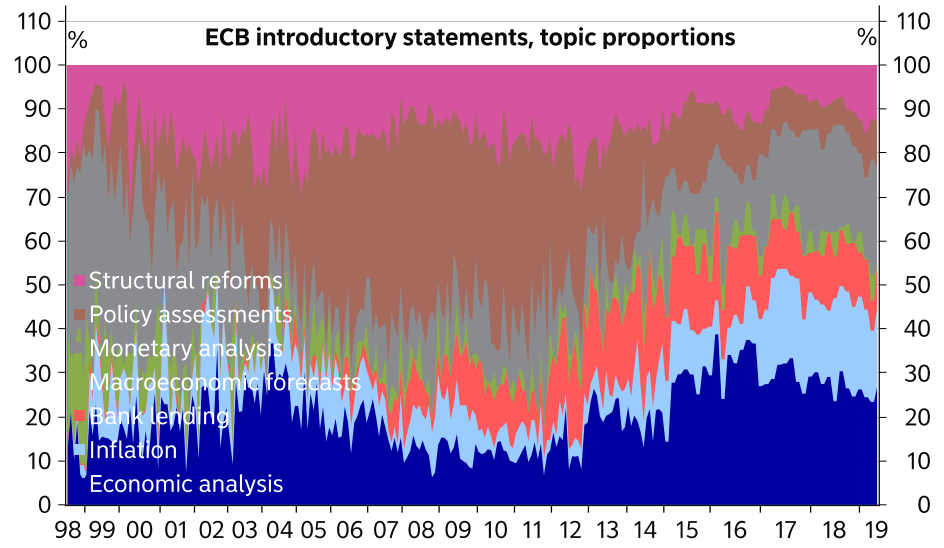

Source: Macrobond and Nordea

Same bank, same meeting, different text sources, different topics

Not much can be concluded from the topic distribution of the ECB's accounts alone, as there has not been much variation over time due to the short time horizon. However, we believe it is worth pointing out that the model picks up remarkably different topics from the accounts than from the introductory statements, suggesting that the introductory statements that are delivered right after the policy meeting do not reflect the discussion held at that meeting.

TOPIC DISTRIBUTION - ECB (INTERACTIVE CHART)

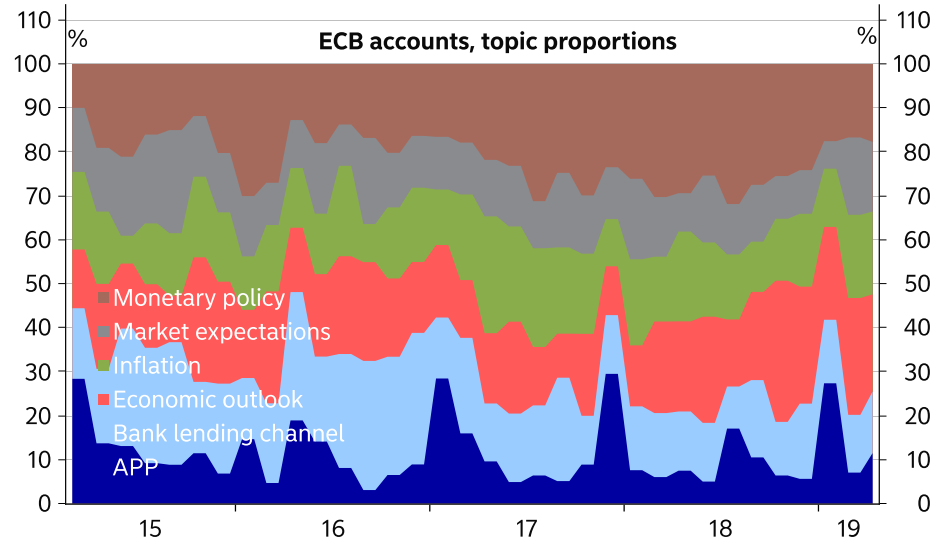

Source: Macrobond and Nordea

The Hawk-o-Meter: Raise, cut or stay still?

The perennial question for central bank watchers has long been: will they raise rates, cut rates or stay still? For anyone doubting their own interpretation, we find that an AI model can quickly and reliably provide a second opinion. We proudly present: our Hawk-o-Meters!

FED: 'THERMOSTAT' HAWK-O-METER

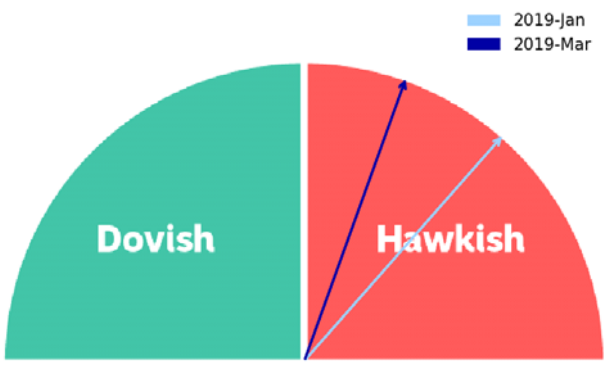

Source: Macrobond and Nordea

RIKSBANK: 'THERMOSTAT' HAWK-O-METER

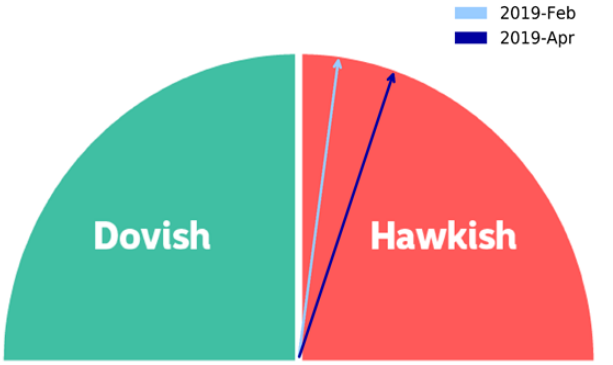

Source: Macrobond and Nordea

The 'Thermostat' Hawk-o-Meters indicate the most probable rate change...

In this section, we detail how our Hawk-o-Meters can be used for financial market participants. Interested readers can find an extensive description of how our Hawk-o-Meters are constructed in the Technical Appendix.

The 'Thermostat' Hawk-o-Meters above indicate the most likely next course of action from a specific central bank meeting, and compares that with the most likely rate outcome from the previous meeting. In other words, it provides a prediction that is the weighted average of the potential rate changes and their respective probabilities. In that way the prediction compares to the rate change implied by a three-month future.

From the Fed's March minutes, the Hawk-o-Meter predicted unchanged rates at the May and June meetings with 69% probability, but the prediction was 10 bp, down from 22 bp in January, as the model saw some likelihood of rate hikes. Thus, our Hawk-o-Meter supports the general view of a less hawkish Fed at the March meeting. The prediction of the model, however, remains in contrast with the market pricing of rate cuts.

For the Riksbank, the change in the Hawk-o-Meter in February was insignificant. This was well in line with pricing and the consensus perception. In April, however, the Riksbank's policy decision emerged as a big, dovish surprise with a massive impact on the bond market. When the minutes were released, the Hawk-o-Meter still showed a massive overweight for an "unchanged" policy stance, but it also pointed to a hawkish change in the policy discussion. This can, on the one hand, be explained by a flaw in our algorithm, as it is not trained against outcomes of policy paths or QE programmes. On the other hand, it could also be a consequence of changed gearing between the policy discussion and the policy decision. It is interesting to note that an unsupervised, algorithmic reading of the minutes (not taking the rate path or QE into account) actually pointed to a more hawkish Riksbank, while the policy action was much more dovish. Perhaps this can help explain the big surprise in the market, while highlighting the big challenge of predicting policy actions from the Riksbank.

FED – PREDICTION OF THE HAWK-O-METER

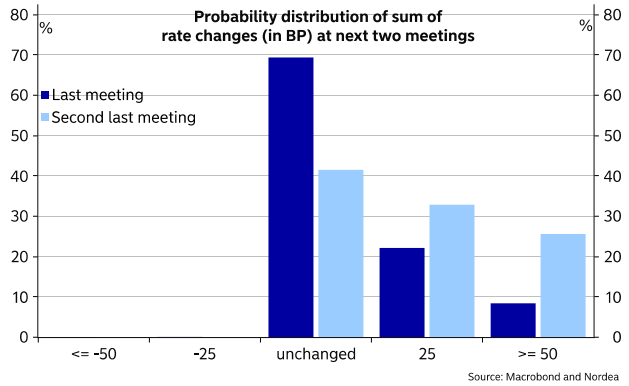

RIKSBANK – PREDICTION OF THE HAWK-O-METER

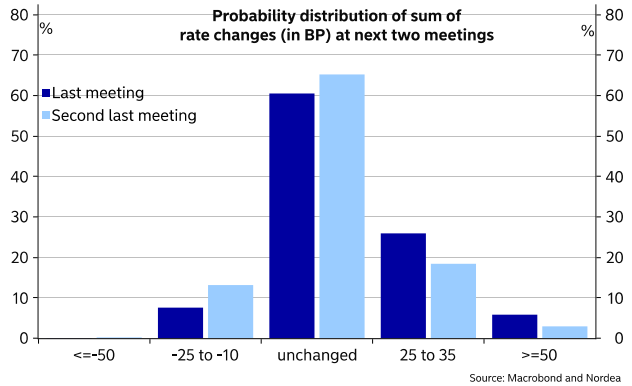

...while the 'Timing' Hawk-o-Meters show if any action should be expected at the next two meetings

The 'Timing' Hawk-o-Meters show how likely central banks are to take a certain policy action at a specific meeting. For example, from the minutes of the March FOMC meeting, the 'Timing' Hawk-O-Meter points to unchanged rates in May and June with 69% probability, much higher than in the January minutes. The Riksbank 'Timing' Hawk-o-Meter, in contrast, predicted unchanged rates for the two next meetings, both in February (65% probability) and April (61% probability), indicating its very cautious approach to policy normalisation.

Measuring the tone of the ECB

To interpret the texts published by the ECB, we take an alternative approach and construct a sentiment score. Given the short history of the ECB's monetary policy accounts and the limited length of the introductory statements, these texts do not lend themselves to machine learning approaches for measuring tone, as we did for the Fed and Riksbank.

The sentiment score marks the tone of the ECB...

Therefore, we take a more hands-on and intuitive approach to quantifying the dovishness or hawkishness of the ECB's communication. For this indicator, we compute a score by counting the number of times hawkish words and dovish words appear in a text, and weighing them against one another. This list of hawkish and dovish words has been predetermined by us. The lower the score, the more dovish the text.

...letting us objectively compare the signals coming from the texts

Having a score for both types of texts – the introductory statements and the accounts – allows for an interesting comparison of their respective signals. From the charts below, it is clear that the sentiment score has tracked movements in the ECB policy rate well, turning dovish in times of rate cuts and hawkish in times of rate hikes. For the more recent period when the rate has not moved, from mid-2016, the sentiment indicator has picked up a hawkish tone from the introductory statements almost throughout that entire period, whereas the score for the accounts shows that the discussion in the meetings was actually more mixed and the signal has alternated between hawkish and dovish.

It turns out that your text source matters

On the one hand, the differing signals mean that there is different information communicated through these channels, and there is a benefit from reading both. On the other hand, there is a risk that the opposing signals end up just being noise and that the public will not be able to distinguish the true message from the signals.

ECB – SENTIMENT SCORE FOR INTRODUCTORY STATEMENTS

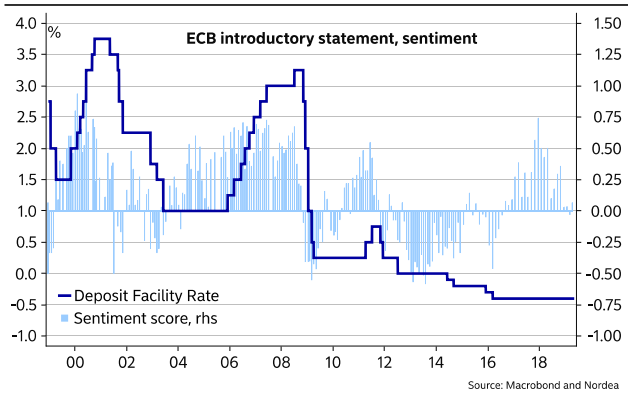

ECB – SENTIMENT SCORE FOR ACCOUNTS

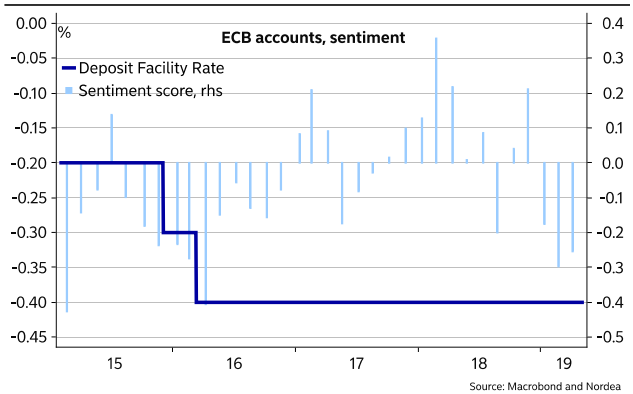

The sentiment score also proves to be a good indicator of the tone of the minutes from the Fed and the Riksbank.

FED – SENTIMENT SCORE

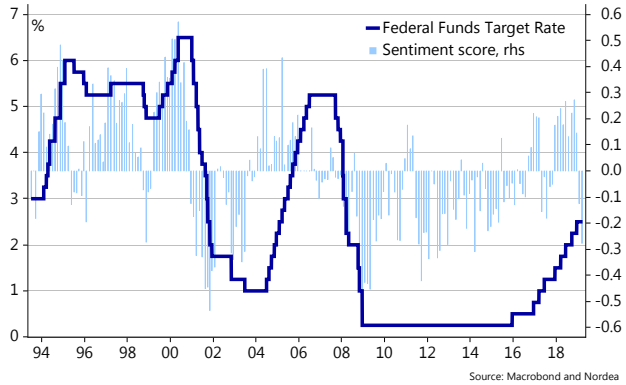

Source: Macrobond and Nordea

RIKSBANK – SENTIMENT SCORE

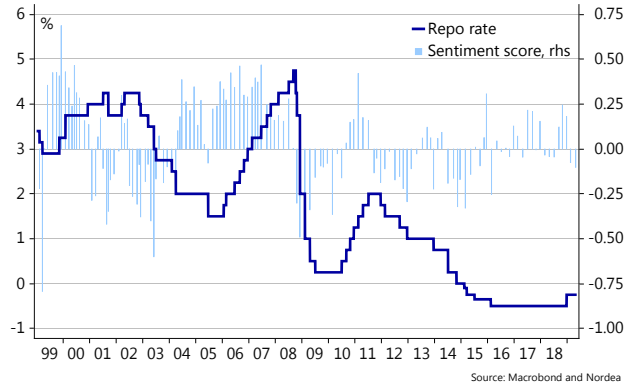

Source: Macrobond and Nordea

Uncertainty over central bank policymaking

Increased uncertainty makes the appropriate policy action harder to decide and prone to more frequent changes. We build an 'uncertainty index' that offers an important cross-check for eg the Hawk-o-Meter presented in the previous section.

Increased uncertainty complicates decision making

During times of uncertainty, for instance when prices in financial markets are very volatile or economic conditions are changing rapidly, the appropriate course of policy action can be hard to find. When trends change, it might take some time to disentangle shocks from trends and assess the new situation. Agents put off spending and investment, and central banks reassess their policies as well. It is useful to detect when such shifts may occur, and when uncertainty in the policy is larger.

Uncertainty can be measured by the changes between minutes

We build an index of lexical innovation, which can be used to identify changing economic and financial conditions and serve as a proxy for the different central bankers' implicit assessment of uncertainty. The degree of uncertainty offers an important cross-check for eg the Hawk-o-Meter presented in the previous section. We use a deep learning method, word mover's distance, to measure the extent of similarity between one set of minutes and the previous one, to capture the change in the central banks' perceptions of the environment, as reflected in their policy discussions. This gives us an implicit assessment of the uncertainty surrounding the outlook and decisions presented by the central banks. In the uncertainty measure, higher values equal greater uncertainty.

FED: MEASURE OF POLICY UNCERTAINTY

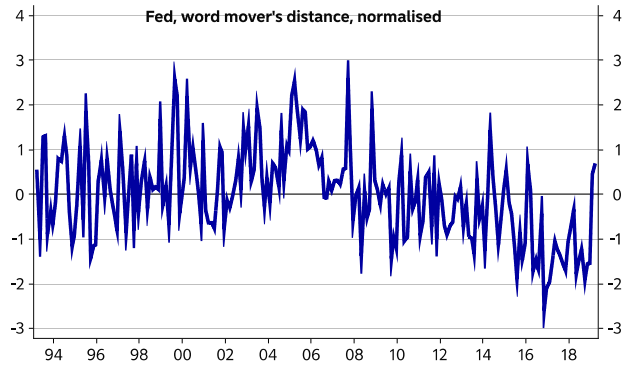

Source: Macrobond and Nordea

ECB: MEASURES OF UNCERTAINTY

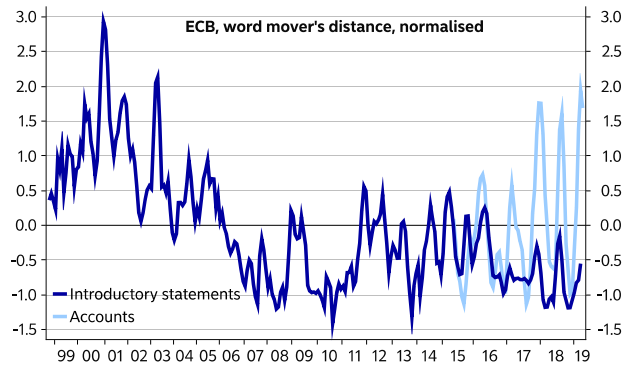

Source: Macrobond and Nordea

RIKSBANK: MEASURE OF UNCERTAINTY

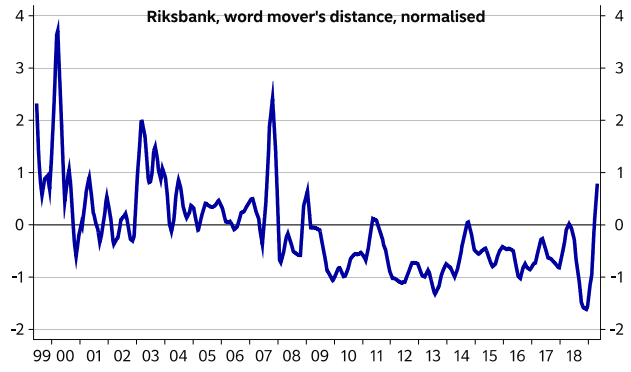

Source: Macrobond and Nordea

The first thing to note from the uncertainty indices is that periods of heightened uncertainty are highly correlated across countries. The second is that uncertainty is a different concept from eg economic expansions or downturns. For example, the Fed perceived much higher uncertainty regarding the appropriate course of policy action in the years prior to the financial crisis starting in 2007-08 than when the crisis erupted. A third observation is that policymakers and outside observers might perceive uncertainty quite differently. For example, financial market participants have been quite surprised by changes in the Riksbank's policymaking since 2013, while the Riksbank has been very certain of the appropriate course of action.

The quickest of quant reports

Using web scrapers and machine learning models can give a preliminary analysis of new minutes from a central bank meeting in – excuse the pun – just a couple of minutes. Below, we present a teaser of what upcoming quant reports could look like. We will publish through Twitter @NordeaMacroLab, which can also be seen through Bloomberg.

A blast from the past – how have the topics evolved over time?

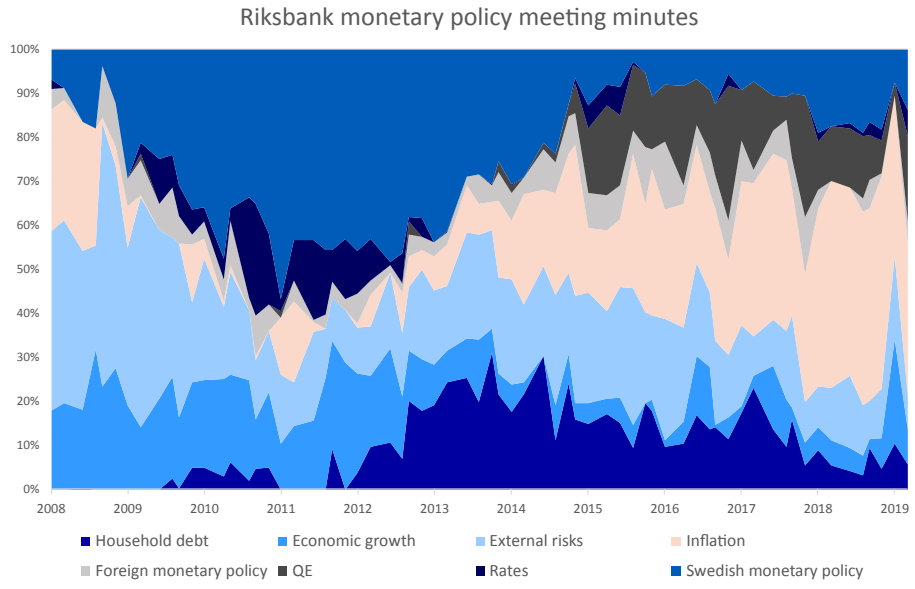

Source: Macrobond and Nordea

How has the focus changed compared with the previous minutes?

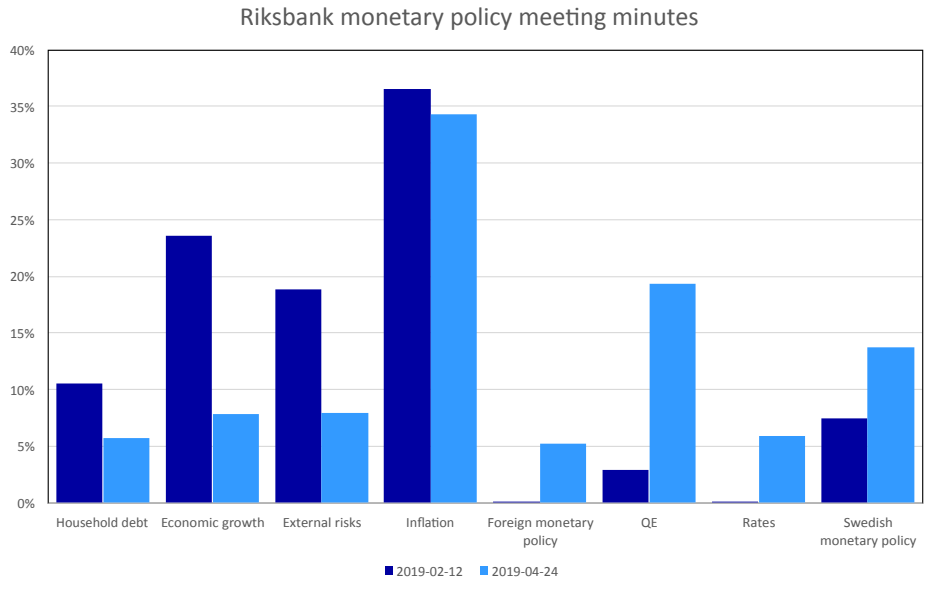

Source: Macrobond and Nordea

Among hawks and doves – the sentiment score captures the tone of the text; lower values mean more dovish tones

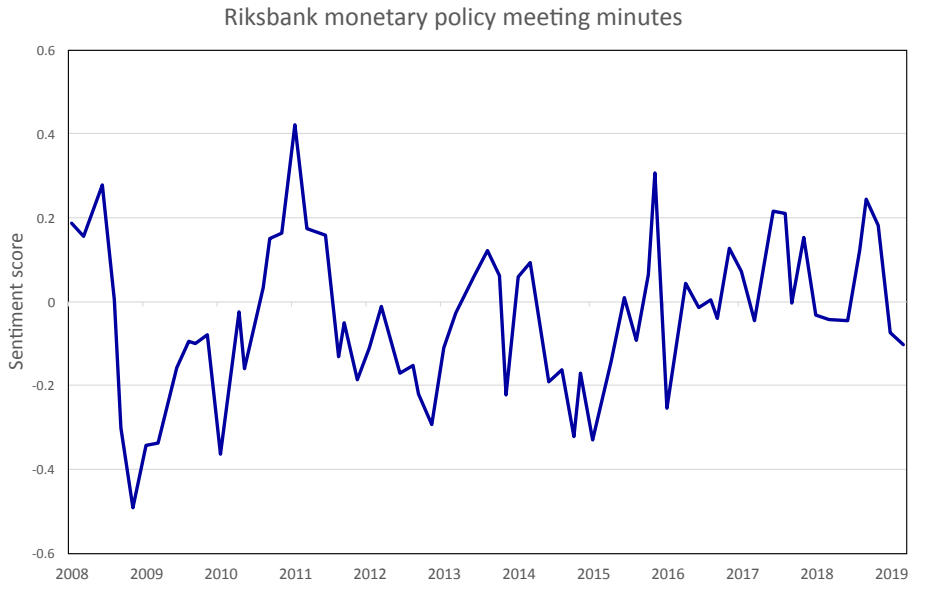

Source: Macrobond and Nordea

What will happen next? The Hawk-o-Meter might hold the answer

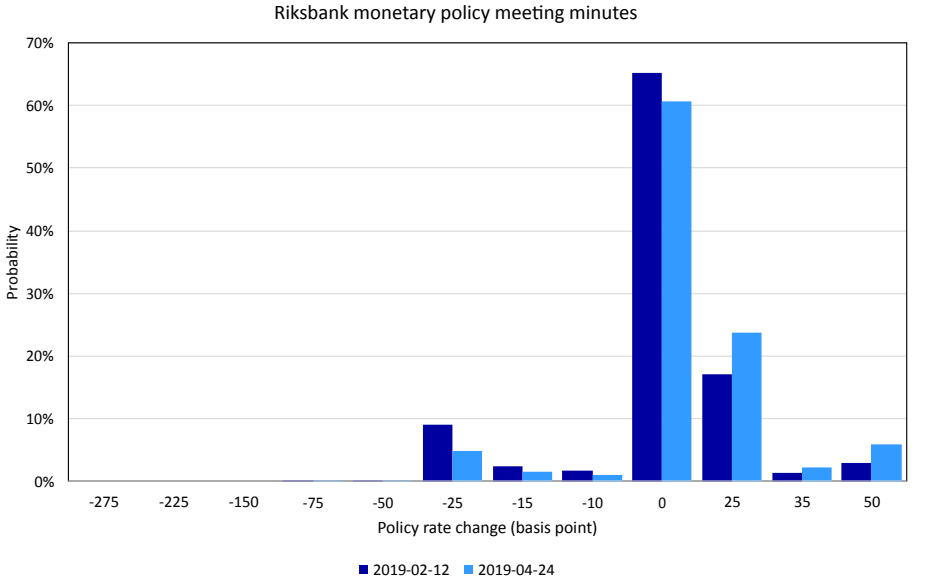

Source: Macrobond and Nordea

Technical Appendix

The techniques used in this paper are various forms of Natural Language Processing (NLP), which turn text into quantifiable data.

Text sources

We use all the monetary policy minutes for the years 1993-2019 for the Fed and 1999-2019 for the Riksbank, excluding minutes from extraordinary meetings. For the ECB we use both the opening statements for the press conferences during the years 1998-2019, and the "monetary accounts" for the years 2015-2019. All texts were retrieved from the respective central bank's homepage.

Topic Modelling - Latent Dirichlet Allocation(LDA)

A common approach to topic modelling is to use Latent Dirichlet Allocation (LDA), first introduced by Blei et al. in 2003. The model has many applications in a variety of fields, from investigating the presence of mental health articles in newspapers, to finding recommendations for articles to read, or studying the history of scientific ideas. The model is also very useful for uncovering the semantic structure of large collections of documents (in the area of language processing called corpus).

We can think of the documents in a corpus as a being generated by several unobservable, or latent, topics. Each topic is in turn represented by an unobservable distribution of words. By using LDA we wish to find the latent distributions of both the topics and the words.

A topic is a vector of probabilities for words

A topic in LDA means a vector with probabilities for each word in the corpus. By manually studying the words with the highest probabilities, one can often put an appropriate name to the topic, but it is essential to understand that the naming of a topic, as is done in this report, is arbitrary and not always straight-forward.

Only the number of topics is assigned beforehand

The number of topics in a corpus is determined beforehand, and the optimal number is far from obvious. The distribution of words in each topic is, however, decided by the training algorithm, as is the distribution of the topics in each document. Thus, LDA can be said to be unsupervised.

Clever and off-the-shelf algos explain LDA's popularity

In LDA, two nested probability distributions are assumed to generate the words in the topics and the topics in the documents. One of the distributions is a Dirichlet distribution, and thus the “D” in LDA. The model can be trained (or fitted) to the corpus by the use of different numerical methods applied to approximations of the rather complex maximum-likelihood function. The fact that these approximations and numerical methods have been done in some very clever, and far from obvious, ways is an important reason of the popularity of LDA. Please see the original articles for details.

Few central bank documents a challenge

There are a number of challenges to applying LDA to central bank minutes or speeches. Assigning the number of topics is one. We have chosen a number of topics that in our judgement created the most meaningful topics when looking at the word distribution in each topic and how the topics were represented in the documents over time. The length and number of documents also play a crucial role in getting comprehensible topics. A rule of thumb is that more documents produce a better model. But long documents may not produce comprehensible topics, as they are likely to contain a mixture of many topics and it may be difficult for LDA to separate them. A solution to this can be to split the document into chapters or sections. In the case of central bank minutes or speeches, one can say in general that the problem is that there are too few documents for an optimal LDA estimation, while the length of the documents is more easily handled.

The Hawk-o-Meter: A snazzy contextual word count

Our Hawk-o-Meters are contextual word counts based on the topic models. Since the topic models allow us to filter out all paragraphs that are highly likely to be about inflation, for instance, we are also able to separate words such as "high" when the word is about inflation from when it is about unemployment, which would have somewhat different implications.

The Hawk-o-Meter weights

To figure out which words are important in which contexts, we roll a "window" in which we evaluate all words in all contexts. The window is the realised change in the policy rate over the coming two meetings. Words that appear frequently in windows with rate

hikes get a high positive weight. Words that appear frequently in windows with rate cuts get a high negative weight. Finally, words that appear frequently in windows with no rate change are weighted with zero. In the end, all words in all contexts are assigned one weight summing up how that has been used ahead of policy changes in the past.

For example, the Hawk-o-Meter for the March FOMC meeting was calculated by first letting the topic model determine the topic distribution for each topic and then count all words in all topics, multiply them by the weights and aggregate up. The result is one number which can be positive (hawkish) or negative (dovish) and says something about what the Fed will do at the coming two meetings.

We take it one step further and do some old-fashioned econometrics on top of the number: a probit model (a type of regression where the dependent variable can take only discrete values), which basically scales the Hawk-o-Meter to predict rate changes in basis points. The probit model also shows the distribution around the predicted outcome.

Sentiment analysis

The tone in the minutes of a new meeting was measured by a simple approach using word counts. Two lists of hawkish and dovish words were manually constructed, containing all forms of a word. Examples of hawkish words are "increase", "raising", "rise", "positive", "hike" and "tightening". The dovish list includes words such as "decreasing", "lower", "uncertainty", "weak", "negative", "cut" and "expansive". Each sentence was then categorised into one of two categories based on the frequency of certain words. The first category contains words such as "inflation", "monetary policy" and "growth", the second includes words such as "uncertainty", "unemployment", "debt" and "risks". The sentences regarding monetary policy or inflation will treat words such as "raising" or "increase" as hawkish words, while sentences discussing eg household debt, uncertainty or risks will most likely treat words like "increase" as a reason for keeping lower rates and an expansive monetary policy, therefore being dovish. All sentences that could not be categorised based on the word counts were placed in the first category.

When the sentences are categorised, words from the hawkish and dovish lists are counted and stored in a hawkish score and a dovish score. Sentences in the first category will give one point to the hawkish score if the word appears in the hawkish list, and dovish words give dovish points. The opposite goes for the second category: if a word appears in the hawkish list it will give one point to the dovish score, and vice versa. When all words are counted, the sentiment score of each of the minutes is calculated.

Lexical Innovations, Word Mover's Distance (WMD)

WMD, developed by Kusner et al (2015), is a way to measure how similar a certain text is to another one, in this case the monetary policy minutes. In order to do this a measure of “distance” between two documents needs to be defined. This is done by first measuring the distance between every word in one document to very word in another document. The distance between two words is measured by simple vector metrics. But first every word needs to be assigned to a vector. This in turn is done by model called “word2vec” that is based on a technique called “word embedding”. The technique applies a simple neural network algorithm to a corpus of documents in order to decide which words are more often used together. One can say that word2vec creates a representation of a word with a vector of relationships to other words. This makes it possible to link semantically similar words to each other even if they are grammatically very different. The technique is called “word embedding”. The minimum distance that the “meaning” of all words in one document have to “travel” in order reach the “meaning” of all words in another document is the simplified explanation of WMD.

We normalise our WMD to have zero mean and a variance of one as the levels have no economic interpretation.

DISCLAIMER

Nordea Markets is the commercial name for Nordea's international capital markets operation.

The information provided herein is intended for background information only and for the sole use of the intended recipient. The views and other information provided herein are the current views of Nordea Markets as of the date of this document and are subject to change without notice. This notice is not an exhaustive description of the described product or the risks related to it, and it should not be relied on as such, nor is it a substitute for the judgement of the recipient.

The information provided herein is not intended to constitute and does not constitute investment advice nor is the information intended as an offer or solicitation for the purchase or sale of any financial instrument. The information contained herein has no regard to the specific investment objectives, the financial situation or particular needs of any particular recipient. Relevant and specific professional advice should always be obtained before making any investment or credit decision. It is important to note that past performance is not indicative of future results.

Nordea Markets is not and does not purport to be an adviser as to legal, taxation, accounting or regulatory matters in any jurisdiction.

Nordea Bank Abp, Satamaradankatu 5, FI-00020 NORDEA, Finland, domicile Helsinki, Business ID 2858394-9

Further information on Nordea available on www.nordea.com