ASSESSMENT OF INFLATIONARY PRESSURES USING NEWSPAPER TEXT ANALYSIS

Mirko Đukić

Available at www.nbs.rs

Economic Research and Statistics Department

NATIONAL BANK OF SERBIA

Belgrade, Kralja Petra 12

Telephone: (+381 11) 3027 100

Belgrade, Nemanjina 17

Telephone: (+381 11) 333 8000

www.nbs.rs

Assessment of inflationary pressures using newspaper text analysis

Mirko Đukić

Abstract: In this paper we present the results of a dictionary-based text analysis of articles from the economic sections of four Serbian daily newspapers, caried out to estimate if those articles contain useful information for assessing inflationary pressures. We analyzed 117,113 economic articles in total for the period 2007–2022, by counting terms related to inflation and price rises and price falls, or counting texts containing those terms. Measures of newspaper inflation sentiment, obtained in this way, were found to be highly correlated with inflation, mainly driven by the periods of large inflation swings. During the period of stable inflation, the correlation is significantly lower. Causal relationship clearly goes from the newspaper inflation sentiment to inflation, and simple inflation models with the sentiment as an explanatory variable beat benchmark AR model in the out-of-sample forecast. We conclude that newspaper inflation sentiment can be used as an indicator of inflationary pressures, especially during periods of high inflation volatility.

Key words: inflation forecasting, text analysis.
JEL Code: C13, C55, E31, E37, E52.

Non-technical summary

Newspapers in Serbia often discuss various topics related to prices and inflation. In this paper we use this textual data to measure newspaper inflation sentiment and check if it can be used as an indicator of inflationary pressures.

We analyzed articles from the economic sections of four Serbian daily newspapers for the period between the beginning of 2007 and June 2022. The sample consists of 117,113 articles, or roughly 650 per month on average.

We quantified the articles using a dictionary-based approach, to obtain measures of inflation sentiment. We first predefined the lists of terms, one related to price rises and the other to price falls, then counted them across months and put those counts in relation to the total number of words in the corresponding months. We also used the so-called Boolean search, in which we counted the number of articles per month that contain at least one of the terms from the lists (in relation to the total number of texts). While the term count measures the intensity of price talk, the text count measures its frequency.

We found measures of inflationary sentiment to be strongly correlated with inflation over the full sample. This is mainly driven by the co-movement during inflation upswings to double-digit figures, preceded with a pick-up in inflationary sentiment in the newspapers. In all of the four such cycles, price-rise related terms jumped over 2.5‰ of total words, while during periods of low inflation they typically stayed below 1‰. During the period of low and stable inflation (mid-2013 to mid-2021), however, the correlation between inflation and inflationary sentiment drops significantly.

Causality between inflation and inflationary sentiment could in principle go both ways, as newspapers write about both past inflation, as well as price changes that will eventually be reflected in inflation. Granger causality test suggests that in all the combinations of inflationary sentiment and CPI, causality goes from the former to the latter, although in some cases it goes both ways.

We also found that a simple inflation model with inflationary sentiment performs better than the benchmark AR model when it comes to out-of-sample forecasting.

Taking everything into account, we concluded that newspaper inflation sentiment can be used as an indicator of inflationary pressures, especially during periods of high inflation volatility, while during stable times it may not perform that well.

Contents

1 Introduction

At the time of writing this paper, global inflation is surging to levels not seen in decades, surprising monetary policy makers month after month. Could central banks have been more prepared for this shock had they given more importance to what newspapers were writing? In this paper we find that, in the case of Serbia, newspaper articles do indeed contain useful information for assesing inflationary pressures, especially during unstable times.

A large number of texts available online – such as newspaper articles, statements, blogs, social media posts, and so on – discuss economic developments in some form. Economic researchers are showing growing interest in exploiting this abundance of textual data for various kinds of analysis, like nowcasting and forecasting economic indicators, or creating measures of some economic concepts that are otherwise not measured well.

Textual analysis typically involves turning texts, as qualitative data, into a numerical measure, using some kind of method – from simple word count to sophisticated machine learning models – and estimating its relationship with observed economic variables. As texts are high-frequency data, correlation with an economic variable can be used to assess current economic trends before the actual data is published.

The early example of using text analysis dates back to 1933, when Cowles classified articles of a specific Wall Street Journal editorial as “bullish”, “bearish”, and “doubtful”, and found that it provided no useful information for forecasting stock prices. On the other hand, in his famous 2007 paper, Tetlock used the psychological Harvard IV-4 dictionary to give articles from a Wall Street Journal column a sentiment score and found it to be a good predictor of one-day-ahead stock market movements. Some researchers (Lucca and Trebbi (2009), Born et al. (2014), Hansen et. al (2018)) studied central bank communication in various ways and found that its sentiment does affect financial markets. Another famous example of text analysis is Baker et al (2016) who counted articles containing at least one word from three categories (“economy”, “policy” and “uncertainty”), and found that such measure of policy uncertainty is a predictor of investment, employment, and production.

Moving closer to our topic, a number of papers examine links between texts and inflation expectations. Carroll (2003), using simple epidemiological model, finds that during periods of intense news coverage of inflation, household expectations adjust more quickly to the expectations of professional forecasters than in the periods of low intensity of inflation-related news. Some papers deal with asymmetry in that respect: Lamla et al. (2012) find that news content sometimes induces bias, overexaggerating negative news, leading to overreaction in inflation expectations; while Drager (2015) finds that only media reports with negative news significantly affect inflation expectations. Larsen et al. (2021) find that news coverage of certain topics, some seemingly unrelated to inflation, have a high predictive power for consumers’ inflation expectations. Some researchers, like Angelico et al. (2021) analyze social media posts. They created a Twitter-based daily measure of inflation expectations (by combining Latent Dirichlet Allocation (LDA) with a dictionary-based approach) and find it to be highly correlated with conventional measures of inflation expectations.

Surprisingly, not many researchers tried to establish a direct link between texts and inflation itself, which is the very topic of our paper. One of the examples is Rambaccussing et al. (2020), who found no evidence that economic policy content in the UK media improves short-term forecast of inflation (as opposed to unemployment and output). On the other hand, Kalamara et al. (2021), by combining counts of terms with supervised machine learning, did find that newspaper articles improve forecast for UK inflation, as well as output and unemployment.

In our paper, we focus on estimating the relationship between newspaper articles and inflation in Serbia. We don't look at inflation expectations because for Serbia it is a survey-based measure, influenced by the current inflation rate, not the other way around, which will be presented later in the paper. We also believe that for practical, policy support purposes, it is more useful to have a model that forecasts inflation directly, rather than indirectly, through expectations.

The link between newspaper texts and inflation comes from the simple fact that in the periods of more frequent and intensive price changes, newspapers write about it more, i.e. what is relevant for our analysis, they use more words related to price changes. These can be mentions of past price changes, or the ones that are about to happen. It is worth noting that prices are a frequent topic in Serbian newspapers. In our sample, 22% of articles from economic sections mention inflation or price changes in some way at least once, and as we will see later on, this topic is more frequently discussed in the periods of high inflation.

We analyzed articles from four Serbian daily newspapers for the period between the beginning of 2007 and June 2022, using a dictionary-based approach and the Boolean search. Concretely, we predefined the list of terms related to price rises and price falls, counted them across the articles, and put those numbers in relation to the total number of words in the corresponding articles. The Boolean search in our case consists of counting the articles that contain at least one word from the lists (relative to the total number of articles in the corresponding months). While the text count (Boolean search) measures how often newspapers talk about prices, the term count measures how intensely prices are discussed.

This way we created monthly series of various measures of newspaper inflation sentiment (NIS): price-rise (related) terms, difference between price-rise and price-fall related terms, price-rise related articles, difference between price-rise and price-fall related articles, as well as HP-smoothed versions of the first two series.

The fact that dictionary-based and Boolean methods take predefined list of terms for search makes them simple to use, but this is sometimes taken as their weakness, as other words are completely ignored[1] (unlike some more complex machine learning methods). We believe, however, that terms related to price changes in both directions are easily identifiable, especially in the Serbian language, making dictionary methods suitable for this kind of analysis. Also, the dictionary approach is more suitable if we want to measure the intensity of the sentiment, and not just clasify texts into topics, for which supervised machine learning techniques can be more appropriate. For instance, Angelico et al. (2021) use LDA to "reduce the noise", that is to extract inflation related tweets, but then use dictionary approach to count terms related to price increases and price declines. In our case, we already reduced the noise

[1] Bholat et al. (2015)

by selecting only economic articles, thus there was no need for machine learning for this purpose.

In the final stage, we looked at the relationship between these various measures of NIS and CPI inflation (y-o-y and m-o-m seasonally adjusted). We found the two classes of series to be highly correlated over the full sample, but not during periods of stable inflation. The high correlation for the full sample is driven by the periods of inflation upswing accompanied by or preceded with a significant increase in the inflationary sentiment in the newspapers.

Granger causality test suggests that in all the combinations of NIS and CPI, causality goes from the former to the latter, although in some cases it goes in both directions. This was important to investigate as newspapers could write about past inflation, and ongoing and future price changes that will eventually be reflected in the inflation figures.

Most models with NIS beat benchmark AR models in out-of-sample forecasts. The ones with the lowest RMSE in the very short term were the model with m-o-m inflation and difference between price-rise and price-fall terms, for nowcast, and the model with HP-smoothed version of the same series, for one month ahead. Also, including NIS in the existing inflation model yields a statistically significant coefficient for the variable of interest, and improves the fit of the model.

It is important to note that the term count outperforms the text count in both correlations and forecasting performance, meaning the intensity of price talk matters for inflation more than its frequency.

Taking everything into account, we concluded that the newspaper inflation sentiment can be used as an indicator of inflationary pressures, with a caveat that it may not perform well during stable times, but rather act as a predictor of large inflation swings.

The paper is structured as follows. First, we present the database of newspaper articles analyzed in the paper. After that we explain the methods used in the text analysis to construct time series of the newspaper inflation sentiment. Finally, we present the results of the estimation of the relationship between NIS and inflation.

2 The database of newspaper articles

The first, most challenging and time-consuming step in our analysis was to create the database of economic articles to be analyzed. We included four Serbian newspapers in the database with sufficiently long history to be exploited for this kind of analysis (see Appendix for details). We downloaded the articles using Rvest package within R software, which was also used for the text analysis in the next stage. The sample covers the period between January 2007 and June 2022, with the exception of one of the newspapers, which have available articles since March 2008. Articles were downloaded directly from available archives on their websites, or indirectly, through other websites that collect news from various sources. The final database analyzed here consisted of 117,113 articles (roughly 650 per month on average) from the economic sections of these newspapers.

Chart 1 **Number of economic newspaper articles in the sample by months**

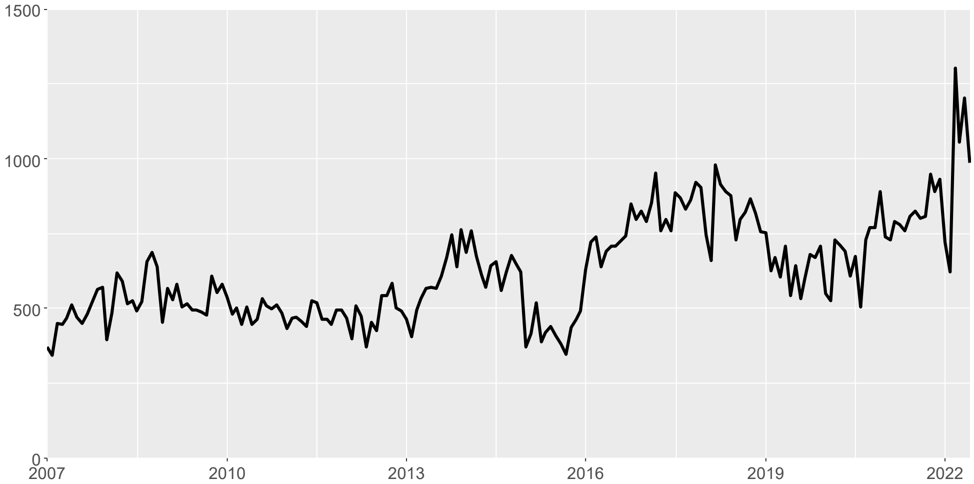

We restricted our analysis to articles from economic sections as classified by the newspapers, for several reasons. One is that including all the articles could add a lot of noise to the series, as occasional big events, such as the covid-19 pandemic, could overshadow other topics, including the ones about inflation. Also, there are words, like “crises” or “depression”, which could have both economic and non-economic meaning. Restricting our analysis only to economic articles would reduce possible mistakes in that respect, although with the terms in our specific analysis that wouldn’t be a big issue.

Figure 1 **Word cloud of economic articles of four Serbian daily newspapers**

In the whole sample, there are over 600,000 different words, out of which we are interested in only a tiny fraction. The word cloud (Figure 1) is a common way of graphically presenting the most common words in the text (the higher the frequency, the larger the font).

Before analyzing, the texts were pre-processed to leave out words that are not directly part of the article (such as "copyrights…", "BONUS video…", and so on). Once this cleaning was done, the articles were ready for the analysis.

3 Text analysis – measuring newspaper inflation sentiment

The text analysis was caried out using a dictionary-based approach, meaning that in the articles we searched for predefined specific Serbian terms related to price rises, on the one hand, and terms related to price falls, on the other hand. A term for the search in this paper is defined as a part of a word ("poskup"), or a shortened phrase ("rast cen"). The algorithm will count any word or a group of words that contain the searched terms.

A term for the search is ideally defined as a part of various types of a desired word. For instance, term "poskup" will pick up nouns in single and plural, in various case forms ("poskupljenje", "poskupljenja", "poskupljenju",…)[2], as well as verbs in different genders and tenses ("poskupeti", "poskupeo", "poskupela", "poskupljuje"…). What is also important is that this term cannot be a part of a word unrelated to prices. When it comes to phrases, the situation is a bit more complex, as only the last word can be defined in a shortened form, whereas for the first one it is necessary to predefine various types of a word ("*rast/rasta/rastu/rastom*"). Table 1 shows the lists of terms that we included in the count.

Table 1 **Lists of terms related to price rises and price falls (in Serbian)**

Searched terms related to price rises	Searched terms related to price falls
poskup	*pojeft*
skuplj	*jeftinij*
rast/rasta/rastu/rastom cen	*pad/pada/padu cen*
povećanje/povećanje/povećanju cen	*sniženje/sniženja/sniženju cen*
viša/više/višoj cen	*niža/niže/nižoj/nižu cen*
visoka/visoke/visokim/visokoj cen	*niska/niske/niskim/niskoj cen*
skok/skoka/skoku/skokovi cen	*deflaci*
digao/digli/dižu/diže cen	*pad inflaci*
inflaci	

Note that we conducted case insensitive search, meaning we treated upper and lower case letters as being the same. Thus, there was no need to transform letters from upper to lower case in the previous stage, as it is often done in this kind of analysis.

The left column in the table 1 includes various combinations of words related to increases (rise, hike, jump, increase, high, higher, inflation) and prices, while the right column combines words related to falls (fall, reduce, decline, lower, low, deflation) and prices. As an example,

[2] "Poskupljenje" is a single Serbian word for a price rise.

in the short article from 2008, based on these predefined terms the algorythm counted 10 terms related to price rises (yellow), and 2 related to price falls (blue):

A litre of petrol to cost 101.6 dinars

Unless the global oil price falls within the next two days, petrol price will go up by 4 dinars, "Blic" found out. According to these estimates, a litre of unleaded petrol would be sold at the price of 101.6 dinars.

The forecast, according to which the price of oil will speed up once it crosses the magical figure of 100 dollars per barrel, is coming true. In recent days, the price of crude oil moved somewhat below its historical peak of 103.05 dollars per barrel (159 litres). Such a high price of the main energy commodity is mainly a result of the weakening of the dollar, but also expectations that members of the OPEC will not increase their daily production. Accordingly, in Serbia another price rise of petrol is expected by the end of this week. As Blic finds out, it is likely that the petrol price increase could amount to four dinars, with a possibility that, if the oil price jumps continue, this could be even stronger. Last price increase took place on 20th February, and the Ministry of Energy confirmed that another one could come on Friday morning.

"After this price increase, I think that within the next month we will not see major jumps in the prices in Serbia. Given that the price of crude oil is near its historical record, it is realistic to expect that in the near term it could go down. However, over the longer horizon, throughout this year oil price rises are expected" – said Nebojša Atanacković from "Oil".

The lower price could be achieved by reducing the excise tax, whichwas increased during the previous price hike. This tax is set by the Ministry of Finance, but, according to earlier statements, its reduction is not expected.

We are aware that this kind of count is not perfect. In this example it included two terms related to price falls, even though they are in conditionals "…unless the global oil price falls in the next two days, the price of petrol will go up by 4 dinars…"), at the same time failing to recognize phrase about price increase in the third sentence ("The forecast, according to which the price of oil will speed up …"). Despite that, we can say that the count clearly captures the inflationary sentiment of this article, and that, given the large number of texts in this analysis (over 650 per month), these occasional failings cannot change the overall tendency in the newspaper sentiment.

Another approach we used in this analysis is the Boolean search, which consists of counting articles that satisfy certain logical conditions, in our case – articles that contain any of the searched terms from Table 1. This approach is less sensitive to extreme values, but its disadvantage is that it neglects the intensity of term usage. Note that, using this approach, a single article can be counted as both price-rise related and price-fall related, as would be the case with the article in our example.

Finally, given that the raw series are quite noisy, we also used smoothed time series using HP-filter with the smallest possible smoothing parameter (lambda=1), just to get rid of the noisiest part of the series, while keeping the short-run fluctuations in the trend.

In total we define six **measures of newspaper inflation sentiment**:

- **Price-rise terms (PRT)**, as a ratio between price-rise term count (PRTC) and total number of words (w_m) in the corresponding months (m), per thousand:

$$\mathrm{PRT}_m = \frac{\mathrm{PRTC}_m}{w_m} \cdot 1000$$

- **Price-change terms (PCT)**, as a difference between PRTC and price-fall term count (PFTC) in relation to the total number of words in the corresponding months (per thousand):

$$\mathrm{PCT}_m = \frac{\mathrm{PRCT}_m - \mathrm{PFTC}_m}{w_m} \cdot 1000$$

- **Price-rise articles (PRA)**, as a ratio between counted price-rise related articles and the total number of articles (a_m) in the corresponding months (in %):

$$\mathrm{PRA}_m = \frac{\mathrm{PRAC}_m}{a_m} \cdot 100$$

- **Price-change articles (PCA)**, as a difference between PRAC and price-fall articles count (PFAC) in relation to the total number of articles in the corresponding months (in %):

$$\mathrm{PCA}_m = \frac{\mathrm{PCA}_m - \mathrm{PRAC}_m}{a_m} \cdot 100$$

- **Price-rise terms – smoothed (PRT-HP)**, as an HP-filter of PRT, with smoothing parameter lambda = 1, and

- **Price-change terms – smoothed (PCT-HP)**, as an HP-filter of PCT, with smoothing parameter lambda = 1.

The next chart shows term-count measures, PRT and PCT. The latter can take negative values as in some periods the words related to price falls exceed those related to price rises.

Chart 2 **Newspaper inflation sentiment based on term count** (share of total monthly number of words, in ‰)

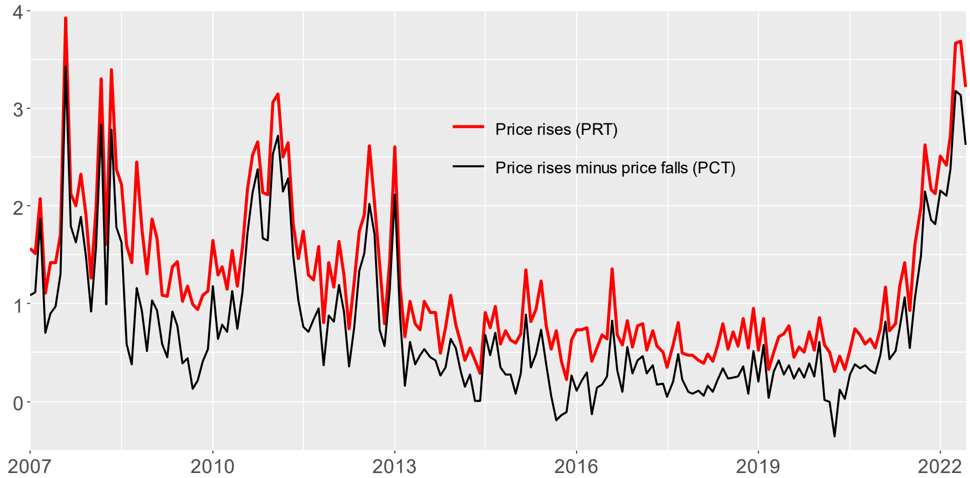

Although these time series are very noisy, one can clearly observe four major cycles in the inflationary sentiment during the analyzed period. Both series peak in 2008, 2010, 2012, and 2022, at over 2.5‰ in case of PRT, and 2‰ in case of the PCT. In the period 2013-2021 the series were rather stable, most of the time between 0 and 1 per thousand.

When it comes to the article count, we can also observe four major cycles, although less pronounced than in the case of the term count. The share of articles with at least one mention of a price increase varies between 10% of total in stable times, to 25–30% at its peaks. The series of differences peaks at around 15%, while mostly hovering between 0–5% in stable times (2013–2021).

Chart 3 **Newspaper inflation sentiment based on text count** (as a share of total monthly number of articles, in %)

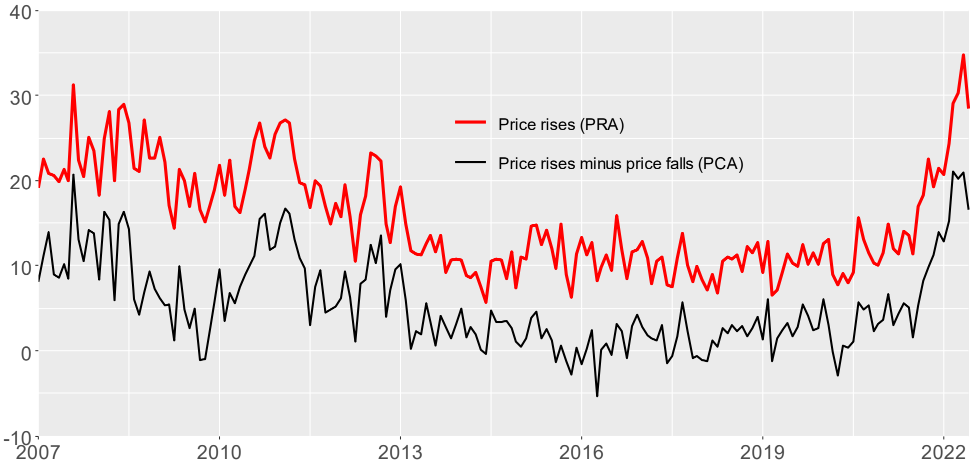

4 Relationship between newspaper inflation sentiment and inflation

There are multiple ways in which the frequency of price-related terms and articles could be related to inflation itself. For instance, a rise in inflation, once it’s published, could trigger increased media interest in inflation figures and individual price changes that drove the rise. On the other hand, newspapers could be writing about ongoing and future price changes that will eventually be reflected in the inflation figures. Another option is that inflationary sentiment in the newspapers is a reflection of inflation expectations that are themselves a driver of inflation movements.

Chart 4 shows that the time series of price-rise related terms, despite being very noisy, is clearly related to inflation. (For relationship between inflation and other measures of NIS see Appendix).

Chart 4 **Inflation and newspaper inflation sentiment**

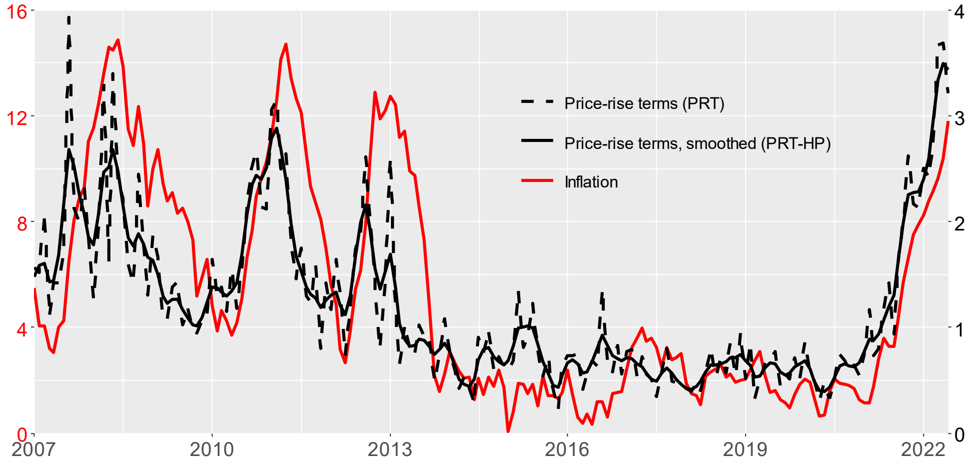

Inflation in the observed period had four major upswings, with double-digit peaks in 2008, 2011, 2012, and 2022. All these inflation cycles were accompanied by or preceeded with a significant increase in the frequency of price-rise related terms in the newspapers to over 2.5 per thousand words.

Also, during the eight-year period of inflation stability (mid-2013 to mid-2021), PRT remained comparably low, never exceeding 1.5‰, although it seems that during that period the two series occasionally diverged. Like, for instance, in 2015 when the newspaper mentions of price-rises was not followed by the inflation pick-up, or in early 2017 when the newspapers failed to predict a rise in inflation.

We also checked how measures of NIS correlate with the seasonally adjusted monthly inflation rate. In this case, the analyzed series would be fully consistent – monthly price increase with monthly term count.

Here, it's worth noting that all of the term-count series do not have seasonality, so there was no need to seasonally adjust them. Also, all of the inflationary sentiment measures, as well as two inflation series are stationary, according to Dickey-Fuller test.

Table 2 shows statistically significant correlations in all the cases over the full sample. The most strongly correlated series are y-o-y inflation and smoothed series of PRT on the 4th lag. Measures of inflation sentiment show somewhat weaker correlations with the monthly measure of inflation.

Table 2 **Peak correlations between measures of NIS and inflation for the period 2007–2022** (lags in brackets)

	terms		terms smoothed		articles	
	PRT	PCT	PRT-HP	PCT-HP	PRA	PCA
y-o-y inflation	0.75 (2)	0.71 (3)	0.86 (4)	0.83 (4)	0.76 (3)	0.72 (4)
m-o-m inflation s.a.	0.64 (0)	0.62 (0)	0.62 (0)	0.62 (0)	0.58 (0)	0.55 (0)

In line with our visual inspection, over the period of stable inflation (mid-2013 – mid-2021) correlations are much weaker (Table 3). In this case, monthly inflation rates show somewhat stronger correlation with the measures of inflationary sentiment.

Table 3 **Peak correlations between measures of NIS and inflation for the period mid-2013 – mid-2021**
(lags in brackets)

	terms		terms smoothed		articles	
	PRT	PCT	PRT-HP	PCT-HP	PRA	PCA
y-o-y inflation	0.22 (0)	0.26 (0)	0.23 (0)	0.26 (0)	0.14 (0)	0.23 (3)
m-o-m inflation s.a.	0.30 (0)	0.37 (0)	0.25 (0)	0.31 (0)	0.22 (5)	0.17 (0)

Here, it's important to note that the inflation figures in Serbia are published on the 12th day of a month for the previous month, so even if the correlations are only contemporaneous (with no lags), we can say that inflationary sentiment measures precede inflation figures, and therefore hold some predictive power.

However, before drawing such a conclusion, we need to establish the direction of the causality. Based on the Granger causality test (Table 4), measures of inflationary sentiment cause inflation in all the combinations, with causality going both ways in some of the cases. (See Appendix for detailed results).

Table 4 **Direction of Granger causality between measures of NIS and inflation*** (full sample)

	terms		terms smoothed		articles	
	PRT	PCT	PRT-HP	PCT-HP	PRA	PCA
y-o-y inflation	↙ ↗	↙	↙ ↗	↙ ↗	↙	↙ ↗
m-o-m inflation s.a.	↙	↙	↙ ↗	↙ ↗	↙	↙

*↙ means causality from inflationary sentiment to inflation, ↗ the other way around, while – means no Granger causality in either direction

Impulse response function from a VAR(7) model with inflation and PCT is a nice illustration of the relationship. As shown on Chart 5, inflation clearly reacts to inflationary sentiment in newspapers. The peak response of inflation to a 1‰ increase in PCT is 4%, which is roughly in line with observations on Chart 2 (increase in PRT from 0–0.5‰ to 2–3‰ is followed by inflation pick-up by 8-10%, from trough to peak).

Chart 5 **Impulse response from a VAR model with y-o-y inflation and PCT**

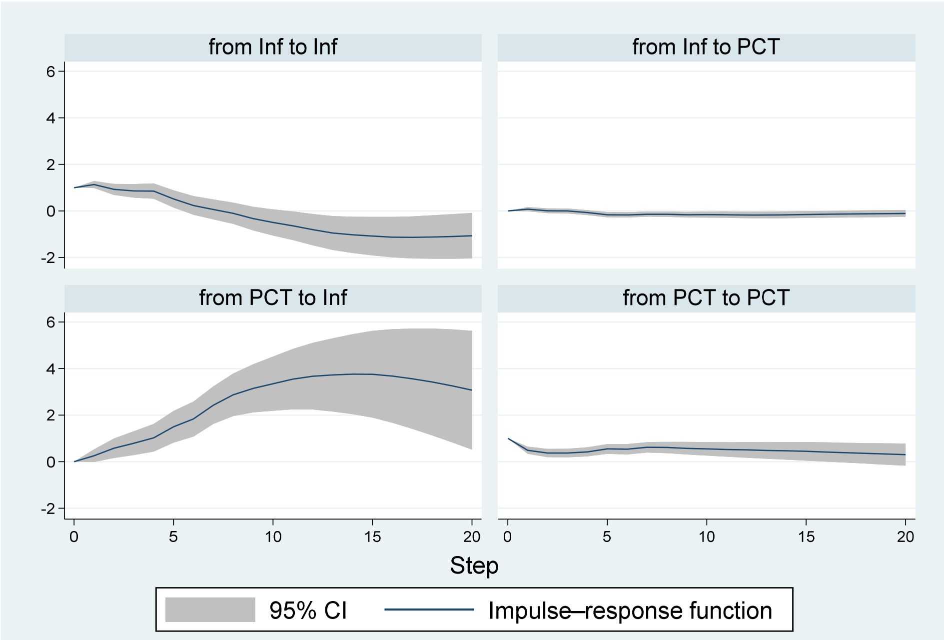

So far, we have concluded that inflationary sentiment is a leading indicator of inflation. But what could be the reason behind such a relationship? Is it that newspaper reports on ongoing and future price changes are eventually reflected in inflation figures; or could it be that these reports reflect inflation expectations, which are themselves a factor of inflation?

If the latter was the case, we should see inflation expectations causing inflation. However, Granger causality test tells a different story, showing that inflation expectations are Granger-caused by inflation itself, and not the other way around (Appendix).

Impulse response function from a VAR(2) model with inflation, PCT and household inflation expectations (Exp) show that inflation expectations react to PCT, as well as to inflation itself. We can also observe that the inclusion of expectations in the model, doesn't change our conclusion that PCT causes inflation, and not the other way around (Chart 6).

Chart 6 **Impulse response from a VAR model with y-o-y inflation, PCT, and inflation expectations**

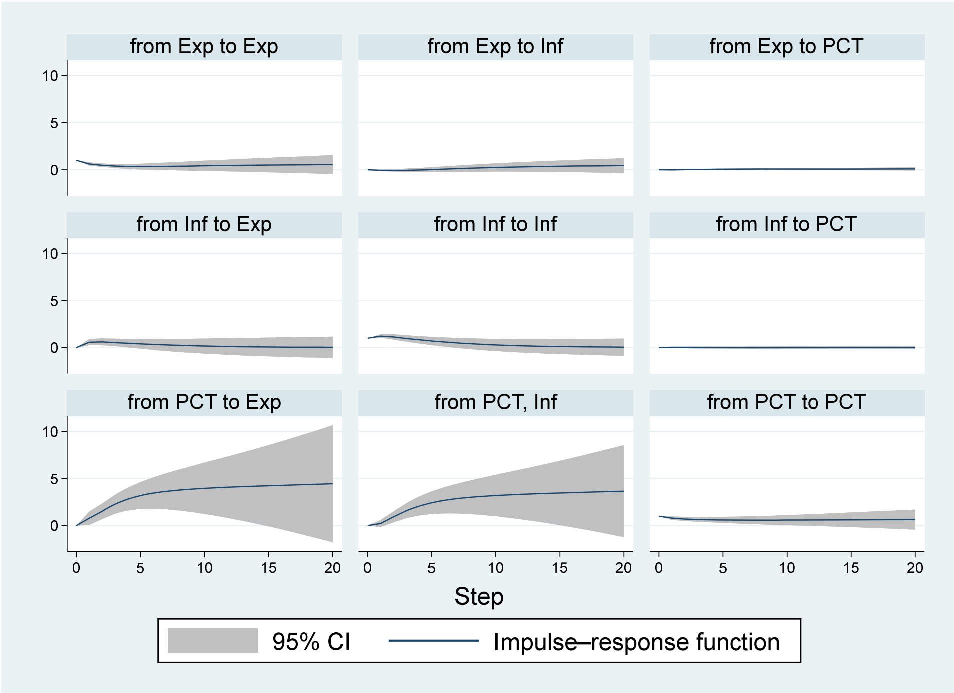

Now, could we exploit this relationship for inflation forecasts? To make a verdict, we did a standard out-of-sample forecast exercise with models including various inflationary sentiment measures and two inflation measures, comparing them to AR(1) models for y-o-y and m-o-m s.a. inflation (Table 5). We tested VAR(2) models, as well two-equation models, where one is linking inflation to various NIS measures, and the other one being an AR(2) equation for NIS:

$$Inf_m = \alpha_1 \cdot Inf_{m-1} + \alpha_2 \cdot NIS_{m-1} + \varepsilon_m$$

$$NIS_m = \beta_1 \cdot NIS_{m-1} + \beta_2 \cdot NIS_{m-2} + \gamma_m$$

To make it comparable to simulations with y-o-y inflation, in case of m-o-m inflation we calculate RMSE for cumulative forecasting errors. See Appendix for details.

The RMSEs of various models are shown in Table 5.

Table 5 **Out-of-sample RMSE for different inflation models with NIS measures (leads in columns)**

	nowcast	1	2	3	4	5	6
Modes with y-o-y inflation							
AR(1) - benchmark	0.613	0.613	0.929	1.172	1.384	1.556	1.739
PRT model	0.567	0.549	0.770	0.958	1.145	1.307	1.516
PCT model	0.558	0.548	0.770	0.968	1.160	1.303	1.499
PRA model	0.582	0.569	0.843	1.062	1.253	1.411	1.574
PCA model	0.601	0.563	0.824	1.054	1.249	1.414	1.609
PRT VAR	NA	0.662	1.007	1.286	1.583	1.893	2.194

	nowcast	1	2	3	4	5	6
Modes with y-o-y inflation							
PCT VAR	NA	0.658	0.994	1.275	1.576	1.880	2.189
PRT-HP model	0.593	0.654	0.963	1.204	1.447	1.583	1.742
PCT-HP model	0.586	0.641	0.956	1.240	1.520	1.652	1.804
Modes with m-o-m inflation							
AR(1) - benchmark	0.396	0.396	0.699	1.016	1.345	1.649	1.969
PRT model	0.324	0.338	0.522	0.704	0.890	1.048	1.258
PCT model	0.322	0.342	0.540	0.748	0.960	1.141	1.372
PRA model	0.334	0.338	0.527	0.712	0.879	1.004	1.188
PCA model	0.347	0.342	0.546	0.766	0.974	1.156	1.400
PRT VAR	NA	0.350	0.552	0.764	0.986	1.176	1.413
PCT VAR	NA	0.346	0.536	0.728	0.932	1.111	1.336
PRT-HP model	0.329	0.262	0.413	0.582	0.763	0.891	1.057
PCT-HP model	0.328	0.256	0.422	0.629	0.830	0.969	1.153

We can see that the models with NIS improve forecasting performance relative to benchmark AR(1) models in almost all the cases, with the exception of VAR models with y-o-y inflation. Models with monthly inflation perform better than the ones with y-o-y inflation. The lowest RMSE is achieved for models linking monthly inflation to some of the term count measures: PCT for nowcast, PCT-HP for one lag, and PRT beyond the 1st lag.

The reason we used ARs as benchmark models is that the NBS short-term inflation forecasts are made on disaggregated inflation data, as a combination of model-based approach (for some components), judgement and external information on various price changes. There is an NBS medium-term projection model described in Đukić, Momčilović, Trajčev (2010), but we don't have an ambition to use newspaper sentiment for the medium-term forecasts.

Finally, we checked if NIS brings some additional information to models with usual factors of inflation. For that purpose we used ARDL model for the headline inflation in Serbia from Ivković, Jakovljević, Miletić (2022), that includes euro area inflation, agricultural product prices, wages and exchange rate as explanatory variables. In that model we introduced PCT as another variable, and found that it is statistically significant, both in the short and long run (Table 6). Therefore, we can say that NIS holds explanatory power of inflation, even when we control for main factors of inflation. As this is a single equation model, it cannot be used for forecasting purposes, so that's why we didn't use it as a benchmark model in the out-of-sample forecast exercise.

Table 6 **Estimated results of the ARDL inflation model**

Conditional Error Correction Regression

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	-0.632643	0.375823	-1.683354	0.0991
LOG(PRICES_SA(-1))*	-0.082318	0.027850	-2.955764	0.0049
LOG(EA_PRICES_SA)**	0.199081	0.135368	1.470666	0.1482
LOG(AGRIC_PRICES)**	-0.003742	0.008988	-0.416369	0.6791
LOG(WAGES_SA(-1))	0.018141	0.020716	0.875701	0.3857
LOG(NEER)**	-0.011635	0.021442	-0.542639	0.5900
PCT**	0.016931	0.002373	7.133902	0.0000
DLOG(WAGES_SA)	-0.193639	0.062357	-3.105336	0.0033
VES2012Q4	0.035339	0.005320	6.642341	0.0000

* p-value incompatible with t-Bounds distribution.
** Variable interpreted as Z = Z(-1) + D(Z).

Levels Equation
Case 2: Restricted Constant and No Trend

Variable	Coefficient	Std. Error	t-Statistic	Prob.
LOG(EA_PRICES_SA)	2.418436	1.214257	1.991700	0.0524
LOG(AGRIC_PRICES)	-0.045461	0.137099	-0.331593	0.7417
LOG(WAGES_SA)	0.220374	0.257840	0.854692	0.3972
LOG(NEER)	-0.141345	0.321418	-0.439756	0.6622
PCT	0.205681	0.097416	2.111374	0.0402
C	-7.685328	2.595410	-2.961123	0.0048

EC = LOG(PRICES_SA) - (2.4184*LOG(EA_PRICES_SA) -0.0455
*LOG(AGRIC_PRICES) + 0.2204*LOG(WAGES_SA) -0.1413*LOG(NEER)
+ 0.2057*RAZLIKA -7.6853)

F-Bounds Test — Null Hypothesis: No levels relationship

Test Statistic	Value	Signif.	I(0)	I(1)
			Asymptotic: n=1000	
F-statistic	37.18888	10%	2.08	3
k	5	5%	2.39	3.38
		2.5%	2.7	3.73
		1%	3.06	4.15
			Finite Sample: n=55	
Actual Sample Size	55			
		10%	2.226	3.241
		5%	2.617	3.743
		1%	3.543	4.839

5 Conclusion

The purpose of the analysis presented in this paper was to investigate if newspaper articles in Serbia contain information that could be used as an indicator of inflationary pressures. We found that they indeed do.

The text analysis was carried out by applying a dictionary-based approach and Boolean search on four Serbian newspapers, for the period between the beginning of 2007 to June 2022. We first predefined two lists of terms for the search – one related to price rises, and the other one related to price falls – then we counted those terms, as well as articles containing at least one of the terms from the lists, by months. Using that information, we created several measures of newspaper inflation sentiment: price-rise terms, price-change terms, price-rise articles, price-change articles, HP-smoothed price-rise terms, and HP-smoothed price-change terms.

These measures were found to be highly correlated with inflation over the full sample. The correlation is stronger with y-o-y inflation with a certain lag (around 0.8), than with m-o-m inflation (around 0.6). The relationship between inflation and the sentiment is striking during periods of high inflation volatility. For instance, in all four cycles when inflation exceeded 10%, price-rise related terms jumped over 2.5‰ of total words, while during periods of low inflation they tipically remained below 1‰. On the other hand, during stable times this link seems to be much weaker. Thus, restricting the sample to the period of low and stable inflation (mid-2013 to mid-2021), the correlation between NIS and inflation drops to 0.2–0.3, way lower than for the full sample.

It was also important to establish the direction of causality, which could in principle go both ways, as newspapers write about about both past inflation, as well as price changes that will eventually be reflected in inflation. Granger causality test suggests that in all the combinations of NIS and CPI, causality goes from the former to the latter, although in some cases it goes both ways. This suggests that NIS holds predictive power for inflation. Impulse response function from VAR model with price-rise terms and inflation suggests a one-way causal relationship, where a 1‰ jump in price-rise terms predicts a 4% rise in inflation. This is roughly in line with observations from Chart 4, where we can see that upswings in inflation of 8–10 percentage points were accompanied by jumps in price-rise related terms of 2–2.5‰.

A number of papers treat textual information as a reflection of inflation expectations, which themselves are a factor of inflation. For that reason, we investigated whether the causal relationship between the sentiment and inflation works through the expectations. This notion is disputed by the Granger causality test and VAR model with the sentiment, inflation and inflation expectations. Expextations were found to lag inflation movements, while the relationship between the sentiment and inflation doesn't change significantly after including expectations in the model. Therefore, we conclude that the main reason behind the causality is the information content in the newspapers about price changes that are eventually reflected in the inflation figures.

Next, we found that various combinations of NIS measures and the two inflation measures in most cases perform better against the benchmark AR inflation models when it comes to out-of-sample forecasts. Although correlations were stronger with the y-o-y inflation, the models with m-o-m inflation showed better forecasting performance. In the very short term, the best performing models were the one with m-o-m inflation and difference between price-rise and price-fall terms, for nowcast, and the model with HP-smoothed version of the same series, for one month ahead.

It is worth noting that in each stage of the analysis, the term count measures outperformed the text count measures, meaning the intensity of price talk is more relevant than its frequency.

Finally, we included price-rise terms in a quarterly headline inflation model published in Ivković, Jakovljević and Miletić (2022) and found that the coefficient for NIS is statistically significant, even when most relevant factors of inflation are included in the model.

Taking everything said into account, we conclude that newspaper inflation sentiment is a good indicator of inflationary pressures, with a caveat that it may not be the case in stable times. Its usefulness may, therefore, come mainly in sounding alarm bells about possible major inflation upswings, such as the one observed globally since late 2021.

Appendix

Out-of-sample forecast

This exercise is used to check forecasting performance of the models on the history. We run forecasts with small inflation models with various measures of NIS (PRT, PCT, PRA, PCA, PRT-HP and PCT-HP), and autoregresive equations for NIS:

$Inf_m = \alpha_1 \cdot Inf_{m-1} + \alpha_2 \cdot NIS_{m-1} + \varepsilon_m$

$NIS_m = \beta_1 \cdot NIS_{m-1} + \beta_2 \cdot NIS_{m-2} + \gamma_m$

In this “horse-race” we also included VAR(2) models with Inf and NIS, while as a benchmark models we used AR(2) models for inflation.

We carried out the exercise in the usual manner: we first estimate models for the sample that ends at a certain moment in history (we begin with June 2014), forecast the inflation using that model for 0 to 6 months ahead, and then calculate the deviation of the forecast from the actual values of that variable for all the leads. We then extend the estimation sample month by month, and do the same excercise again, until we reach the last month in history. Finally, we calculate root-mean-squared error, as a root of average squared deviation of forecast to the actual values for all the samples, for 0 to 6 leads. RMSE for lead *l* would be calculated in the following way:

$$\mathrm{RMSE}_l = \sqrt{\frac{\sum_{m=t+1}^{endhist}(\mathrm{Inf}_{m+l}^{fcast} - \mathrm{Inf}_{m+l})^2}{endhist - m}}$$

where t is the last period of the estimation sample, and endhist the last month of the full sample (June 2022).

When it comes to monthly inflation models, we cumulated the errors, thus making them comparable to y-o-y forecast errors:

$$\mathrm{RMSE}_l = \sqrt{\frac{\sum_{m=t+1}^{endhist}(\sum_{i=1}^{l}(\mathrm{Inf}_{m+i}^{fcast} - \mathrm{Inf}_{m+\mathrm{i}}))^2}{endhist - m}}$$

Here, it's important to explain how we calculated the nowcasting error (0 leads), since literally speaking it is not an out-of-sample forecast. We did it by forecasting one-period ahead inflation with actual, realized NIS measure for *t*+1 (which makes it out-of-sample). The idea is to simulate the situation at the end of a month where we have indicators of newspaper sentiment, but we still don't have the inflation figure for that month. This is different from out-of-sample forecast for one month ahead, where we run forecast not with an actual NIS, but with a model-based prediction.

Chart A1 **Inflation and indicators of newspaper inflation sentiment**

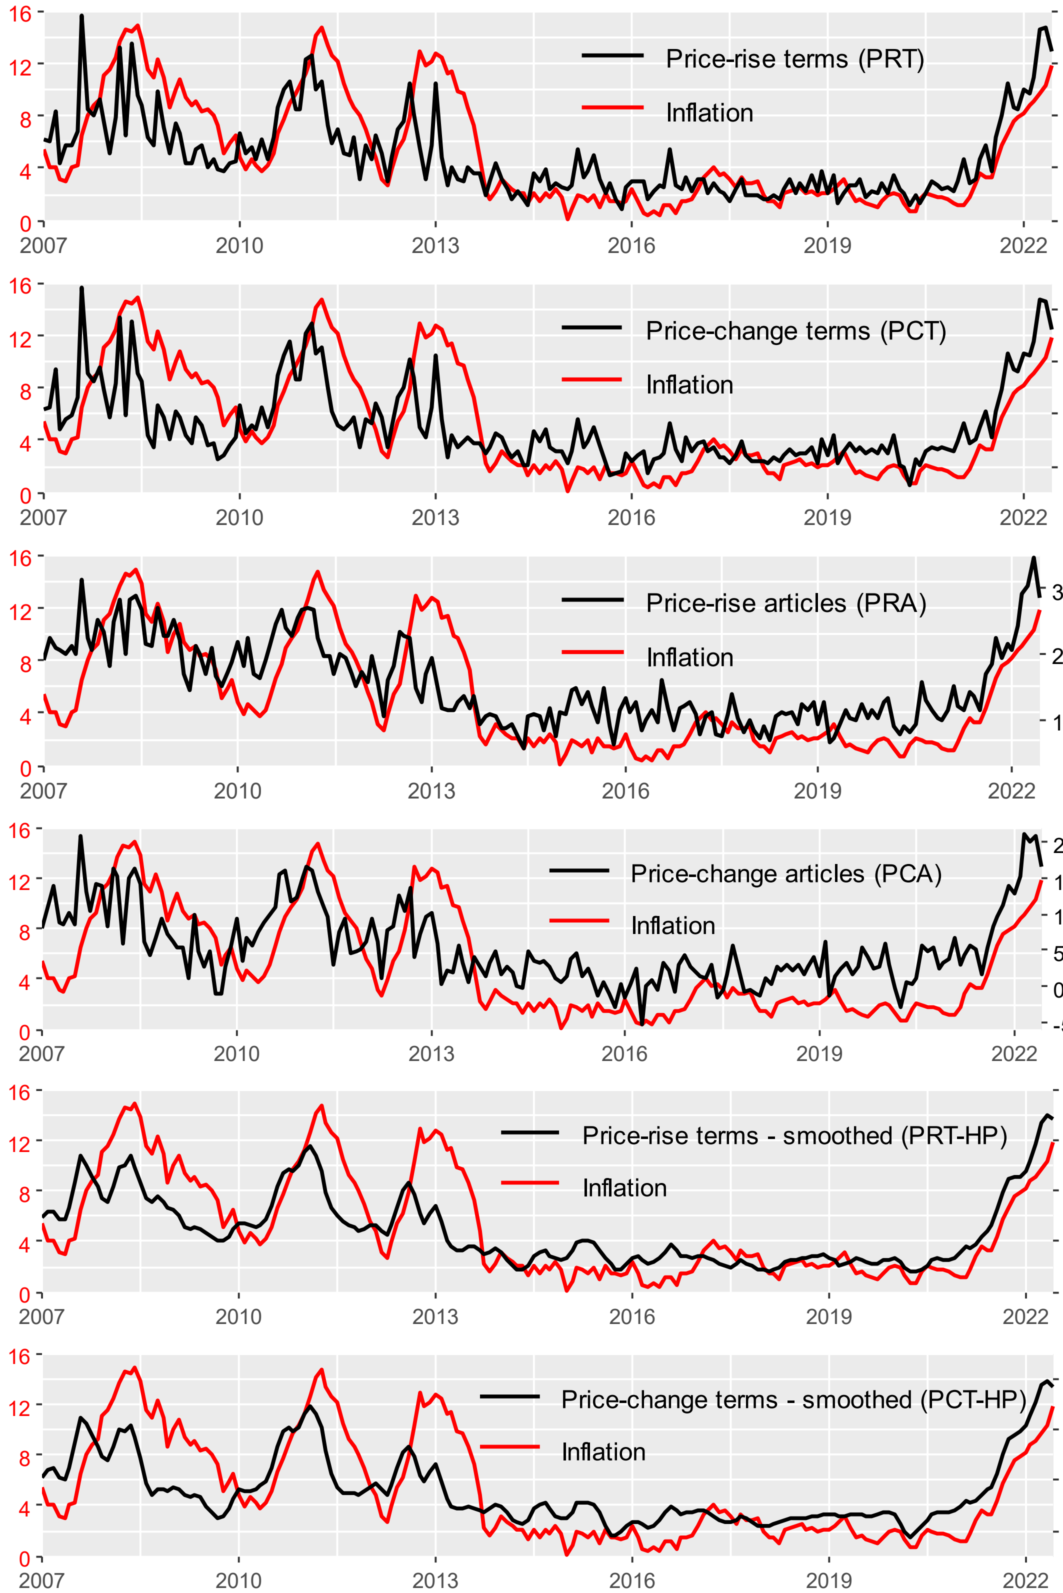

Chart A2 **Monthly inflation and indicators of newspaper inflation sentiment**

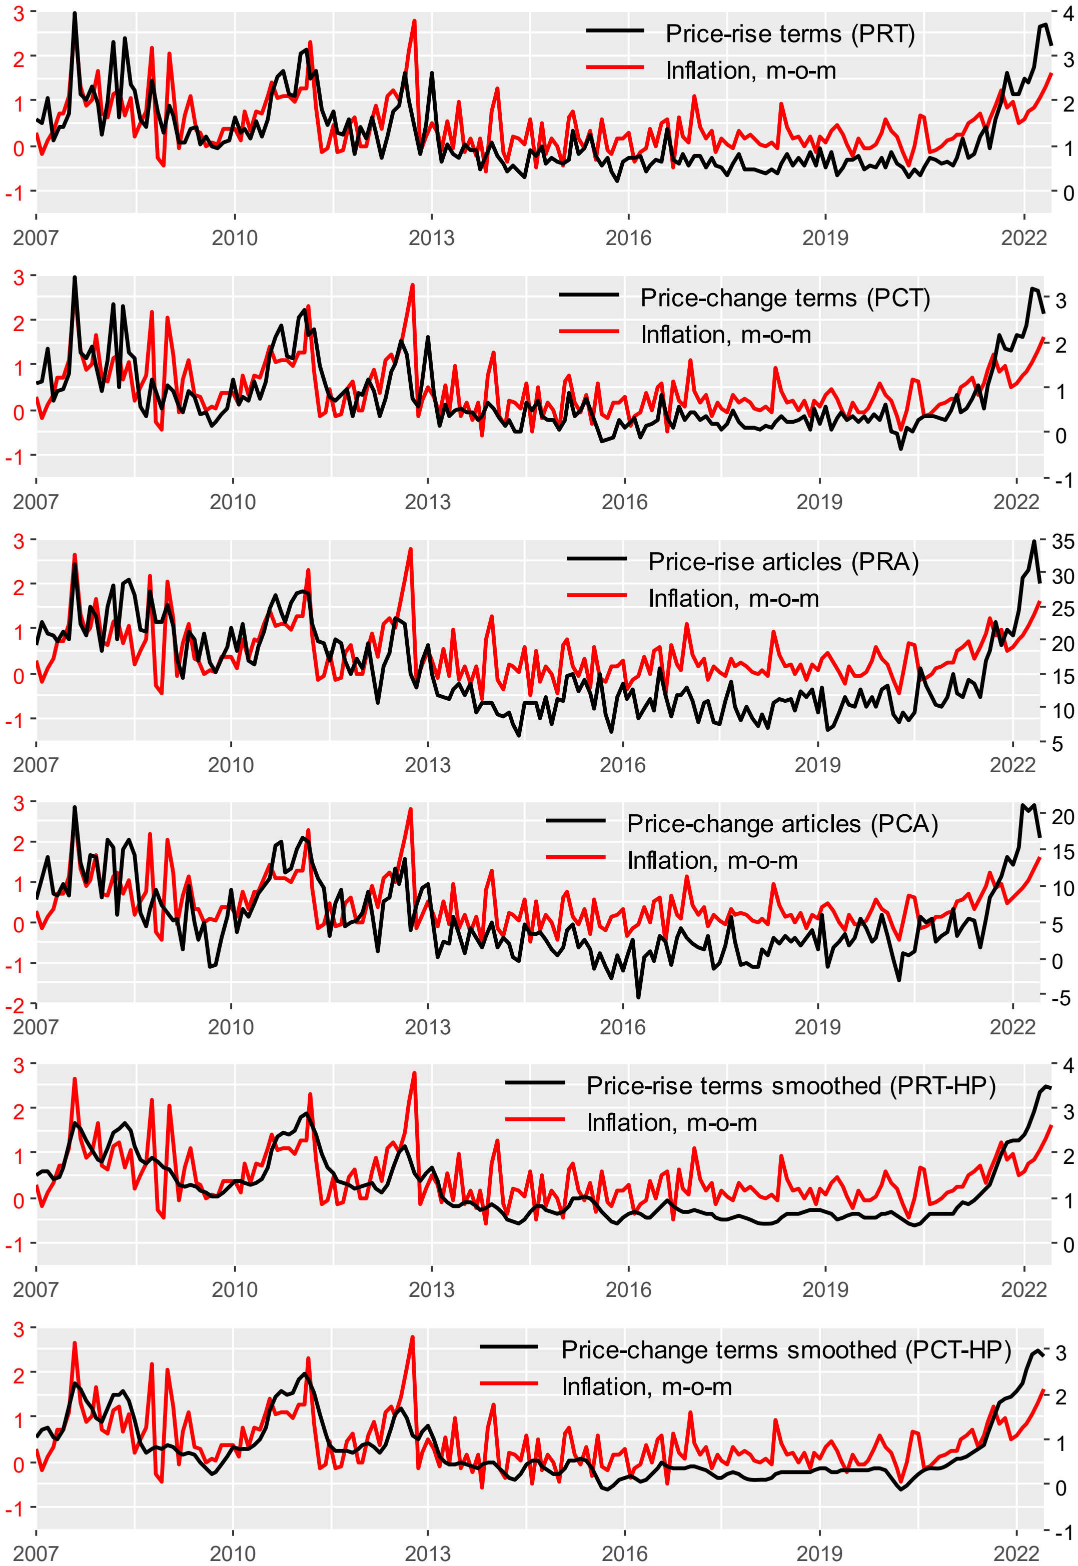

Table A1 **Granger causality between measures of NIS and inflation ** (full sample)**

	terms		terms smoothed		articles	
	PRT	PCT	PRT-HP	PCT-HP	PRA	PCA
y-o-y inflation	5.95017* 3.04455*	11.4523* 0.98914	18.5593* 5.84182*	19.2880* 4.70734*	5.10133* 1.73623	6.73472* 2.56426*
m-o-m inflation s.a.	10.8278* 2.31504	2.31504* 0.80398	11.8014* 5.08440*	12.1729* 3.90030*	10.2410* 2.36838	11.2649* 1.95995

* means statistical significance at the 5% level, meaning there is causality

** first term in the field is F statistics for causality from NIS to inflation, and second term the other way around

Table A2 **Granger causality test between inflation and inflation expectations**

From expectaions to inflation: F = 0.35442 (p=0.7023)
From inflation to expectaions: F = 15.1907* (p=0.0000)

The database of articles

The database includes four Serbian newspapers: Politika, Večernje Novosti, Danas and Blic. For three of them (Politika, Danas, and Blic), we found available articles since January 2007, while for Novosti the sample begins with March 2008. While some of the newspapers (Politika and Danas) have publicly available archives of articles on their websites, others (V. Novosti and Blic) don't, so we had to download their articles through other websites that collect news from various sources. We used only articles from economic sections.

References

Angelico C., Marcucci J., Miccoli M. & Quarta F., (2021). “Can we measure inflation expectations using Twitter?,” Temi di discussione (Economic working papers) 1318, Bank of Italy.

Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. The quarterly journal of economics, 131(4), 1593-1636.

Bholat D., Hans S., Santos P. & Schonhardt-Bailey C. (2015). Text mining for central banks, Handbooks, Centre for Central Banking Studies, Bank of England, number 33.

Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. The Quarterly Journal of Economics, 118(1), 269-298.

Cowles A. (1933). Can Stock Market Forecasters Forecast? Econometrica, 1(3), 309–324.

Dräger, L. (2015). Inflation perceptions and expectations in Sweden – Are media reports the missing link?. Oxford Bulletin of Economics and Statistics, 77(5), 681-700.

Dräger, L., & Lamla, M. J. (2012). Updating inflation expectations: Evidence from micro-data. Economics Letters, 117(3), 807-810.

Đukić M., Momčilović J. & Trajčev Lj., 2010. Medium-term projection model of the National Bank of Serbia, Working papers 17, National Bank of Serbia.

Hansen, S, McMahon m, and Tong M. (2018) The long-run information effect of central bank narrative. Working paper, Oxford University and Bank of England.

Kalamara, E., Turrell, A., Redl, C., Kapetanios, G., & Kapadia, S. (2020). Making text count: economic forecasting using newspaper text. Staff Working Paper No. 865, Bank of England.

Ivković A., Jakovljević S., Miletić M. (2022). Estimation of the impact of global and domestic factors on inflation in Serbia. Working Papers Bulletin II – March 2022, National Bank of Serbia.

Larsen, V. H., Thorsrud, L. A., & Zhulanova, J. (2021). News-driven inflation expectations and information rigidities. Journal of Monetary Economics, 117, 507-520.

Lucca D. & Trebbi F., (2009). “Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements,” NBER Working Papers 15367, National Bureau of Economic Research, Inc.

Rambaccussing, D., & Kwiatkowski, A. (2020). Forecasting with news sentiment: Evidence with UK newspapers. International Journal of Forecasting, 36(4), 1501-1516.

Tetlock. P. (2007). Giving Content to Investor Sentiment: The Role of Media in the Stock Market. The Journal of Finance, 62(3), 1139–1168.