# (Stage-Based) Identifying the Effect of War on Inflation*

*Acauã Brochado*†

*National Treasury of Brazil and GSEFM/Goethe University Frankfurt*

July 19, 2024

## Abstract

Identifying the causal effect of a macroeconomic shock that hits the whole country, or the whole world, at the same time, is still a challenge. Identifying what is moving inflation is especially important to guide central banks in their monetary policy decisions. A big war, potentially, can impact inflation simultaneously across the globe, and a novel method may be able to clearly identify its effect. I apply the Stage-Based Identification method on inflation data for countries in the Eurozone, regions in USA, regions in Brazil, and Länder in Germany to estimate the impact on inflation of Russia invading Ukraine in 2022. The results indicate that, given the state of the world at the time, the war had a positive and economically relevant impact, between 1 and 2 percentage points, on headline inflation in most countries and subnational regions studied. But in Europe, the great spike in inflation observed between 2021 and 2022 was only partially explained by it.

**Keywords:** inflation, policy effect, causality, war, regional effects
**JEL Classification:** C01, E31, E32, E65, R10

* I thank my supervisor, Alexander Ludwig, and, of course, the co-authors Christian Alemán, Christopher Busch and Raül Santaeulàlia-Llopis, for kindly providing the base code for running SBI. I also thank my co-supervisor Georg Dürnecker, and the colleagues from the chair of Public Finance and Macroeconomics Dynamics, and from the Frankfurt Quantitative Macro Group (FQMG) for useful comments. As usual, all errors are my own. Research funding from the State of Hessen is gratefully acknowledged.

† brochado@econ.uni-frankfurt.de - PhD candidate at the Graduate School of Economics, Finance, and Management – GSEFM / Goethe University Frankfurt. Analyst at the National Treasury of Brazil (on leave). Opinions are my own and do not represent institutional views.

## 1. Introduction

It is usually tempting to consider the specific method one is using as being "the best", or "the correct" one, across all methods available in the literature. Experienced analysts, however, know that every method have strengths and weaknesses, and that it is the specific task at hand that can better define which method (or methods) is (are) the most appropriate. With this work I aim to apply a novel method and introduce it to the literature that identifies causality of shocks on inflation, by challenging the notion that the war (invasion of Ukraine by Russia in February 2022) caused (most of) the big spike in inflation observed in that year. I evaluate the shock across many countries and estimate aggregate as well as subnational regional effects.

Empirically identifying the causal effect of a macroeconomic shock that affects at the same time the whole country, or the whole world, is still a challenge. Most empirical macroeconomic exercises suffer from at least one of the many econometric fragilities. There is hardly a case where control groups can be found, and synthetic ones are always subject to criticism. Parallel pre-trends are also questionable in many cases. Controlling for confounding factors is hard because important features of the economy may not be observable. Ad hoc restrictions in estimation models, as done, for example, in the VAR literature, can hardly create a consensus. And last, but not least, potential endogeneity is almost always present. Ramey (2016) goes through most of these issues in macroeconometrics, and discusses the most recent tools available, its advantages and fragilities. The challenge seems so hard that she ends one section in a very discouraging note: *"… we simply do not have enough information to produce estimates with any great certainty."* (Ramey, 2016)

However difficult it may be, the economy is moving, and agents need to act. Thus, to guide them, our job is to produce the best estimates we can.

Identifying what is moving inflation is especially important to guide central banks in their monetary policy decisions. Broadly speaking, optimal monetary policy crucially depends on what are the characteristics of the inflation process at hand. Is it a short-lived shock or is it long lasting? Does it come from excess demand or from a supply shock?

A big war, for example, can potentially impact inflation relevantly and simultaneously across the globe. In general, it also triggers a heated debate across the society (media, politics, and academia) on the questions mentioned above, precisely those that are challenging for empirical macroeconomists.

In this paper I contribute to this discussion bringing a novel method that may be able to clearly identify the effect of a single and simultaneous global shock on inflation. Here I focus on the invasion of Ukraine by Russia in February 2022, which was commonly cited as one of the important drivers of the big spike in world inflation on that year, following a decade-long disinflationary trend.

To identify the effect of the shock, I apply the Stage-Based Identification method (SBI) developed by Alemàn, Busch, Ludwig and Santaeulàlia-Llopis, from now on just ABLS, (2023).

The SBI method has, of course, a limitation. Though the identification of the effect is theoretically very robust, in practice, the convergence of the code and the precision of the estimated policy effect relies on a minimum level of similarity between the outcome paths across the subjects. That is, since the method normalizes individual outcome paths to a reference path, the original paths must not be too disparate (ABLS, 2023).

It turns out that looking at inflation across the globe we see a fairly common trend following the onset of the COVID pandemic in the beginning of 2020, which is subdued inflation, and also a common upward trending inflation after one year or so. This upward trend however presents different timings across countries, and the peaks (or trend reversals) are also heterogenous. Figure 1 shows headline inflation in 25 selected advanced and emerging economies across the world and a vertical line marking the moment of the invasion. Since the shock occurred at a point in time where most of the series were already rising, a question that comes naturally is whether the war really had an impact over inflation, and what its size was.

***Figure 1 - Inflation Across the Globe - 12-month %***

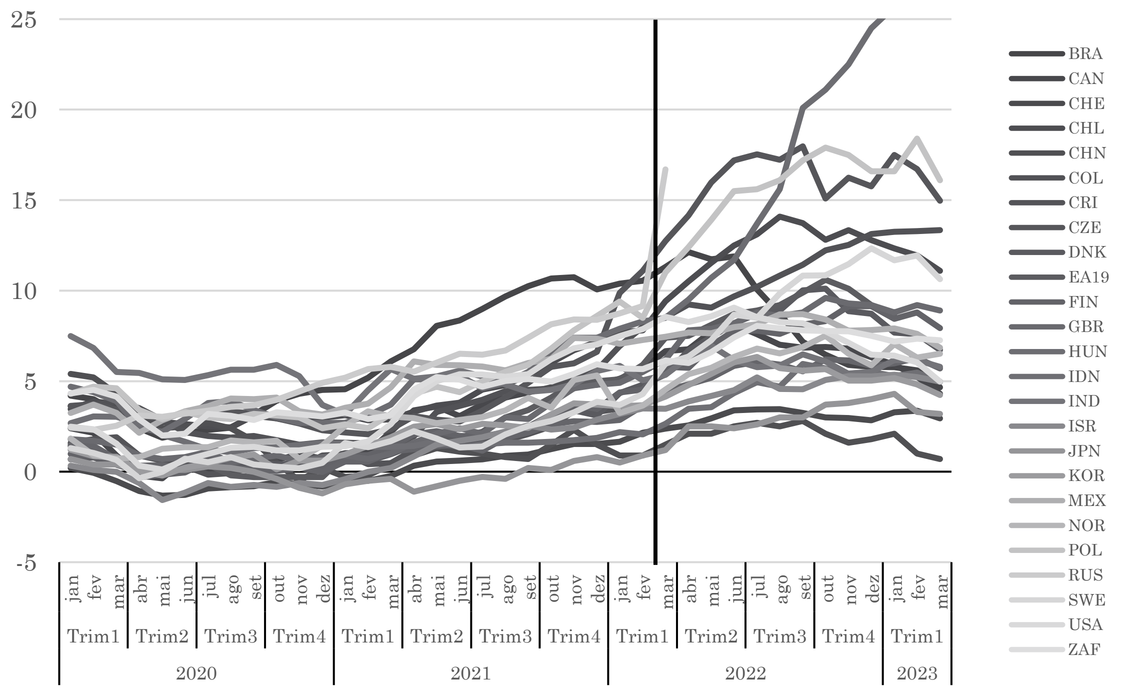


Source: OECD, own elaboration. Note: 12-month percentage variation of headline price index. The vertical black line represents the moment of the invasion of Ukraine by Russia.

Despite the theoretical soundness of the SBI model, such a practical implementation can raise questioning about the degree to which other factors could get in the way of the correct identification of causality and its size. With that in mind, in this paper I use inflation data for countries in the Eurozone, regions in the United States, regions in Brazil, and Länder in Germany. Consequently, then, I work with a broad spectrum of the trade-off between the similarity in the underlying inflation trends and the heterogeneity needed for identification.

The results indicate that, given the state of the world at the time, the war had indeed a positive impact in headline inflation in most countries and regions studied. The effect estimated from the four samples in this study range from close to 1 to a little above 2 percentage points. The estimates are statistically significant for two cases (Brazilian regions and German Länder). For the Eurozone countries and US regions the estimate cannot be considered statistically significant, considering a 90% confidence interval.

Comparing these estimations with the variation of total headline inflation, it is possible to say that in Brazil and in the United States other factors pulled inflation down between 2021 and 2022, despite the positive (and big for Brazil) impact of the war. On the contrary, in the Eurozone and in Germany, other factors pushed inflation up (significantly) more than the estimated impact of the war did.

In this paper, after this introduction, I briefly review some recent literature on identification of macroeconomic shocks, determinants of the recent inflation surge, and causality. After that I focus on Stage-Based Identification and give an overview of the method. In the following section I quickly describe the data, after which I present the empirical strategy pursued. Next, I describe and comment the results before I wrap-up and conclude. In the appendix I present details of the data and an additional exercise.

## 2. Literature Review

Inflation and monetary policy have been extensively studied in economics. Here I will focus on the most recent literature that deals with two specific topics inside this broad field: i) identifying causality across macroeconomic variables; and ii) identifying the causes of the recent surge in inflation, among them, the Russian-Ukrainian war, and the resulting food and energy price crisis. I will also touch on the recent econometric literature on causality.

### 2.1. Identifying the effect of macroeconomic shocks

As a macroeconomic shock is generally a treatment that affects a whole country, or the whole world, most standard methods that rely on the crosssectional dimension are not applicable for the matter of identifying causality. On the other hand, as noted by Ramey (2016), time series approaches must rely on numerous assumptions to impose some degree of exogeneity of the shock used for estimations. Ramey (2016) also claims that there is a growing difficulty in ever more sophisticated models to identify a monetary policy shock, which may result simply from better monetary policy practices and better understanding of it by the agents. She ends that section in an unencouraging note: *"… we simply do not have enough information to produce estimates with any great certainty"* (Ramey, 2016).

Still, brave scholars keep on searching for new ways of measuring all types of relevant macroeconomic parameters. In the following subsection I give a brief overview of some papers that study the recent inflation surge. In the next, I briefly comment on the state-of-the-art methods for causality identification.

### 2.2. Effect of war and supply shocks on inflation

The big spike of inflation in 2022 prompted a renewed interest in causes of inflationary pressures and identification of effects of shocks on inflation.

Ball et al. (2022), for instance, try to understand what determines the rise in inflation in the pandemic times by decomposing it into a core component and deviations of headline from the core. They then study what determined the deviations away from the core using three types of shocks, one being rises in energy prices. They estimate that, from the total 6.9 percentage points increase in 12-month inflation between September 2021 and December 2022, the energy price shock contributed with 2.7 percentage points. This is almost 40% of the increase observed. It is worth noting, however, that their measure of energy price shock is the difference between energy price inflation and median inflation. In this sense, they consider all changes in energy prices that occurred along the period, and do not do a specific event study. With a similar indirect measure of supply shocks, which is a combination of deviations of energy and food prices from headline "core" inflation, Benigno and Eggertson

(2023) build a NK-model with a non-standard Phillips curve and arrive at a very different result. They conclude it was a labor shortage that caused the majority of the climb in inflation in this early 2020's. The exercise estimates the supply shock to have contributed with around 0.2 percentage points to the movement of inflation between 2021 and 2022. They also conclude that: *"… the supply shocks (…) contribute to the initial surge of inflation in the second and third quarters of 2021, but the main reason for the ongoing inflation during 2022 is the tight labor market"* (Benigno and Eggertson, 2023). In this paper I contribute to this discussion by focusing on the single event of the invasion of Ukraine by Russia. I also move away from any decomposition of inflation or imposed exogeneity assumptions.

Hall et al. (2023) use a VAR approach to study the evolution of inflation in three monetary zones (U.S., U.K. and Eurozone) and their interdependencies. They try to identify the effects of several shocks on inflation, introducing a method that would circumvent the usual caveats of the standard Cholesky decomposition by solving the VAR backwards, which they claim can better capture the shocks. Their work however also does not focus only on the war, but instead use oil prices as one of the variables and calibrate a hypothetical shock of 50% to build impulse response analysis. Here I deviate from this approach by completely avoiding the VAR estimation and its caveats. I also stick to the observed shock, evaluated in the observed state of the world.

Aastveit et al. (2023) use a SVAR approach to study how the effect of oil price shocks on actual and expected inflation depends on underlying demand and supply conditions, over short and long horizons. Since their data set goes only until 2019, they do not include the recent event triggered by the war in 2022, but they study four major periods of oil price shocks. Their conclusions indicate that, conditional on other demand and supply states, in some types of events (three of the four cases they study) an oil shock does have a significant effect on both actual and expected inflation, but the effect fades to zero already after three months. Here I contribute both by studying the shock observed in 2022 and, more importantly, avoiding the restrictions necessary for a SVAR estimation.

Financial history literature also links war to inflation, for example, in studies about the hyperinflation cases between the great wars of Germany, Austria, Hungary and Poland (curiously close to the current war). Don-Siemion (2021) uses a combination of time-series structural breaks estimations with narrative analysis. She concludes that the military events are a relevant cause of the Polish hyperinflation in that period. I provide here a modern case estimated with a different method with a result that goes in line with those.

### 2.3. Empirical identification of causality

Economic literature has been increasingly worried about soundly identifying causality effects from shocks and policy interventions. This work relates to that literature mostly through the empirical strategies used on settings close to natural experiments and the ones that use synthetic control methods. ABLS (2023) provide a comprehensive but quick overview of the most recent literature on these models. Here I briefly summarize their overview.

To deal with events close to natural experiments, a part of the literature relies on difference-in-difference methods to get a counterfactual path that serves as control (see ABLS (2023) for a list of references). These, in general, depend on the existence of one untreated region or a staggered policy implementation. This is not the case for the SBI method, which can deal with a shock that affected all individuals in the sample at the same time. The key to SBI is that the individuals must differ in stages at the time of the shock, as will be explained in the next section.

This literature usually also depends on the so called "parallel trends assumption", which basically requires that the paths before the shock differ, possibly, only in level, but not in slope or curvature. SBI also does not require that, as it uses a (parsimonious) mapping of paths before the shock, that is, it projects one path into the other, thereby making them comparable instead of assuming it. The only assumption is that, absent the shock, the normalization done with data before treatment would also work for the later periods.

The synthetic control methods, on their side, require the existence of some untreated individuals that could be used as reference to build a synthetic counterfactual for the treated individuals. This is also not the case of SBI, because it can identify the effect of a shock that hit all individuals of the sample. The identification comes from the different stages in which the shock hit each individual.

Together these characteristics make SBI a well-suited candidate to study the effects on inflation caused by the war, that potentially affected every country in the world (or every region in a country) at the same time, while inflation was also maybe not evolving in parallel across all countries (regions). In the next section I describe the method.

## 3. The SBI Estimation

Broadly speaking, to identify a causal relation between a treatment and an outcome we need essentially two things: i) two groups of individuals, treated and non-treated; and ii) similar characteristics of individuals across the two groups before the treatment.

When economists want to study the effects of policies or shocks that already happened in reality, however, these two conditions are rarely met. In many cases that the policy or shock did not affect all individuals in the sample, it is hard to guarantee that, before the event, the treated were sufficiently similar to the untreated. On the other hand, in many cases that the individuals of the sample could be considered similar among each other, they all receive the treatment at the same time, and there is no control group available.

ABLS (2023) develop a method that deals with both these shortcomings, also when they occur together. Their method can uncover causality of an event that affects all individuals in the sample at the same time, even if individuals are different[1] before the treatment. Their idea is to transform the time of the shock, which is the same for every individual, into the stage at which each individual was when the shock hit. Hence, they called the method "Stage-Based Identification" (SBI). Here I summarize the main ideas of their method.

Imagine a hypothetical situation in which we have two regions in a country. Initially their price levels are stable and inflation hovers around zero in both regions. After a period of high economic growth, inflation starts to pick up. For some reason[2] inflation in Region 1 grows faster and higher than in Region 2. At a certain moment both regions are hit by a shock (reducing inflation in this case). Panel (a) in Figure 2[3] illustrates this situation.

***Figure 2 – SBI normalization illustration***

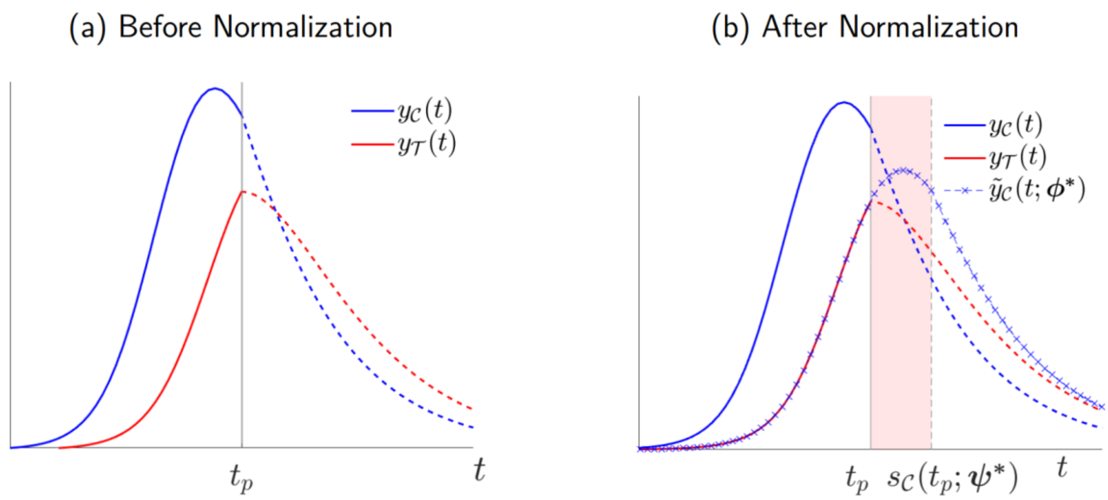


Source: ABLS, 2023.

Visual inspection of panel (a) shows that the parallel pre-trends assumption does not hold. And we only have these two regions, so there is no control group.

ABLS (2023) idea is to explore the apparent similar underlying process of the outcomes in both regions, even though parameters may be significantly different. They map the path of the outcome (in this case, inflation) of one

[1] Up to a certain degree. Individuals actually have to be different before the treatment, but not too much, as will be briefly explained in this section.
[2] But we are not interested on that.
[3] Taken from ABLS (2023).

region into the path of the other region. This can be seen in panel (b) of the graph above, where the blue dashed line with crosses is the inflation of Region 1 mapped into the path of inflation in Region 2. Then, the **stage** at which the shock hit Region 1 is now different than the stage at which it hit Region 2. Moreover, with the normalization, both regions can be considered sufficiently similar before the treatment (identical in this hypothetical example), and there is an overlapping window inside which Region 1 was not treated and Region 2 was. This allows for a clear identification of the causal effect of the shock inside the window. The causal effect of the event is then simply the difference between the blue dashed line with crosses and the red dashed line inside the window, in panel (b) of the graph in Figure 2. The graph in Figure 3 zooms in into the window, in panel (c), and shows only the policy effect, which is the difference between the two lines, in panel (d). The magenta horizontal line marks the final effect of the shock.

***Figure 3 – SBI normalization zoom and policy effect***

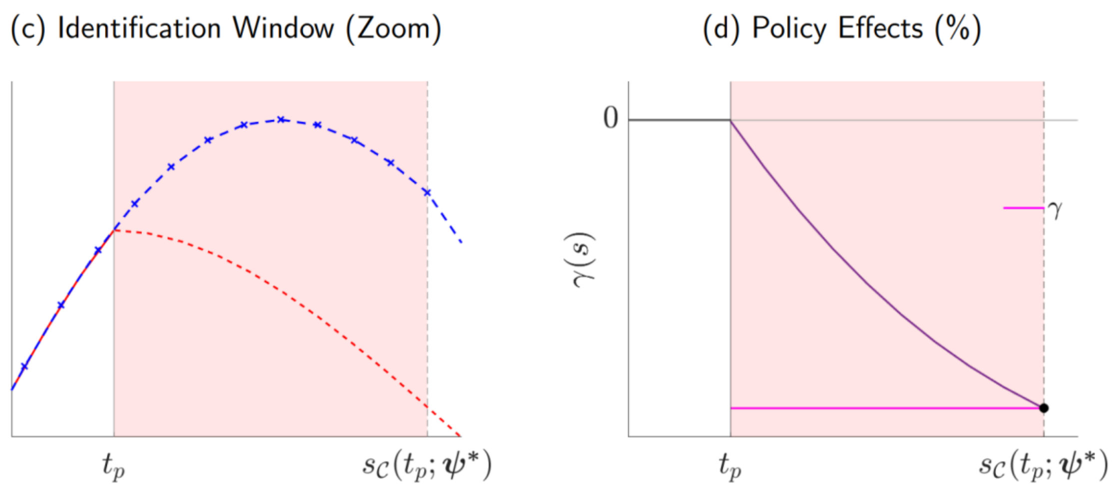


Source: ABLS, 2023.

For the mapping, the authors propose a parsimonious approach, mapping the outcome paths using a composite polynomial function with at most four parameters. More specifically, the parameters adjust the time into stage by speed and (maybe) acceleration, and adjust the outcome for level and slope. Parameters of the polynomials are estimated by minimizing the distance between the paths in the two regions.

For practical purposes, they also suggest smoothing the data before the mapping is performed.

The authors provide proofs and tests showing the method gives the same results irrespective of the choice of the reference region used as reference for mapping. For instance, if we map Region 2 into Region's 1 path instead, the theoretical result will be the same.

The authors also perform a series of tests to assess the robustness of the method in many dimensions. The method correctly captures a zero effect of a shock if it theoretically has no effect.

The key assumption for the method to deliver consistent causality identification is also not so restrictive. The method requires that the mapping parameters, that is, the coefficients of the polynomial in the composite function, do not change if the policy is implemented. Since the coefficients are estimated using only pre-shock data, the only challenge to the method will be in cases where the policy (or shock) was announced before happening, the agents acted differently because of the announcement, and additionally the actions were heterogeneous across the regions.

In practice, the method works well if the paths of the outcome variable have roughly the same "shape" across regions, but differ in some characteristics, whether in level, slope, stage, speed, or all of them. That is, the outcome paths must be different, but not too much[4].

Another practical requirement is that the paths before the shock cannot be completely flat, because in that case the stages cannot be defined.

To emphasize, the biggest advantages over other methods available to identify and estimate causal relations are that SBI: i) does not need a control group (not even a synthetic one); ii) does not rely on trends being parallel

[4] For example, it could be hard to map an exponential path into a sinusoidal one with a simple 4-parameter polynomial.

before treatment; and iii) is robust to confounding factors and endogeneity (up to a certain degree).

To properly work, that is, to deliver a reliable normalization (mapping) and an identification window, SBI requires that: i) there is a somewhat defined "shape" of the outcome time series before the treatment; ii) this "shape" has to show some heterogeneity across individuals; but iii) this heterogeneity cannot be too large. In other words, to apply SBI, the time series of the outcome variable must have some similar underlying process, but they cannot be exactly equal before treatment.

If those conditions are satisfied, it is likely that the method will successfully map one region outcome into the other, define an overlapping window and uncover a soundly identified causal effect of the event on the outcome, conditional on all other factors that were present at the time. In that sense, the method does not uncover the "pure" effect of the shock alone, but rather the effect of the shock in that particular situation.

To use the method for inference, the authors propose a bootstrap routine to estimate confidence intervals around the estimated shock/policy effect.

ABLS (2023) demonstrate the method and apply it, for illustration, to three quite different cases, but int this paper I apply it with the objective of answering a specific question. I do that, however, using a handful of different samples of the same outcome variable, which is inflation, across many countries (and regions) around the world. In the following section I describe the inflation data used in the exercise.

## 4. The data

The object of this study is one of the most followed economic variables through time and space: inflation. So, there is little need to digress about its definitions and characteristics. The main results illustrated in the next sections were performed using the most common, official, headline 12-month inflation rate measure, that is, the percentage change in the headline price index between the reference month and the same month in the previous year.

I perform the exercise using data for four different groups of individuals (samples). This is motivated by two reasons. First, the SBI method requires an *a priori* unknown balance between similarity and heterogeneity across individuals. Second, and more importantly, measuring the effect of the war on inflation is interesting for every country and region. Therefore, I use inflation data from different combinations of countries, or regions inside different countries. Specifically, I selected groups of individual countries/regions that arguably cover a relevant part of the spectrum of the "inflation stage heterogeneity". They also differ in the potential confounding factors basket that could influence the outcome variable[5].

The first group is composed by 17 countries in the Eurozone. Figure 4 shows inflation in the countries of the Eurozone, and the aggregate of the Eurozone (EUR) in red. They show a pattern of inflation flat in 2020 and upward-sloping from 2021 onwards. Apart from the three Baltic countries, heterogeneity is seemingly smaller than in the comparison across the globe saw in Figure 1.

For the 9 US regions, Figure 5 shows a similar pattern, but with the beginning of the upward trend and the peak occurring sooner. The heterogeneity is smaller than across the Eurozone countries.

The 16 metropolitan regions[6] in Brazil show again roughly a similar pattern in Figure 6, but it further differs from the Eurozone, with a well-defined trough in mid-2020 and a sharper and sooner reversal in mid-2022. It is however also noticeable that the heterogeneity seems a little higher than what is observed across the regions in the USA.

---

[5] For example, in an exercise using advanced and emerging countries across the world we could probably cite as potential confounding factors: i) monetary policy, ii) (sovereign) fiscal policy, iii) institutions, iv) development stage, v) local fiscal policy, and vi) income. But in the comparison focused on the German Länder, maybe only the last two would remain.

[6] The statistical institute (IBGE) does not collect inflation data in all 27 Brazilian states, but only in the biggest 16 metropolitan regions.

Last, the 16 Länder[7] in Germany also show a similar pattern in Figure 7, but with a less accelerated rise, and a later reversal. It is also noteworthy that heterogeneity is even smaller than in the Brazilian regions.

***Figure 4 - Inflation in Euro Area – 12-month %***

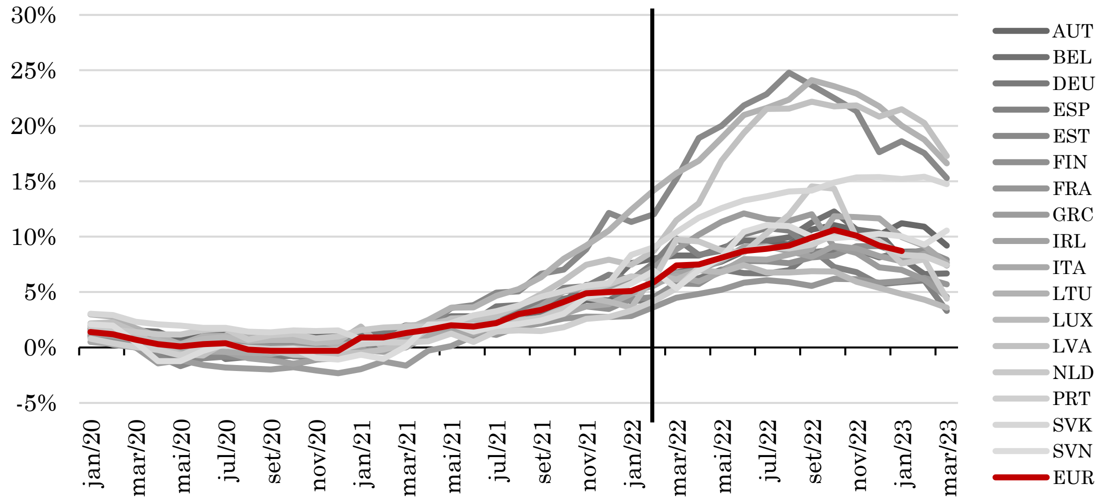


Source: OECD, own elaboration. Note: 12-month percentage variation of headline price index. The vertical black line represents the moment of the invasion of Ukraine by Russia.

***Figure 5 - Inflation in US-regions - 12-month %***

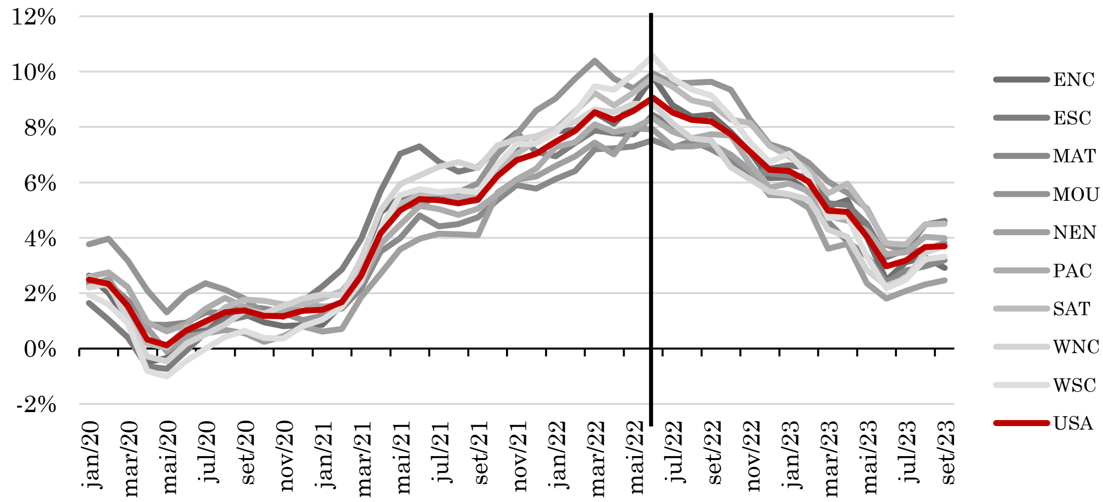


Source: OECD, own elaboration. Note: 12-month percentage variation of headline price index. The vertical black line represents the moment of the invasion of Ukraine by Russia.

[7] Länder is the German term that denotes the states (and city-states) of the German federation.

***Figure 6 - Inflation in Brazilian Regions - 12-month %***

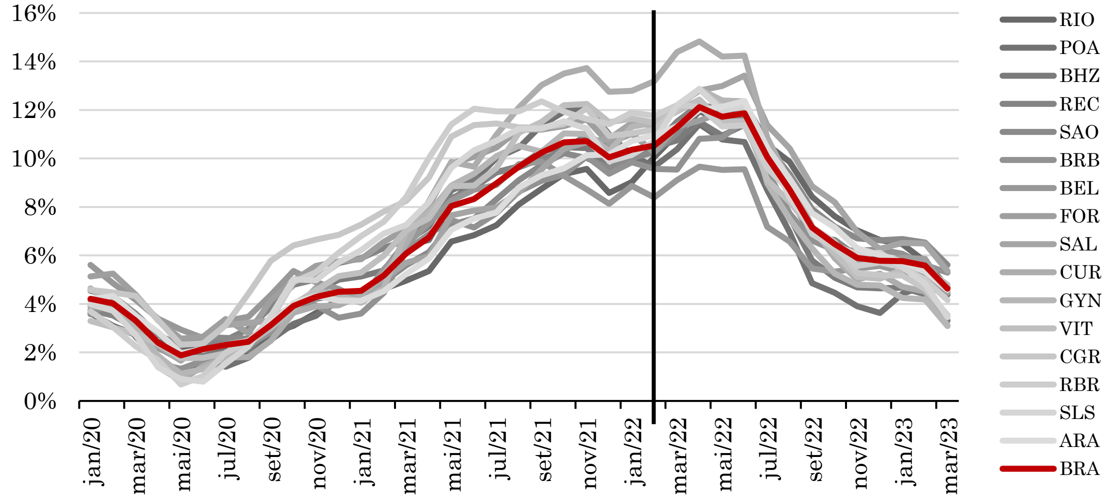


Source: IBGE, own elaboration. Note: 12-month percentage variation of headline price index. The vertical black line represents the moment of the invasion of Ukraine by Russia.

***Figure 7 - Inflation in German Länder - 12-month %***

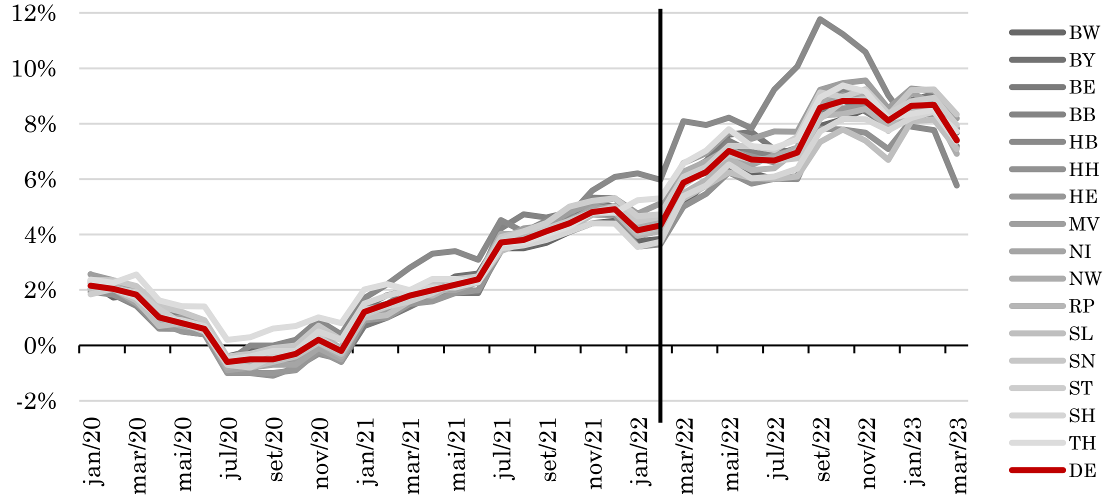


Source: Destatis, own elaboration. Note: 12-month percentage variation of headline price index. The vertical black line represents the moment of the invasion of Ukraine by Russia.

Though the groups show slightly different movements, it is clear that there is an upward trend that begins somewhere between mid-2020 and mid-2021. This trend lasts until at least mid-2022. That is, the shock hit inflation in the middle of an upward trend. Eyeballing an effect is hard.

In this work I focus the analysis on the period between January 2020 and March 2023. The starting point takes advantage of the fact that the world was still experiencing a subdued inflation period following the Great Financial Crisis (GFC) of 2008-2009 and the Euro Crisis of 2011-2012 and experienced a further disinflationary period in the beginning of the 2020 pandemic, which was shared by virtually all countries in the world (or in the sample). This helps to ensure that all paths have some similarity in their underlying evolution, at least in the first stages of the period studied.

The next section describes the empirical strategy I use in this setting to implement the Stage-Based Identification (SBI) method and uncover the effect of the war on inflation across the globe.

## 5. Empirical Strategy

As described in ABLS (2023) and summarized above, given a pair of regions (or countries[8]), the method first uncovers the region for which the path of the outcome variable in question leads the other. This is done endogenously by the method, after mapping one outcome path into the other, comparing the normalized paths and the resulting stages when the shock occurred. The leading region will then be considered the control subject by the method, while the lagging region will be the treated one. This is because the shock occurs in a later **stage** for the leading region, therefore opening an identifying window (overlap) inside which the leader was not treated but the other region was. The effect of the shock is then calculated simply by comparing the normalized paths inside the identification window.

In a hypothetical case with only two regions and no noise in the data, this is roughly as simple as it sounds. But in practice, with noisy data and multiple regions in a sample, other steps are necessary.

Next, I briefly describe each step. More details of the application for each sample and the results will follow in the next section.

[8] For the sake of readability, from now on I will use mostly the term "regions", even though I may be referring to sovereign countries that are part of a supranational group.

### 5.1. Smoothing

To begin with, before mapping one path into another, the authors suggest that the outcome series be smoothed, when we are not interested in the noisy (or cyclical) variance around the path, but in the underlying path itself. The smoothing can be done in different ways. The authors use for their three examples a Chebyshev polynomial smoothing and propose as alternatives, for example, moving averages, HP filters, and cubic splines. Here I follow the advice and smooth the data, but with a more parsimonious strategy, smoothing the pre-shock data with standard (monomial bases) polynomials of low degrees.

### 5.2. Finding the leading region

The discovery of the leader is done endogenously by the model in a pairwise comparison, but in this exercise, I must discover the leader among many regions. Because of that, as an intermediate step, I run the normalization partially, only up to the part that calculates the overlap interval, for each region in the sample.

Based on the algorithm developed by ABLS (2023), to find the overall leader I compare each region with the aggregate[9] and find the resulting overlapping windows. I define the overall leader as the one with the biggest (positive) overlapping interval with respect to the aggregate.

### 5.3. Model choices

The SBI normalization is fundamentally based on the procedure that maps one region's outcome into the other. This mapping also is subject to choices by the researcher. Here I follow ABLS (2023) and use a parsimonious mapping function, choosing a linear time-stage mapping ($\varphi_0 + \varphi_1 t$) and a linear

[9] The Eurozone, USA, Brazil, or Germany in our four samples.

(additive + proportional) level adjustment ($\omega_0 + \omega_1 y$). So, the composite mapping function that defines the normalized outcome has the following format:

$$\tilde{y}(t) = \omega_0 + \omega_1 y(\varphi_0 + \varphi_1 t) \tag{1}$$

### 5.4. Two types of measures

The effect of the shock is estimated by comparing a control region and a treated region. If there are only two regions to begin with, there is no confusion about that. But in a setting in which the available sample has multiple regions, this can be done in more than one way.

To estimate a measure of the effect of the shock on a whole country, using many regions inside that country, the strategy proposed by ABLS (2023) is to find an overall leading region, which will serve as reference, and group all other regions into an aggregate called "Rest-of-…". The estimated effect can then be considered as the aggregate effect on the country, in relation to what happened with the leading (untreated) region.

On top of that measure, however, one could also be interested in the individual regional effects inside a country. For that, each region can be separately compared to the overall leader. In this case we would have a different estimated effect of the shock for each region, except for the leader itself. These regional effects could then be used to answer other questions, for example, whether the regions geographically closer to Ukraine and Russia were more affected by the war.

In this paper I proceed with the first type of measure, following ABLS (2023), leaving the second type for future work.

### 5.5. Building the "Rest-of-…" aggregate

To uncover the effect of the shock according to the first measure described above, one needs to re-aggregate all regions except the leader. We call that the "Rest-of-…" region. This should be done using appropriate weights, which in the case of headline inflation would be the original regional weights used by the statistical offices to compose the aggregate headline inflation index from the regional subindexes.

The calculated outcome for the "Rest-of-…" region is then smoothed with the same procedure used for all individual regions.

### 5.6. Mapping the outcome paths

With the leader (control) and the "Rest-of-…" (treated) defined, the next step is to map (normalize) one outcome path into the other.

The mapping consists of minimizing the distance between the two (smoothed) paths before the shock. In an empirical setting, this is generally done by minimizing the sum of squared differences between the observed outcome from the two regions, at each point in time (stage). The minimization is then performed numerically. In some cases, however, it is possible to avoid the numerical solution. This happens when the number of coefficients in the mapping function equals the number of coefficients in the smoothing functions that represent the outcome data.

In this work I choose a mapping function composed by four coefficients $(\omega_0, \omega_1, \psi_0, \psi_1)$, and I smooth the pre-shock data with a polynomial function. If the chosen "optimal" degree of the smoothing polynomial turns out to be three, resulting in a functional form with four parameters[10], then an analytical closed-form solution is available. As described in the following sections, this turned out to be the case for three of the samples.

The analytical solution for the case of a 3rd-degree polynomial is presented by ABLS (2023) in their Appendix.

---

[10] $\beta_0 + \beta_1 t + \beta_2 t^2 + \beta_3 t^3$

With the data appropriately normalized, the next step is to finally calculate the effect of the shock.

### 5.7. Calculating the effect of the shock

Here I depart slightly from ABLS (2023). Their proposed measure for the policy effect is a ratio of integrals between the two normalized paths. It is, therefore, a relative measure of a cumulative effect, suitable for a flow variable.

In this study the outcome variable is headline inflation, which is not exactly a flow variable. Instead, we are interested in the simple difference between inflation with and without the shock (either along the path inside the overlap or at the end of it). Hence, the effect of the shock I discuss below is defined as the simple difference between the treated region outcomes and the normalized outcomes of the control region[11], for each point inside the overlapping identification window. The effect is estimated as follows:

$$\gamma(t) = y_T(t) - \widetilde{y_C}(t) \simeq \widetilde{y_T}(t) - y_C(t) \tag{2}$$

where $\gamma(t)$ is the estimated effect, $y_i$ is the outcome variable, T and C refer, respectively, to treated and control, and the tilde (~) indicates the normalized series.

### 5.8. Bootstrapping

To advance into inference and have a say on whether the estimated policy effect is statistically significant, we perform a bootstrap that constructs 90% confidence intervals around the measure described above.

In the next section I present and discuss the results for most of these steps.

[11] And vice-versa for the reverse mapping, which, as explained in ABLS (2023), produce fundamentally the same results.

# 6. Results

In this section I present the results of the SBI normalization, and the resulting estimated effects of the war shock on inflation, across the four above-mentioned samples of regions (countries). I start performing all the steps described above for the countries in the Eurozone. After that I summarize the same steps to the other three samples.

## 6.1. Countries in the Eurozone

The countries in the Eurozone form a sample particularly interesting for this exercise. They share important factors that influence inflation, mainly the common monetary policy, but also other institutional frameworks. And, of course, they are geographically very close to the war, and some are considered the most affected by the war (outside Ukraine, of course).

### *6.1.1. Smoothing*

I ran the smoother for all regions, using, as described above, a polynomial smoothing with monomial bases. I identified what I call here the "optimal" degree of the polynomial as the degree for which the adjusted-$R^2$ reduced for the first time, for each region. In other words, I find the first local[12] maximum adjusted-$R^2$ for each region. Since I want to smooth the data from all regions with the same functional form, I use the simple average of those "optimal" degrees, rounded to the lower integer.

The "optimal" smoothing degrees for the Eurozone turned out to be 3.

After this first smoothing step, I proceed to find the overall leader.

[12] I increase the degree, evaluate the adjustment, and pick the degree immediately before the first drop in adjusted R-squared.

### 6.1.2. *Finding the leading region*

To find the region that is considered as the control inside the identification window, I do the following.

In a first step, I run a part of the SBI normalization algorithm for each region against the aggregate[13] of the sample. The step produces the estimated overlapping windows of each region against the aggregate. The region for which the overlap is biggest is then considered the overall leader and will be the control of the second step, described in the next section.

In practice, a couple of important extra actions were taken. Firstly, a very large overlap for one region could indicate that the algorithm did not work properly in that case. In this work I discard those regions. Second, even if we were agnostic about the size of the interval and kept everyone at first, the version of the SBI algorithm used here still has some limitations and cannot calculate the effect of the shock when the overlap is big enough to go beyond the end of the sample. Hence, here I restrict the analysis to a subsample of regions for which the algorithm was capable of producing final results.

For the countries in the Eurozone, the results in Table 1 show a wide range of values for the estimated overlapping identification windows, including more than a couple extreme values. Considering the limitations of the algorithm, I am left with a subsample of mid-range overlapping identification windows from 6 countries: ESP, EST, FRA, FIN, IRL and LVA. I named this subsample "EUm" (m from "mid-range").

From those 6 countries, the "overall" leader is Ireland (IRL), with an overlap of +5.9 months against the aggregate Eurozone.

[13] For countries in the Eurozone, the outcome of the aggregate is the headline inflation index for the whole Eurozone itself.

***Table 1 – Eurozone countries identification windows – 1st step***

| Country | Overlap (months) |
|---|---|
| GRC | -11,3 |
| AUT | -9,6 |
| LUX | -8,4 |
| SVN | -7,5 |
| EST | -3,1 |
| LVA | -2,7 |
| ESP | 2,7 |
| FIN | 2,8 |
| FRA | 4,2 |
| IRL | 5,9 |
| ITA | 15,4 |
| DEU | 18,0 |
| PRT | 21,6 |
| LTU | 29,4 |
| SVK | 36,8 |
| NLD | 62,7 |
| BEL | 63,7 |

Note: Months between the date of shock in EUR and the normalized date (stage) of shock in each country. Shaded lines represent the "extreme" values, that is, those which the algorithm could not handle.

### *6.1.3. Building the "Rest-of-…" aggregate*

After finding out which region is the leader, I reaggregate all other regions into a group called "Rest-of-EUm". For the reaggregation I use the original weights applied by the respective statistical offices in the composition of the aggregate index from the regional sub-indexes. I exclude the overall leader and recalculate the "Rest-of-…" aggregate as a simple weighted average of the other regions in the sample.

The weights used are briefly described in the Appendix I – Weights for regional subindexes.

### *6.1.4. Stage-Based Mapping*

So far, we found the leader, which will serve as the control, and constructed the "Rest-of-EUm" aggregate, which will be the treated subject. The next step of SBI normalization is to map the path of the outcome variable from one region into the other. The specifics of the mapping procedure and some example plots of the original series and the normalized (mapped) series are presented below.

As mentioned above, for the Eurozone countries the chosen degree of the smoothing polynomial is three (four parameters when including the intercept). Since the mapping function is composed of four coefficients, a closed-form analytical solution for the mapping is available. See the Appendix from ABLS (2023) for a detailed derivation.

The plot in Figure 8 - panel (a) shows the original and smoothed data before the normalization. The plot in panel (b) shows the original path and data for "Rest-of-EUm" (red triangles) and the normalized data for IRL (blue circles). In panel (b), as expected from the analytical solution, the smoothed paths align perfectly.

***Figure 8 – Mapping of Ireland into Rest-of-Eurozone***

Panel a – Original and smoothed Panel b - Normalized

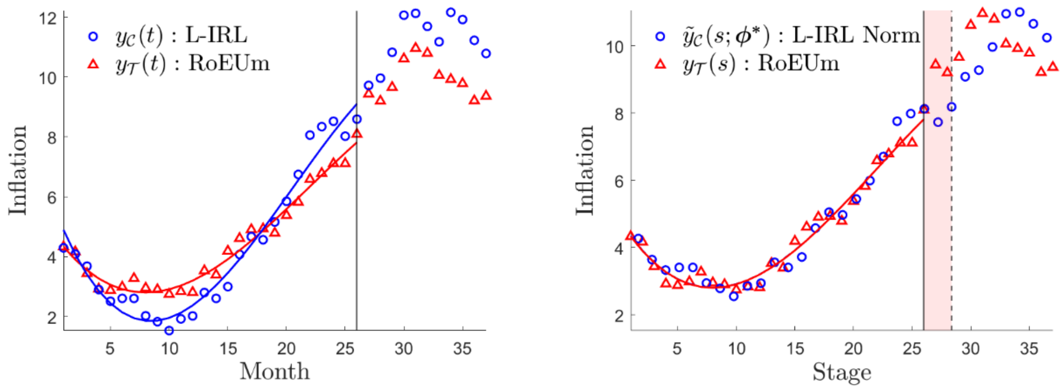


Note: Panel a: circles and triangles depict the observed data and lines represent the smoothed paths with a 3rd-degree polynomial. Panel b: the same, but data and smoothed path for IRL (blue) are normalized to the same stage as RoEUm according to the SBI mapping function. The black solid line shows the date of the shock (same as the stage for RoEUm) and the black dashed line represents the stage when the shock hit IRL. The pink shaded area, in between, marks the overlapping identification window.

### *6.1.5. Estimating the effect of the shock*

The effect of the shock on inflation can be estimated from those normalized paths simply as the difference between the data inside the identification window. In what follows I zoom in the identification window and show both the normalized data and the difference between them.

The effect of the war on inflation in the countries of the Eurozone (EUm) can be seen in Figure 9. In panel (a), the plot depicts the observed inflation for the "Rest-of-EUm" aggregate and the normalized data for IRL. After the shock, inside the identification window (shaded area), the data for the treated subject (Rest of EUm) goes above the normalized data for the control (IRL) by roughly 1 percentage point. This difference is depicted in panel (b).

***Figure 9 – Effect of war on inflation in the Eurozone***

Panel a – Normalized data

Panel b - Difference

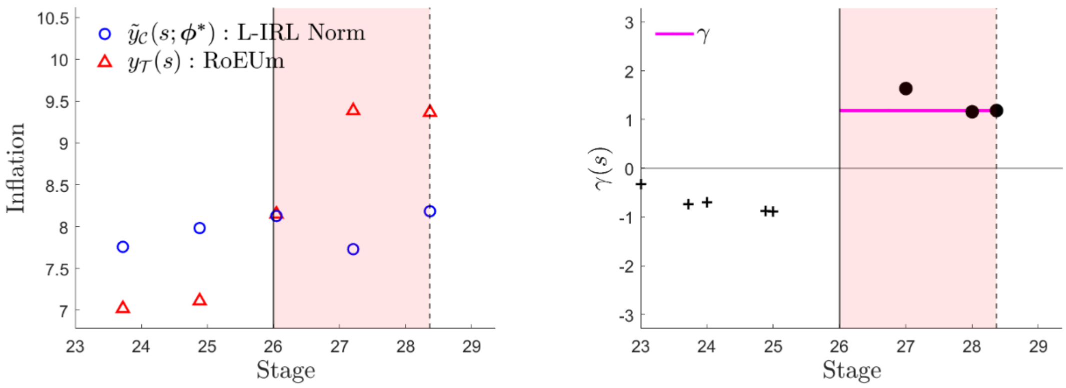


Note: Panel a: circles and triangles depict the observed and normalized data. Panel b: in the y-axis $\gamma(s)$ measures the effect of the shock in percentage points. The crosses represent the differences between data points before the shock. The black points represent the differences inside the identification window (shaded area). The magenta solid line shows the difference in the last period of the window.

### *6.1.6. Bootstrapping and inference*

For the bootstrapping, as proposed by ABLS (2023), I assume that the deviations from the smoothed paths are normally distributed, with zero mean and variances estimated from the data. I draw 1000 samples from the error distribution, add them to the smoothed paths and perform the SBI normalization for each one of the 1000 samples.

The graph in Figure 10 shows the results of the simulations for the EUm subsample. It is possible to say that, just before the shock, the difference between the (normalized) paths was zero. After the shock the effect seems to go up to around 1 percentage point, which is quite close to the point-estimate obtained with the actual data. Nevertheless, the confidence intervals indicate that the effect is not statistically different from zero.

***Figure 10 – Effect of war on inflation in the Eurozone – bootstrap***

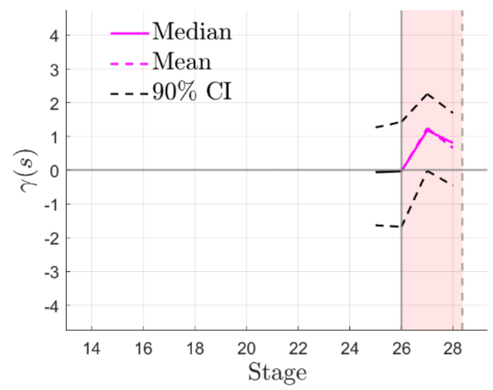


Note: In the y-axis, $\gamma(s)$ is the estimated effect of the shock, that is, the difference between the normalized paths, in percentage points. The black solid line represents the median of the difference between the paths before the shock. The magenta lines represent the median and mean of the effect after the shock (solid and dashed, respectively). The dashed black lines indicate the 5% and 95% centiles of the simulations, that is, the 90% confidence interval.

### *6.1.7. Placebo test*

To try and address the worry that the estimated effect could be a spurious result of the particular choice for the date of the shock, I run a placebo test, changing precisely that. That is, I feed the code with a date of the shock different than the actual. Around a date different than the actual shock, we expect that the method estimates an effect close to zero.

Note that this placebo test is different than the one described in ABLS (2023). There they put the SBI method through a placebo test in a controlled environment, that is, they test model generated data with a policy and without a policy. In the real-life case at hand, of inflation and war, we do not have data for which there was no shock from which we could run SBI and check if the estimate is close to zero. The placebo test will therefore still be influenced by the shock if we take placebo dates around the actual date.

Here I simply feed the code with a date of the shock two months earlier than the actual, that is, I ask SBI to calculate what was the effect on inflation of the placebo shock that happened in December 2021. I run with the placebo date all steps of the algorithm described above, including the smoothing, the leader-finding step, up to the bootstrapping.

***Figure 11 – Placebo test of the effect on inflation in the Eurozone***

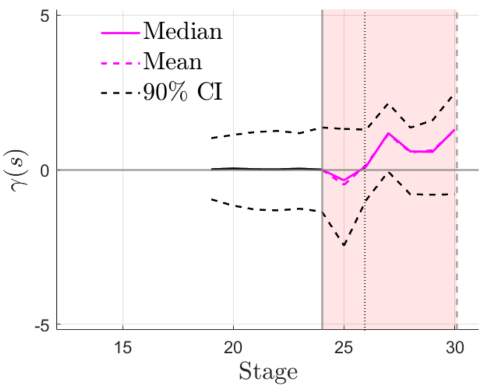


Note: In the y-axis, $\gamma(s)$ is the estimated effect of the shock, that is, the difference between the normalized paths, in percentage points. The black solid line represents the median of the difference between the paths before the shock. The magenta lines represent the median and mean of the effect after the shock (solid and dashed, respectively). The dashed black lines indicate the 5% and 95% centiles of the simulations, that is, the 90% confidence interval. The dotted line depicts the actual date of the shock.

Figure 11 shows the results of the bootstrapping estimates for this placebo test. The median/mean lines show that the estimated effects were seemingly close to zero in the first two months after the placebo shock, from December 2021 until February 2022, after which the estimates grow to close to one as in the main estimation described in the previous section.

## 6.2. Brazilian Regions, US Regions, and German Länder

In this section I describe the results of performing the same steps using inflation data for US regions, Brazilian regions, and German Länder.

### *6.2.1. Smoothing*

The "optimal" degrees for the smoothing polynomial found by the method in each sample are listed in Table 2.

***Table 2 – Degree of smoothing polynomial***

| | USA | Brazil | Germany |
|---|---|---|---|
| **Degree:** | 3 | 3 | 6 |

### *6.2.2. Finding the leading region*

The calculated overlapping windows for US regions, Brazilian regions and German Länder, shown in Table 3, panels A, B and C, respectively, are smaller than the ones estimated for the European countries, and gradually decreasing in dispersion. This is very likely due to the, unsurprisingly, more homogenous paths of the data, when we compare these three samples against the Eurozone. Looking only to the inflation paths, regions in USA show less heterogeneity than the countries in the Eurozone, but more so than the regions in Brazil and Germany, which is the less heterogeneous sample. There is however still sufficient variability, even inside Germany, for SBI to proceed, as will be shown in the following subsections.

***Table 3 – US, Brazilian and German identification windows***

*Panel A – USA*

| Region | Overlap |
|---|---|
| MOU | -7,1 |
| NEN | -2,3 |
| PAC | -1,9 |
| SAT | -1,1 |
| MAT | -0,4 |
| ENC | 1,1 |
| WSC | 1,5 |
| WNC | 3,1 |
| ESC | 5,4 |

*Panel B – Brazil*

| Region | Overlap |
|---|---|
| SAL | -4,2 |
| ARA | -3,2 |
| RIO | -2,0 |
| SAO | -1,0 |
| BRB | -0,2 |
| VIT | -0,1 |
| CUR | -0,1 |
| REC | 0,3 |
| GYN | 0,6 |
| BHZ | 1,0 |
| POA | 1,5 |
| FOR | 1,9 |
| CGR | 2,2 |
| RBR | 2,4 |
| SLS | 2,5 |
| BEL | 2,7 |

*Panel C – Germany*

| Region | Overlap |
|---|---|
| TH | -4,0 |
| HB | -1,0 |
| BE | -1,0 |
| HE | -0,3 |
| NW | 0,0 |
| BY | 0,1 |
| RP | 0,2 |
| BB | 0,2 |
| BW | 0,2 |
| NI | 0,3 |
| SH | 0,4 |
| HH | 0,6 |
| MV | NA |
| SL | NA |
| SN | NA |
| ST | NA |

Note: Months between the actual date of shock and the normalized date (stage) of shock in each region after normalization against the aggregate. A positive value means the region is in a stage ahead of the aggregate. Shaded lines indicate regions for which the subsequent steps of the code did not run, so they were taken out of the sample[14]

### *6.2.1. Building the "Rest-of-…" aggregate*

After finding out which region is the leader, I reaggregate all other regions into a group called "Rest-of-…". The weights used for the four samples in this exercise are briefly described in the Appendix I – Weights for regional subindexes.

### *6.2.2. Stage-Based Mapping*

The next step of SBI normalization is to map the path of the outcome variable from one region into the other. For the US and Brazilian regions, the "optimal" degree of the smoothing polynomial turns out to be three, so the analytical solution is also chosen, which shows on the perfect fit between the lines before the shock in panels B1 and B2 of Figure 12. For Germany, even though the degree of the polynomial is 6, the numerical mapping also returned a very precise mapping, shown in panel B3.

### *6.2.1. Estimating the effect of the shock*

Figure 13 zooms in the identification window to show the effect of the war on inflation in US, Brazilian and German regions. The panels in column A show a zoom-in into the identification windows and depicts, for each sample, the control region (Ro---) data and the normalized data for the leading region. Panels in column B show the simple difference between the two regions, for each point in time (stage) inside the identification window.

The results show that the estimated effects reach almost 1 percentage point in the USA, around 2.5 p.p. in Brazil and a little above 1 p.p. in Germany.

[14] In "Appendix II – Trimmed sample for German Länder", I run the same experiment, changing only the start of the sample, discarding the first two months and starting the sample in March 2020. With this small change, the "optimal" degree is again three and the code runs smoothly to the end. Results do not change significantly.

***Figure 12 – Mapping of USA, Brazil, and Germany***

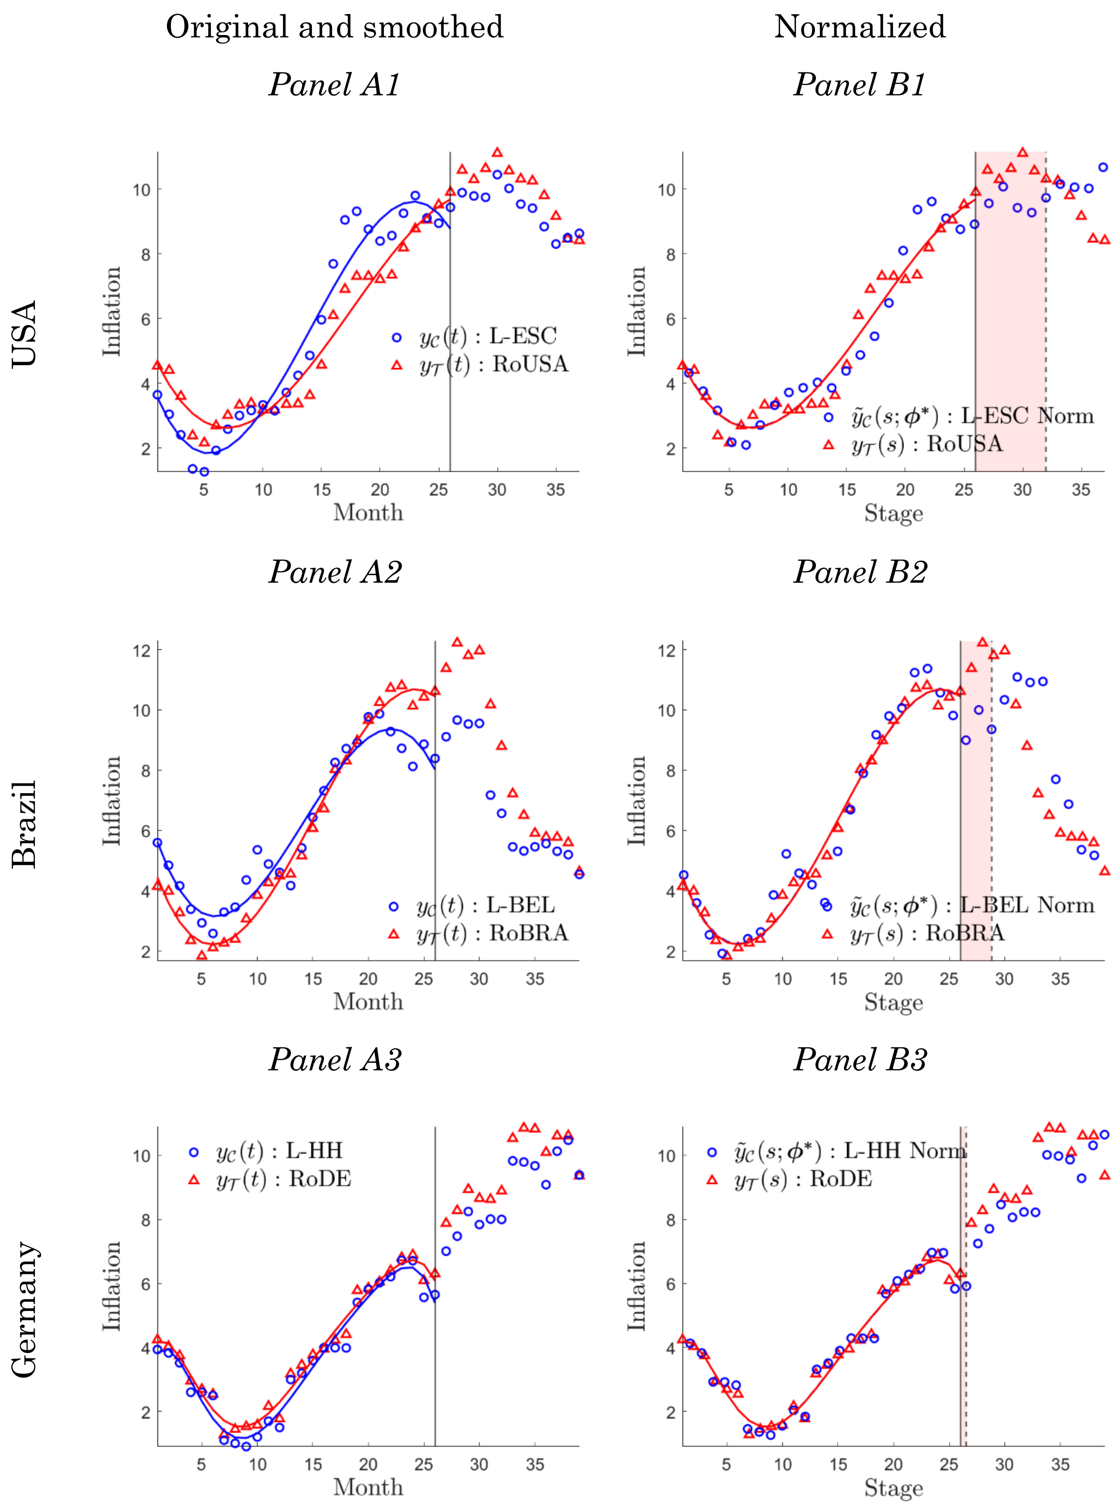


Note: Panel A: circles and triangles depict the observed data and lines represent the smoothed paths. Panel B: the same, but data and smoothed path for the leading (L) region (blue) are normalized to the same stage as that of the "Rest-of-..." aggregate (Ro---) according to the SBI mapping function. The black solid line shows the date of the shock (same as the stage for Ro---) and the black dashed line represents the stage when the shock hit the leading (L) region. The pink shaded area, in between, marks the overlapping identification window.

***Figure 13 – Effect of war on inflation in Brazil, USA and Germany***

Normalized data | Difference

*Panel A1*

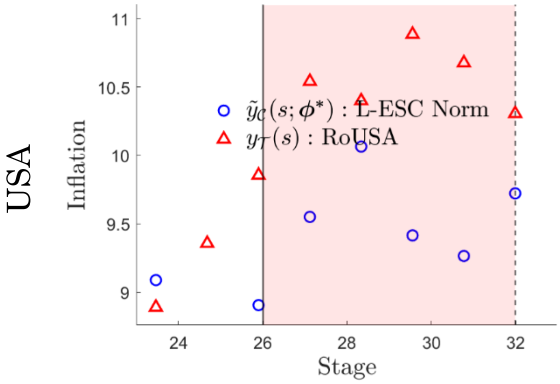


*Panel B1*

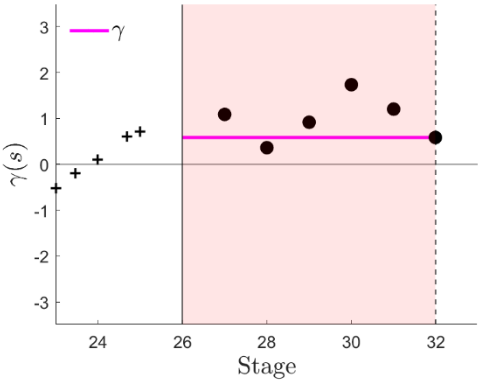


*Panel A2*

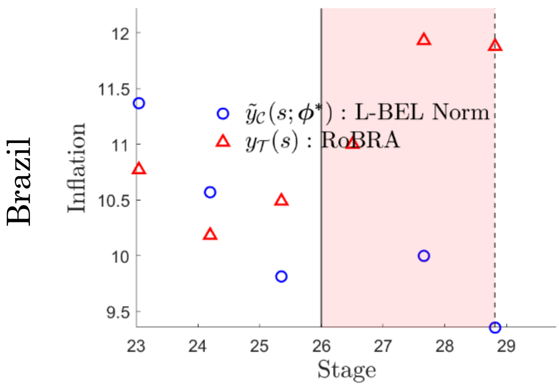


*Panel B2*

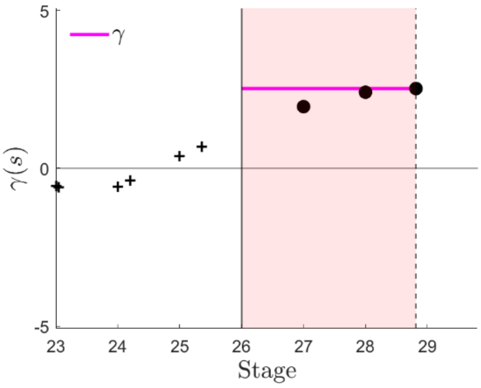


*Panel A3*

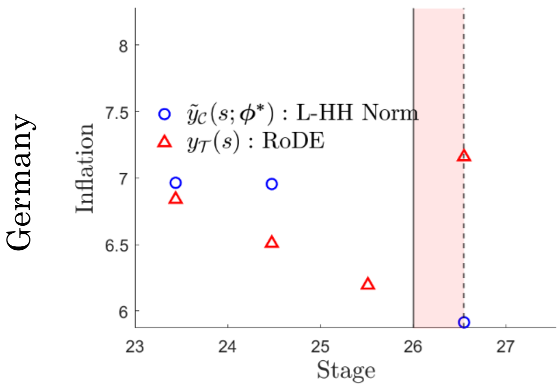


*Panel B3*

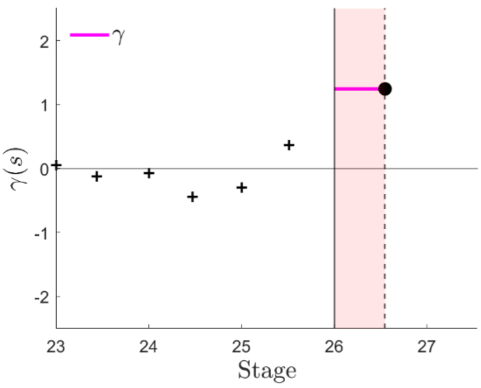


Note: Panel A: red triangles and blue circles depict the observed and normalized data, respectively. Panel B: in the y-axis $\gamma(s)$ measures the effect of the shock in percentage points. The crosses represent the differences between data points before the shock. The black points represent the differences inside the identification window (shaded area). The magenta solid line shows the difference in the last period of the window.

*6.2.2. Bootstrapping and inference*

Figure 14 illustrates the results of the bootstrap simulations of a 90% confidence interval (CI) for the effect of the shock in US and Brazilian regions (Panels A and B, respectively). For USA, the central measures are close to 1 percentage points, and for Brazil, they are close to 2.5 p.p., both close to the point estimates shown above. The confidence intervals indicate that the estimate is not statistically different than zero for the USA, while the Brazilian case is borderline significant.

***Figure 14 – Effect of war on inflation – bootstrap***

*Panel A – USA* *Panel B – Brazil*

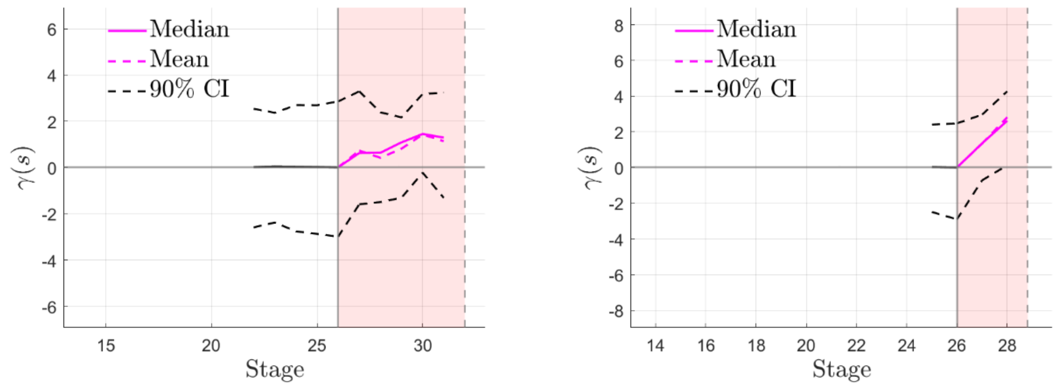


Note: In the y-axis, $\gamma(s)$ is the estimated effect of the shock, that is, the difference between the normalized paths, in percentage points. The black solid line represents the median of the difference between the paths before the shock. The magenta lines represent the median and mean of the effect after the shock (solid and dashed, respectively). The dashed black lines indicate the 5% and 95% centiles of the simulations, that is, the 90% confidence interval.

For the German Länder subsample, the algorithm could not produce the plots, but the confidence intervals were calculated and are as follows. While the point estimate (Figure 13) indicate that the effect of the war on Germany was approximately 1.2 percentage points, the mean and median estimates of the bootstrapping procedure are around 2.5 percentage points. The lower and upper bounds of the 90% confidence interval are, respectively, 0.3 and 3.8 percentage points, indicating that the effect was positive and statistically significant.

### *6.2.3. Placebo tests*

I run the placebo tests with the placebo shock date two months before the actual. The results are depicted in Figure 15. For USA the central estimates detach from zero already in the first period of the overlap interval, but less than what is seen after the actual shock. For Brazil the estimates remain close to zero until the actual date and then start rising.

***Figure 15 – Placebo tests in USA, Brazil and Germany***

*Panel A – USA*

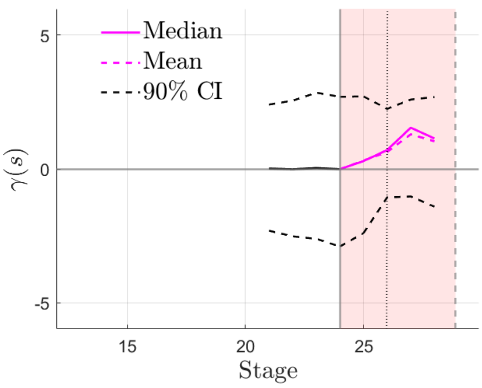


*Panel B – Brazil*

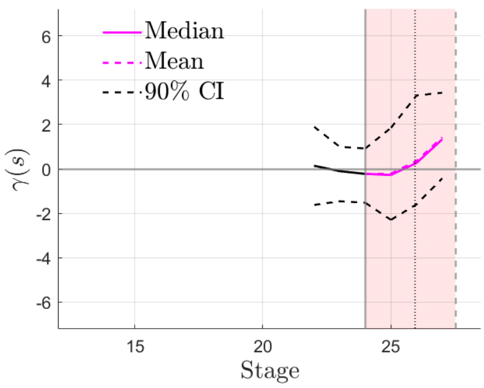


Note: In the y-axis, $\gamma(s)$ is the estimated effect of the shock, that is, the difference between the normalized paths, in percentage points. The black solid line represents the median of the difference between the paths before the shock. The magenta lines represent the median and mean of the effect after the shock (solid and dashed, respectively). The dashed black lines indicate the 5% and 95% centiles of the simulations, that is, the 90% confidence interval. The dotted line depicts the actual date of the shock.

## 6.1. Individual regional effects

Instead of rebuilding the "Rest-of-the-country" data to get an aggregate effect, it is also possible to run the SBI normalization for each region. I keep all specifications of the model the same as detailed in the previous sections and use the same leader already found for each country. With that we can uncover the effects of the shock on inflation inside each individual region, apart from the leading one. The results are described in Table 4.

In general, estimates are positive and do not stray too far from the aggregate estimate (also included in the last line for comparison), but there is a sizeable heterogeneity. Whether the differences between regions are statistically significant still needs to be assessed but considering the 90% confidence intervals from the bootstrapping we can have a hint. For the Eurozone, the two extremities (LVA and EST) have an estimated effect of the shock on inflation of +5 and -5 p.p. respectively, which fall outside the confidence interval simulated for the aggregate analysis. For the United States the confidence interval is larger and there are no extreme estimates, so no individual region falls outside it. In Brazil, only two regions (ARA and SAL) fall below the interval with sizeable negative estimates (-1.7 and -3.3), which warrant further investigation. The estimates for the German Länder all fall inside the 90% confidence interval of the aggregate effect.

***Table 4 – Individual regional effects***

| **Eurozone** | | **United States** | | **Brazil** | | **Germany** | |
|---|---|---|---|---|---|---|---|
| **Region** | **Effect** | **Region** | **Effect** | **Region** | **Effect** | **Region** | **Effect** |
| LVA | 5.0 | WSC | 2.6 | GYN | 3.0 | BE | 2.1 |
| FIN | 2.0 | ENC | 2.1 | CUR | 2.8 | HE | 2.1 |
| FRA | 1.2 | MAT | 1.1 | REC | 2.8 | NW | 1.6 |
| ESP | 0.5 | MOU | 1.1 | SAO | 2.6 | BY | 1.3 |
| EST | -5.0 | WNC | 0.8 | BRB | 2.3 | BB | 0.9 |
| IRL | - | SAT | 0.4 | BHZ | 2.1 | RP | 0.9 |
| | | PAC | 0.0 | RIO | 1.9 | SH | 0.9 |
| | | NEN | -0.8 | VIT | 1.2 | BW | 0.7 |
| | | ESC | - | SLS | 1.1 | NI | 0.7 |
| | | | | CGR | 1.0 | HH | - |
| | | | | FOR | 1.0 | | |
| | | | | RBR | 0.9 | | |
| | | | | POA | 0.8 | | |
| | | | | ARA | -1.7 | | |
| | | | | SAL | -3.3 | | |
| | | | | BEL | - | | |
| Mean | 0.8 | Mean | 0.9 | Mean | 1.2 | Mean | 1.2 |
| Median | 1.2 | Median | 0.9 | Median | 1.2 | Median | 0.9 |
| RoEUm | 1.2 | RoUSA | 0.6 | RoBRA | 2.5 | RoDE | 1.2 |

Note: Impact of shock on 12-month headline inflation, estimated for each individual region, in percentage points. Shaded cells show estimates that fall outside the 90% confidence interval simulated for the aggregate. The table also shows the simple average and median of the regional estimates in each country. For comparison, the last line shows the estimates for the aggregate effect calculated in the previous section.
Source: Own calculations.

Apart from reinforcing the robustness of the estimates for the aggregate performed before, this information can also be used as input to other studies that may, for example, relate regional characteristics to the individual estimated effects. I leave that for future studies. Next, I summarize the main results and put them into perspective with a brief discussion.

### 6.2. Results summary

Taking all results together, the evidence points to inflation being positively affected by the invasion of Ukraine by Russia in February 2022. The positive effect has an economically relevant size of around 1 or 2 percentage points, depending on the sample, which compares to yearly inflation levels swinging roughly inside the 5% to 10% range between 2021 and 2022. Figure 16 shows inflation for these two years and highlights the identified effect of the war in 2022, in the four samples.

It is then possible to say that in Brazil and the USA other factors pulled inflation down between 2021 and 2022, despite the positive (and big for Brazil) impact of the war. On the contrary, in the Eurozone and in Germany, other factors pushed inflation up (significantly) more than the estimated impact of the war did.

The estimated values are statistically different from zero for two of the samples studied in this paper (Brazilian and German regions), but not significant (with a 90% confidence interval) for the Eurozone countries nor for US regions.

In particular, for the German Länder, the 90% confidence interval lies between +0.3 and +3.8 percentage points. For the Brazilian regions the confidence interval goes from +0.9 to +5.0 percentage points. As for the countries in Europe, the interval contains zero and it is not possible to sustain that the effect is statistically significant.

***Figure 16 – Inflation in 2021 and the effect of war in 2022***

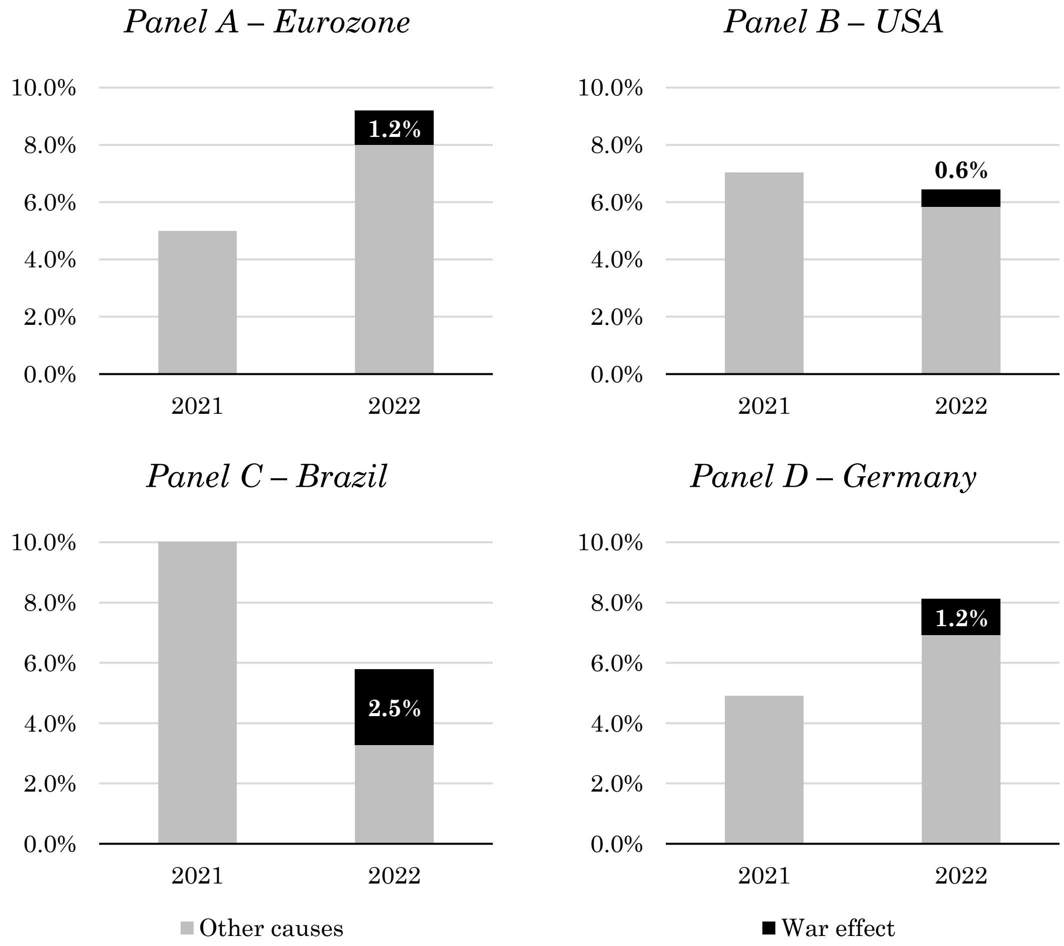


Note: Gray bars represent headline inflation for 2021 and headline inflation excluding the effect of the war in 2022. Black bars represent the estimated effect of the war, so that the total height of the bar is equal to total headline inflation in 2022.

These estimates compare with the values already found in the literature of 0.2, 1.2 and 2.7 percentage points, estimated by Ball et al. (2022), Hall et al. (2023), and Benigno and Eggertsson (2023), respectively. These estimates however do not refer specifically to the shock of the invasion of Ukraine by Russia in February 2022, but to the broader supply and energy crises observed in 2021/2022. They also refer only to the effect on US inflation.

In this study I also estimate subnational regional effects. To the best of my knowledge this is the first paper to do so. Individual regional effects are broadly in line with the aggregate estimates. Heterogeneity is apparent, but the significance of the difference across regions is yet to be assessed.

In conclusion, on top of the undeniable difficulties faced by many sectors of the world economy caused by the war, not to mention the humanitarian costs, this study indicates that the war did push inflation up. But in Europe, the great spike in inflation observed between 2021 and 2022 was only partially explained by the war. These results are summarized in Table 5.

***Table 5 – Results summary***

| | **Eurozone countries*** | **US regions** | **Brazilian regions** | **German Länder*** |
|---|---|---|---|---|
| 2021 Inflation | 5 | 7 | 10,1 | 4,9 |
| 2022 Inflation | 9,2 | 6,5 | 5,8 | 8,1 |
| Leader | Ireland | East South Central | Belém | Hamburg |
| Window (months) | 2,4 | 6,0 | 2,7 | 0,6 |
| Impact of shock (p.p.) | *1,2* | *0,6* | 2,5 | 1,2 |
| 90% C.I. | *- 0,8 to +1,8* | *- 1,1 to +3,8* | +0,9 to +5,0 | +0,3 to +3,8 |

* Subsamples for Eurozone countries and German Länder, as described in previous sections.
Note: Inflation values are for the full standard aggregate country (group of countries), with no exclusions. Impact of shock estimated with the respective working subsamples. Estimates for which the 90% confidence interval contains the zero (not statistically significant) are shaded and italic.
Source: OECD, BLS, IBGE, Destatis. Own calculations.

Of course, these results must be taken with a grain of salt, as usual. To begin with, in this paper I perform the exercise with only a subset of the samples for the Eurozone countries and for the German Länder. Further work needs to be done to extend the analysis to all countries and regions in these samples. Second, the effect is identified (and should be consistently estimated) inside the overlapping window, which is only a couple to a handful of months long, depending on the sample. It is possible that a relevant part of the shock only impacts inflation with a lag greater than the time span captured by the windows estimated above. The window of identification in the comparisons performed in this study, however, goes up to six months in the case of the US regions, which is not so far from most estimates of relevant time horizons for price adjustments[15] (see, for example, Aastveit et al., 2023). Finally, it is relevant to remember that the results estimated by the SBI method are conditional on the specific situation, that is, if the same shock would occur when, for example, economic growth, monetary policy etc. where in a very different stance, the effect could potentially change significantly (see, again, Aastveit et al., 2023).

## 7. Conclusion

With this work I apply a novel method and introduce it to the literature that identifies causality of shocks on inflation, by challenging the notion that the war (the invasion of Ukraine by Russia in February 2022) caused (most of) the big spike in inflation observed in that year. To the best of my knowledge this is the first work that systematically and comprehensively evaluates this shock across many countries and estimates aggregate as well as subnational regional effects.

To identify the effect of the shock, I apply the Stage-Based Identification method (SBI) developed by Alemàn, Busch, Ludwig and Santaeulàlia-Llopis – ABLS – (2023). The SBI method normalizes individual outcome paths to a reference path by transforming time into stage. If the shock hit the individuals at different stages, there is a window inside which one region was treated and the other was not. For that to work, the method requires a mix between similarity and heterogeneity, that is, the original paths must be different to allow for identification, but not too much, so normalization works (ABLS, 2023).

Looking at inflation across the globe we see a fairly common trend following the onset of the COVID pandemic in the beginning of 2020, which is subdued inflation, and also a common upward trending inflation after one year or so.

[15] "*For instance, an economic activity shock, that increases oil prices by 10 percent, elicits an extremely persistent response in both expected and actual inflation. In contrast, when the economy is hit by shocks to supply, consumption demand, or inventory demand, we find that both expected and actual inflation initially increase but the effect typically dies out after a quarter. This highlights the importance of identifying the underlying oil market shocks when analyzing the effect of the oil price on actual inflation and expected inflation*." (Aastveit et al., 2023)

This upward trend however presents different timings across countries, and the peaks (or trend reversals) are also heterogenous.

Despite the theoretical soundness of the model, a practical implementation in such a setting can raise questions about the degree to which other factors could get in the way of the correct identification of causality and its size. With that in mind, in this paper I use inflation data for countries in the Eurozone, regions in the United States, regions in Brazil, and Länder in Germany. Consequently, I have a broad spectrum of the trade-off between the similarity in the underlying inflation trends and the heterogeneity needed for identification.

The results indicate that, given the state of the world at the time, the war had a positive impact on headline inflation in most countries. The effect estimated from the fourthree samples in this study has an economically relevant size of around 1 or 2 percentage points, depending on the sample, which compares to yearly inflation levels swinging roughly inside the 5% to 10% range between 2021 and 2022. The estimates are statistically significant for two cases (Brazilian regions and German Lãnder). For the Eurozone countries the estimate cannot be considered statistically significant.

Comparing these estimations with the variation of total headline inflation, it is possible to say that in Brazil and the USA other factors pulled inflation down between 2021 and 2022, despite the positive (and big for Brazil) impact of the war. On the contrary, in the Eurozone and in Germany, other factors pushed inflation up (significantly) more than the estimated impact of the war did.

One dimension that is key to the SBI method and can also be explored in this setting is the regional heterogeneity. To the best of my knowledge this is the first paper to evaluate the subnational regional dimension with this shock. For this I implement the SBI normalization across all regions of each sample, individually paired against the leading one of the respective countries. This uncovers individual estimated effects that are broadly in line with the estimates for the aggregate ones but show sizeable heterogeneity.

Significance of the heterogeneity is yet a matter of research but comparing the regional effects with the bootstrapping 90% confidence intervals for the aggregate, only two countries in the Eurozone and two regions in Brazil fall outside the simulated range. Apart from reinforcing the robustness for the aggregate effects estimates performed before, these regional effects could be used to answer other questions, for example, whether the regions geographically closer to Ukraine and Russia were more affected by the war, or what characteristics of consumption baskets and productive structures correlate with the regional heterogeneity observed here.

Additional tests over the exact specification of the normalization procedure would also help impose robustness on the results found in these exercises.

Another extension would be to apply the same exercise in yet a more localized data set, using a consumer price index collected in half a dozen urban areas inside the city of São Paulo – SP, in Brazil, by Fipe. This will give another combination of similarity and heterogeneity, even less prone to critics about confounding factors, since they are all collected inside the same municipality.

Finally, the estimation could be made more robust if done with a higher frequency data set, resulting in more data points inside the identification overlapping window. This could be done for Brazil, for example, by using at least three additional data sets. First, there is a mid-month "preview" of the official monthly price index published by the statistical office (IBGE), which would result in a time series with a 15-day periodicity. There is also another consumer price index published every 10 days by Fundação Getúlio Vargas (FGV). And, finally, it would be interesting to use a series, also collected by FGV, which tracks the official inflation index with daily surveys.

To conclude, the results from all these exercises may call attention in academia and policy institutions for this novel tool that can identify causal macroeconomic effects of shocks. The paper should also contribute to policy makers, offering an answer to a common question (what part of current inflation was caused by the war?), but derived from a different method.

## 8. References


Aastveit, K. A., Bjørnland, H. C., Cross, J. L. (2023). Inflation Expectations and the Pass-Through of Oil Prices. *The Review of Economics and Statistics*; 105 (3): 733–743; doi: https://doi.org/10.1162/rest_a_01073

Alemán, C., Busch, C., Ludwig, A., Santaeulàlia-Llopis, R. (2023). Stage-Based Identification of Policy Effects. Mimeo. https://alexander-ludwig.com/wp-content/uploads/2023/10/paper.pdf

Ball L. M., Leigh, D., Mishra, P. (2022). Understanding U.S. Inflation During the COVID Era. Working Paper 30613, National Bureau of Economic Research. https://www.nber.org/papers/w30613

Benigno, P., Eggertsson, G. B. (2023). It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve. Working Paper 31197, National Bureau of Economic Research. https://www.nber.org/papers/w31197

Don-Siemion, T. (2021) 'We'll give up our blood but not our gold': money, debt, and the balance of payments in Poland's Great Depression. PhD thesis, London School of Economics and Political Science. http://etheses.lse.ac.uk/4395/

Eurostat (2018). *Harmonised Index of Consumer Prices (HICP): Methodological Manual*. Luxembourg: European Union. https://ec.europa.eu/eurostat/documents/3859598/9479325/KS-GQ-17-015-EN-N.pdf/d5e63427-c588-479f-9b19-f4b4d698f2a2?t=1547028935000

Hall, S., Tavlas, G., Wang, Y. (2023). Drivers and spillover effects of inflation: The United States, the euro area, and the United Kingdom. *Journal of International Money and Finance*, Volume 131, https://doi.org/10.1016/j.jimonfin.2022.102776

IBGE (2019). Atualização das estruturas de ponderação a partir da POF 2017-2018. Nota técnica 02/2019. https://ftp.ibge.gov.br/Precos_Indices_de_Precos_ao_Consumidor/Sistema_de_Indices_de_Precos_ao_Consumidor/Notas_Tecnicas/snipc_nota_tecnica_2019_02.pdf

Ramey, V. A. (2016). Chapter 2 - Macroeconomic shocks and their propagation. In: *Handbook of macroeconomics*, Eds: John B. Taylor, Harald Uhlig, Volume 2, Pages 71–162. Elsevier, . https://doi.org/10.1016/bs.hesmac.2016.03.003

Statistisches Bundesamt – Destatis (2019). Hintergrundpapier zur Revision des Verbraucherpreisindex für Deutschland 2019. Wiesbaden. https://www.destatis.de/DE/Presse/Pressekonferenzen/2019/HGG_VPI/Statement_HGG_VPI_PDF.pdf?__blob=publicationFile

## Appendix I – Weights for regional subindexes

### *1. Countries in the Eurozone*

For Europe we use the country weights calculated by Eurostat (2018) according to which "*… each country gets a weight that corresponds to its share of consumption expenditure in the total of the group*." The weights are updated annually according to that methodology. The values used in this paper are the following:

***Table A.I.1 – Weights for Eurozone countries***

| | **2020** | **2021** | **2022** | **2023** |
|---|---|---|---|---|
| AUT | 3,5% | 3,4% | 3,3% | 3,5% |
| BEL | 3,8% | 3,8% | 4,1% | 4,0% |
| DEU | 27,7% | 29,3% | 28,4% | 28,0% |
| ESP | 11,8% | 10,6% | 11,1% | 11,1% |
| EST | 0,2% | 0,2% | 0,3% | 0,3% |
| FIN | 1,9% | 2,0% | 2,0% | 1,9% |
| FRA | 20,3% | 20,3% | 20,6% | 19,8% |
| GRC | 2,2% | 2,1% | 2,2% | 2,3% |
| IRL | 1,5% | 1,6% | 1,5% | 1,5% |
| ITA | 17,1% | 16,4% | 16,5% | 16,8% |
| LTU | 0,5% | 0,6% | 0,6% | 0,6% |
| LUX | 0,3% | 0,3% | 0,3% | 0,4% |
| LVA | 0,3% | 0,3% | 0,3% | 0,3% |
| NLD | 5,3% | 5,6% | 5,4% | 5,7% |
| PRT | 2,3% | 2,3% | 2,3% | 2,5% |
| SVK | 0,8% | 0,9% | 0,8% | 0,9% |
| SVN | 0,4% | 0,4% | 0,4% | 0,5% |
| **EUR** | **100,0%** | **100,0%** | **100,0%** | **100,0%** |

Source: Eurostat, HICP - country weights (prc_hicp_cow), downloaded from: https://ec.europa.eu/eurostat/data/database. More precisely the [COWEA] subsample downloaded from: https://ec.europa.eu/eurostat/databrowser/view/PRC_HICP_COW__custom_7158022/default/table?lang=en

## *2. Regions in Brazil*

The weights for Brazilian metropolitan regions in the headline inflation index are derived from the "monetary income available to families". These are currently taken from the 2017-2018 Family Budget Survey[16].

The weights used in this paper are thus the following:

***Table A.I.2 – Weights for Brazilian regions***

| Abbreviation | Region | Weight |
|---|---|---|
| ARA | Aracaju | 1,0% |
| BEL | Belém | 3,9% |
| BHZ | Belo Horizonte | 9,7% |
| BRB | Brasília | 4,1% |
| CGR | Campo Grande | 1,6% |
| CUR | Curitiba | 8,1% |
| FOR | Fortaleza | 3,2% |
| GYN | Goiânia | 4,2% |
| POA | Porto Alegre | 8,6% |
| RBR | Rio Branco | 0,5% |
| REC | Recife | 3,9% |
| RIO | Rio de Janeiro | 9,4% |
| SAL | Salvador | 6,0% |
| SAO | São Paulo | 32,3% |
| SLS | São Luís | 1,6% |
| VIT | Vitória | 1,9% |
| **BRA** | **BRASIL** | **100,0%** |

Source:
https://ftp.ibge.gov.br/Precos_Indices_de_Precos_ao_Consumidor/Sistema_de_Indices_de_Precos_ao_Consumidor/Notas_Tecnicas/snipc_nota_tecnica_2019_02.pdf

[16] Pesquisa de Orçamento Familiar – POF, in Portuguese.

### *3. German Länder*

For the German Länder, the weights are based on the relative shares of private consumption expenditures for the German Länder (Destatis, 2019). The weights used in this paper are, therefore, the following:

***Table A.I.3 – Weights for German Länder***

| Abbreviation | Region | Weight |
|---|---|---|
| BB | Brandenburg | 2,6% |
| BE | Berlin | 3,9% |
| BW | Baden-Württemberg | 14,1% |
| BY | Bayern | 16,8% |
| HB | Bremen | 0,8% |
| HE | Hessen | 7,7% |
| HH | Hamburg | 2,4% |
| MV | Mecklenburg-Vorpommern | 1,6% |
| NI | Niedersachsen | 9,4% |
| NW | Nordrhein-Westfalen | 21,7% |
| RP | Rheinland-Pfalz | 5,1% |
| SH | Schleswig-Holstein | 3,6% |
| SL | Saarland | 1,2% |
| SN | Sachsen | 4,4% |
| ST | Sachsen-Anhalt | 2,4% |
| TH | Thüringen | 2,3% |
| **DE** | **Deutschland** | **100,0%** |

Source:
https://www.destatis.de/DE/Presse/Pressekonferenzen/2019/HGG_VPI/Statement_HGG_VPI_PDF.pdf?__blob=publicationFile

## Appendix II – Trimmed sample for German Länder

For this run of the experiment, I use a slightly trimmed version of the sample for the regions in Germany, which discards the first two months of 2020. I call the sample "Dt" (from Deutschland trimmed). The results are close to the ones obtained using the whole sample.

The trimming of the sample allows for a much more parsimonious smoothing function. The "optimal" degree for the smoothing polynomial in this case is 3, so the SBI algorithm can map the paths using the closed-form analytical solution described in the Appendix of ABLS (2023). With this, all overlapping intervals can be calculated. The results are listed below. The estimated windows are very close to the ones estimated with the original sample. There is almost no change in order, except for the four regions that had no values and are now ordered in between the others. Hamburg (HH) is still the overall leader.

***Table A.II.1 – Overlapping intervals for German Länder – trimmed sample***

| **Region** | **Overlap (months)** |
|---|---|
| TH | -3,01 |
| HB | -2,08 |
| BE | -0,88 |
| MV | -0,80 |
| HE | -0,58 |
| ST | -0,35 |
| SN | -0,28 |
| SL | -0,24 |
| NW | 0,02 |
| RP | 0,15 |
| BY | 0,18 |
| BB | 0,21 |
| NI | 0,30 |
| BW | 0,61 |
| SH | 0,71 |
| HH | 0,79 |

Source: Own calculations.

The mapping between Hamburg (HH) and "Rest-of-Germany" (RoDt) is illustrated in the plot below. The quality of the mapping, due to the analytical solutions, is clear in panel (b).

***Figure A.II.1 – HH and RoDt mapping***

Panel a – Original and smoothed | Panel b - Normalized

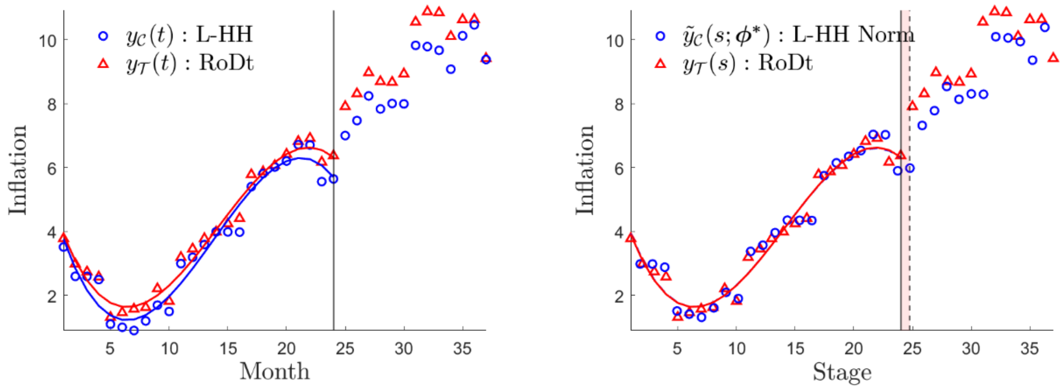


Note: Panel a: circles and triangles depict the observed data and lines represent the smoothed paths with a 3rd-degree polynomial. Panel b: the same, but data and smoothed path for HH (blue) are normalized to the same stage as RoDt according to the SBI mapping function. The black solid line shows the date of the shock (same as the stage for RoDt) and the black dashed line represents the stage when the shock hit HH. The pink shaded area, in between, marks the overlapping identification window.

Zooming in the identification window the plot below illustrates the effect of the shock.

As happened with the original sample, the bootstrapping could not produce plots, but the values are very similar. The median and mean of the simulated effects are both approximately 1.5 percentage points, close to the 1.6 point-estimate. This compares with the 1.2 point-estimate for the original sample. The 90% confidence interval goes from +0.3 to +3.0 percentage points, implying also a positive and statistically significant effect of the war on inflation in Germany.

*Figure A.II.2 – Effect of war on inflation in Germany - trimmed*

Panel a – Normalized data

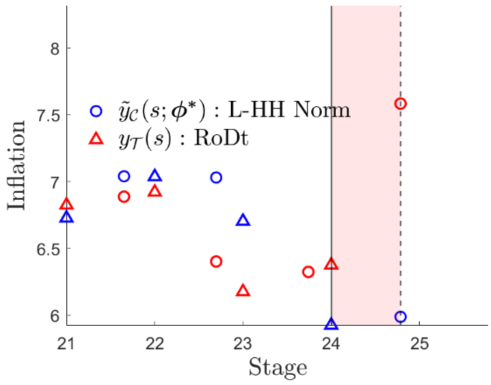


Panel b - Difference

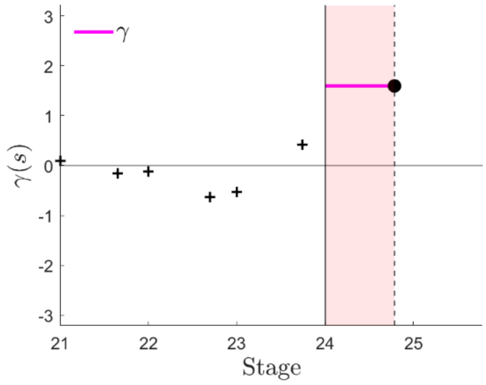


Note: Panel a: circles and triangles depict the observed and normalized data. Panel b: in the y-axis $\gamma(s)$ measures the effect of the shock in percentage points. The crosses represent the differences between data points before the shock. The black points represent the differences inside the identification window (shaded area). The magenta solid line shows the difference in the last period of the window.