Huber, Florian; Schreiner, Josef

Article

Are Phillips curves in CESEE still alive and well behaved?

Are Phillips curves in CESEE still alive and well behaved?

Florian Huber, Josef Schreiner[1]

This paper estimates Phillips curve relationships using nonparametric vector autoregressions for four Central, Eastern and Southeastern European (CESEE) economies. The novel feature of our model, which builds on Bayesian additive regression trees, is that it allows for unveiling possible asymmetries with respect to the size and sign of structural shocks. We simulate how unexpected movements in the unemployment rate impact inflation measures across the countries under consideration. We provide evidence that the reactions of inflation to labor market shocks are highly asymmetric: Small shocks trigger no statistically significant response of inflation whereas larger shocks induce strong, significant and persistent reactions for all countries in our sample. When focusing on differences between positive and negative unemployment shocks, we find that benign shocks lead to stronger price reactions than adverse movements in unemployment rates. These results all highlight substantial nonlinearities in the dynamic relationship between unemployment rates and inflation rates.

JEL classification: E31, E32, E50
Keywords: Phillips correlation, Bayesian vector autoregressions, business cycle shocks, asymmetries

Inflation is currently the most pressing topic on economic policymakers' agendas in Central, Eastern and Southeastern Europe (CESEE) and, in fact, around the globe. Among the most frequently cited reasons for accelerating inflation are the rebound of demand following COVID-19-related lockdowns, combined with emerging demand-supply mismatches, the strengthening of households' net financial wealth during the pandemic, adverse weather conditions (e.g. droughts) in some parts of the world and, more recently, the economic consequences of Russia's war against Ukraine (and primarily its impact on energy and food prices). Policymakers in CESEE have acted proactively and decisively to tame surging inflation by raising monetary policy rates to historic levels. As of early 2023, they have been successful in putting a break on ever-increasing headline inflation rates, helped by a rebalancing of European energy demand and the associated decrease of energy prices. Core inflation rates, however, kept on rising unabatedly. At the same time, CESEE labor markets are in full swing, as the COVID-19 pandemic has not left any lasting scars: Once the pandemic-related restrictions were lifted, the region's labor markets quickly returned to practically full employment amid tight labor supply, occasional skill mismatches and accelerating wage growth.

The question arises of how much (if any) of the 2021–2022 price surge in CESEE can be attributed to labor market tightness. This question is usually addressed through the lens of the Phillips curve, a concept that generally postulates a negative relationship between measures of economic slack (in our case, the unemployment rate) and inflation rates, meaning that tighter labor markets cause inflation rates to rise. In CESEE, however, there has been a visible disconnect

[1] *Oesterreichische Nationalbank, Central, Eastern and Southeastern Europe Section, josef.schreiner@oenb.at; Paris Lodron University of Salzburg, Department of Economics, florian.huber@plus.ac.at. Opinions expressed by the authors of studies do not necessarily reflect the official viewpoint of the OeNB, the Eurosystem or the Paris Lodron University of Salzburg.* *The authors would like to thank Fabrice Kreuzbichler (Paris Lodron University of Salzburg) for excellent research assistance.*

between labor market and price developments for at least the past decade. While this has caused quite some confusion among economists analyzing the CESEE region, it fits well into empirical evidence gathered for many advanced economies that found a flattening of the Phillips curve since the mid-1990s (see e.g. Kuttner and Robinson, 2010; IMF, 2013).

The motivation of this paper is, therefore, to examine the Phillips curve in a sample of CESEE countries and to see whether it is still alive and well behaved. This broader topic has drawn considerable attention in economic literature over the past years (see e.g. Stock and Watson, 2019; Del Negro et al., 2020). Present literature deals with this question by looking at parametric econometric models that take a strong stance on the nature of any nonlinearities the Phillips curve might exhibit. To circumvent introducing strong prior assumptions on the functional relationship between prices and real economic activity, we use the nonparametric multivariate time series model originally developed in Huber and Rossini (2022) and Huber et al. (2023). This model uses Bayesian additive regression trees (BART, see Chipman et al., 2010) to handle structural breaks, changing trends and any form of nonlinearities in the conditional mean.

Empirically, we are interested in how shocks to the unemployment rate impact inflation as measured by the Harmonised Index of Consumer Prices (HICP). The flexibility of our model allows us to focus on whether shocks to the unemployment rate trigger a nonproportional reaction of inflation or whether positive shocks trigger different inflation reactions than negative shocks. By decomposing core inflation into a cyclical and a noncyclical component, we can then assess whether labor market shocks have the potential to shift trend inflation or whether variations in core inflation are purely driven by the reaction of the transitory component.

This paper is structured as follows: Section 1 describes the data and shows some descriptive statistics. Section 2 introduces the econometric framework we employ, briefly discusses the prior setup and outlines our estimation strategy. Section 3 presents our empirical results, including the impulse response of different measures of inflation to unemployment shocks. Section 4 discusses the results. Finally, section 5 puts the results into context and applies them to rationalize some stylized facts of CESEE's recent inflation history; it elaborates on some policy implications and further research questions and concludes the paper.

1 Data description

In our analysis, we concentrate on the link between unemployment rates and HICP inflation in four CESEE EU member states that conduct an independent monetary policy: Czechia, Hungary, Poland and Romania. We collected quarterly data over the period from Q1 00 to Q2 22. The dataset includes the following series: annual percentage changes in HICP inflation, core inflation (defined as overall HICP inflation excluding energy and unprocessed food), cyclical and noncyclical items within core inflation, real GDP (seasonally and working day adjusted), the nominal compensation per employee (whole economy, seasonally and working day adjusted) and nominal unit labor costs (whole economy, per person, seasonally and working day adjusted) as well as the unemployment rate (EU Labor Force Survey methodology), the respective policy rate, the three-month money market rate and ten-year government bond yields in percent.

The distinction between cyclical and noncyclical core inflation follows the methodology used in Lian and Freitag (2022). In particular, core inflation items (on the two-digit level of the HICP classification or – in areas that include energy – on the three-digit level) are split into cyclical and noncyclical components based on their average correlation with a simple HP-filtered output gap. Cyclical components include rentals for housing; maintenance and repair of dwellings; furnishings, household equipment and routine household maintenance; health; transport services; recreation and culture; restaurants and hotels; and processed food including alcohol and tobacco. Noncyclical components include clothing and footwear; water supply and miscellaneous services relating to the dwelling; purchase of vehicles; communications; education; and miscellaneous goods and services. Individual cyclical and noncyclical components are aggregated using their country-specific item weights in the consumption basket.

Chart 1 shows the unemployment rate and annual changes in HICP inflation in the four CESEE countries in our sample. In the period under review, price developments in CESEE were characterized by a broad-based disinflation trend that lasted approximately up to 2005, reflecting economic stabilization after the early years of transition, increased competition (especially at the international level), a monetary policy shift away from exchange rate stabilization toward inflation targeting in many countries and – later on – a stronger reform momentum in the run-up to EU accession. Unemployment rates trended up moderately but basically did not move too much over the first five years in our time series. As always, there is some variation across countries. Disinflation, for example, was especially pronounced in Romania, as the country had experienced a period of very high price growth after its currency reform and the elimination of most price controls in the late 1990s.

In the boom years following the 2004 EU enlargement round, prices trended up again, reflecting buoyant (partly credit-fueled) domestic demand and record-high GDP growth as well as tightening labor markets amid continuing emigration. Consequently, the unemployment rates also declined in all countries in our sample except Hungary, where pronounced macroeconomic imbalances and high fiscal deficits weighed on growth and the labor market.

The crisis of 2008 and the subsequent years put an end to this phase and sent prices on a downward trend. This trend – temporarily interrupted between 2011 and 2013, when oil prices climbed to above USD 100 per barrel – culminated in a period of deflation around 2015 and 2016. Up until the pandemic and the recent price boost, inflation rates only increased very moderately, hovering between 2% and 4%. This is even more striking as the four CESEE countries under review experienced a period of swift economic expansion between 2014 and 2019 after economic imbalances and crisis legacies from 2008 were finally cleaned up. In this boom period, unemployment rates embarked on a remarkable downward path (from the elevated levels they had reached during the Great Recession) and reached historically low levels on the eve of the pandemic. Even the disruptions caused by the COVID-19-related lockdowns did not persistently alleviate labor market tightness. Unemployment rates only increased moderately in 2020 and early 2021 and then quickly returned to their pre-pandemic levels.

Chart 1

Development of HICP inflation and unemployment rates

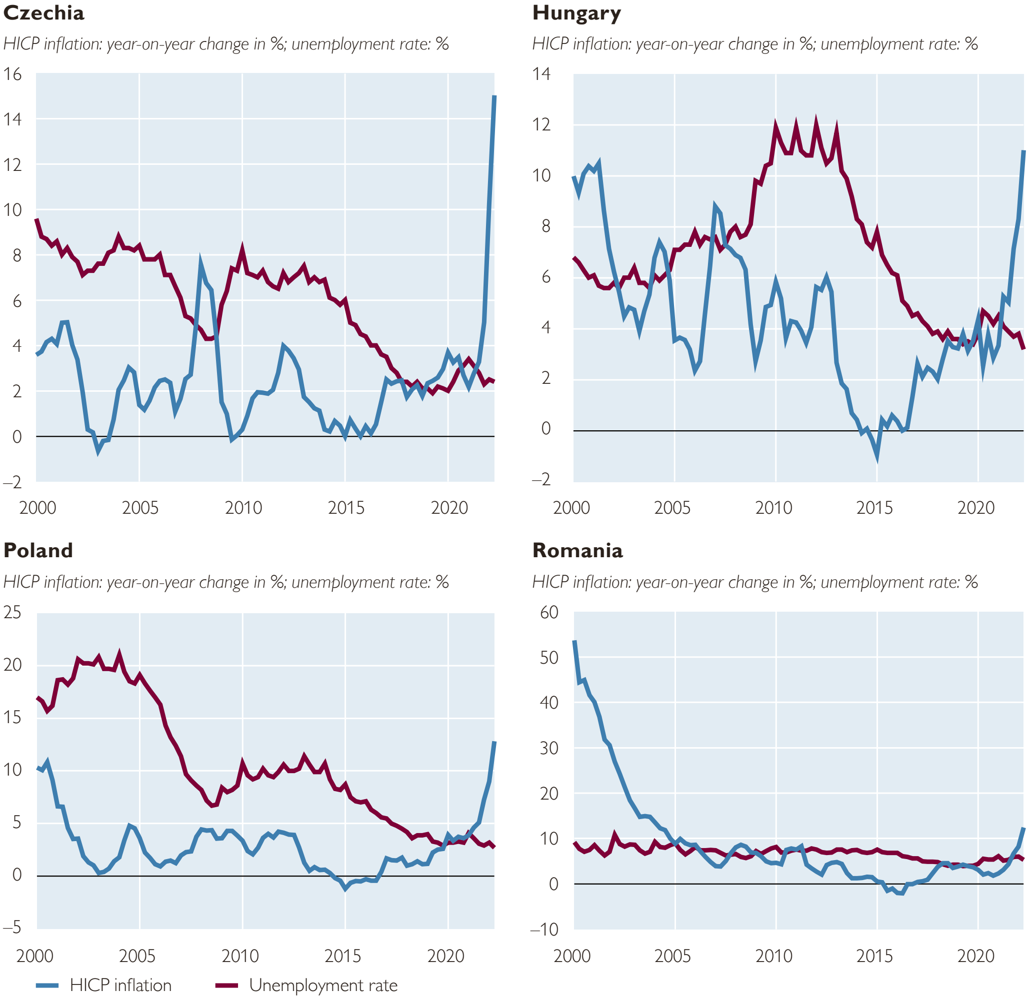

Source: Eurostat.

It is already clear from the remarks above that unemployment and price trends in the countries under review disconnected at least for certain periods. Chart 2 shows the correlation between HICP inflation and the unemployment rate both over the entire observation period and for a moving window of 20 quarters. Over the whole period, correlation coefficients range from –0.25 in Czechia to 0.43 in Romania.

For different time frames, correlations show some distinct patterns across countries: large and negative correlation coefficients in the years surrounding the economic downturn after the financial crisis, large and positive correlation coefficients for the deflationary period around 2015, and large and negative correlation coefficients for the boom years preceding the pandemic. At the most recent end of our sample, correlations weakened notably again.

Chart 3 shows the development of core inflation and its cyclical and noncyclical components. Core inflation very much mimics the dynamics of headline inflation throughout the sample and over the entire observation period. However, it was

Chart 2

Correlation between HICP inflation and unemployment rates

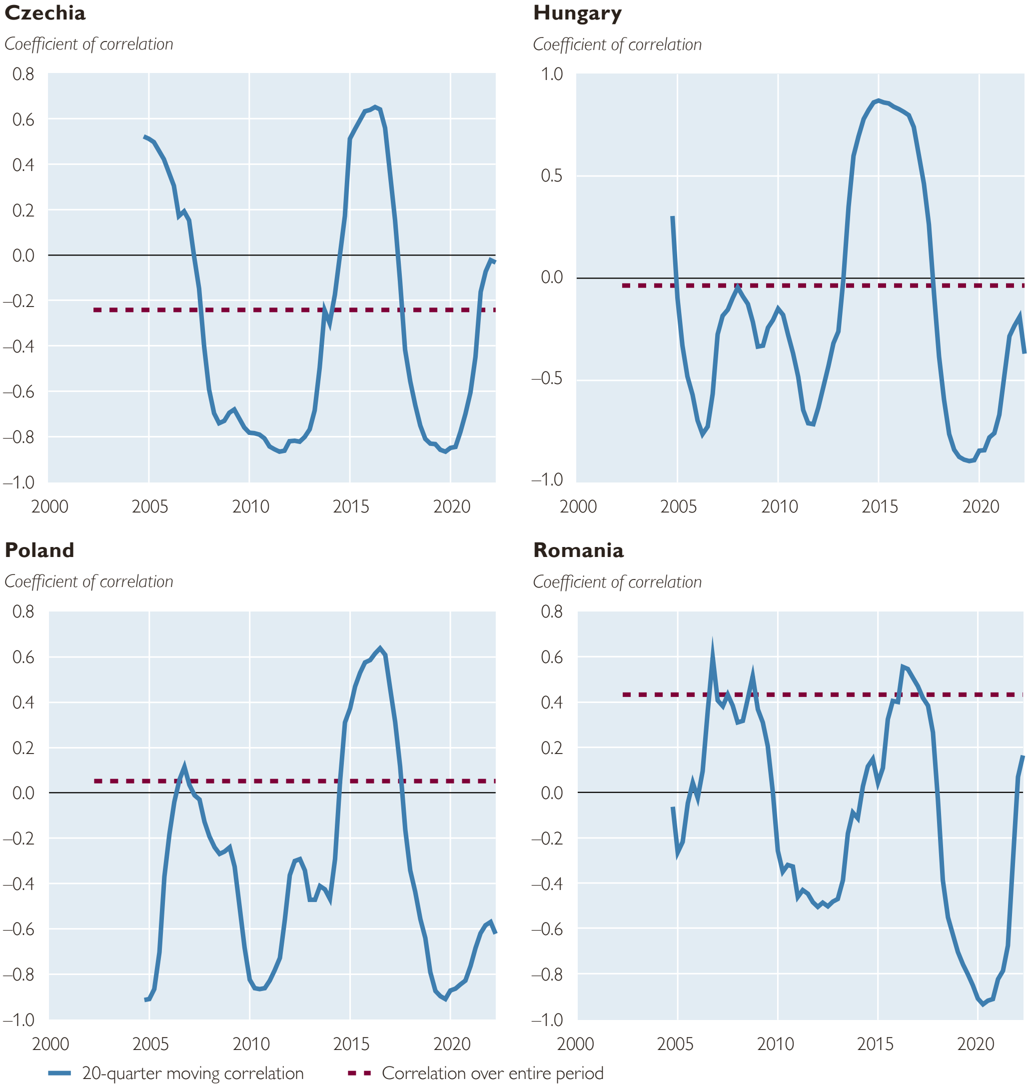

Source: Authors' calculations.

generally lower than headline inflation except for the disinflationary (or deflationary) period between 2013 and 2017 and a brief period during the COVID-19 pandemic.

Furthermore, it is clearly visible that cyclical core inflation generally outpaced noncyclical core inflation across all four countries, sometimes by a large margin. The only exceptions to this pattern are once again the disinflationary (or deflationary) period around 2015 and – at least in Poland and Czechia – the years of the COVID-19 pandemic. This suggests that structural price trends generally had a disinflationary effect in CESEE, very much echoing the discussions on secular stagnation that made headlines throughout much of the 2010s.

Chart 3

Development of core inflation and its cyclical and noncyclical components

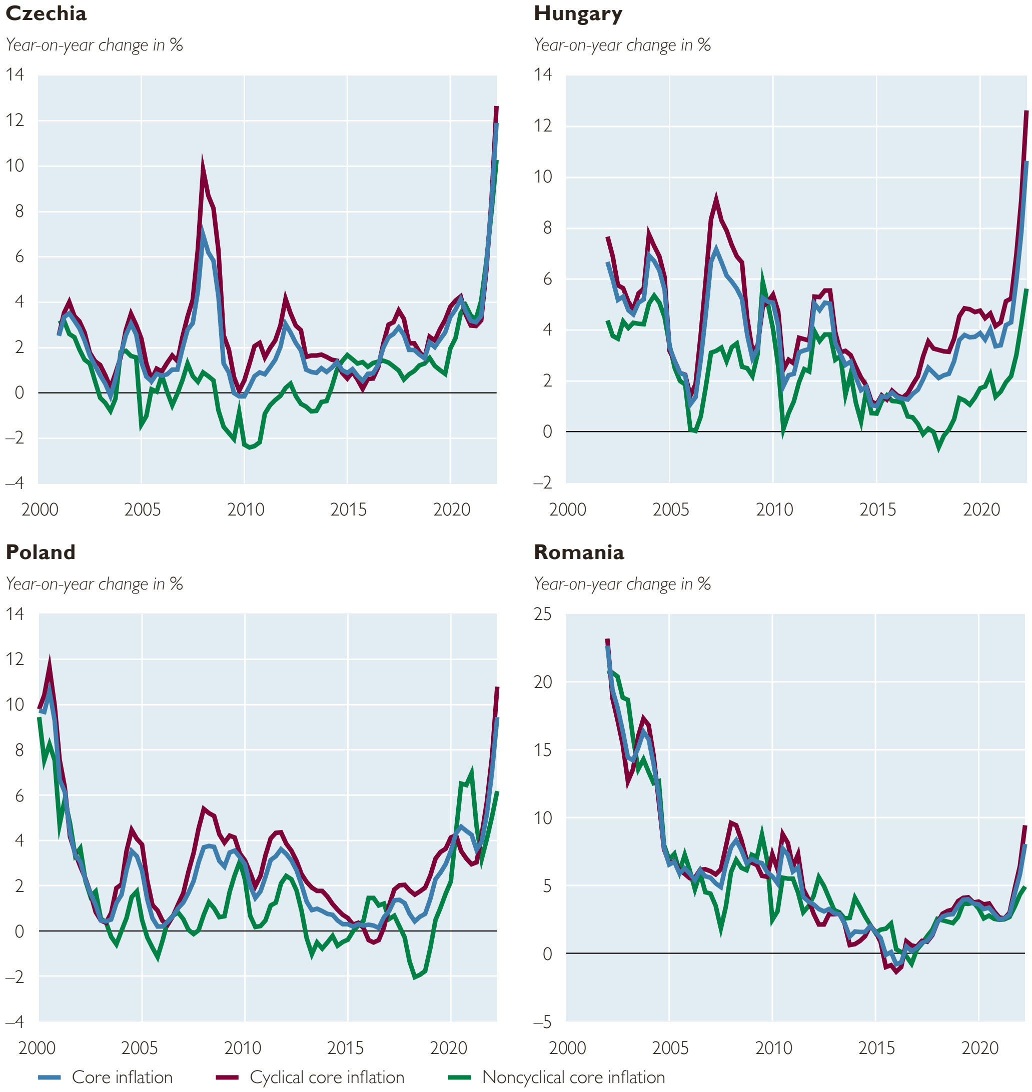

Source: Eurostat, authors' calculations.

2 Econometric framework

Our flexible econometric model combines BART with vector autoregressions (VARs) along the lines of the BART-VAR developed in Huber and Rossini (2022) and Huber et al. (2023). Linear models such as standard VARs are not capable of producing asymmetries in impulse responses with respect to the size or the sign of a structural shock of interest. If the researcher is interested in these forms of asymmetries, nonlinear multivariate time series models such as regime-switching VARs (see Sims and Zha, 2006; Huber and Fischer, 2018), time-varying parameter VARs (see Primiceri, 2005; Koop et al. 2009; Korobilis, 2013) or smooth transition models (see Hauzenberger et al. 2021) are possible approaches to an analysis. However, all these models take a strong prior stance on how nonlinearities are captured, and their assumptions might be inconsistent with the data. In this paper, our approach is nonparametric and does not rely on specific assumptions of the relationship between the vector of endogenous variables y_t, which is $M \times 1$ dimensional, and $x_t = (y'_{t-1}, \dots, y'_{t-p})'$, which stores the lags of the endogenous variables. The model is given by

$$(1) \qquad y_t = F(x_t) + \varepsilon_t, \ \ \varepsilon_t \sim N(0, \Sigma)$$

where F denotes an unknown function that takes x_t as input and returns $F(x_t) = (f_1(x_t), \dots, f_M(x_t))'$ as output. The equation-specific functions $f_j(x_t)$ $(j = 1, \dots, M)$ allow for different functional relations across the elements in y_t. Finally, we let $\mathcal{E}_t$ denote a Gaussian shock term with error covariance matrix $\sum$. The matrix $\sum$ can be made time varying but, given the recent forecasting evidence in Clark et al. (2022), we leave this possibility aside and focus on the homoscedastic case for reasons of simplicity.

The unknown equation-specific functions f_j are approximated using BART. The BART approximation of f_j is:

$$(2) \qquad f_j \approx \sum_{s=1}^{S} g(X|\mu_s, T_s),$$

where we let $X = (x_1, \dots, x_T)'$ denote a $T \times Mp$ full data matrix of regressors and each function g is a tree function that depends on the tree structure T_s and terminal node parameters μ_s. The intuition behind the tree functions is as follows. The tree structure is made of a sequence of decision rules of the form $\{x_t < c\}$ or $\{x_t \geq c\}$ and thus decomposes the input space into several disjoint subsets. These forms of decision rules are applied iteratively and, after testing all these splitting rules, we reach a terminal node. Each terminal node is associated with a terminal node parameter μ_s. The terminal node parameter is then the predicted value of the corresponding regression.

It is worth illustrating this concept by means of a simple example that sets $S=M=1$. Suppose that $x_t=t$ and t runs from 1 to T. Let $c=30$ and the number of terminal nodes be equal to 2. In this trivial case the tree simply splits the sample in two: The first part includes all observations from the beginning of the sample up to 29, whereas the second part includes observations 30 to T. For each of these samples, we then simply estimate the mean over the samples. These means are then the terminal node parameters. The corresponding predictions are then given by:

$$E(y_t) = \begin{cases} \mu_1 \ if \ \ t < 30 \\ \mu_2 \ if \ \ t \geq 30 \end{cases}.$$

Since this specification is very simple (it models through a single structural break), the question arises whether it would pay off to allow for more complex tree structures. Chipman et al. (2010) discuss this possibility but argue for a model that uses very simple tree functions and – instead of taking one single tree – sum over many simple trees. This is what we do in equation (2). Instead of just using a single tree that implies a single structural break, it would be possible to use many (i.e., S) trees. In such a case, the joint model will be able to fit more complex patterns in the data while minimizing the risk of overfitting.

We carry out model estimations precisely along the lines suggested in Huber and Rossini (2022). Our approach is Bayesian and we use the benchmark priors proposed in Chipman et al. (2010). Posterior simulation is carried out using an equation-by-equation algorithm that simulates the terminal node parameters, tree structures and error covariances using a Metropolis-within-Gibbs sampler. Further information can be found in Huber and Rossini (2022) or Clark et al. (2022).

3 Empirical results

This section discusses the impulse responses of the unemployment rate and different price measures to shocks to the unemployment rate. We focus on an unemployment shock defined along the lines of Del Negro et al. (2020). This implies that we rank unemployment first and then consider a Cholesky decomposition of the error covariance matrix Σ. Notice that this approach leaves open the question of whether changes in the unemployment rate are driven by demand- or supply-side shocks. One way to address this question would be to use identified shock measures. As we are interested in estimating nonlinear Phillips curves and high-frequency instruments are not readily available for the CESEE countries, we leave this possibility aside.

We consider three different shocks to the unemployment rate: a one-standard deviation (weak), a five-standard deviation (medium) and a ten-standard deviation (strong) shock.

Chart 4 shows the impulse responses of the unemployment rate, illustrating the magnitude and the evolution of the respective shock. In all four countries under review, a weak shock leads to an immediate increase of the unemployment rate by 0.4 to 0.5 percentage points.[2] The shock fades out rather quickly and becomes statistically insignificant after 5 to 7 quarters.

A medium shock leads to an immediate increase of the unemployment rate by 1.8 percentage points in Czechia, 2.1 percentage points in Hungary and 2.5 percentage points in Poland and Romania. The unemployment rate returns to its initial level quickly at first, and then gradually. The shock ceases to be statistically significant after 9 quarters in Romania, after 10 quarters in Poland and after 12 quarters in Hungary. In Czechia, the shock delivers significant increases in the unemployment rate even after 20 quarters.

A strong shock raises the unemployment rate by 3.7 percentage points in Czechia, 4.2 percentage points in Hungary, 5 percentage points in Poland and 5.1 percentage points in Romania. Again, the unemployment rate returns to its initial level quickly at first, and then gradually. The shock fades out after 10 quarters in Romania, after 12 quarters in Hungary and after 13 quarters in Poland. In Czechia, the shock is still statistically significant after 20 quarters and thus more long lived.

Chart 5 depicts the reaction of the unemployment rate to a strong positive and a strong negative shock. The impulse responses are largely symmetrical for Hungary and Romania, except for the first few quarters when negative shocks fade out somewhat more quickly than positive shocks. More variation can be observed in Czechia and Poland. In both countries, negative shocks impact the unemployment rate more strongly than positive shocks do. The difference between the two shocks

[2] *The figures for the responses of selected variables to a shock in the unemployment rate reported in this section refer to the median estimates. The confidence intervals are depicted in the respective charts.*

Chart 4

Unemployment rate: impulse responses to different positive unemployment shocks

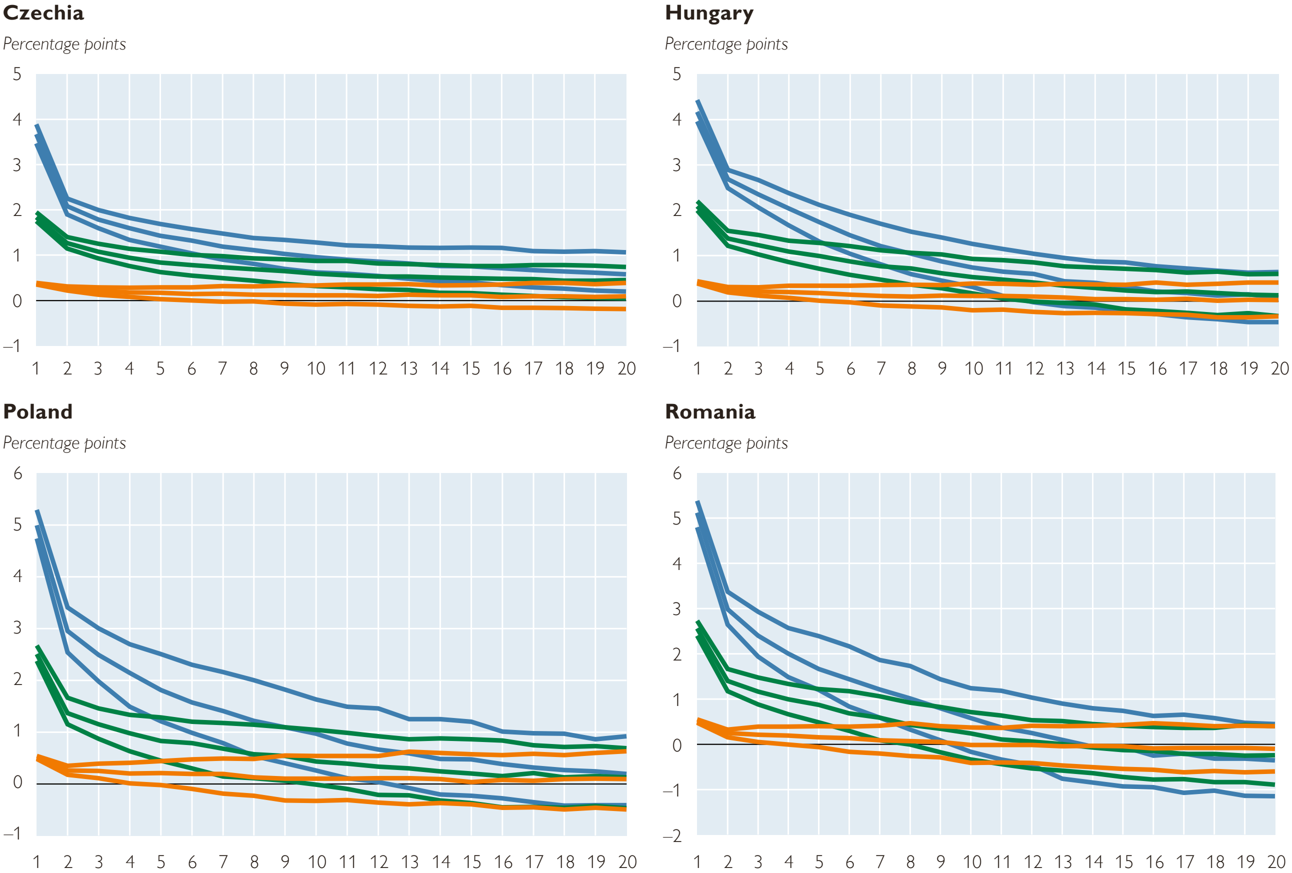

Source: Authors' calculations.

Note: Median, 16th and 84th credible intervals of the posterior distribution of the dynamic responses to a strong (blue line), medium (green line) and weak (orange line) unemplomyent shock.

reaches a maximum of about 0.4 percentage points after 4 and 5 quarters, respectively, and – in the case of Czechia – continues to be observed even after 20 quarters.

Chart 5

Unemployment rate: impulse responses to strong positive and strong negative unemployment shock

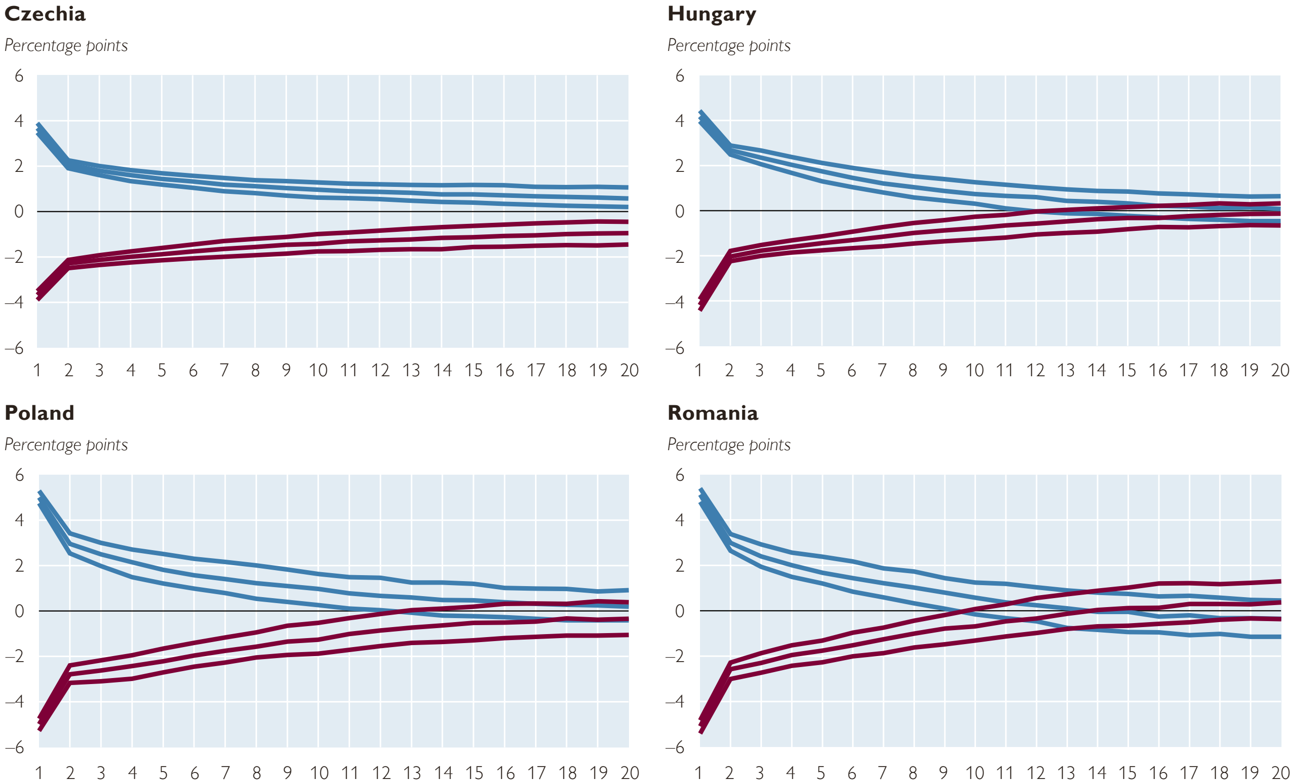

Source: Authors' calculations.

Note: Median, 16th and 84th credible intervals of the posterior distribution of the dynamic responses to a strong positive (blue line) and a strong negative (red line) unemplomyent shock.

Next, our focus is on how HICP inflation responds to unexpected movements in the unemployment rate. This exercise sheds light on whether inflation dynamically reacts to movements in the unemployment rate and whether these movements are consistent with economic theory.

The price reactions to shocks of different sizes are shown in chart 6. The most important information from this chart is that it takes quite a substantial shock to the unemployment rate to trigger a significant reaction in HICP inflation. Small shocks do not translate into any significant price reactions. While medium-sized shocks reduce inflation across the countries under review, responses are very moderate in some cases and tend to fade out in most cases. A medium-sized shock reduces inflation by a maximum of 1.2 percentage points in Czechia (after 17 quarters), 0.7 percentage points in Hungary (after 6 quarters), 0.6 percentage points in Poland (after 10 quarters) and 0.9 percentage points in Romania (after 7 quarters). Only in Czechia, the response remains significant after 20 quarters. In Hungary, the shock becomes insignificant after 17 quarters, in Poland after 13 quarters and in Romania after 16 quarters.

Only strong shocks substantially reduce inflation and produce significant reactions even after 20 quarters. The largest effects can be observed in Czechia and Romania, where the shock reduces inflation by a maximum of 1.6 percentage points after 15

Chart 6

HICP inflation: impulse responses to different positive unemployment shocks

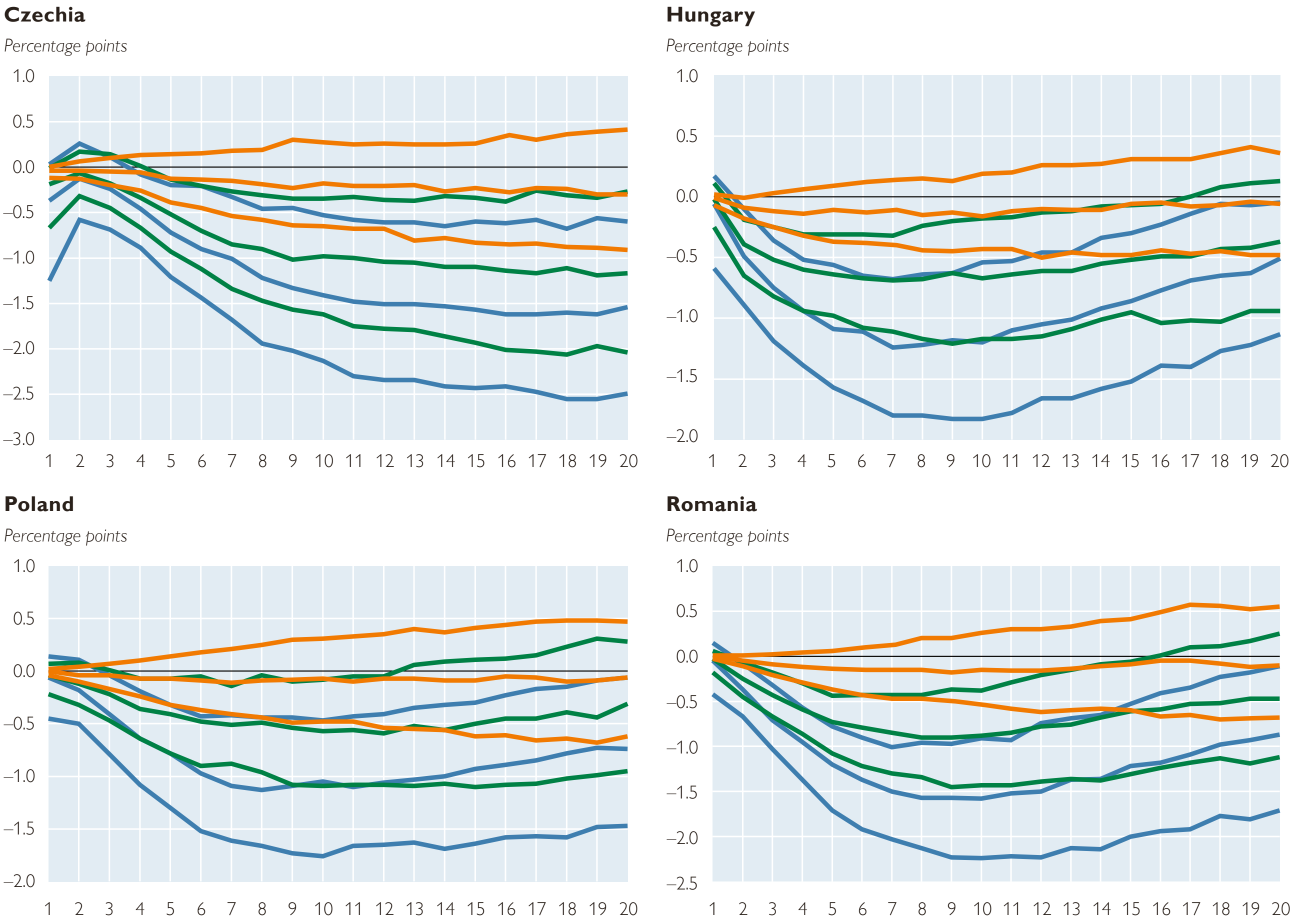

Source: Authors' calculations.

Note: Median, 16th and 84th credible intervals of the posterior distribution of the dynamic responses to a strong (blue line), medium (green line) and weak (orange line) unemplomyent shock.

and 8 quarters, respectively. The impact is somewhat weaker in Hungary and Poland but remains substantial at a maximum of 1.2 and 1.1 percentage points, respectively, after 7 quarters.

Chart 7 shows the impulse responses of HICP inflation to a strong positive and a strong negative shock. Only for Hungary we find responses that are symmetrical, indicating that positive and negative shocks impact prices to a similar extent. For the other countries, responses differ somewhat. In Czechia, a positive unemployment shock has a much weaker impact on HICP inflation than a negative shock. This means that an increase in the unemployment rate reduces inflation by a lesser margin than a decrease in the unemployment rate drives inflation up. This difference reaches a maximum of 1.1 percentage points after 12 quarters and remains substantial even after 20 quarters. The picture is similar for Poland, although the differences are not quite as pronounced and reach a maximum of only 0.4 percentage points after 5 quarters. In Romania, a negative shock produces a stronger response up until 10 quarters; after that, a positive shock tends to impact on inflation somewhat more strongly.

Chart 7

HICP inflation: impulse responses to strong positive and strong negative unemployment shock

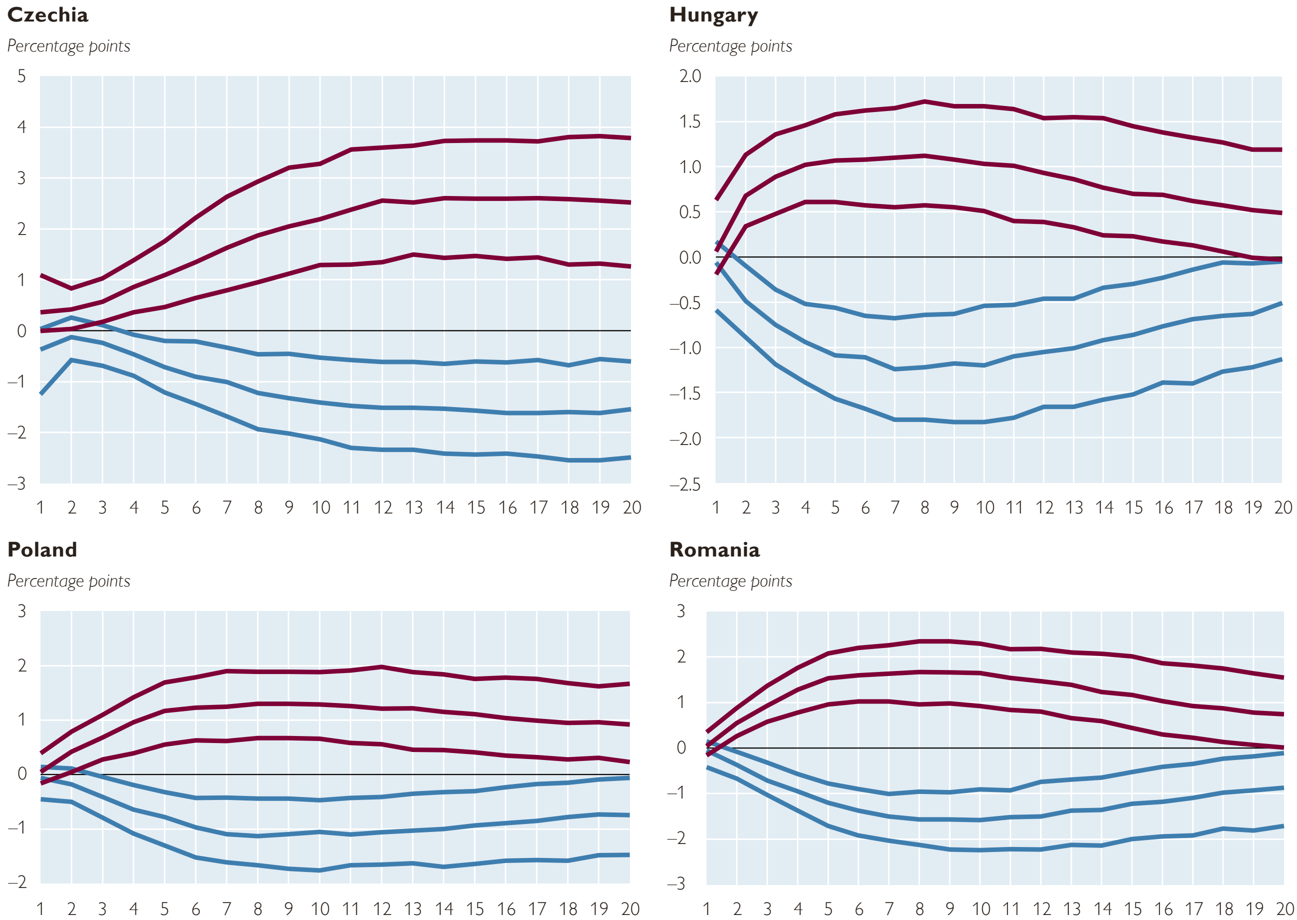

Source: Authors' calculations.

Note: Median, 16th and 84th credible intervals of the posterior distribution of the dynamic responses to a strong positive (blue line) and a strong negative (red line) unemplomyent shock.

Chart 8 shows the impulse responses of core inflation (HICP inflation excluding energy and unprocessed food) to different positive shocks. As is the case with headline inflation, only medium and strong shocks result in statistically significant changes to core inflation. A medium-sized shock reduces core inflation by a maximum of 0.9 percentage points in Czechia (after 12 quarters), 0.5 percentage points in Hungary (after 6 quarters), 0.4 percentage points in Poland (after 6 quarters) and 0.6 percentage points in Romania (after 6 quarters). Only in Czechia, the response remains significant after 20 quarters. In Hungary, the shock becomes insignificant after 17 quarters, in Poland after 12 quarters and in Romania after 16 quarters.

A strong shock to the unemployment rate reduces core inflation by a maximum of 1.3 percentage points in Czechia (after 12 quarters), 0.9 percentage points in Hungary (after 8 quarters), 0.8 percentage points in Poland (after 8 quarters) and 1.1 percentage points in Romania (after 8 quarters). The effects of the shock remain statistically significant after 20 quarters. With that, shocks to core inflation produce effects that are comparable to the effects of shocks to headline inflation in

Chart 8

Core inflation: impulse responses to different positive unemployment shocks

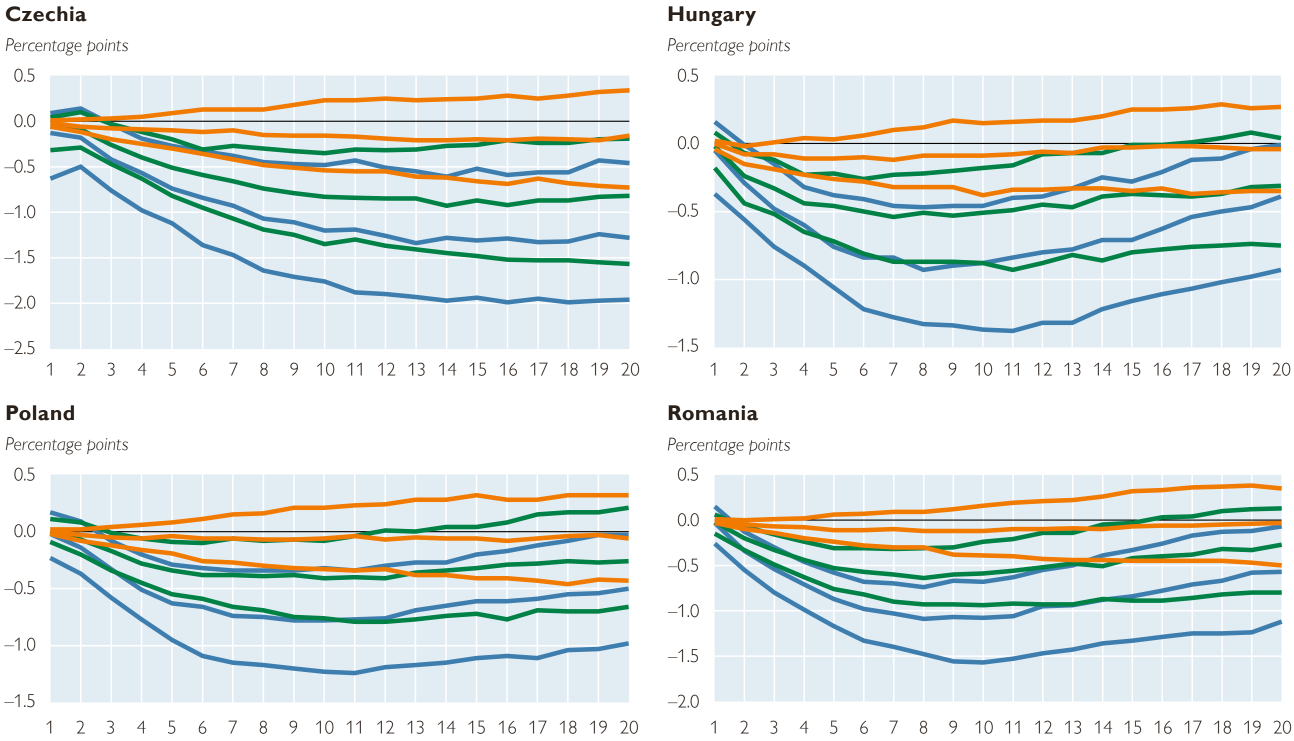

Source: Authors' calculations.

Note: Median, 16th and 84th credible intervals of the posterior distribution of the dynamic responses to a strong (blue line), medium (green line) and weak (orange line) unemplomyent shock.

terms of timing and persistence; their quantitative effects are generally lower, however.

We now turn to the question of whether there are any asymmetries with respect to the sign of the shock to core inflation. Chart 7 shows the impulse responses of core inflation to strong positive and negative shocks to the unemployment rate. As is the case with headline inflation, impulse responses are largely symmetrical for Hungary (even though the negative shock impacts somewhat more strongly on core inflation over the first 3 quarters). In Czechia, core inflation responds more strongly to negative shocks than to positive shocks throughout the observation period, with the difference reaching a maximum of 0.8 percentage points (after 14 quarters). The same is true for Poland, but the difference in the two responses only climbs to 0.2 percentage points (after 5 quarters). In Romania, the impact of the negative shock outpaces the impact of the positive shock up until the 10th quarter, after which the two shocks become largely indistinguishable.

Chart 9

Core inflation: impulse responses to strong positive and strong negative unemployment shock

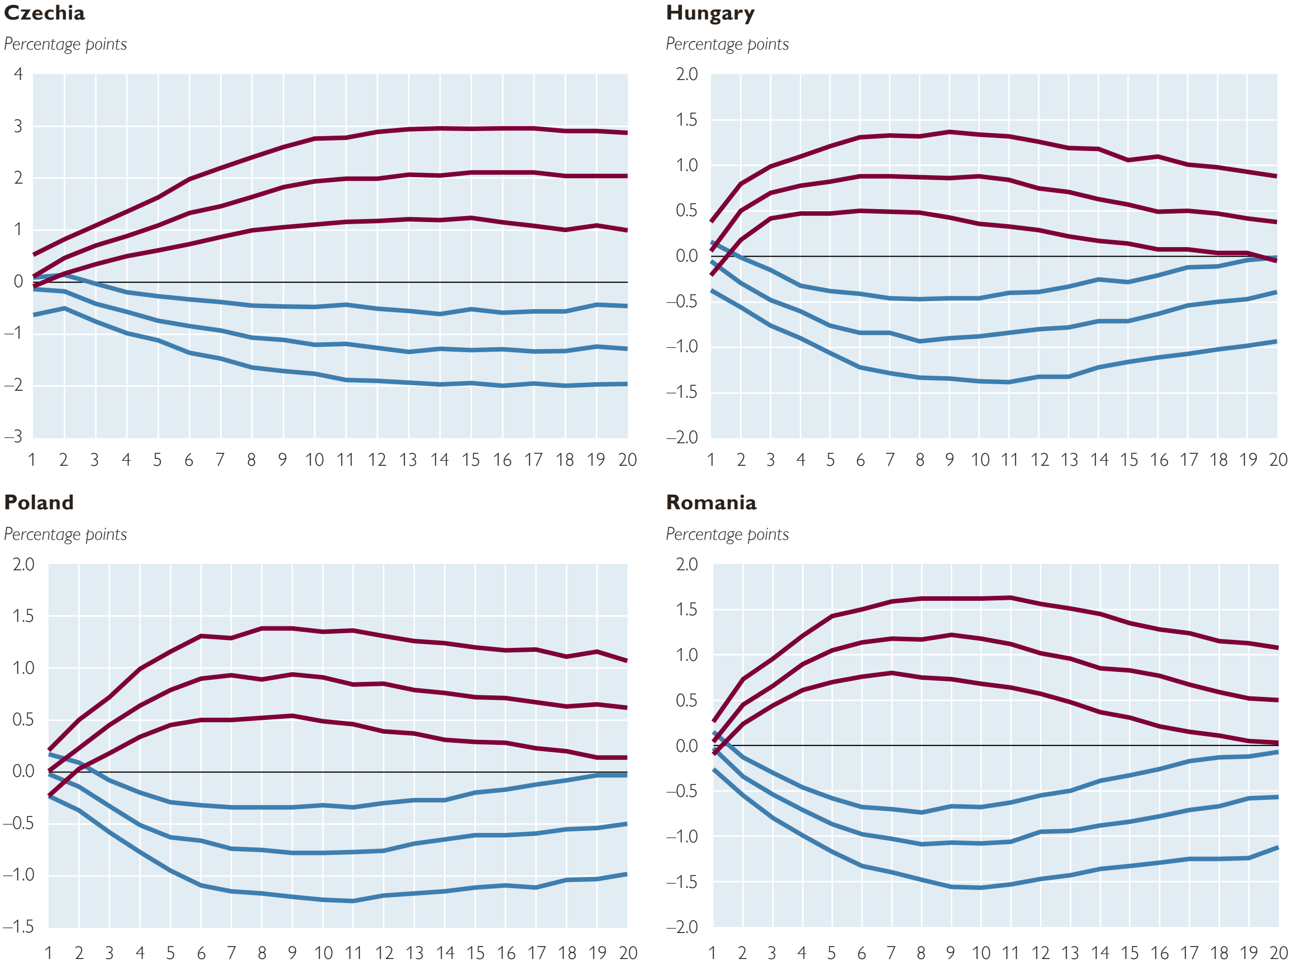

Source: Authors' calculations.

Note: Median, 16th and 84th credible intervals of the posterior distribution of the dynamic responses to a strong positive (blue line) and a strong negative (red line) unemplomyent shock.

To improve our intuition of what drives dynamics in core inflation, we decompose core inflation into a cyclical and a noncyclical component. Chart 10 depicts the reaction of cyclical core inflation to different unemployment shocks. Again, it takes at least a medium-sized shock to produce statistically significant results.

A medium-sized shock reduces cyclical core inflation by a maximum of 0.9 percentage points in Czechia (after 12 quarters), 0.6 percentage points in Hungary (after 5 quarters), 0.5 percentage points in Poland (after 9 quarters) and 0.9 percentage points in Romania (after 11 quarters). The effects remain significant even after 20 quarters in Czechia, but fade away after 15 quarters in Hungary, after 11 quarters in Poland and after 15 quarters in Romania.

Large shocks reduce cyclical core inflation by a maximum of 1.3 percentage points in Czechia (after 12 quarters), 1.1 percentage point in Hungary (after 9 quarters), 0.9 percentage points in Poland (after 8 quarters) and 1.4 percentage points in Romania (after 11 quarters). The effects remain statistically significant even after 20 quarters in all countries under observation.

Chart 10

Cyclical core inflation: impulse responses to different positive unemployment shocks

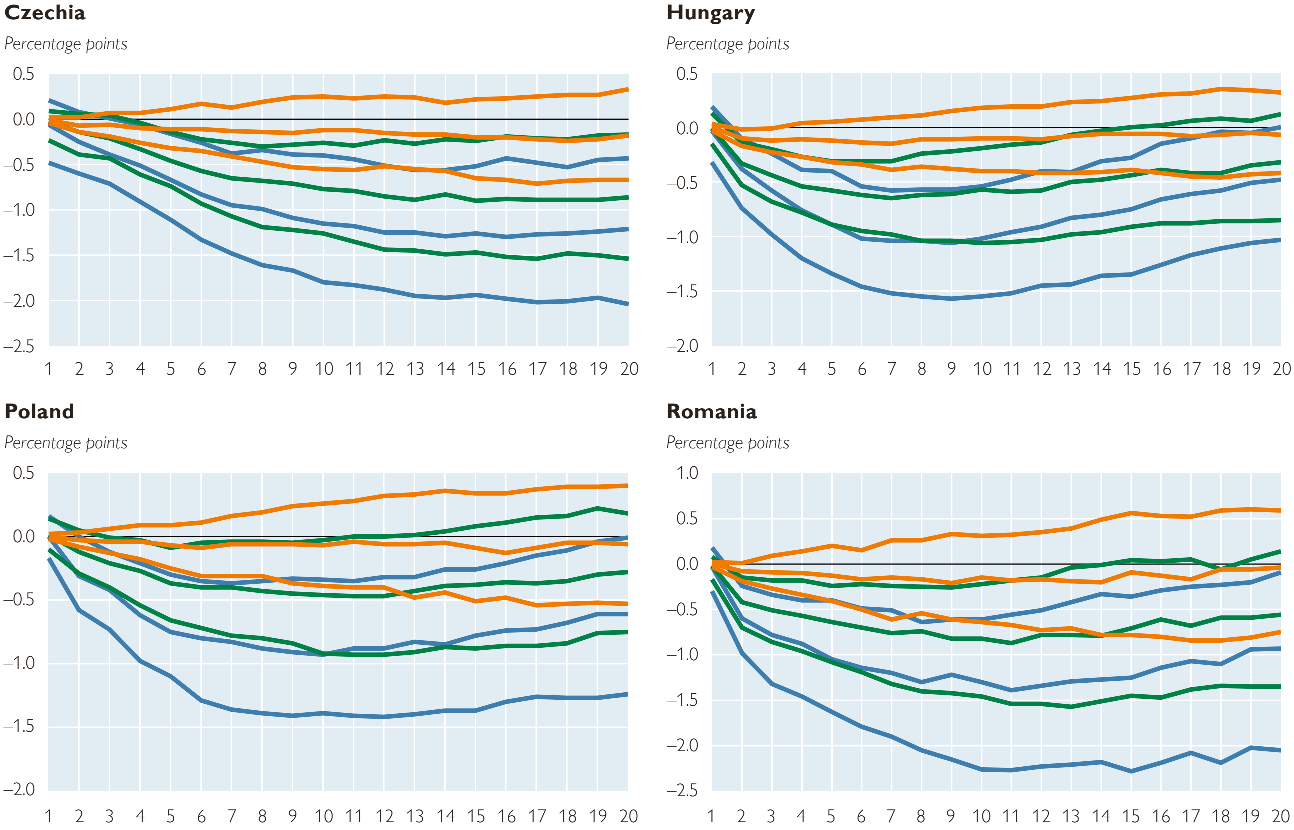

Source: Authors' calculations.

Note: Median, 16th and 84th credible intervals of the posterior distribution of the dynamic responses to a strong (blue line), medium (green line) and weak (orange line) unemplomyent shock.

Chart 11 shows the response of cyclical core inflation to strong negative and strong positive shocks to the unemployment rate. The responses are largely symmetrical for Hungary and Romania, even though in Hungary, the negative shock impacts somewhat more strongly on cyclical core inflation over the first 5 quarters. In Czechia, cyclical core inflation responds more strongly to negative shocks than to positive shocks throughout the observation period, with the difference reaching a maximum of 0.8 percentage points (after 10 quarters). The same is true for Poland, but the difference in the two responses only climbs to 0.3 percentage points (after 6 quarters).

Chart 11

Cyclical core inflation: impulse responses to strong positive and strong negative unemployment shock

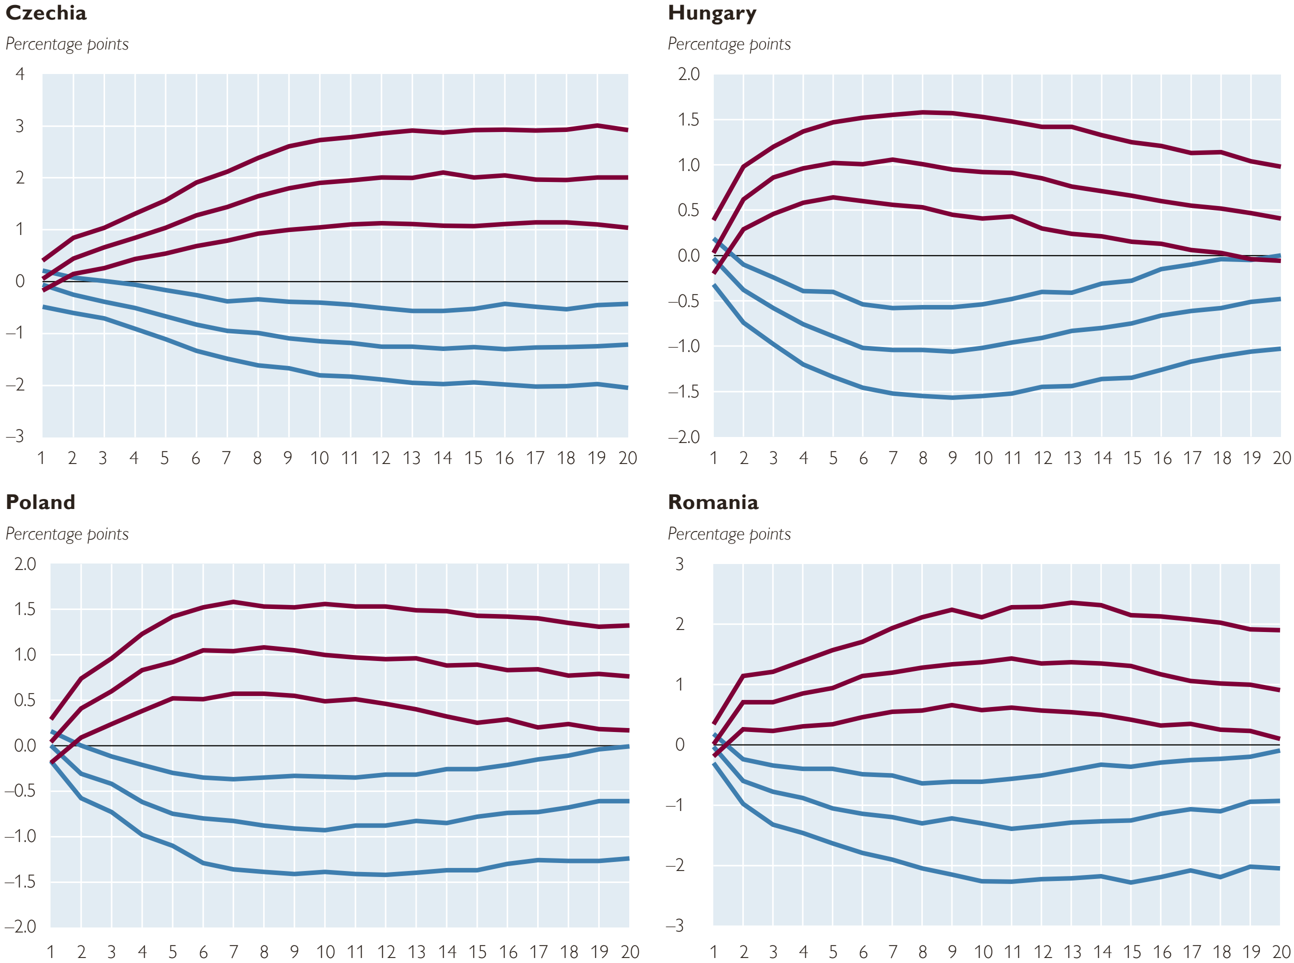

Source: Authors' calculations.

Note: Median, 16th and 84th credible intervals of the posterior distribution of the dynamic responses to a strong positive (blue line) and a strong negative (red line) unemplomyent shock.

Chart 12 shows the reaction of the noncyclical part of core inflation to shocks to the unemployment rate. Two things stand out: First, only large shocks significantly reduce noncyclical core inflation for an extended time span. Second, even with strong shocks it takes quite some time to produce a significant negative effect on noncyclical core inflation (between 4 quarters in Czechia and 7 quarters in Romania).

Strong shocks reduce noncyclical core inflation by a maximum of 1.2 percentage points in Czechia (after 11 quarters), 0.5 percentage points in Hungary (after 9 quarters), 0.5 percentage points in Poland (after 7 quarters) and 0.9 percentage points in Romania (after 13 quarters). The effects of the shock remain statistically significant even after 20 quarters in all countries but Poland, where the effect fades after 16 quarters.

Chart 13 shows the responses of noncyclical core inflation to strong negative and positive shocks to the unemployment rate. We observe largely symmetrical effects in Hungary and Romania. In Czechia, noncyclical core inflation responds

Chart 12

Noncyclical core inflation: impulse responses to different positive unemployment shocks

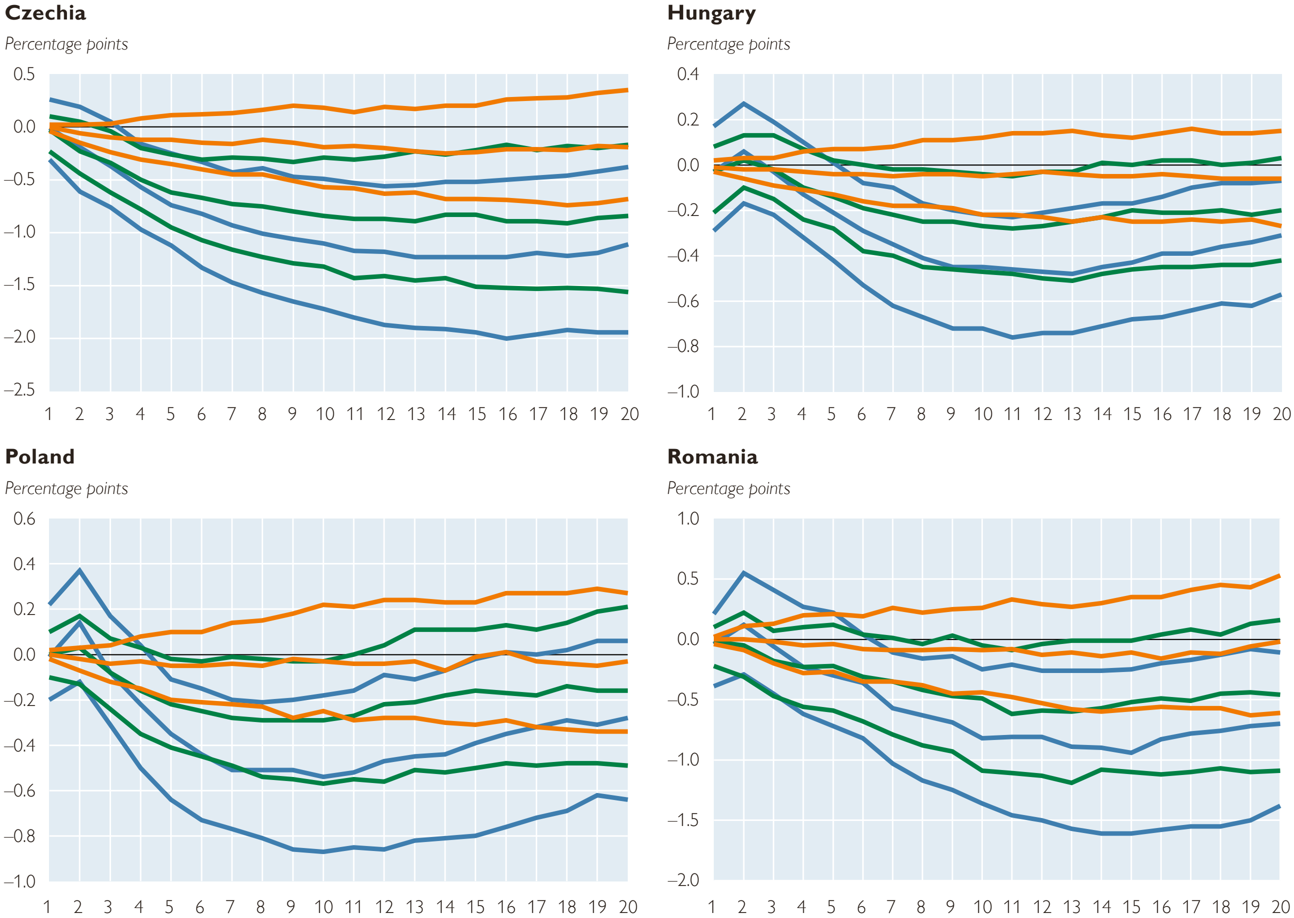

Source: Authors' calculations.

Note: Median, 16th and 84th credible intervals of the posterior distribution of the dynamic responses to a strong (blue line), medium (green line) and weak (orange line) unemplomyent shock.

more strongly to negative shocks than to positive shocks throughout the observation period, with the difference reaching a maximum of 0.7 percentage points (after 12 quarters). The same is true for Poland, but the difference in the two responses only climbs to 0.3 percentage points and reaches its maximum already after 2 quarters. After 10 quarters, the two responses are largely indistinguishable.

Chart 13

Noncyclical core inflation: impulse responses to strong positive and strong negative unemployment shock

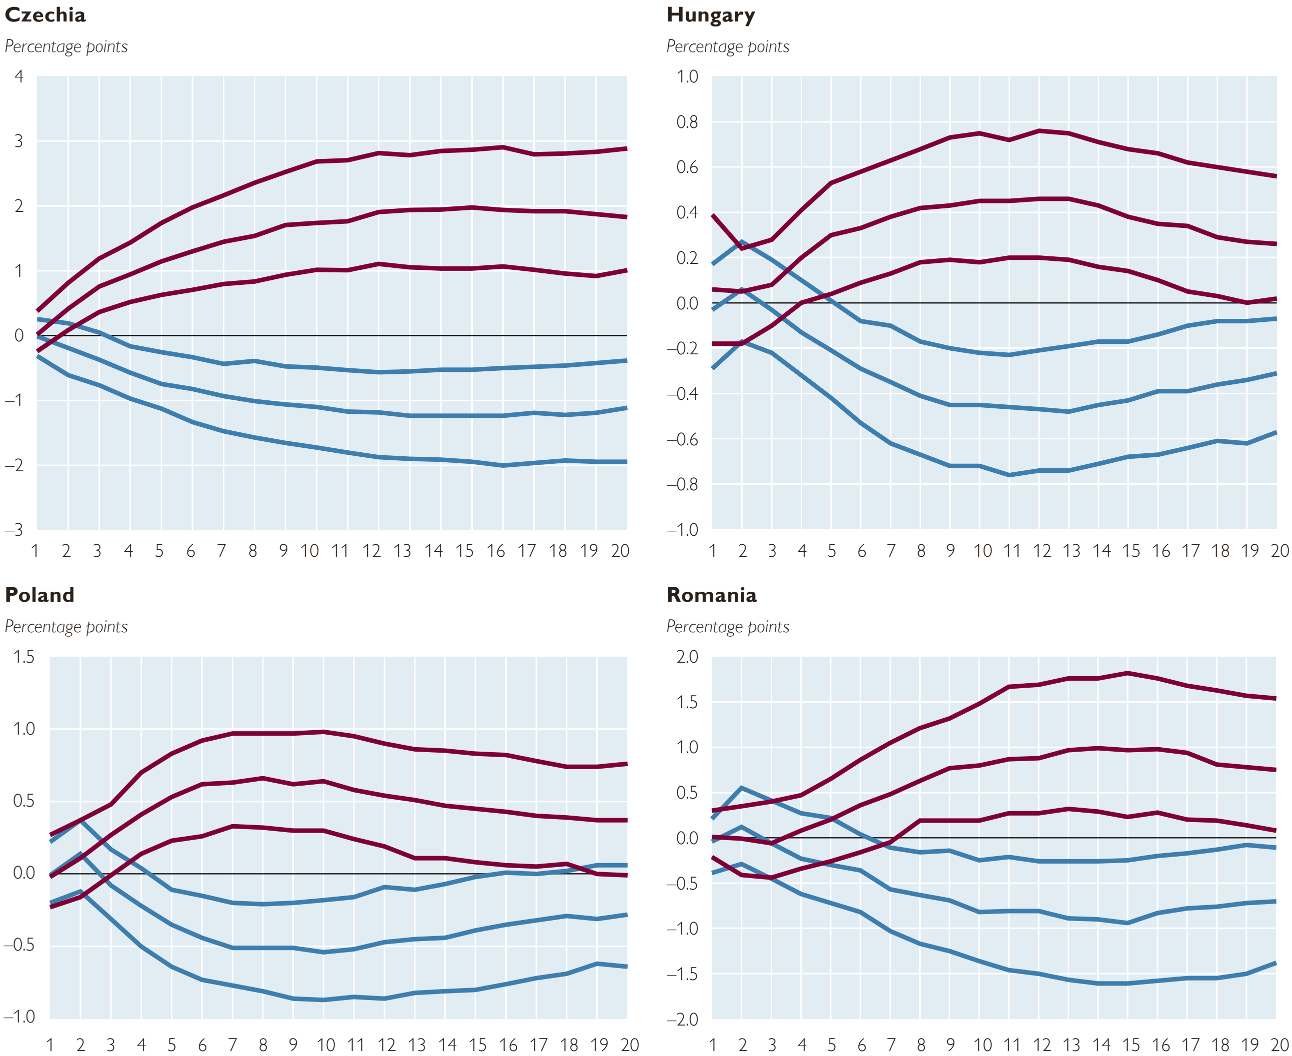

Source: Authors' calculations.

Note: Median, 16th and 84th credible intervals of the posterior distribution of the dynamic responses to a strong positive (blue line) and a strong negative (red line) unemplomyent shock.

4 Discussion of results

Our paper aims to find out whether the Phillips curve is still alive and well behaved in CESEE. Our results for a sample of four CESEE countries suggest the following: The Phillips curve is still alive, but it may be somewhat sleepy. It takes quite a push to wake it up. Only a strong shock to the unemployment rate induces a significant, broad and lasting effect on price growth in the countries in our sample. Our results point toward substantial nonlinearities in the relationship between real economic activity and inflation.

We find that an increase in the unemployment rate by a magnitude of ten standard deviations lowers headline HICP inflation by a maximum of 1.1 to 1.6 percentage points, depending on the individual country. The impact is somewhat weaker for core inflation, where a strong shock induces a decline by a maximum of between 0.8 and 1.3 percentage points. Within core inflation, cyclical components react much more strongly than noncyclical components (0.9 to 1.4 percentage points vs. 0.5 to 1.2 percentage points). Usually, the maximum impact is reached after about 10 quarters and weakens again toward the end of the observation period. Across individual countries, the strongest effects can be observed for Czechia, followed by Romania and – with quite some margin – Hungary and Poland.

Compared with the reactions to strong shocks, those to medium-sized shocks (five standard deviations) are roughly half as strong and usually lose statistical significance after some time at least in Hungary, Poland and Romania. In Czechia, even medium shocks produce significant results for 20 quarters and beyond and impact somewhat more strongly on inflation than in the other CESEE countries.

The fact that only strong shocks induce substantial, broad and lasting reactions constitutes one important nonlinearity in the Phillips curves of the countries under review. Another important nonlinearity can be found in the way inflation reacts to shocks with different signs. At least in the case of Czechia and Poland, inflation reacts much more strongly to negative shocks than to positive shocks. This means that a reduction in the unemployment rate drives up inflation to a stronger extent than an increase in the unemployment rate lowers it. This finding is robust across different inflation measures. The situation is somewhat more heterogeneous in Romania. Here, negative shocks tend to produce stronger results than positive shocks. The differences, however, are less pronounced, they are often restricted to certain time periods and they are found predominantly for the broader aggregates of headline and core inflation. Only in Hungary, impulse response functions are largely symmetrical for both positive and negative shocks.

5 Conclusions

Our estimations show that the Phillips curve is alive and well in CESEE. However, it displays some nonlinearities that are vital in understanding the impact of labor market developments on inflation over the past ten years. It takes a substantial shock to the unemployment rate to trigger a notable and sustained move in the inflation rate. Furthermore, it takes around two and a half years until such a shock reaches its maximum impact. This should introduce quite some inertia in the Phillips curve relationship in cases where the unemployment rate changes steadily but only slowly. Exactly such a setting was observed in the years preceding the pandemic, when gradual improvements in the unemployment rate stretched out over many years and did not initiate a strong increase in inflation.

Nevertheless, some inflationary pressure was probably accumulated. This pressure finally started to contribute to price rises after post-pandemic and war-related disruptions led to a regime change from a low inflation to a high inflation environment. In such a situation, price changes in individual subsegments increasingly affect other subsegments and a limited change in relative prices tends to translate into a stronger generalized inflation momentum (see BIS, 2022). Transitioning back from such a high inflation regime can be very costly once it becomes entrenched. Our research suggests that – at least in some CESEE countries – this transition could be made even more costly by the fact that inflation reacts more weakly to a loosening than to a tightening of the labor market. Against this backdrop, classic macroeconomic demand management, including demand management by means of monetary policy, would be called upon to act particularly strongly and decisively to keep inflation in check.

On top of that, it remains unclear how strongly economic policy could contribute to labor market loosening in the CESEE countries under consideration. In CESEE, several factors keep labor markets tight even when macroeconomic demand conditions are weakening: (1) The production factor labor is particularly heavily utilized; (2) labor supply is adversely affected by demographic headwinds such as population aging, skill mismatches, emigration and – in some countries – cross-border commuting for work; and (3) catching-up related positive growth differentials vis-à-vis Western Europe and a structural shift of the economy toward labor-intensive services keep labor demand high. CESEE labor markets have therefore operated (almost) at full capacity for much of the past decade and labor shortages have become chronic. While there are some possible remedies for this situation – e.g. automation – labor markets will probably remain tight at least over the medium term, thereby limiting the functioning of the Phillips curve relationship. Against this backdrop, future research might explicitly address the functioning of the Phillips curve in an environment of structurally tight labor markets and put a spotlight on special nonlinearities related to unemployment rates that are near the "zero lower bound."

References

Bank of International Settlements (BIS). 2022. Inflation: a look under the hood. Annual Economic Report 2022. Chapter 2.

Chipman, H. A., E. I. George and R. E. McCulloch. 2010. BART: Bayesian additive regression trees. The Annals of Applied Statistics 4(1). 266–298. March.

Clark, T. E., F. Huber, G. Koop, M. Marcellino and M. Pfarrhofer. 2022. Tail forecasting with multivariate Bayesian additive regression trees. International Economic Review. Forthcoming.

Del Negro, M., M. Lenza, G. E. Primiceri and A. Tambalotti. 2020. What's Up with the Phillips Curve? Brookings Papers on Economic Activity. 301–357.

Hauzenberger, N., M. Pfarrhofer and A. Stelzer. 2021. On the effectiveness of the European Central Bank's conventional and unconventional policies under uncertainty. Journal of Economic Behavior & Organization 191. 822–845.

Huber, F. and M. M. Fischer. 2018. A Markov switching factor-augmented VAR model for analyzing US business cycles and monetary policy. Oxford Bulletin of Economics and Statistics 80(3). 575–604.

Huber, F. and L. Rossini. 2022. Inference in Bayesian additive vector autoregressive tree models. The Annals of Applied Statistics 16(1) 104–123.

Huber, F., G. Koop, L. Onorante, M. Pfarrhofer and J. Schreiner. 2023. Nowcasting in a pandemic using non-parametric mixed frequency VARs. Journal of Econometrics 232(1). 52–69.

International Monetary Fund (IMF). 2013. The Dog That Didn't Bark: Has Inflation Been Muzzled or Was It Just Sleeping? World Economic Outlook. Chapter 3.

Koop, G., R. Leon-Gonzalez and R. W. Strachan. 2009. On the evolution of the monetary policy transmission mechanism. Journal of Economic Dynamics and Control 33(4). 997–1017.

Korobilis, D. 2013. Assessing the transmission of monetary policy using time-varying parameter dynamic factor models. Oxford Bulletin of Economics and Statistics 75(2). 157–179.

Kuttner, M. and T. Robinson. 2010. Understanding the Flattening of the Phillips Curve. In: The North American Journal of Economics and Finance 21. 110–125.

Sims, C. A. and T. Zha. 2006. Were there regime switches in US monetary policy? American Economic Review 96(1). 54–81.

Stock, J. H. and M. W. Watson. 2019. Slack and Cyclically Sensitive Inflation. NBER Working Papers 25987. National Bureau of Economic Research, Inc.

Lian, W. and A. Freitag. 2022. Inflation Dynamics in Advanced Economies: A Decomposition into Cyclical and Non-Cyclical Factors. IMF Working Paper WP/22/91.

Primiceri, G. E. 2005. Time varying structural vector autoregressions and monetary policy. The Review of Economic Studies 72(3). 821–852.