BIS Working Papers

No 1159

The 'plucking' model of the unemployment rate floor: Cross-country estimates and empirics

by Jing Lian Suah

Monetary and Economic Department

January 2024

JEL classification: E24, E31, E32, E52

Keywords: Plucking model; Unemployment rate; Non-linear Phillips curve; Threshold effects

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

The ‘Plucking’ Model of the Unemployment Rate Floor: Cross-Country Estimates and Empirics

Jing Lian Suah

Bank for International Settlements

January 2024

Abstract

The unemployment rates (u-rates) of 19 economies (10 advanced, and 9 emerging) demonstrate properties consistent with the plucking model. That the amplitude of expansions and subsequent contractions are unrelated, but that the deeper the contraction, the greater the subsequent expansion. The plucking model, which suggests that the u-rate hovers at or above a theoretical floor, has implications for the unemployment-inflation trade-off as well as shock propagation mechanisms, including the effects of policy shocks. This paper does three things. First, building on existing empirics, it demonstrates a straightforward way to estimate the u-rate floor based on identified peaks in the business cycle and interpolation methods. Second, it analyses the empirical relationship between the u-rate and core inflation, and the effect of a binding u-rate floor on this. Third, it analyses the threshold effects of the u-rate gap on the propagation of macroeconomic shocks, with special attention given to interest rates, using a threshold panel local projections model. The paper finds that: (i) the u-rate hovers at or above the floor and converges towards the floor after each downturn; (ii) the relationship between core inflation and the u-rate weakens when the u-rate is further from the floor; and (iii) the propagation of interest rate, price and output shocks display threshold effects, while exchange rate and debt shocks do not.

JEL Codes: E24, E31, E32, E52

Keywords: Plucking model; Unemployment rate; Non-linear Phillips curve; Threshold effects

† The views expressed in this paper are those of the author and do not necessarily represent those of the Bank for International Settlements (BIS), the Central Bank of Malaysia or the Government of Malaysia. The author would like to thank, without implication, James Yetman for critical review and feedback, as well as Ilhyock Shim, Tirupam Goel, Benoit Mojon and seminar participants at the BIS and Hong Kong Institute for Monetary and Financial Research (HKIMR) for comments. All data used in this study are open access. Replication codes and data vintages are freely available at github.com/suahjl/global-plucking.

1 Introduction

1.1 Motivation

Various studies have demonstrated empirical evidence of the plucking properties underlying Friedman (1993)'s plucking model of the business cycle — (i) that the amplitude of the current downturn is unrelated to that of the preceding expansion, and (ii) the deeper the preceding contraction, the larger the succeeding expansion. In essence, the plucking model posits that output operates at or below some unobserved ceiling, analogous to its potential, and is occasionally 'plucked' down by shocks, before rebounding. In the labour market, this means that the unemployment rate (u-rate) hovers at or above a floor, corresponding to the output ceiling, and is occasionally 'plucked' up by shocks. The business cycle in this view is asymmetric.

In recent literature, Dupraz et al. (2019) demonstrated the plucking properties in US monthly u-rate data, and further proposed a computational, real-time method to identify peaks and troughs that closely match those identified by the National Bureau of Economic Research (NBER), whose methodology is described in Hall et al. (2003). Kohlscheen et al. (2023) analysed the GDP growth rates, and u-rates of 12 advanced economies (AEs), while Hartley (2021) analysed GDP per capita of 169 economies, and Suah (2023) analysed the real GDP growth rates of about 80 economies. All of these studies identified properties consistent with the plucking model in the economies studied.

This paper examines a mixed sample of 19 economies (10 advanced, and 9 emerging).[1] Figure 1 shows that the plucking properties are present in the u-rates for these economies. Specifically, the greater the amplitude of a contraction (peak-to-trough change in the u-rate), the larger the decline in u-rate in the subsequent expansion (trough-to-peak change in the u-rate) (subfigure (a)); and the peak-to-trough change in the u-rate during a contraction is unrelated to the trough-to-peak change in the u-rate during the preceding expansion (subfigure (b)).[2] These empirical regularities demonstrate the plucking properties in a cross-country setting covering a range of advanced and emerging economies, and motivate the

[1] Australia, Brazil, Chile, China, France, Germany, Hong Kong SAR, India, Indonesia, Italy, Japan, Malaysia, Mexico, Singapore, Korea, Thailand, the United Kingdom, the Philippines, and the United States of America; these economies were selected based on the following criteria — (i) coverage of major advanced, and emerging economies (EEs); (ii) key regional economies from the Asia-Pacific region (including ASEAN); and (iii) sufficient u-rate data at quarterly frequency

[2] This set of relationships hold even when the average rate of change in the u-rate is used in lieu of amplitude in figure A.1, indicating that base effects throughout business cycle phases play little role. The faster the u-rate rises during a contraction, the faster the u-rate falls in the expansion that follows (subfigure (a)), and how fast the u-rate rises during a contraction is largely unrelated to how fast the u-rate falls in the preceding expansion (subfigure (b)). In the latter, the slight downward slope is driven by a select few outliers.

application of the plucking model to the u-rate.

Figure 1

(a) Amplitude of the u-rate during the subsequent expansion against the current contraction

(b) Amplitude of the u-rate during the subsequent contraction against the current expansion

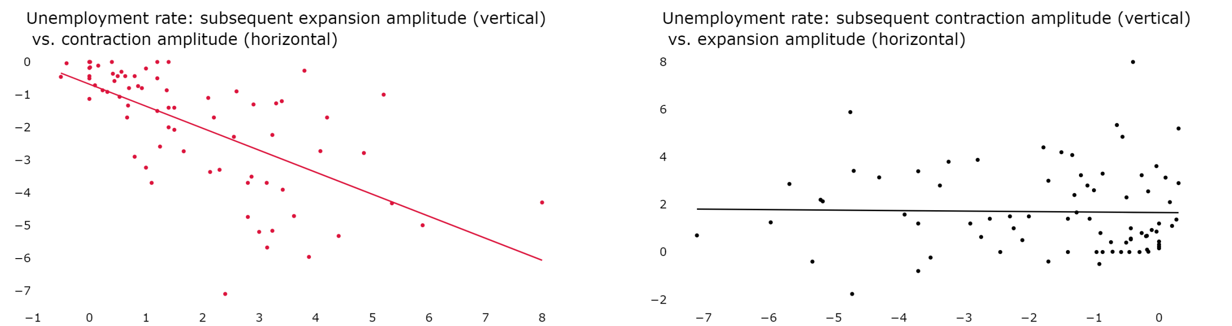

Subplot (a) shows a scatterplot of the change in the u-rate during a contraction (horizontal axis) against that of during the subsequent expansion (vertical axis). Subplot (b) shows the reverse: the change in the u-rate during an expansion (horizontal axis) against that of during the subsequent contraction (vertical axis). In both subfigures, each observation represents unique expansion-contraction (or contraction-expansion). A version without the outliers is available in the appendix in figure A.2. Another version using quarterly changes in the unemployment rate is also available in the appendix in figure A.1.

This empirical evidence in favour of the plucking model leads to further research questions. First, can the model a useful guide to policy, akin to models based on other latent variables (natural interest rates, potential output, structural unemployment rates)? Second, what are the implications of the plucking model for macroeconomic stabilisation policy?

1.2 Literature review

Regarding the first question, there is an existing literature that attempts to estimate the so-called u-rate floor (or equivalently output ceiling) based on the plucking model. Kim & Nelson (1999) modelled the asymmetry in the plucking model described by Friedman (1993) as a Markov-switching process, estimated with a Kalman Filter on US GDP and u-rate data. However, when ported to a broader set of economies, De Simone & Clarke (2007) demonstrated that Kim & Nelson (1999)'s approach often generated estimates that trace actual output, even following major shocks. Rather, output ought to be plucked down far below its ceiling in these instances, rather than both output and the ceiling moving down together. More recently, Fontanari et al. (2020) applied a revised method from Okun (1962) to US data, and produced a highly asymmetric estimate of the output gap, which is not unlike the implications of the plucking model. Similarly, Suah (2023) applied a peak-trough identification algorithm modified from Dupraz et al. (2019) to fit real GDP, labour force and capital, supplemented with interpolation techniques and an augmented production function, to Malaysian data to estimate an output ceiling and hence the output gap. The resultant output gap is also asymmetric, and stable during downturns. A commonality between these asymmetric output gaps, as noted in Fontanari et al. (2020), is their independence from inflation during the estimation process (ex ante). In contrast, building a measure of potential output based on the assumption that there is a non-inflation accelerating rate of resource allocation tends to generate relative symmetry of the output gap across the business cycle. The same applies to the estimates of the NAIRU (non-inflation accelerating rate of unemployment), as in Cusbert (2017) and Ruberl et al. (2021). Suah (2023) showed that the plucking model estimate of the output gap fitted inflation data for Malaysia well, despite no ex ante assumption of such relationships. Suah (2023) further highlighted some policy pitfalls of symmetric measures, whose reliance on time series filters often result in estimates prone to large retrospective revisions, adding to the aforementioned issues observed in De Simone & Clarke (2007), and hence offer little reliability for policy decisions that must be made in real time. Galbraith (1997) had also argued against using the NAIRU in policy, including its lack of empirical support. Estimating the u-rate floor based on the plucking model is feasible but documented attempts are scarce. There are two guiding objectives in this estimation. First, the methodology needs to generate stable estimates for policy usability. Second, the u-rate floor is to be estimated without reference to inflation (ex ante), leaving the ex post relationship with inflation as a means of examining the model.

Regarding the second question, the typical purpose of estimates of economic slack (whether output gaps or u-rate gaps) in policy questions is primarily for informing strategies to smooth economic fluctuations in response to shocks. For central banks, a large part of this is to manage unemployment-inflation trade-offs. The plucking view, taken at face value, suggests that policy stands to minimise social welfare losses asymmetrically by attempting to bring output quickly, and as close as possible, back to the ceiling (or the u-rate back to the floor), as discussed in Dupraz et al. (2019). Note that this implication for policy depends on the idea that inflation and economic slack are insensitive to each other, so that there is no unemployment-inflation trade-off, at least when the economy experiences slack.

On this, however, the existing empirical evidence across countries is mixed. The literature tends to agree on the existence of an inflation-unemployment trade-off, but disagree on its strength. Hazell et al. (2022) estimated that the Phillips curve slope was small (near-zero), and had remained small over time, using US state-level data. Wellmann (2023) further replicated this finding for the eurozone using country-level data in euro area. In contrast, Furuoka et al. (2007) (error correction model estimates with national-level data), Furuoka & Harvey (2015) (empirical estimates of a Neo-Keynesian Phillips curve), and Tang & Lean (2007) (error correction model estimates) demonstrated relatively a strong Phillips curve relationship in Malaysia.[3] This is related to research on sacrifice ratios, such as in Ball (1994), Cecchetti & Rich (2001), Daniels et al. (2005), and Bowdler (2009). These papers suggest that the inflation cost of output stabilisation policy evolves over time, and varies across countries (due to, for instance, capital account openness and central bank independence). In the context of economic slack, when the u-rate is at the floor, hence the supply constraint binds, does the slope of the observed Phillips curve steepen, raising the inflation cost of output stabilisation for policymakers? If so, then the degree of policy aggressiveness changes according to the position of the u-rate relative to the floor. Addressing this question requires two steps: first, find out if there is such a non-linearity in the unemployment-inflation link that is tied to the u-rate gap; second, subsequently find out if macroeconomic propagation is different when the u-rate floor binds.

1.3 Outline

To answer the aforementioned questions, this paper is structured as follows. Section 2 provides an overview of the data. Section 3 estimates the u-rate floor based on the methodology in Dupraz et al. (2019) and Suah (2023) on a panel of 10 advanced, and 9 emerging economies. Section 4 provides descriptive analysis of the relationship between inflation and the u-rate when the u-rate floor binds. Section 5 then analyses the threshold effects of the closing of the plucking u-rate gap on the propagation of macroeconomic shocks, including those due to monetary policy. Section 6 concludes.

2 Data

This paper uses macroeconomic time series for 19 economies (10 advanced, and 9 emerging), sourced from respective national authorities via CEIC data (a proprietary data compiler), as listed in table 1.[4] Special attention is given to the Asia-Pacific region.

[3]For context, Malaysia is an upper-middle income emerging economy with an export-oriented mixed sector structure.

[4]In various parts of the analyses, the economies may be further segregated into major advanced economies (all advanced economies excluding Singapore, Hong Kong SAR, and Korea), major emerging economies, newly industrialised Asian economies (NIEs; Singapore, Hong Kong SAR, and Korea), and the ASEAN-4 emerging economies.

Table 1: List of economies included

Advanced Economies (AEs)	Emerging Economies (EEs)
United States	Mexico
Germany	Chile
France	Brazil[5]
Italy	China[1]
United Kingdom	India[3,5]
Japan	Malaysia[5]
Australia	Thailand
Singapore	Indonesia[3, 6]
Korea	Philippines[3,4]
Hong Kong SAR[2]	

[1] Excluded from panel estimation of Phillips curve equations and threshold panel local projections analysis, owing to its unique economic structure relative to the rest of the world.
[2] Excluded from panel estimation of Phillips curve equations and threshold panel local projections analysis due to the absence of published core CPI data.
[3] The u-rate is seasonally adjusted by the author using X13 (as only non-seasonally adjusted data is available).
[4] Excluded from threshold panel local projections analysis due to unavailability of comparable non-financial sector credit (private sector debt) statistics.
[5] Sensitivity analyses excluding Brazil, India, and Malaysia are available.
[6] Excluded from panel estimation of Phillips curve equations, and threshold panel local projections analysis, as data are published half-yearly.

Table 2 shows the variables included in the analysis. Section 3 deploys only the de-seasonalised u-rates to estimate the u-rate floors. Section 4 uses the estimated u-rate floors from section 3, along with actual u-rate, core inflation, expected inflation and the real effective exchange rate (REER) to analyse the unemployment-inflation link. Section 5, which analyses propagation of shocks with a threshold panel local projections model, uses the estimated u-rate gaps, real GDP growth, core inflation, expected inflation, credit to the private non-financial sector, the REER and the short-term interest rate.

Table 2: List of variables included

No.	Variable
1	Unemployment rate
2	Core CPI inflation
3	12-month ahead expected inflation (computed as the rolling weighted 12-month average of the median forecasts from Consensus Economics)
4	Real GDP growth
5	Real effective exchange rate (REER)
6	Credit to the private non-financial sector
7	Short-term interest rate

3 Estimating the unemployment rate floors

The estimation framework is based on the approach in Suah (2023), which uses the peak-trough identification algorithm in Dupraz et al. (2019). A major benefit of this algorithm is its computational simplicity, and the ability to call peaks and troughs in a timely manner, avoiding the long lags, for example, from the NBER approach in Hall et al. (2003). The estimation framework consists of the following steps.

First, we identify the peaks and troughs using Dupraz et al. (2019)'s peak-trough identification algorithm with the tolerance threshold X tuned to match major cycles of the respective economies. In this application, X is modified to be a multiplier of the standard deviation of the first (quarterly) difference of the u-rate of the corresponding economy for computational tractability across economies.[5] Second, the peaks are interpolated to estimate the u-rate floor using a quadratic spline.[6] Third, the end points are extrapolated using the average rate of change in the estimated u-rate floor in the neighbouring peak-to-peak phase. Finally, we adjust the estimates of the floor where necessary to ensure that they are non-negative and no higher than the observed u-rate. Table 3 provides the choices of X and the implied percentage points in the u-rate by economy. In general, economies whose u-rates tend to fluctuate with greater magnitude tend to have higher thresholds to avoid misclassifying 'jitters' as turning points.

[5] The original application uses a straightforward percentage points in the u-rate, which is reasonable for a single-economy application. Suah (2023) demonstrated that the tolerance threshold might be economy-specific, through a comparison between Malaysian and US data.

[6] In the event of insufficient nodes or observations, a linear spline interpolation is used instead. If the number of nodes or observations are still insufficient, then linear interpolation is used.

Table 3: Tolerance thresholds (multiplier of the standard deviation of the u-rate)

Economy	X	**Percentage points**
United States	1	1.7
Germany	0.3	0.68
France	0.5	0.88
Italy	0.4	0.65
United Kingdom	0.3	0.72
Japan	0.25	0.3
Australia	1.2	2.15
Singapore	2	1.41
Korea	0.8	0.41
Hong Kong SAR	0.4	0.66
Mexico	1	0.83
Chile	2	3.33
Brazil	1.4	3.65
China	1	0.27
India	0.8	0.92
Malaysia	0.35	0.19
Thailand	0.18	0.12
Indonesia	0.7	0.76
Philippines	1.5	3.42

Figures 2 to 5 show the quarterly u-rates and the estimated u-rate floor in different groups of economies. In most major AEs (figure 2), the u-rate has recovered in the aftermath of the COVID-19 pandemic to be close to the floor (with the exception of Italy). In major EEs (figure 3), the u-rates remain above the floor, except in India and Mexico, but the u-rate gap has closed substantially since COVID-19 everywhere. In the Asian AEs excluding Japan (figure 4), the u-rate gap closed towards the end of 2022 (earlier in the case of Korea), likely reflecting active labour market policies during the COVID-19 crisis. Moreover, prior to the crisis, unlike some major AEs, the Asian NIEs entered the COVID-19 crisis with their u-rates at their respective floors. Finally, the ASEAN-4 (figure 5) were similar to the Asian NIEs, with u-rates at the floor both prior to and by the end of the COVID-19 pandemic.

Figure 2

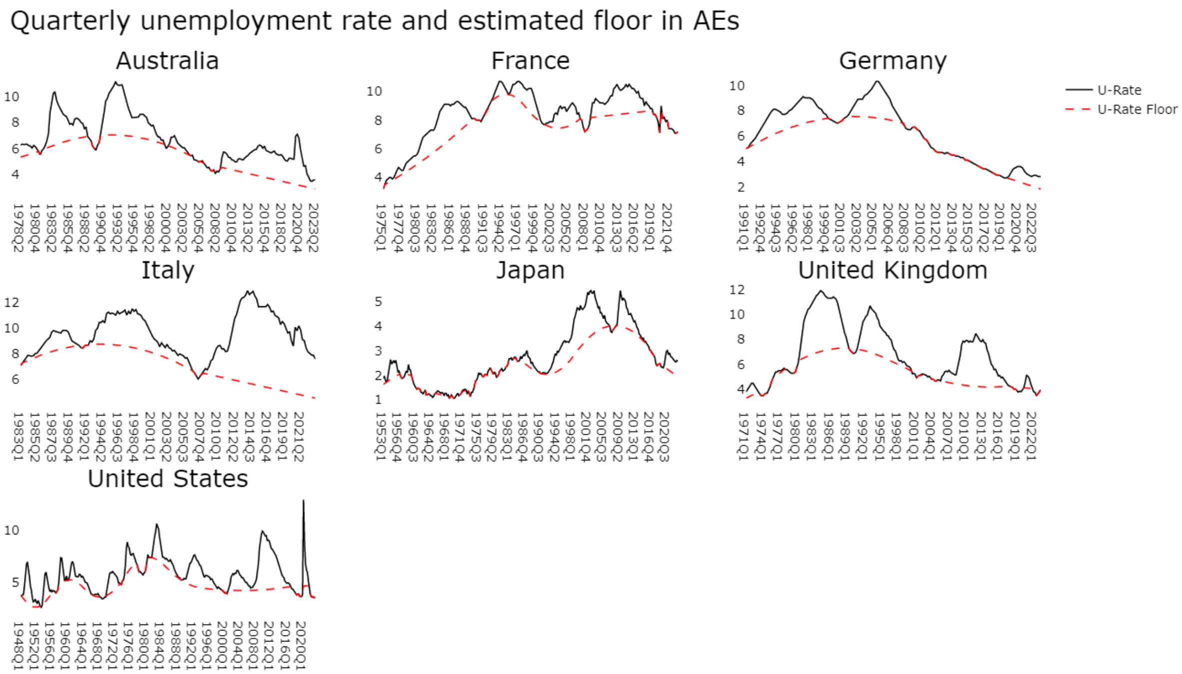

Figure 3

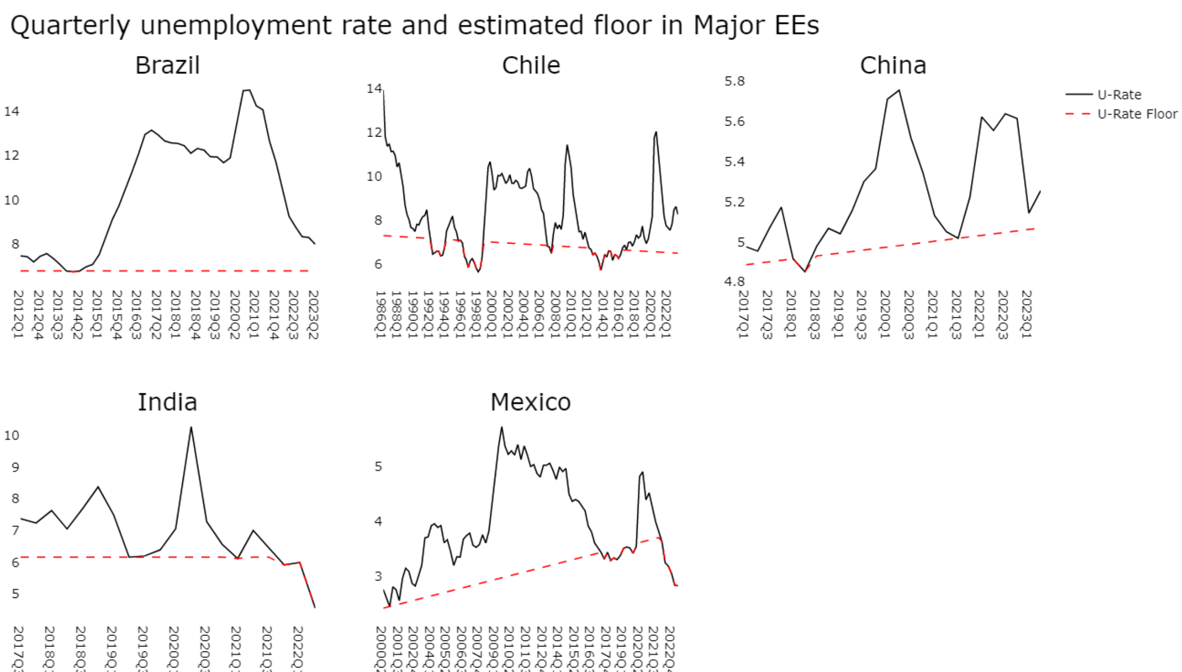

Figure 4

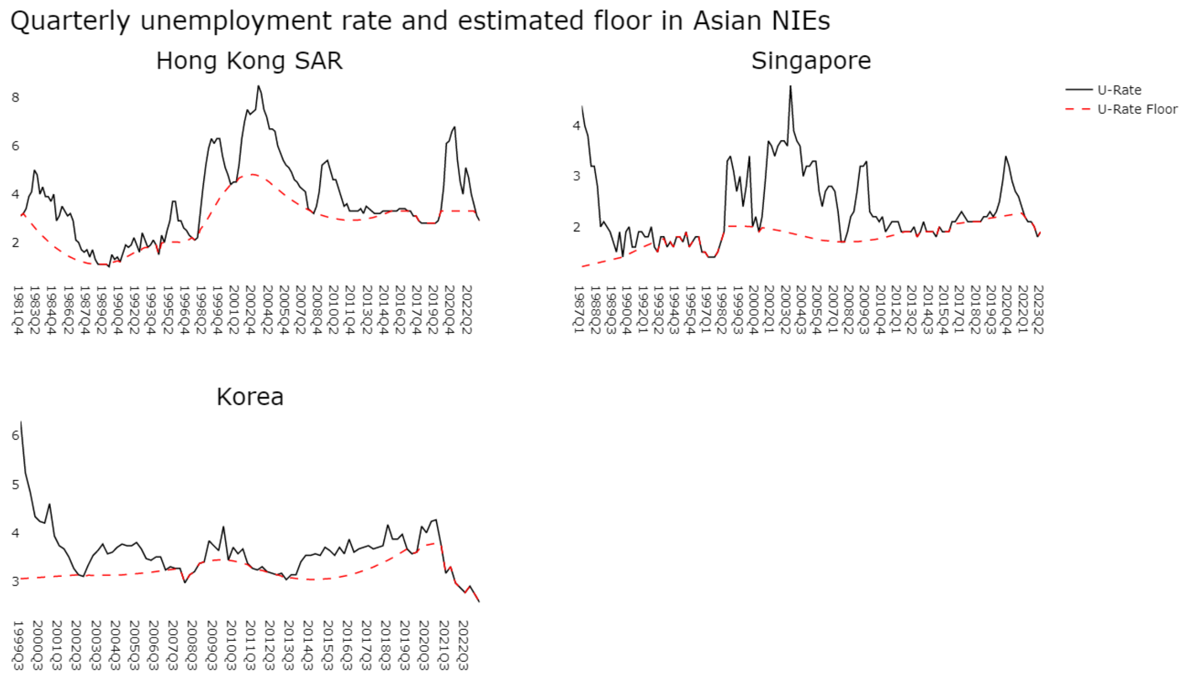

Figure 5

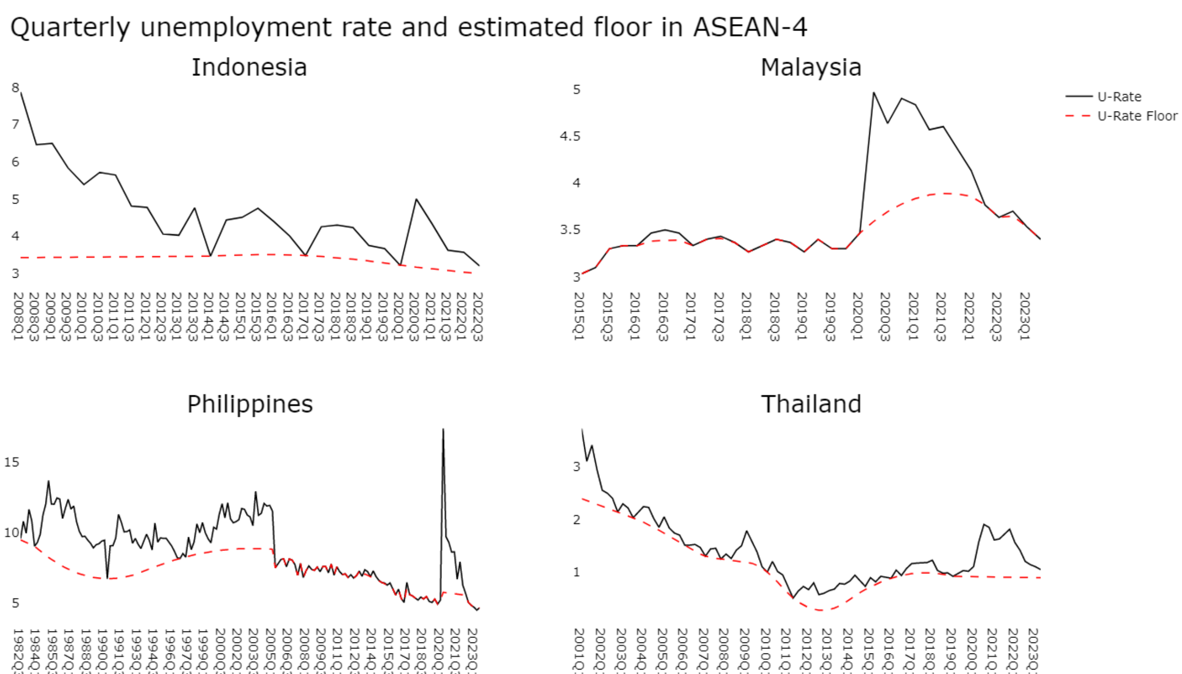

3.1 Comparison of the NAIRU and the plucking u-rate floor models using US data

Putting the presence of the plucking properties shown in section 1 aside, the u-rate floor in the plucking view and the non-inflation-accelerating rate of unemployment (NAIRU) are not necessarily antithetical. The former describes a theoretical maximum given resource constraints while the latter describes an allocative equilibrium that, given underlying frictions, generates a stable inflation rate. However, they each have their own testable implications. Here, we compare them based on US data, given the ready availability of NAIRU estimates.

There are a plethora of methodologies in estimating the NAIRU, not unlike potential output. Section 1 earlier discussed Cusbert (2017) and Ruberl et al. (2021), which are based on the Phillips curve framework. Figure 6 plots the actual u-rate, estimated u-rate floor and the Congressional Budget Office (CBO)'s published estimate of the natural rate of unemployment.[7] It shows that, with the exceptions of the late 1970s to 1980s and around 2020, the the NAIRU remained above the u-rate floor. Figure 7 then tabulates observations between 1949 Q1 and 2023 Q2 where (i) the u-rate was at or below the NAIRU versus where (ii) the u-rate was at the estimated u-rate floor. Comparing the two illustrates asymmetry in overheating and underheating episodes, with 20.2% of observations at the u-rate floor, compared to 49.3% falling below the NAIRU.[8] Thus, while both concepts are not antithetical in a vacuum, the plucking properties shown in figures 1 and A.1 appear to reject the symmetry implied by the NAIRU estimates.

[7]Extracted from the Federal Reserve Economic Data (FRED) portal.

[8]As the u-rate floor cannot be above the observed u-rate, the conceptual equivalent for economic tightness in the plucking model is when the u-rate floor binds.

Figure 6

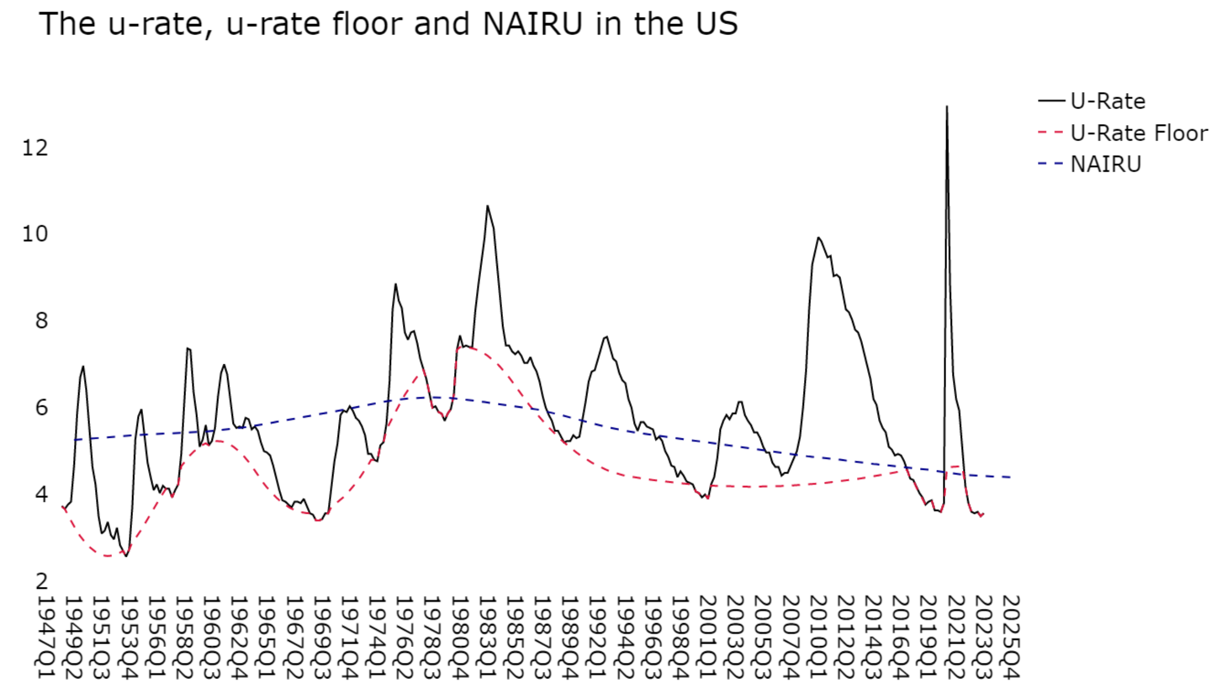

Figure 7

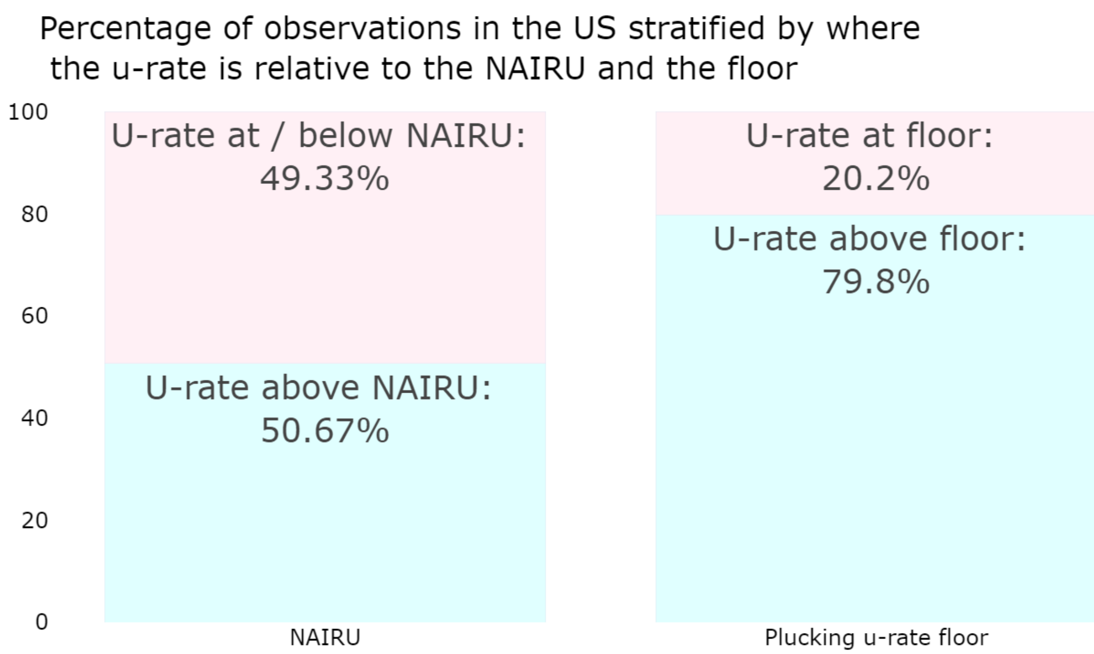

Further, our estimates of the u-rate gaps relative to the floor are visually similar to the output gaps estimated by Fontanari et al. (2020) with a revised approach based on Okun (1962), in that the economy experiences slack most of the time. This partly reflects two key methodological similarities between the approaches. First, estimation abstracts from inflation; hence neither the estimated potential nor the floor are constrained by inflationary behaviour. Second, neither uses time series filters as part of the estimation process. By contrast, in previous approaches by De Simone & Clarke (2007), which uses a state-space Markov-switching model, estimates of the u-rate floor (and output ceiling) often traced the actual u-rate (output level), even after large shocks. However, if the u-rate floor (or analagously the natural u-rate, or potential output) is a structural representation of efficient resource allocation, then large or sudden fluctuations are likely to be uncommon. As demonstrated by Suah (2023), they are partly an artefact of applying two-sided symmetric time series filters, which result in large retrospective revisions once turning points are included in the input data.[9] By contrast, our approach is relatively robust to the inclusion of large shocks.[10] Instead, the estimated u-rate floor is a slow-moving asymmetric trend that the u-rate converges to in the long run, with economic shocks plucking the u-rate above the floor.

3.2 Supply shocks, labour supply and the u-rate floor in US data

If the u-rate floor in the plucking view is to represent the level of unemployment corresponding to a theoretical maximum of economic output conditional on resource constraints, then it is a measure of the production frontier. Conceptually, negative supply shocks should raise the u-rate floor. A general supply shock, such as an oil price shock, should thus be accompanied by an increase in the floor, and an increase in labour supply should be accompanied by a decline in the floor. The US, following on from the previous subsection, offers a good case study to examine this given the availability of data back dating back to 1949, with NAIRU estimates published by the CBO. We measure oil price shocks using the "net oil increase" variable in Hamilton (1996), later referenced as the "max-oil" variable when adapted to uncertainty shocks in Jackson et al. (2019). The max-oil variable takes the maximum of zero and the growth in oil prices relative to the highest point over the previous year (4 quarters), shown in equation 1. Labour supply is proxied by the labour force participation rate (LFPR).

[9] Fernald (2015) also noted this issue.

[10] See, for example, figures B.1 to B.4 in the appendix, which illustrate the effect of restricting input to the pre-COVID period, resulting in limited retrospective revision in the u-rate floor estimates of most economies. Calibration parameters for the pre-COVID vintages are shown in table B.1.

$$\widehat{P_{it}^{oil}} = \max\{0, \frac{P_{it}^{oil} - \max\{P_{i,t-1}^{oil},\ P_{i,t-2}^{oil},\ P_{i,t-3}^{oil}, P_{i,t-4}^{oil}\}}{\max\{P_{i,t-1}^{oil},\ P_{i,t-2}^{oil},\ P_{i,t-3}^{oil}, P_{i,t-4}^{oil}}\} \tag{1}$$

Figure 8 plots the u-rate floor, the CBO's estimate of the NAIRU, year-on-year wage inflation and the max-oil variable for the US since 1947. Figure 9 overlays the same variables against the LFPR instead of max-oil. We can make several observations. First, despite a lack of oil price shocks in the 1950s, the u-rate floor rose against a backdrop of declines in the LFPR. Second, the increase in the u-rate floor and wage inflation in the 1970s were preceded by the two largest oil price shocks since 1947. Third, in the 2000s, a combination of smaller oil price shocks and a sustained decline in labour supply corresponded to a turnaround in the u-rate floor. Finally, during the COVID-19 pandemic, the LFPR declined steeply and rapidly converged towards pre-pandemic rates, accompanied by an increase in the u-rate floor (before a subsequent decline in the u-rate floor). Examination of US macro data visually appears to corroborate supply shocks and related factors as key determinants of the u-rate floor. We leave further investigation of this to future work.

Figure 8

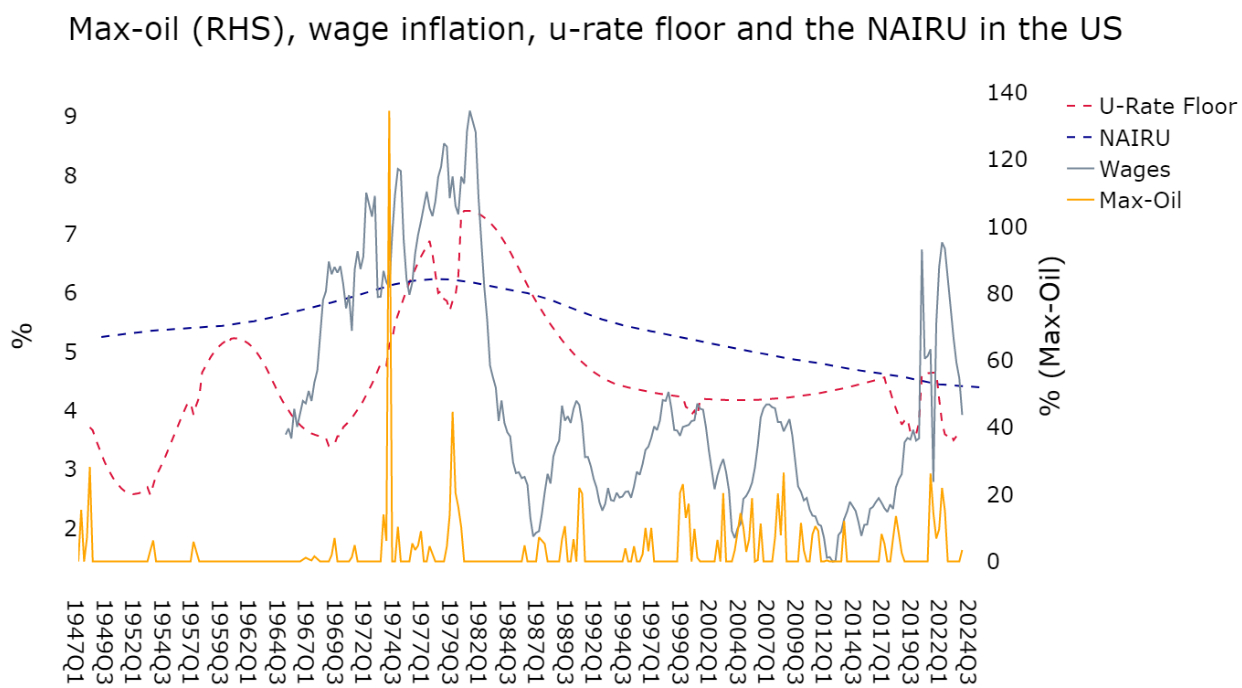

Figure 9

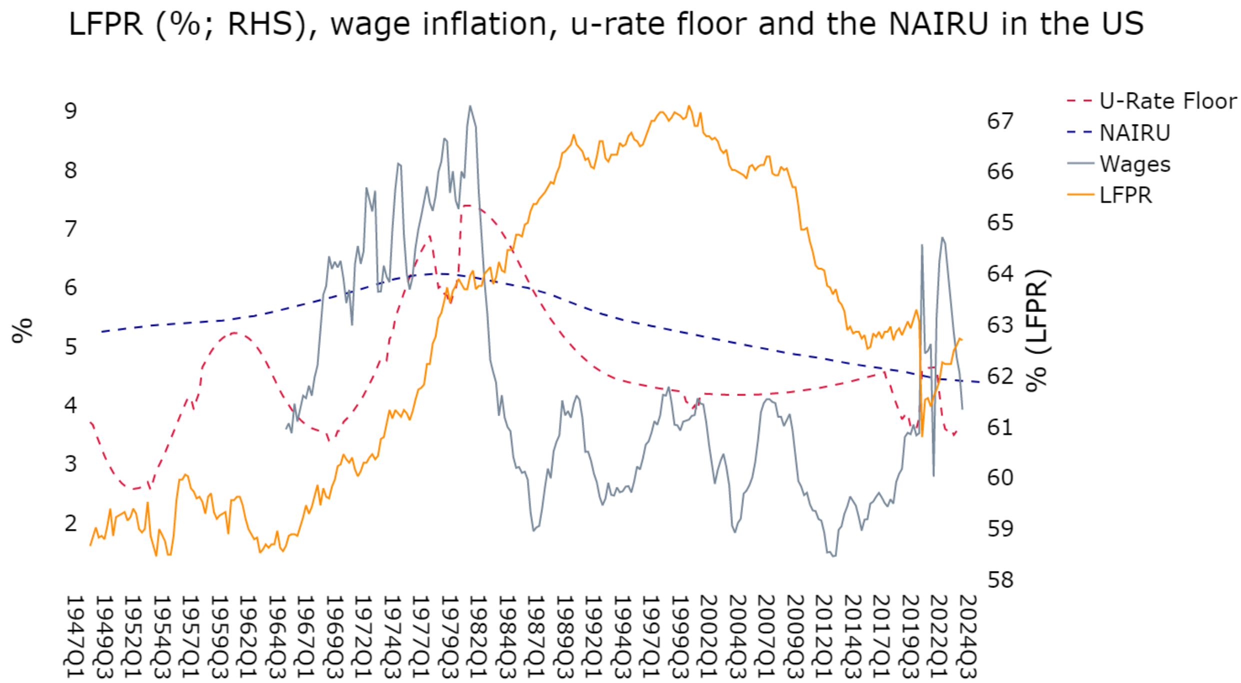

4 The empirical Phillips curve

Having established the properties of the plucking model and its ability to match key business cycle characteristics, a natural next step is to examine how it changes policy calculus. If economic slack is an asymmetric concept, what does that mean for the underlying parameters of the economy? To answer this, we first analyse the observed relationship between inflation and the u-rate, which encapsulates the perennial unemployment-inflation trade-off problem for economic policymakers, especially central banks. When the u-rate is at the floor, does it change the slope of the observed Phillips curve systematically? Does the size of the gap matter as well?

4.1 Methodology

This paper estimates fixed effects (FE) linear regressions, specified to resemble backward and forward-looking Phillips curves. The baseline models in equations 2 to 5 exclude u-rate gaps from the analysis. Core inflation π_{it}, defined as the year-on-year (YoY) percentage change in core CPI, is the dependent variable for economy i in period t. In equation 2, the set of explanatory variables include the YoY change in the u-rate u_{it}, the weighted average of the 12-month-ahead inflation expectations by professional forecasters $\mathbb{E}_t \pi_{t,t+4_{it}}$, and lagged core inflation π_{t-1}, while α_i refers to the fixed effect for economy i. Equation 3 additionally includes the YoY growth in the REER z_{it} to reflect exchange rate effects on the domestic prices, especially for the small open economies in the sample. Equations 4 and 5 are variations of equations 2 and 3 that further include the US u-rate as an explanatory variable to proxy for global economic developments.[11]

$$\pi_{it} = \alpha_i + \beta_1 u_{it} + \gamma \mathbb{E}_t \pi_{t,t+4_{it}} + \lambda \pi_{t-1} + \varepsilon_{it} \tag{2}$$

$$\pi_{it} = \alpha_i + \beta_1 u_{it} + \gamma \mathbb{E}_t \pi_{t,t+4_{it}} + \lambda \pi_{t-1} + \delta z_{it} + \varepsilon_{it} \tag{3}$$

$$\pi_{it} = \alpha_i + \beta_1 u_{it} + \gamma \mathbb{E}_t \pi_{t,t+4_{it}} + \lambda \pi_{t-1} + \mu_1 u_{it}^{USA} + \varepsilon_{it} \tag{4}$$

$$\pi_{it} = \alpha_i + \beta_1 u_{it} + \gamma \mathbb{E}_t \pi_{t,t+4_{it}} + \lambda \pi_{t-1} + \delta z_{it} + \mu_1 u_{it}^{USA} + \varepsilon_{it} \tag{5}$$

Equations 6 to 9 then expand the analysis to include the estimated u-rate gaps from section 3, defined as the distance between the actual u-rate and the estimated u-rate floor. In equations 6 and 7, the u-rate gap u_{it}^{gap} and its interaction with the YoY change in the u-rate $u_{it} * u_{it}^{gap}$ are additionally included as explanatory variables.[12] Equations 8 and 9 also include the corresponding variables for the US to account for global economic developments. Similar to the recent analysis by Hazell et al. (2022), and Wellmann (2023), fixed effects account for economy-specific, time-invariant variation, for example due to inflation expectations. The expected inflation variable then explicitly controls for the time-varying component of expectations within respective economies. Lagged inflation is included to accommodate backward-looking price-setting, and persistence in inflation arising from contractual commitments from previous periods. For completeness and comparison against naive estimation, parameters in equations 6 and 7 are also estimated with pooled ordinary least squares (POLS) with the economy fixed effects α_i replaced by a constant α.

$$\pi_{it} = \alpha_i + \beta_1 u_{it} + \beta_2 u_{it}^{gap} + \beta_3 u_{it} * u_{it}^{gap} + \gamma \mathbb{E}_t \pi_{t,t+4_{it}} + \lambda \pi_{t-1} + \varepsilon_{it} \tag{6}$$

$$\pi_{it} = \alpha_i + \beta_1 u_{it} + \beta_2 u_{it}^{gap} + \beta_3 u_{it} * u_{it}^{gap} + \gamma \mathbb{E}_t \pi_{t,t+4_{it}} + \lambda \pi_{t-1} + \delta z_{it} + \varepsilon_{it} \tag{7}$$

[11] For equations 4 and 5, the US is excluded from the panel.

[12] Results from additional analyses including squared terms of the u-rate gap and their interaction terms are available. All reported results in the main text are robust to these alternative parametrisations.

$$\begin{aligned}\pi_{it} = \alpha_i + \beta_1 u_{it} + \beta_2 u_{it}^{gap} + \beta_3 u_{it} * u_{it}^{gap} \\ +\gamma \mathbb{E}_t \pi_{t,t+4_{it}} + \lambda \pi_{t-1} \\ +\mu_1 u_{it}^{USA} + \mu_2 u_{it}^{gap,USA} + \mu_3 u_{it}^{USA} * u_{it}^{gap,USA} + \varepsilon_{it}\end{aligned} \tag{8}$$

$$\begin{aligned}\pi_{it} = \alpha_i + \beta_1 u_{it} + \beta_2 u_{it}^{gap} + \beta_3 u_{it} * u_{it}^{gap} \\ +\gamma \mathbb{E}_t \pi_{t,t+4_{it}} + \lambda \pi_{t-1} + \delta z_{it} \\ +\mu_1 u_{it}^{USA} + \mu_2 u_{it}^{gap,USA} + \mu_3 u_{it}^{USA} * u_{it}^{gap,USA} + \varepsilon_{it}\end{aligned} \tag{9}$$

We are interested primarily in the estimates of β_1 and β_3 in equations 6 to 9, which are the coefficients on the observed u-rate, and the interaction between the u-rate and the u-rate gap, respectively. The estimated slope ω_{it} of the Phillips curve for economy i in period t is a linear function of the u-rate gap, as indicated in equation 10. From recent literature on the Phillips curve, notably Hazell et al. (2022), and Wellmann (2023), we expect the slope β_1 to have a small negative value. That is, when the u-rate accelerates, core inflation decelerates. However, as we include a variety of EEs and AEs, we anticipate some economy-specific heterogeneity.[13] β_3 is of particular interest. Kumar & Orrenius (2016) found that, with US state-level data, wage pressures were larger when a decline in the u-rate placed it below rather than above the historical average. The analysis in this paper develops this possibility further to see if there is a kink in the Phillips curve at the u-rate floor. Is the Phillips curve flatter when the u-rate floor does not bind or, equivalently, is the slope steeper when the floor binds? In this case, we expect $\beta_3 > 0$. This possibility makes conceptual sense if we view the u-rate floor as a theoretical production frontier, where supply constraints bind, so that demand pressures quickly translate into relative price pressures and hence higher inflation.

$$\omega_{it} = \beta_1 + \beta_3 u_{it}^{gap} \tag{10}$$

More generally, is there is point beyond which the inflation-unemployment relationship changes? For example, changes in the slope of the Phillips curve in response to the u-rate gap could be continuous rather than discrete, given that the u-rate gap from the plucking view can only take on non-negative continuous values. A threshold analysis is used to provide a sensitivity check on whether non-linearity in the inflation-unemployment link is present. For this purpose, a model similar to equation 6 and 7 is estimated but with a dummy for when the u-rate gap is below or at a particular threshold $\mathbb{1}\{u_{it}^{gap} \leq \tau\}$, in lieu of the u-rate gap

[13] For example, Furuoka et al. (2007), Furuoka & Harvey (2015), and Tang & Lean (2007) previously demonstrated the presence of a non-negligible Phillips curve slope in Malaysia, although with different methodologies, primarily error correction models (ECMs).

directly. The threshold τ is chosen based on grid-search of all possible values of u-rate gap u_{it}^{gap}. Hence, the approach here searches for τ that best fits the data based on the models, shown in equations 11 (without REER) and 12 (with REER).

$$\pi_{it} = \alpha_i + \beta_1 u_{it} + \beta_2 \mathbb{1}\{u_{it}^{gap} \leq \tau\} + \beta_3 u_{it} * \mathbb{1}\{u_{it}^{gap} \leq \tau\} + \gamma \mathbb{E}_t \pi_{t,t+4_{it}} + \lambda \pi_{t-1} + \varepsilon_{it} \quad (11)$$

$$\pi_{it} = \alpha_i + \beta_1 u_{it} + \beta_2 \mathbb{1}\{u_{it}^{gap} \leq \tau\} + \beta_3 u_{it} * \mathbb{1}\{u_{it}^{gap} \leq \tau\} + \gamma \mathbb{E}_t \pi_{t,t+4_{it}} + \lambda \pi_{t-1} + \delta z_{it} + \varepsilon_{it} \quad (12)$$

4.2 Findings

This subsection will be divided as follows — model estimates including US as part of the panel, model estimates with US variables as controls, the evolution of the inflation-unemployment link, the threshold model estimates and policy discussion.

4.2.1 Model estimates with the US as part of the panel

Subfigures (a) and (b) of figure 10 show the FE estimates of the baseline models without u-rate gaps from equations 2 to 3, respectively. Subfigures (c) and (d) show the POLS estimates of the corresponding equations. As expected, estimates of β_1 are precise, small, and negative. This provides a baseline to operate on when proceeding to the subsequent set of analyses that include the u-rate gap. A sensitivity check, excluding Brazil, India and Malaysia on account of lack of data prior to the GFC, shows quantitatively similar estimates in figure 11.

Figure 10

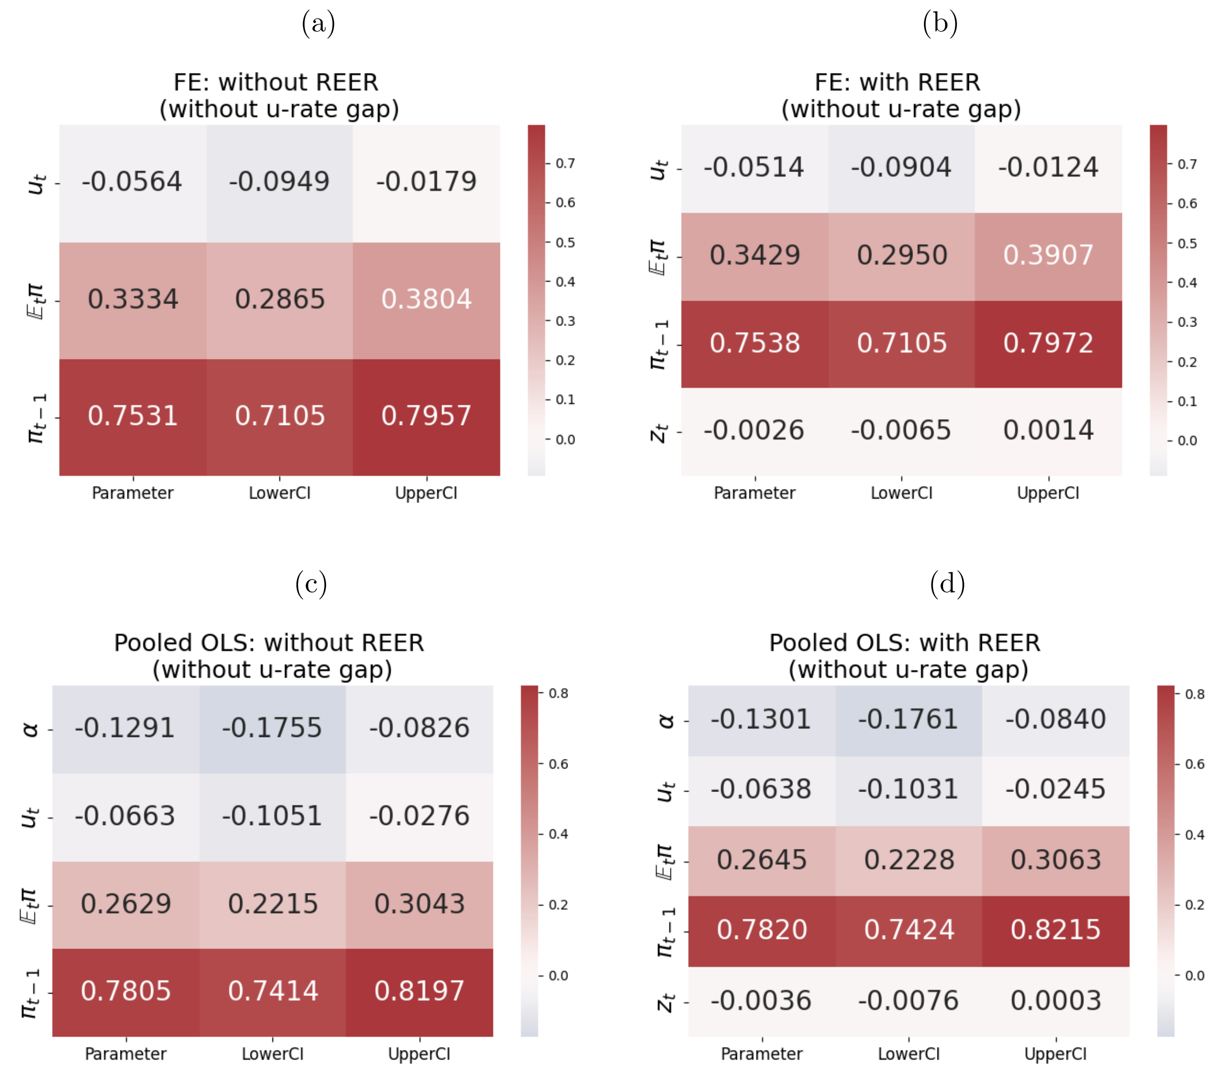

These are coefficient estimates from a fixed effects regression of core inflation on economy fixed effects, the year-on-year (YoY) change in the u-rate, 12-month-ahead expected inflation and the lag of core inflation. The confidence bands refer to 95% confidence intervals, and are computed with heteroscedasticity-robust standard errors.

Figure 11

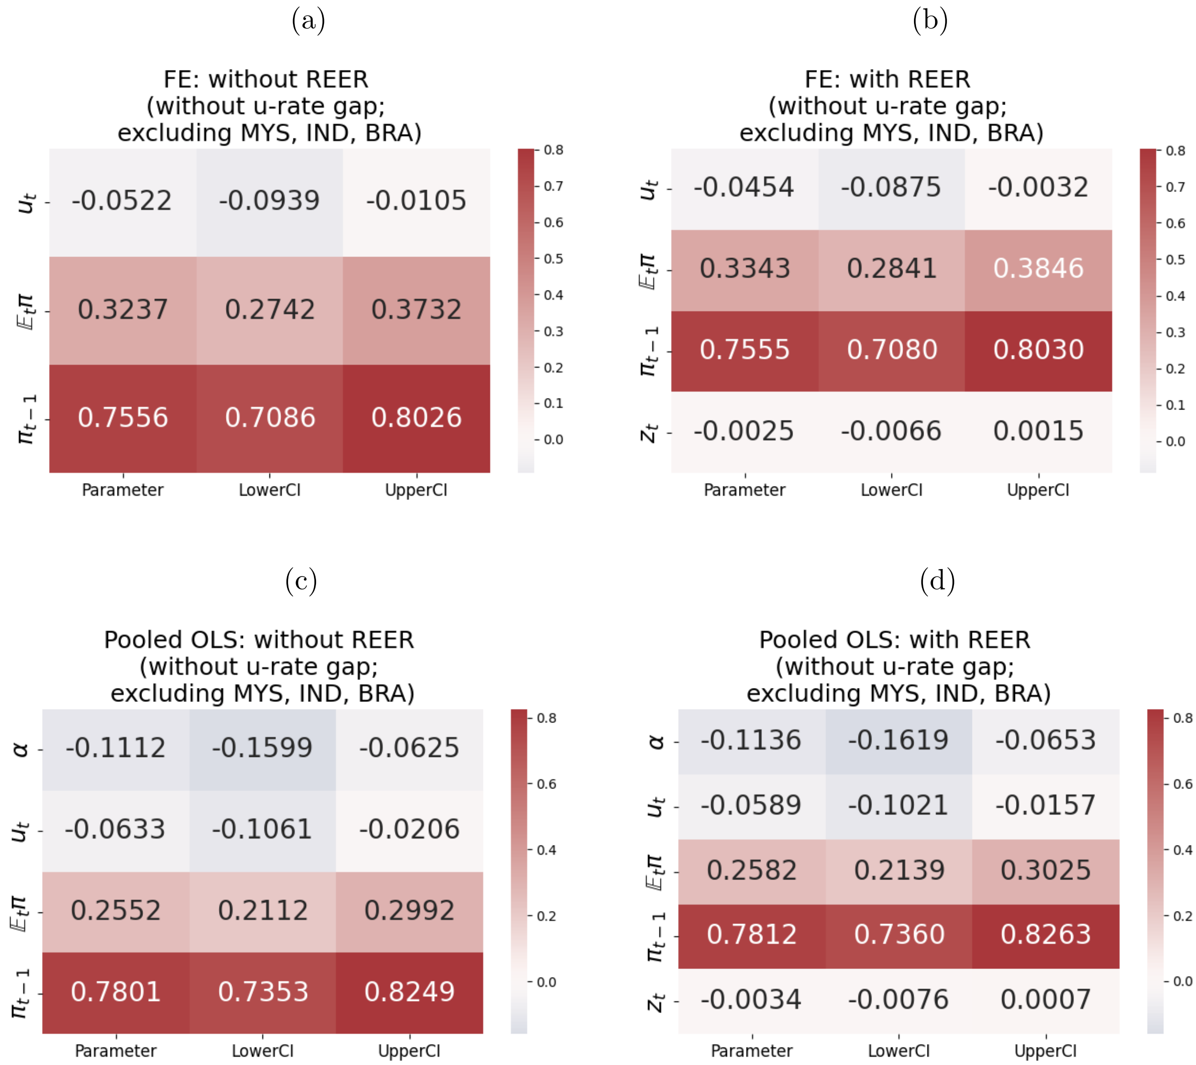

Estimates here further exclude Brazil, India and Malaysia. These are coefficient estimates from a fixed effects regression of core inflation on economy fixed effects, the year-on-year (YoY) change in the u-rate, 12-month-ahead expected inflation and the lag of core inflation. The confidence bands refer to 95% confidence intervals, and are computed with heteroscedasticity-robust standard errors.

Figures 12 shows the FE coefficient estimates, and the corresponding 95% confidence bands, of equation 6 in sub-figure (a), and equation 7 in sub-figure (b). Similar to the baseline specifications in figures 10 and 11, β_1 is small and negative, consistent with the latest literature on the Phillips curve. β_3 is positive and precisely estimated, which suggests that a 1 percentage point increase in the u-rate relative to the floor, i.e., larger slack, flattens the Phillips curve. Subfigures (c) and (d) then show the POLS estimates of equations 6 and 7, respectively. The POLS estimates are similar to their FE counterparts, but with β_1 marginally larger, and β_3 marginally smaller. Excluding Brazil, India and Malaysia results in quantitatively similar precise estimates, as illustrated in figure 13.

Figure 12

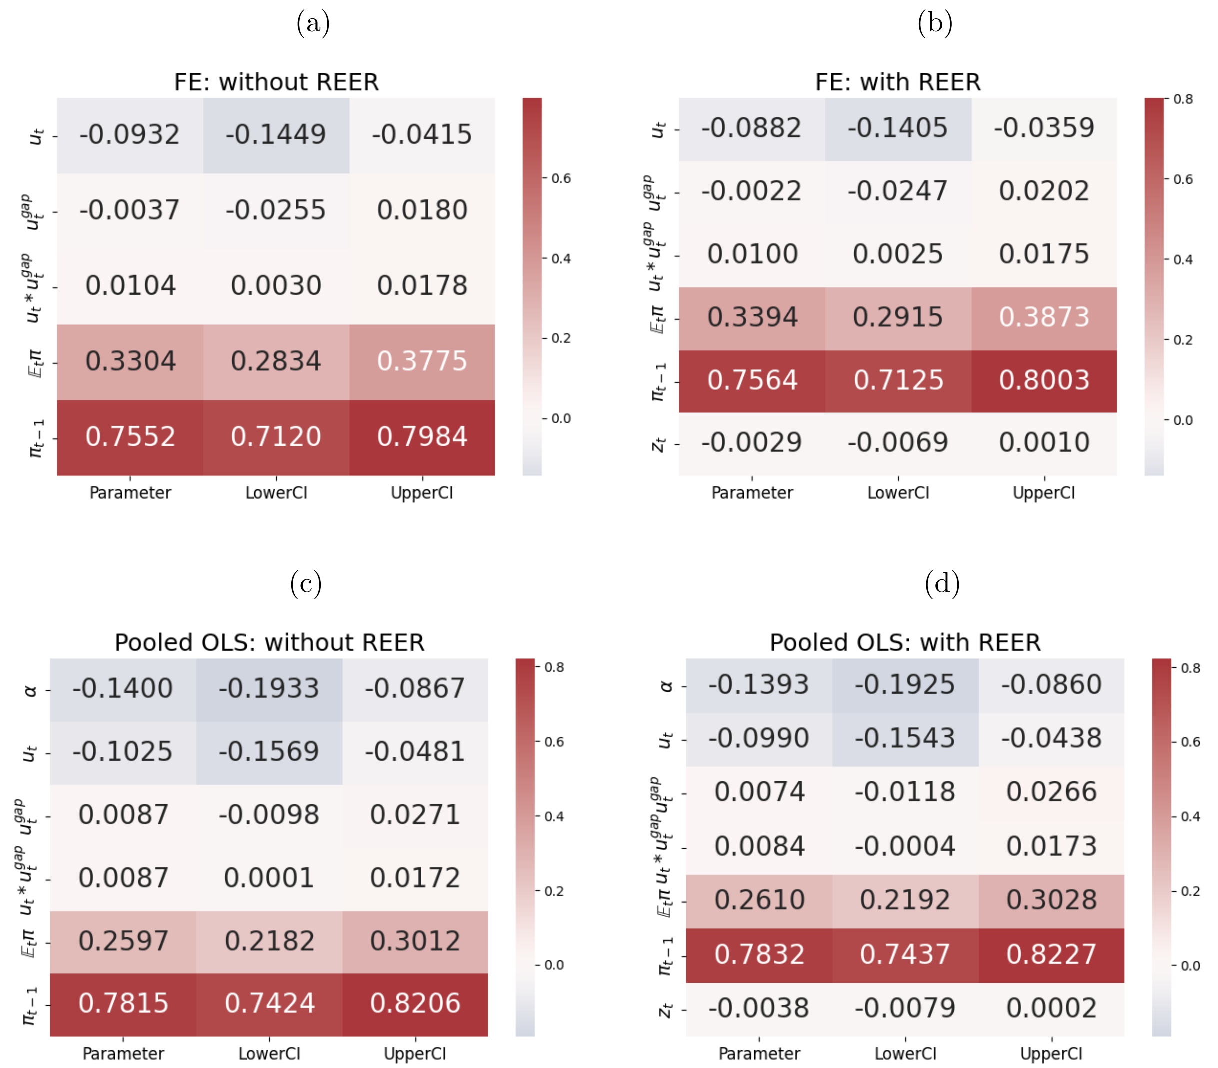

These are coefficient estimates from a fixed effects regression of core inflation on economy fixed effects, the year-on-year (YoY) change in the u-rate, the u-rate gap (takes values of at least 0), an interaction of the YoY change in the u-rate and the u-rate gap, 12-month-ahead expected inflation and the lag of core inflation. The confidence bands refer to 95% confidence intervals, and are computed with heteroscedasticity-robust standard errors.

Figure 13

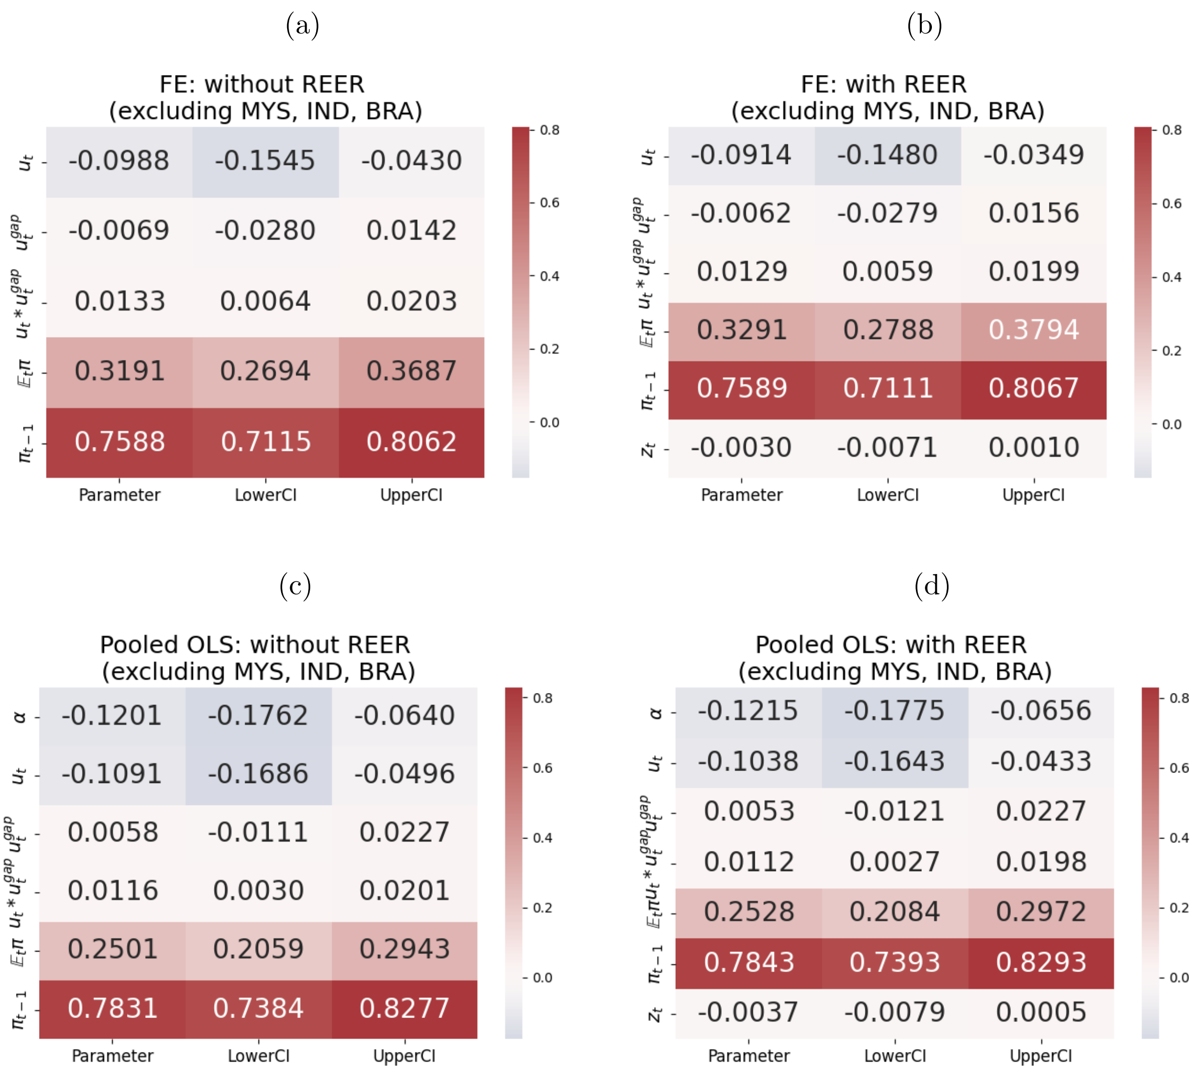

Estimates here further exclude Brazil, India and Malaysia. These are coefficient estimates from a fixed effects regression of core inflation on economy fixed effects, the year-on-year (YoY) change in the u-rate, the u-rate gap (takes values of at least 0), an interaction of the YoY change in the u-rate and the u-rate gap, 12-month-ahead expected inflation and the lag of core inflation. The confidence bands refer to 95% confidence intervals, and are computed with heteroscedasticity-robust standard errors.

4.2.2 Model estimates with US variables as controls

Figures 14 to 17 show the parameter estimates of equations 4, 5, 8 and 9, respectively. These correspond to the earlier figures, except for the inclusion of US u-rates and u-rate gaps as additional explanatory variables. Broadly speaking, the β_1 estimates are now smaller, but remain negative and precisely estimated. Estimates of β_3 are also smaller but remain positive and, when Brazil, India and Malaysia excluded, are precisely estimated. This corroborates the more parsimonious specification that includes the US directly as part of the panel than that includes US variables as explanatory variables for the other economies. As there are minimal differences quantitatively and qualitatively, the more parsimonious specification earlier will be used for subsequent analysis.

Figure 14

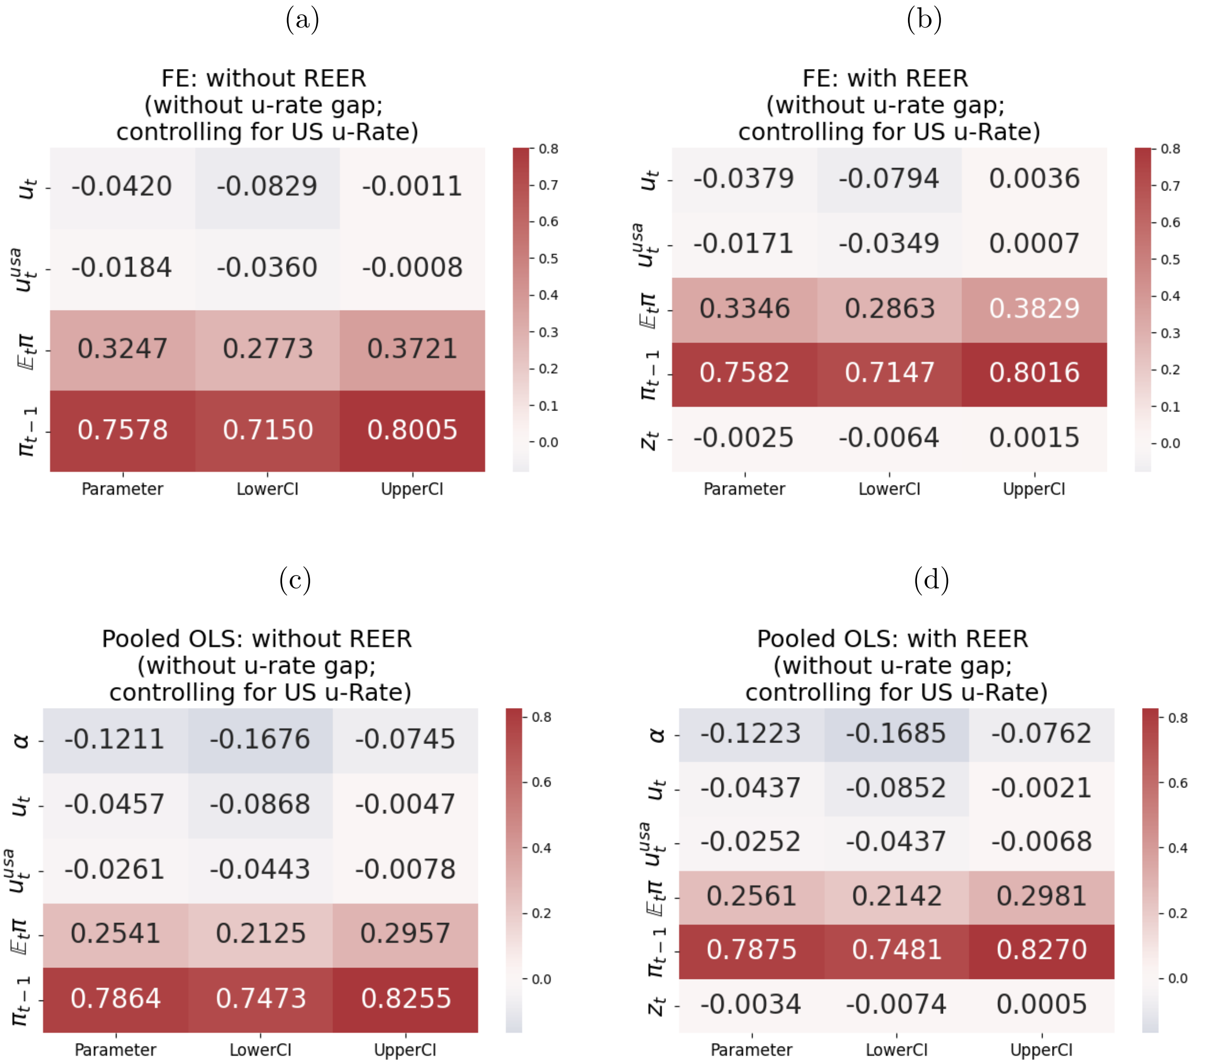

These are coefficient estimates from a fixed effects regression of core inflation on economy fixed effects, the year-on-year (YoY) change in the u-rate, 12-month-ahead expected inflation and the lag of core inflation; and further controls for the YoY change in US u-rate. The US is excluded from the analysis to avoid multicollinearity. The confidence bands refer to 95% confidence intervals, and are computed with heteroscedasticity-robust standard errors.

Figure 15

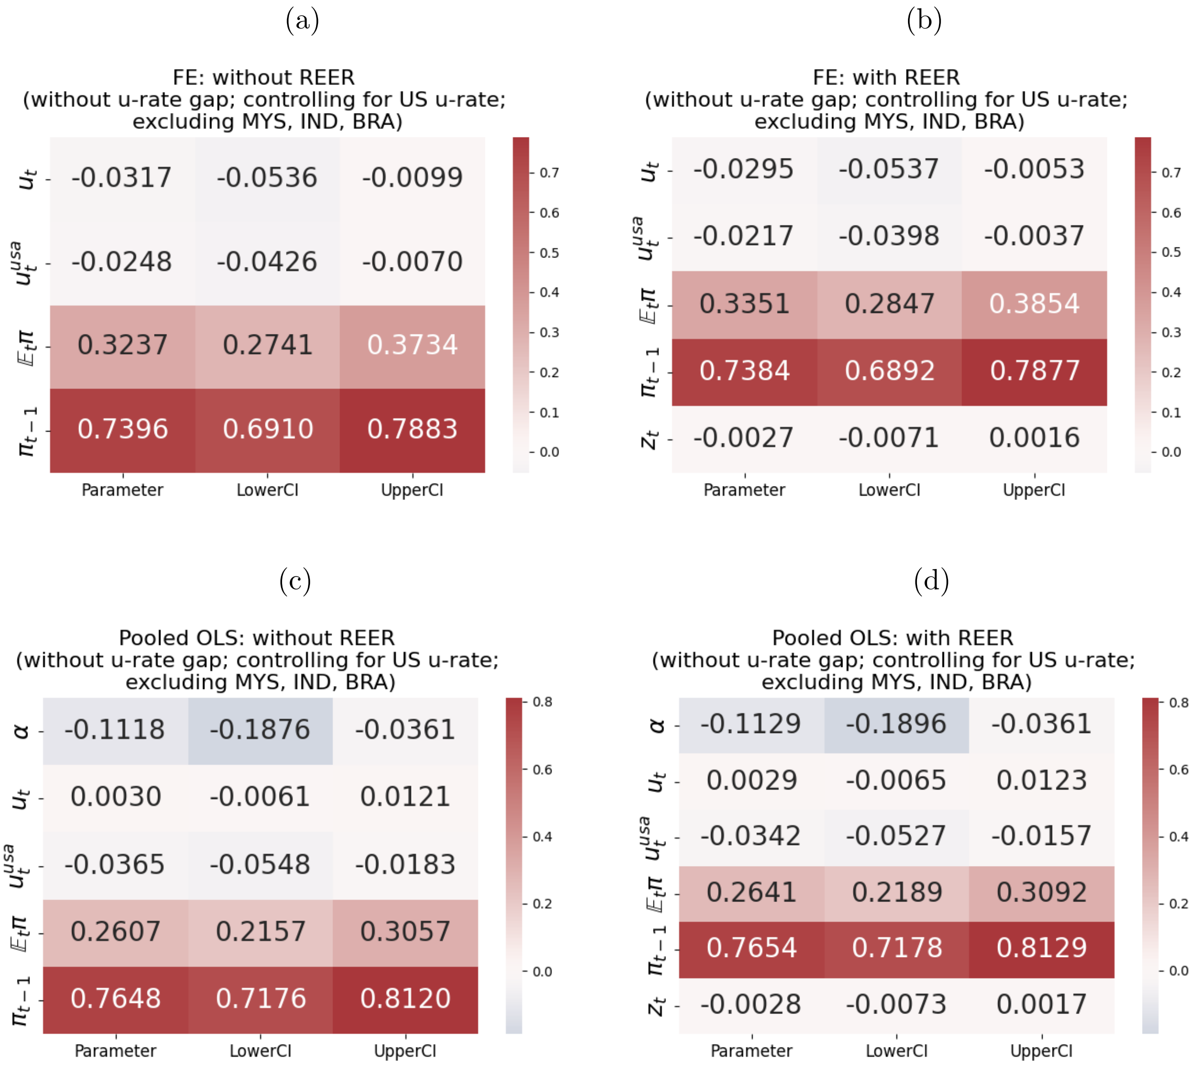

Estimates here further exclude Brazil, India and Malaysia. These are coefficient estimates from a fixed effects regression of core inflation on economy fixed effects, the year-on-year (YoY) change in the u-rate, 12-month-ahead expected inflation and the lag of core inflation; and further controls for the YoY change in US u-rate. The US is excluded from the analysis to avoid multicollinearity. The confidence bands refer to 95% confidence intervals, and are computed with heteroscedasticity-robust standard errors.

Figure 16

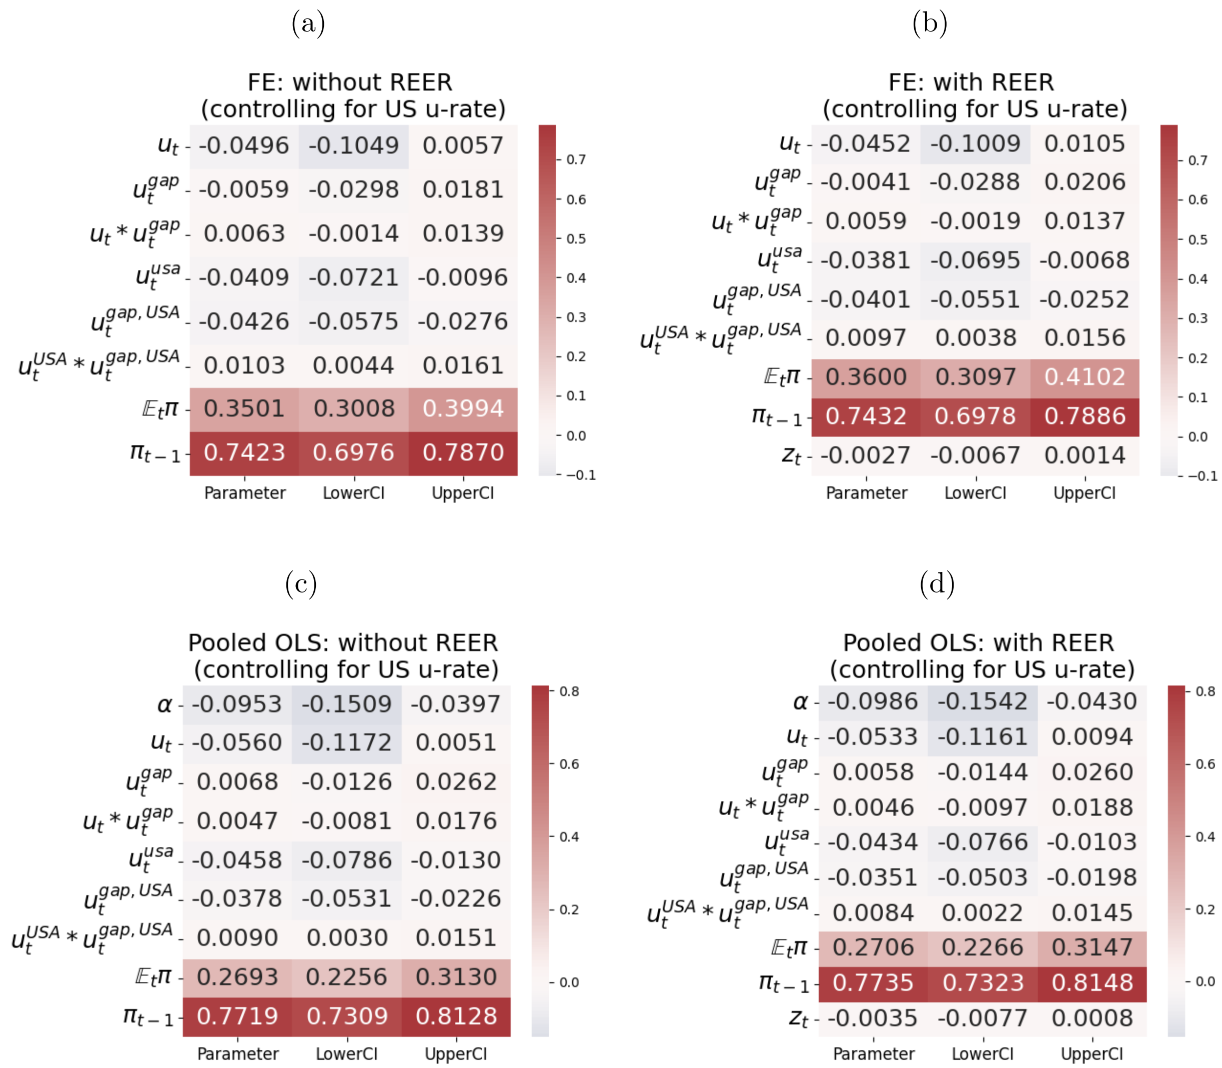

These are coefficient estimates from a fixed effects regression of core inflation on economy fixed effects, the year-on-year (YoY) change in the u-rate, the u-rate gap (takes values of at least 0), an interaction of the YoY change in the u-rate and the u-rate gap, 12-month-ahead expected inflation and the lag of core inflation; and further controls for the YoY change in US u-rate, US u-rate gap, and the interaction between the YoY change in US u-rate and the US u-rate gap. The US is excluded from the analysis to avoid multicollinearity. The confidence bands refer to 95% confidence intervals, and are computed with heteroscedasticity-robust standard errors.

Figure 17

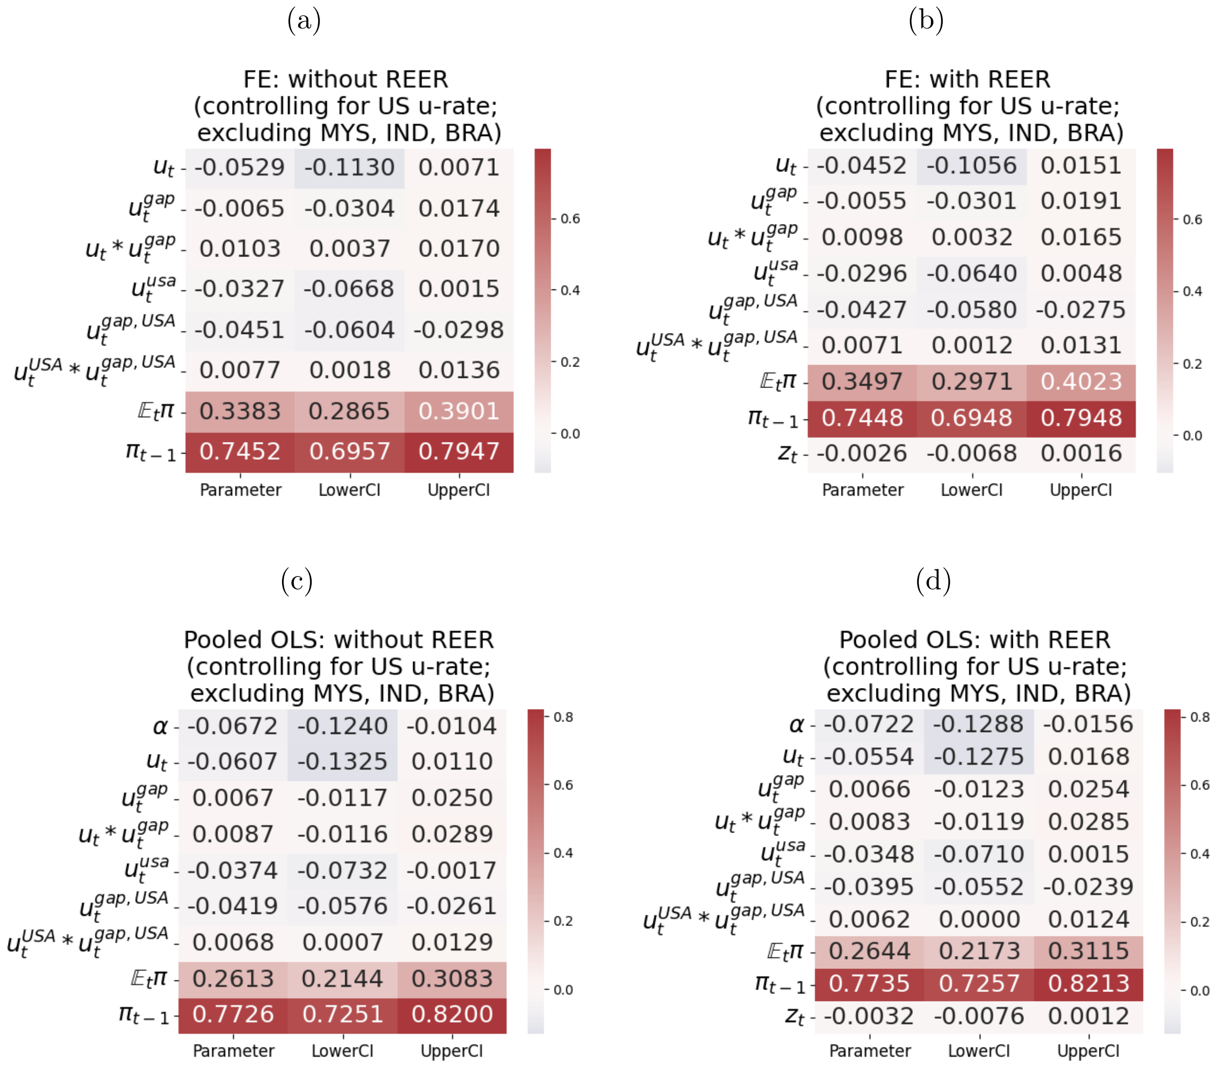

Estimates here further exclude Brazil, India and Malaysia. These are coefficient estimates from a fixed effects regression of core inflation on economy fixed effects, the year-on-year (YoY) change in the u-rate, the u-rate gap (takes values of at least 0), an interaction of the YoY change in the u-rate and the u-rate gap, 12-month-ahead expected inflation and the lag of core inflation; and further controls for the YoY change in US u-rate, US u-rate gap, and the interaction between the YoY change in US u-rate and the US u-rate gap. The US is excluded from the analysis to avoid multicollinearity. The confidence bands refer to 95% confidence intervals, and are computed with heteroscedasticity-robust standard errors.

4.2.3 Evolution of the inflation-unemployment link

By how much does the observed relationship between inflation and the u-rate flatten over time across economies? Is the degree of change meaningful for policy? Figures 18 to 21 plot the estimated slopes of the Phillips curve ω_{it} for the respective economies, stratified by economy group. The solid line indicates the point estimates, and the dashed lines the

95% confidence bands, based on the FE estimates of the model with REER included, and with the US included as part of the panel (figure 12b).[14] Across all economies, the degree of flattening of the Phillips curve ranges primarily from about 20% to 50% when u-rate gaps peak (around the trough of the corresponding business cycle phase). Except for the US, Brazil, and the Philippines, the COVID-19 pandemic did not differ substantially in terms of the degree of flattening in most economies, relative to earlier downturns.

Figure 18

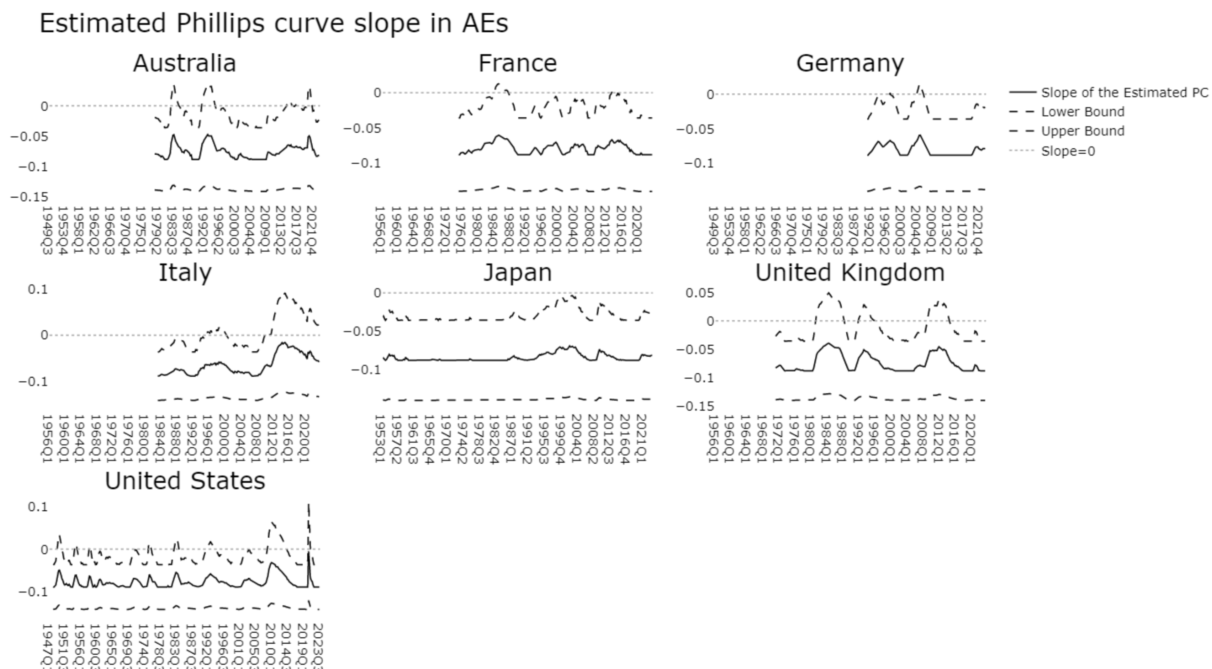

[14]Similar results are obtained when the US u-rate and u-rate gaps are included as explanatory variables for the other economies.

Figure 19

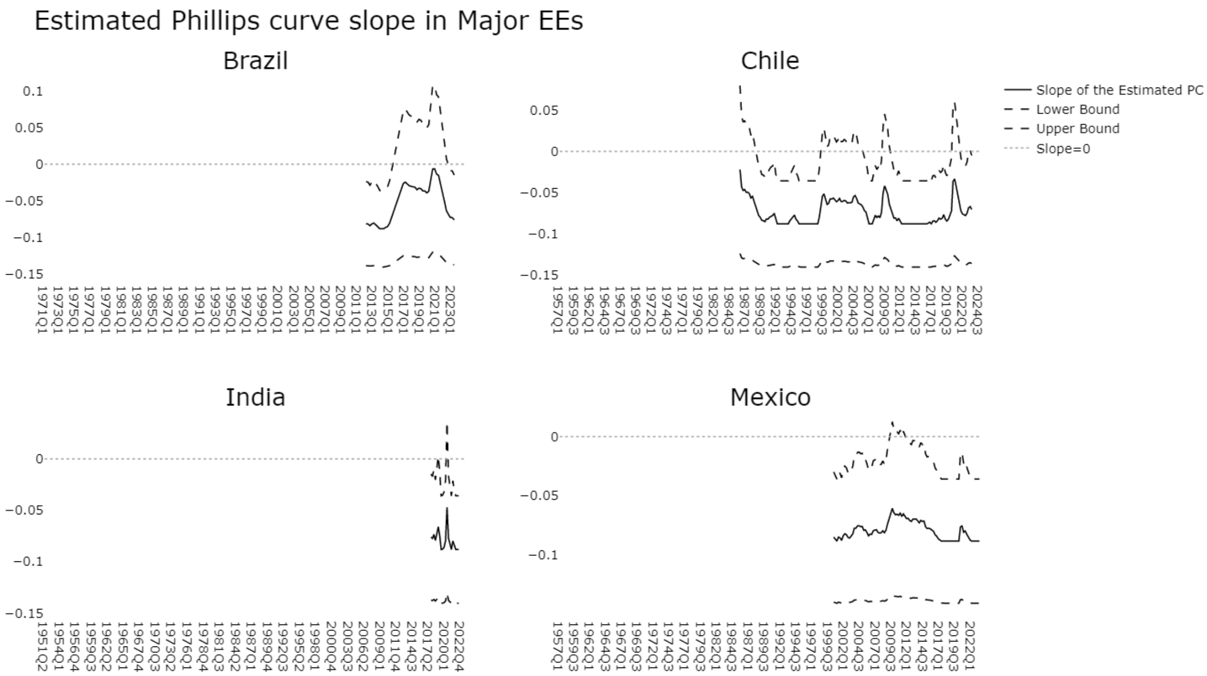

Figure 20

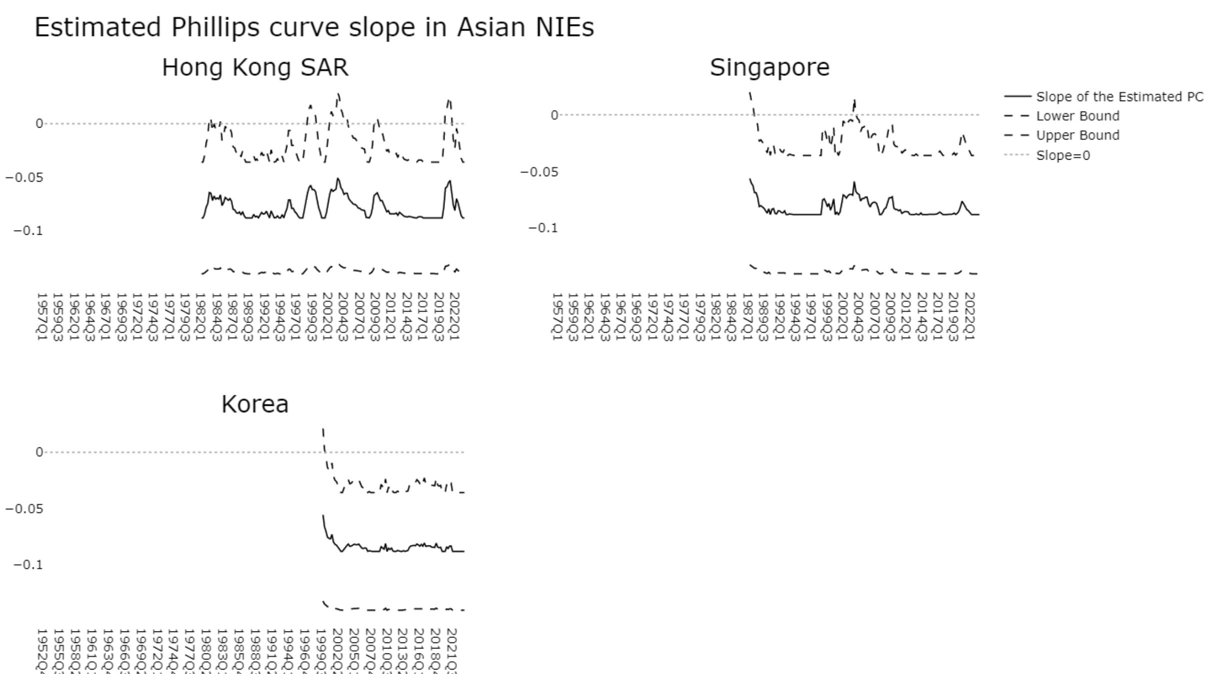

Figure 21

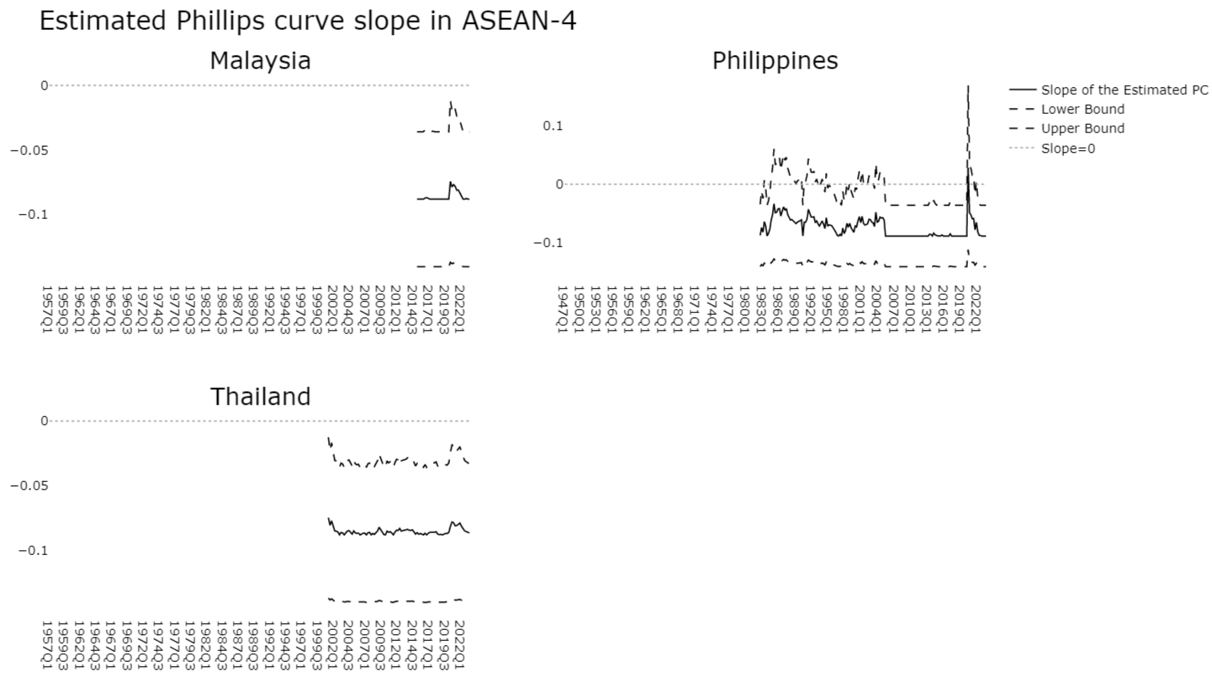

4.2.4 Threshold model estimates

When a binary indicator is used in lieu of the u-rate gap directly, as in equation 11 and 12, estimates indicate that there is a weaker inflation-unemployment link when the u-rate is further from the floor, as shown in figure 22. Subfigure (a) and (b) correspond to the FE estimates, and (c) and (d) the POLS estimates. The identified threshold of $\tau = 1.32$ (the u-rate being 1.32 percentage points from the u-rate floor) is close to the mean (1.50) and median (1.06), and could reflect gravitation towards these measures of central tendency. As in the continuous u-rate gap specification, β_1 estimates in the binary threshold variable model are small, negative and precisely estimated. As the dummy indicator switches on when the u-rate gap is at or below the threshold, the expected sign for β_3 is negative. If the u-rate gap is closer to zero, then the estimates in the previous analyses suggest that the inflation-unemployment link becomes more negative. As expected, the β_3 estimate is negative and precisely estimated. Summing β_1 and β_3 up also tallies with the earlier estimates of $\omega_{it} = 0.1$ when the u-rate gap is at its floor, affirming the finding that the inflation-unemployment link is (1) small and negative, and (2) flattens when the u-rate is further from the floor.

Figure 22

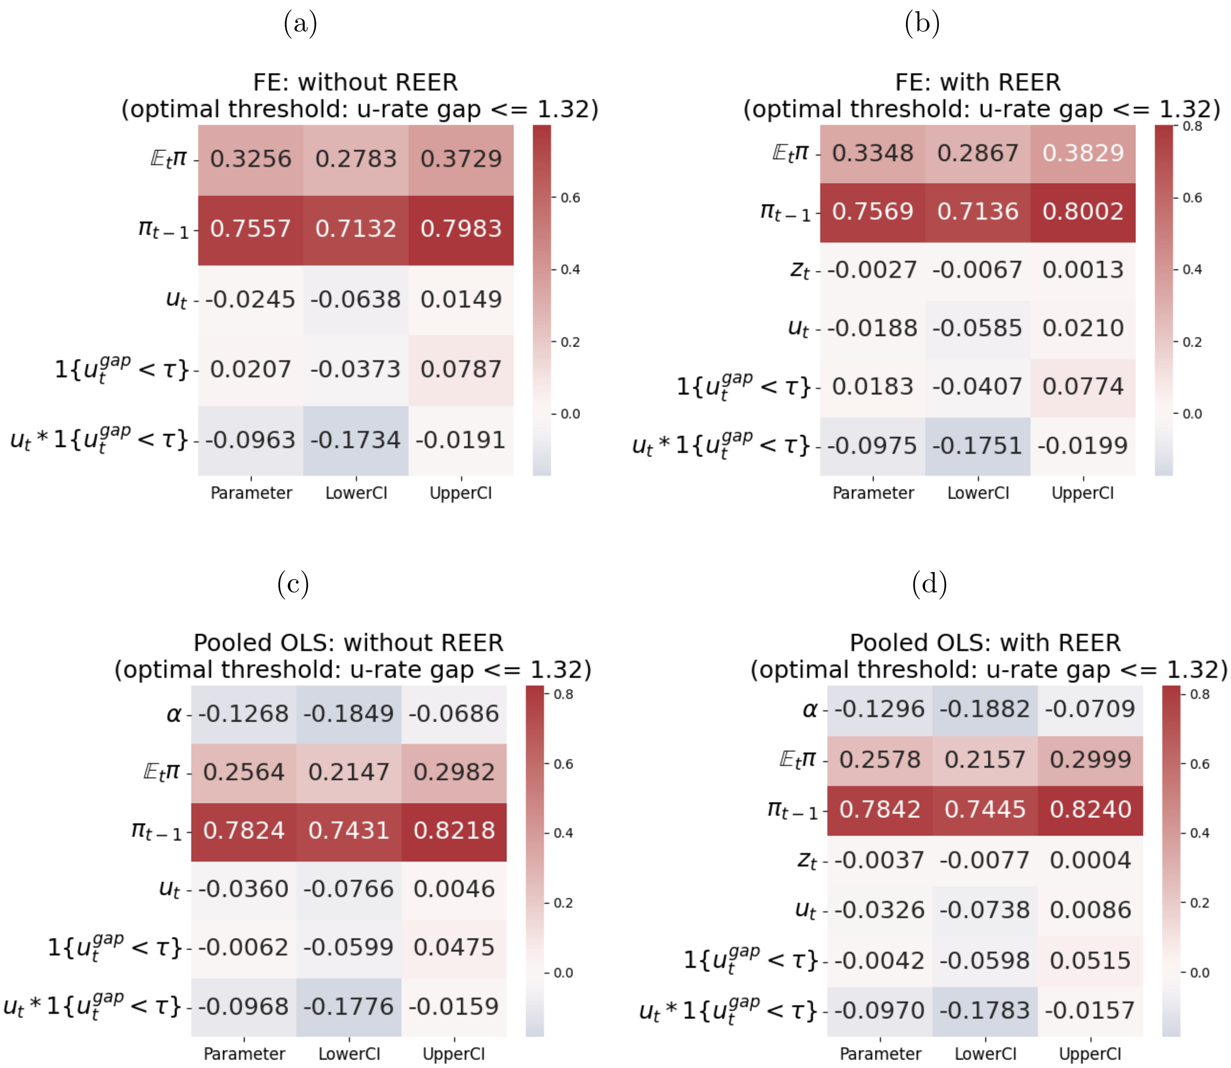

These are coefficient estimates from a fixed effects regression of core inflation on economy fixed effects, the year-on-year (YoY) change in the u-rate, the u-rate gap (takes values of at least 0), an interaction of the YoY change in the u-rate and the u-rate gap, 12-month-ahead expected inflation and the lag of core inflation. The confidence bands refer to 95% confidence intervals, and are computed with heteroscedasticity-robust standard errors.

4.2.5 Alternative model estimates with the NAIRU for the US

How does this compare against conventional measures of the unemployment gap based on the NAIRU? While published and regularly updated, estimates of NAIRU are hard to come by for most economies, the US' CBO publishes quarterly estimates of the NAIRU, as previously shown in figure 6. Confined to the US context, a comparison of model estimates with the plucking view's u-rate gap, and with the symmetric view's NAIRU gap is feasible. We now examine the relative fit of these models.

Figure 23 estimates equation 3 in subfigure 23a with only US data, replaces the u-rate with the NAIRU gap in subfigure 23b, and replaces the u-rate with the plucking u-rate gap in subfigure 23c. For the benchmark interaction model, subfigure 23d estimates equation 7, while the u-rate gap relative to the floor is replaced with the NAIRU gap in subfigure 23e. Figure 24 shows estimates of the same models without REER. In all variants of the baseline specification, estimates are imprecise. However, the magnitude and signs of the estimates agree with the earlier panel estimates. In addition, figure 25 compares the models based on goodness-of-fit measures and indicates that the model that interacts the u-rate with the plucking view's u-rate gap has the lowest corrected Akaike Information Criterion (AICc) and uncorrected AIC, and the highest log-likelihood, compared to the models with the NAIRU gap. Thus, where comparable data is available (e.g., the US), the asymmetric plucking u-rate gap appears to provide better explanatory power of observed inflation dynamics than using the CBO's symmetric NAIRU gap.

Figure 23

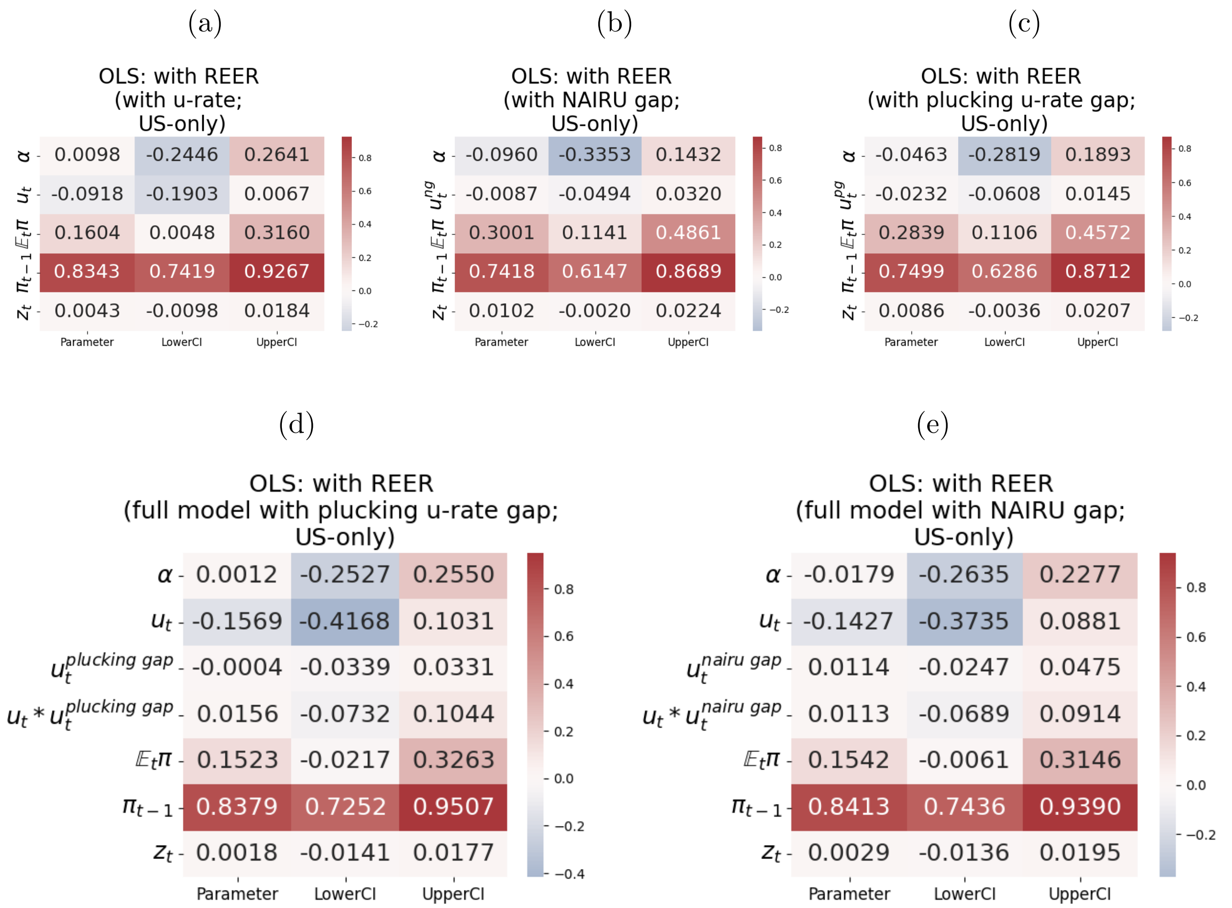

These are OLS coefficient estimates from linear regression of YoY core inflation on an intercept, lagged YoY core inflation, 12-month-ahead expected YoY inflation and YoY growth in REER, using only US observations. Subfigure (a) includes u-rate, subfigure (b) includes the NAIRU gap (difference between the u-rate and NAIRU shown in figure 6), subfigure (c) includes the plucking u-rate gap (difference between u-rate and the floor), subfigure (d) includes the u-rate interacted with the u-rate gap relative to the floor, while subfigure (e) includes the u-rate, interacted with the NAIRU gap.

Figure 24

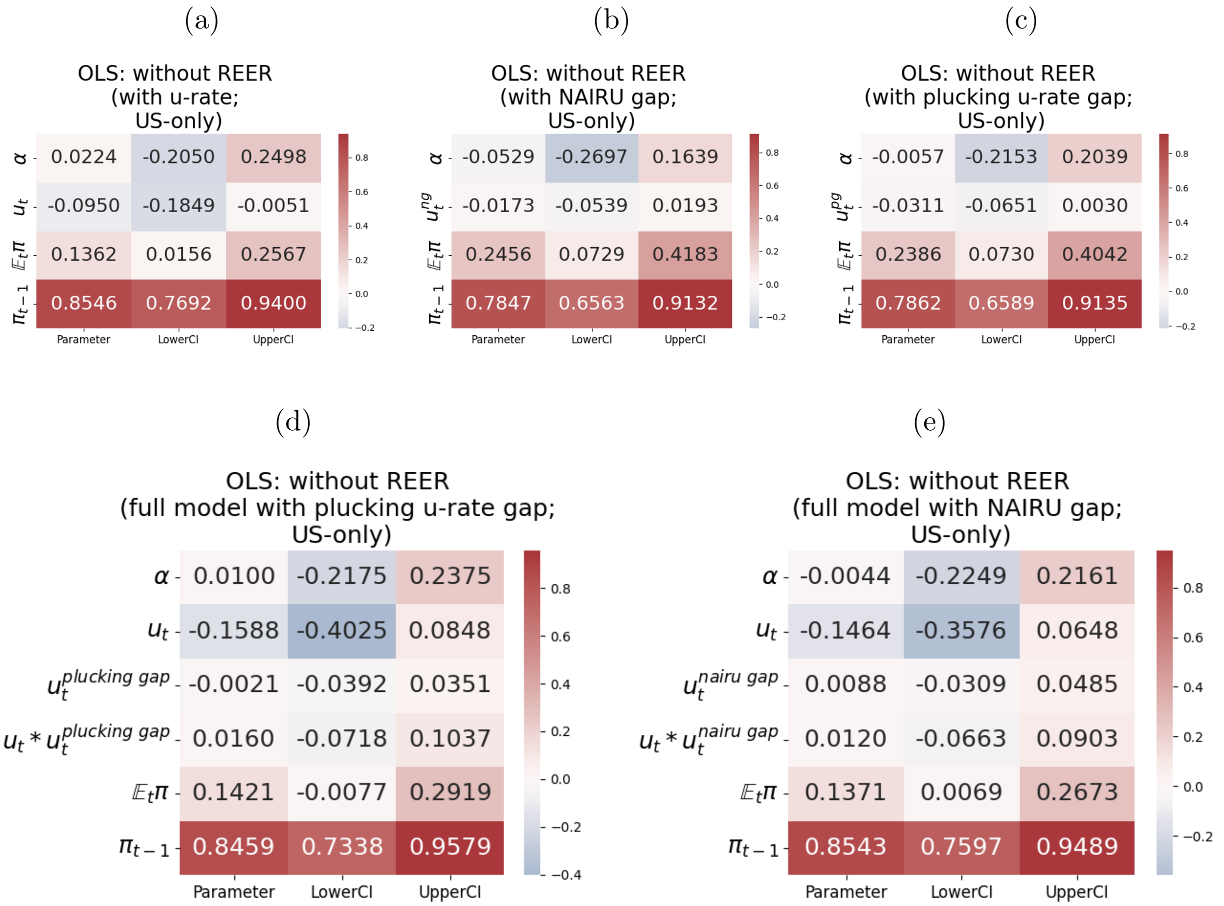

These are OLS coefficient estimates from linear regression of YoY core inflation on an intercept, lagged YoY core inflation, 12-month-ahead expected YoY inflation and YoY growth in REER, using only US observations. Subfigure (a) includes u-rate, subfigure (b) includes the NAIRU gap (difference between the u-rate and NAIRU shown in figure 6), subfigure (c) includes the plucking u-rate gap (difference between u-rate and the floor), subfigure (d) includes the u-rate, interacted with the u-rate gap relative to the floor, while subfigure (e) includes the u-rate, interacted with the NAIRU gap.

Figure 25

(a)

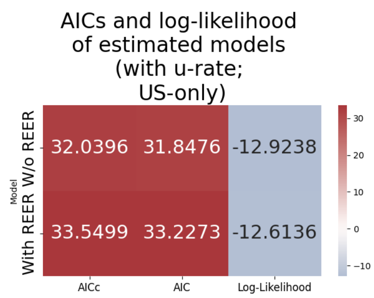

(b)

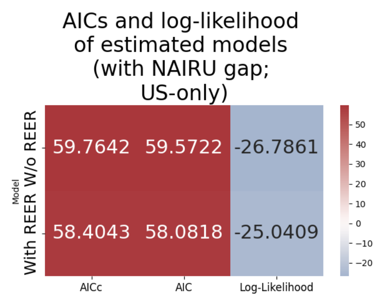

(c)

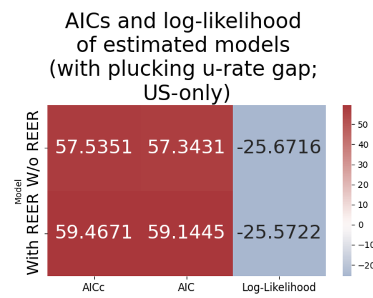

(d)

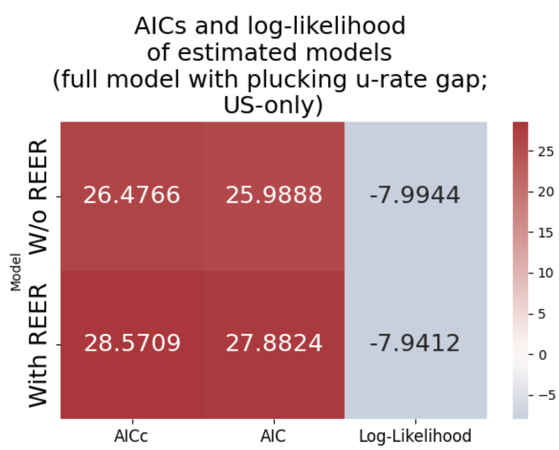

(e)

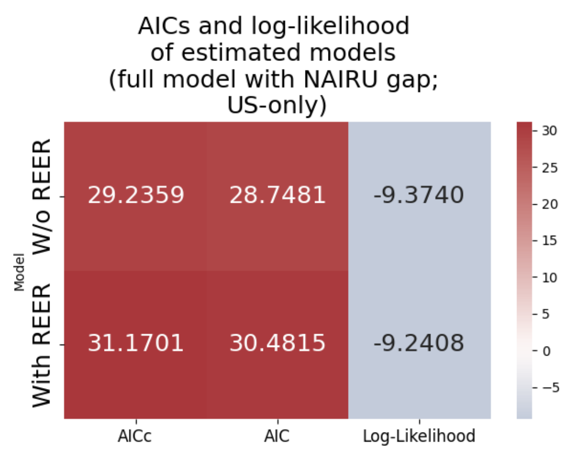

These are AICc, AIC and log-likelihoods of OLS linear regression estimates of YoY core inflation on an intercept, lagged YoY core inflation and 12-month-ahead expected YoY inflation, using only US observations. Row 1 of each subfigure corresponds to estimates without YoY growth in REER, shown in figure 24. Row 2 of each subfigure corresponds to model estimates with YoY growth in REER, shown earlier in figure 23. Subfigure (a) includes the u-rate, subfigure (b) includes the NAIRU gap (difference between u-rate, and NAIRU shown in figure 6), subfigure (c) includes the plucking u-rate gap, subfigure (d) includes the u-rate interacted with the u-rate gap relative to the floor, and subfigure (e) includes the u-rate, interacted with the NAIRU gap.

4.2.6 Policy discussion

What do these results mean for policy conduct? Insofar as the unemployment-inflation trade-off plays into policy decisions, the observed inflation cost of expansionary policy may be larger when the u-rate is already at the floor, and smaller when the u-rate is above the floor. Nonetheless, there are limits to how far the estimates changed. Except for the Philippines, point estimates of the Phillips curve slopes were always negative, and even in Philippines' case, the point estimate was positive only during the peak of the COVID-19 lockdowns in

2020 Q2 but was statistically insignificantly different from zero.[15] The asymmetry of the plucking model is also reflected here: the steepening of the Phillips curve stops once the u-rate hits the floor. On one level, a β_1 estimate of around 0.1 across all model specifications in figure 12 suggests that the unemployment-inflation trade-off is economically small. However, the flattening of the Phillips curve occurs in response to adverse shocks, which widen the u-rate gap. This is when policymakers are likely to want to provide stimulus, but may be reluctant to do so because of the inflationary impact. Our results suggest that the inflationary costs of such stimulus are likely to be small, and hence a cost-benefit analysis of policy would tend to favour more aggressive expansionary policy. Thus, welfare loss minimisation from stabilisation policy during downturns is asymmetric to that of during upturns. The emphasis on asymmetry is further bolstered by our demonstration of the goodness of fit of the different models using US data. Overall, the results suggest that policymakers may consider closing the gap quickly, and as aggressively as possible, given the reduced cost from an inflation acceleration perspective; likewise, policy easing cycles may justifiably be faster than tightening cycles.

5 Propagation of macroeconomic shocks

So what if the Phillips curve is flatter when the u-rate is above the floor? A firmer case to consider separate policy conduct regimes when the u-rate gap is zero, and above zero, requires understanding of the propagation of economic shocks under both regimes. A natural follow-up is simply to quantify empirically if key macroeconomic variables respond to various shocks, including those from policy, differently when the u-rate floor binds than when it does not. This leads us to an analysis of threshold effects.

5.1 Methodology

Subsequent analysis uses a panel threshold local projections with exogenous variables (LP-X) model, estimated with fixed effects. The equivalent functional form for a single entity setting (non-panel) is also detailed, and its properties discussed, in Gonçalves et al. (2023). The original version of the LP was developed by Jordà (2005), which subsequently provided a panel implementation with fixed effects estimators. Plagborg-Møller & Wolf (2021) showed that the two are conceptually equivalent. Montiel Olea & Plagborg-Møller (2021) then subsequently demonstrated the relative robustness of the LP compared to a vector autoregression (VAR), and also highlighted its simplicity in inference, as the LP impulse

[15] The US and Brazil were the only other economies whose Phillips curve slopes flattened close to zero, but remained negative (US: −0.00462 in 2020 Q2; Brazil: −0.00666 in 2020 Q3 and −0.00642 in 2020 Q4).

responses are essentially regression coefficients. Kilian & Kim (2011) also compared the LP against a bias-adjusted bootstrapped VAR. In general terms, LPs avoid the computational, and hence interpretative, complexities of VARs.

Equation 13 shows a panel LP-X model, without allowing for threshold effects. Estimating this equation yields the impulse response functions (IRFs) for y to various shocks X at horizon h. X enters the equation as first differences ΔX. The dependent variable in the model is a difference of the response variable up to h-period forward and the first lag. Value h also corresponds to the specific horizon of the IRF. Vector $\Delta X_{1,i,t}$ contains the other macroeconomic variables that are deemed to affect y contemporaneously, while vector $\Delta X_{2,h,i,t-1}$ contains variables that affect y with a lag, which also includes the lag of y itself. Vector $\Delta Z_{i,t}$ contains the exogenous variables in first differences. All variables on the right-hand side are expressed in first differences. $\varepsilon_{i,h,t}$, and $\alpha_{i,h}$ are the errors, and economy fixed effects, respectively. The IRFs of y to a shock from X_k up to H-periods ahead are then simply the coefficients $\{\beta_{1,X_k,h=0},\ ...,\ \beta_{1,X_k,h=H}\}$ if X_k shocks propagate to y contemporaneously, and $\{\beta_{2,X_k,h=0},\ ...,\ \beta_{2,X_k,h=H}\}$ if X_k shocks propagate to y with a lag. This essentially resembles recursive identification, i.e., equivalent to Cholesky ordering in the corresponding vector autoregression (VAR) approach.

$$y_{i,t+h} - y_{i,t-1} = \alpha_{i,h} + \beta_{\mathbf{1,h}}\mathbf{\Delta X_{1,i,t}} + \beta_{\mathbf{2,h}}\mathbf{\Delta X_{2,h,i,t-1}} + \beta_{\mathbf{3,h}}\mathbf{\Delta Z_{i,t}} + \varepsilon_{i,h,t} \tag{13}$$

Now, the threshold version of the panel LP-X is shown in equation 14. The only difference is that all variables on the right-hand side of the equation, including the exogenous block, are interacted with state variables, which are dummy indicators for when the u-rate is at its floor, and when the u-rate is above the floor, and the state variables on their own are additionally included as controls. These are $S_{1,it} = \mathbb{1}\{u^g_{it} = 0\}$, and $S_{0,it} = \mathbb{1}\{u^g_{it} > 0\}$, respectively. Essentially, when the u-rate is at the floor, $S_{1,it}$ takes value 1, while $S_{0,it}$ takes value 0, and vice versa when the u-rate is above the floor. When the u-rate floor binds ($u^g_{it} = 0$), then the IRFs of y to a shock from X_k up to H-periods ahead are the coefficients $\{\lambda_{1,X_k,h=0},\ ...,\ \lambda_{1,X_k,h=H}\}$ if X_k shocks propagate to y contemporaneously, and $\{\lambda_{2,X_k,h=0},\ ...,\ \lambda_{2,X_k,h=H}\}$ if X_k shocks propagate to y with a lag. Likewise, the IRFs when the u-rate is above the floor are $\{\pi_{1,X_k,h=0},\ ...,\ \pi_{1,X_k,h=H}\}$, and $\{\pi_{2,X_k,h=0},\ ...,\ \pi_{2,X_k,h=H}\}$, respectively.

$$
\begin{aligned}
y_{i,t+h} - y_{i,t-1} = \\
\alpha_{i,h} + \gamma_1 S_{1,it} + \gamma_0 S_{0,it} \\
+\lambda_{\mathbf{1,h}} \mathbf{\Delta X_{1,i,t}} * S_{1,it} + \lambda_{\mathbf{2,h}} \mathbf{\Delta X_{2,h,i,t-1}} * S_{1,it} \\
+\pi_{\mathbf{1,h}} \mathbf{\Delta X_{1,i,t}} * S_{0,it} + \pi_{\mathbf{2,h}} \mathbf{\Delta X_{2,h,i,t-1}} * S_{0,it} \\
+\delta_{\mathbf{1,h}} \mathbf{\Delta Z_{i,t}} * S_{1,it} + \delta_{\mathbf{2,h}} \mathbf{\Delta Z_{i,t}} * S_{0,it} \\
+\varepsilon_{i,h,t}
\end{aligned}
\tag{14}
$$

The endogenous block of the panel threshold LP-X contain 6 variables, described earlier in section 2. Equation 15 shows the ordering of variables, where expected inflation responds to all shocks contemporaneously, and private sector debt reacts to all shocks, except shocks from itself, with a lag. The exogenous block contains 3 variables, the YoY growth rate in brent crude oil prices to control for global commodity price conditions, the YoY growth rate in the global economic policy uncertainty (GEPU) index described in Baker et al. (2016) and the max-GEPU[16] indicator, computed as per Jackson et al. (2019). Controlling for both the GEPU and the max-GEPU serves to differentiate between the state of prevailing uncertainty, and substantial exogenous uncertainty shocks.

$$
\begin{aligned}
\text{Private Debt} \rightarrow \text{Interest Rate} \rightarrow \text{Real GDP} \\
\rightarrow \text{Core Inflation} \rightarrow \text{REER} \rightarrow \text{Expected Inflation}
\end{aligned}
\tag{15}
$$

Of primary interest to this paper are the responses of core inflation, and real output growth to interest rate shocks, as these directly address how policy transmission differs when the u-rate is at or above the floor. We also investigate how the propagation to other shocks varies.

5.2 Findings

Figure 26 shows the estimated IRFs (the full set of IRFs is reported in figure C.1 of the appendix). In figure 26a, a positive short-term interest rate shock (i.e., a rate hike) negatively impacts core inflation after 4 quarters only when the u-rate floor binds, but hovers around zero when the u-rate floor does not bind. This is consistent with the impact on the 12-month-ahead expected inflation by professional forecasters. Meanwhile, the impact of interest rate shocks on real GDP growth is quantitatively small in both regimes. When the u-rate floor

[16]Max-GEPU is the higher of 0 and the percentage difference between the current GEPU and the highest level of GEPU observed in the previous year (4 quarters). Where U_{it} is GEPU in economy i and period t, this is computed as $\widehat{U_{it}} = \max\{0, \frac{U_{it} - \max\{U_{i,t-1},\ U_{i,t-2},\ U_{i,t-3}, U_{i,t-4}\}}{\max\{U_{i,t-1},\ U_{i,t-2},\ U_{i,t-3}, U_{i,t-4}\}}\}$

binds, a contractionary reaction is observed as early as after 2 to 4 quarters, versus 5 to 6 quarters when the u-rate floor does not bind. Tying in with the analysis earlier on the observed Phillips curve slopes in section 4, policy transmits differently when there is economic slack, compared to when there is not.

Some empirical studies of central banks' behaviour have identified variation in their responses broadly consistent with these asymmetries. Taylor & Davradakis (2006) found evidence of non-linear behaviour by the Bank of England (BOE), while Dolado et al. (2004) found that the Fed exhibited non-linear policy setting after 1983, but not before 1979. Rahman & Serletis (2010), from a monetary policy transmission angle, also found non-linear effects of interest rate shocks on oil prices. Findings of this paper thus far, when placed within the context of these studies, suggest an explanation for differences in policy responses depending on whether the u-rate is at the floor or not.

Figure 26

(a)

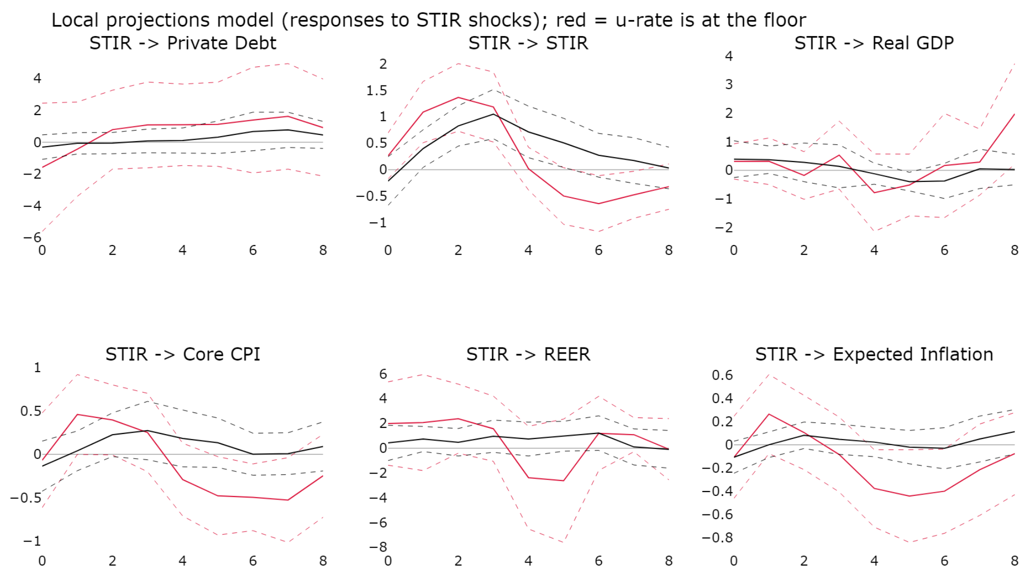

(b)

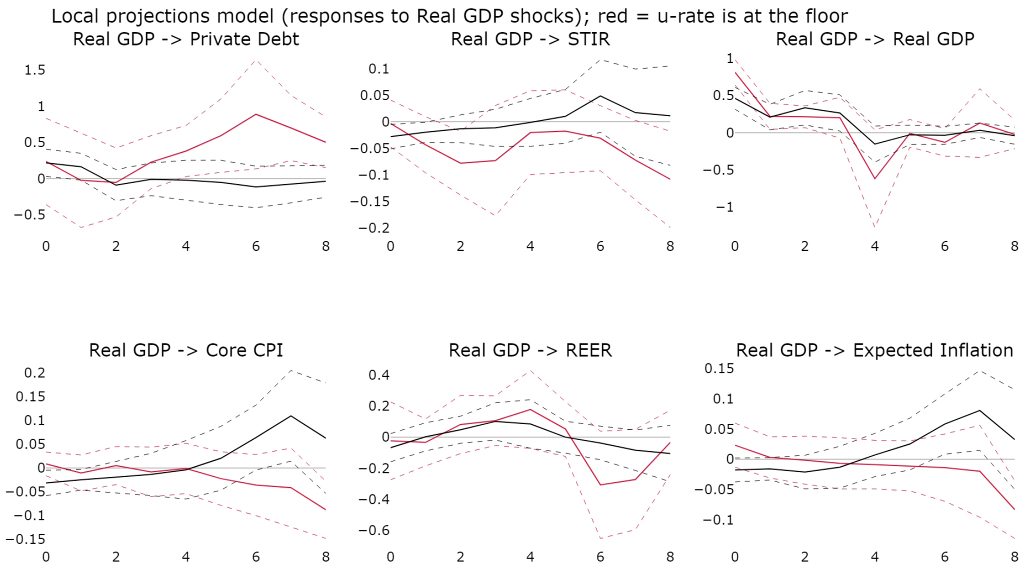

(c)

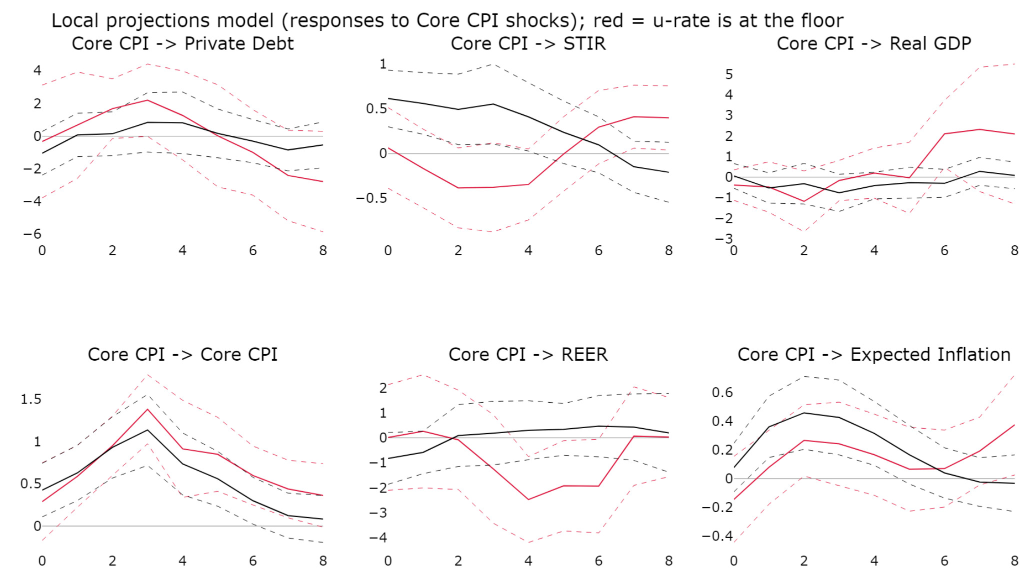

Subfigures show the estimated IRFs of the corresponding variables from select shock variables labelled in the respective titles, which follow the convention "*shock variable* → *response variable*". The panel threshold LP-X model from which the IRFs are estimated from are described in section 5.1. The red lines indicate the IRFs for when the u-rate is at the floor, and black lines for when the u-rate is above the floor. The dotted lines correspond to the respective 95% confidence intervals, estimated with heteroscedastic-robust standard errors. Estimation, and visualisation, are implemented with the `localprojections` package in Python. The full set of IRF estimates is reported in figure C.1 of the appendix.

However, central banks are not concerned solely about monetary policy transmission. Policy decisions are also based on policymakers' understanding of the transmission of non-policy shocks, so as to understand better the superstructure of the economic environment. These also provide evidence in support of the plucking model. In figure 26c, inflation shocks adversely impact output in the short run (after 2 quarters) before growing to more than offset this a few quarters later, when the u-rate floor binds. When the u-rate floor does not bind, the adverse impact of an inflation shock on output growth peaks later and tapers out after 6 quarters. In figure 26b, the impact of a positive output shock on inflation in the medium run is negative when the u-rate floor binds and not otherwise. This supports the interpretation of the u-rate floor (or output ceiling, or potential output) as measures of supply, as in Coibion et al. (2017). In the plucking view, when the u-rate is at the floor, supply is at capacity. Positive output shocks in these periods may have captured increments in the production frontier, hence relaxing capacity. On the other hand, the IRFs estimated for when the u-rate is above the floor may have captured primarily demand-driven inflation responses.[17] These findings are corroborated by previous empirical studies, which found non-linear relationships between output growth and inflation, albeit primarily with thresholds in the magnitude of inflation, e.g., Azam & Khan (2022) on a panel of 27 economies, Munir et al. (2009) on Malaysia, Aydin et al. (2017) on major Organisation of Islamic Cooperation (OIC) economies, and Khan & Ssnhadji (2001) on 140 economies. Eggoh & Khan (2014) further found that the degree of non-linearity is also sensitive to structural factors, including the level of financial development, capital accumulation, and trade openness. The contribution of this paper, within the context of the empirical literature of non-linear inflation-growth nexus, is the possibility of a threshold effect arising from underlying economic slack. For policymakers, these differential state-dependent IRFs suggest that policy response to various shocks, e.g., output, price, or exchange rate, ought to be tailored, depending on whether the u-rate floor binds or not.

6 Conclusion

The plucking model implies that the u-rate hovers at or above a theoretical floor and is occasionally "plucked" up by adverse shocks, before slowly returning in the direction of the floor. This implies that the u-rate behaves asymmetrically over the business cycles: in particular, (i) the amplitude of the u-rate during expansion phases and subsequent contraction phases are largely unrelated; (ii) whereas large u-rate rises during contraction phases are associated with larger subsequent expansions. This paper has demonstrated evidence consistent with

[17] To complete the picture, the reaction of private sector debt growth does not exhibit state-dependency, except for the reaction to REER shocks. However, the impact of an appreciation shock in the REER differs between regimes. When the u-rate floor binds, output growth and inflation react negatively right from when the shock lands. However, when the u-rate floor does not bind, the impact on both is negligible. The relevant IRFs are reported in figure C.1

the plucking model in a sample of 19 economies.

The paper then examined whether, or how, a quantified u-rate floor could be used to inform policy. Specifically, does the unemployment-inflation trade-off depend on where the u-rate is relative to the floor? If so, does it then matter for policy transmission, and the propagation of other shocks, in general?

This paper followed the approach in Suah (2023) and Dupraz et al. (2019), which results in identification of the u-rate floor as a slow-moving asymmetric trend that is relatively stable in response to additional observations, supporting its use in policy decisions that must be made in real time. A close examination of the behaviour of the floor for US data compared with the CBO's NAIRU estimates illustrated that the floor was almost always the lower of the two. Examining the slope of the Phillips curve as a function of the estimated floor, this paper found that the Phillips curve slope flattened as the u-rate moved further from the floor, although it generally remained small and negative. This implies that the cost — in terms of possible higher inflation — from expansionary policy during downturns is likely to be smaller than during other periods, lending support to the idea of aggressive policy easing in response to adverse shocks that result in substantial economic slack. The paper also analysed the IRFs estimated from a panel threshold local projections, which showed that the effects of monetary policy shocks vary depending on whether the u-rate is at or above the floor.

Taken together, the results support the view that bringing a plucking model of unemployment to the data may assist policymakers in macroeconomic surveillance and policy formulation. On this front, there is scope for further research. For instance, while studies on the plucking model have focused on output and unemployment, other variables could exhibit similar characteristics, including financial variables. In addition, most studies of the plucking model are primarily empirical; assessment of the welfare implications of different policy responses in the presence of plucking would require a more analytical approach. This is important to understand the deep interactions between financial and real variables depending on whether the floor binds or not. Finally, there are implications for central bank communication, in particular to fully leverage on the regime shifts in monetary policy transmission and aggressiveness in policy response when the floor binds and otherwise.

References

Aydin, C. et al. (2017), 'The inflation-growth nexus: A dynamic panel threshold analysis for D-8 countries', *Romanian Journal of Economic Forecasting* **20**(4), 134–151.

Azam, M. & Khan, S. (2022), 'Threshold effects in the relationship between inflation and economic growth: Further empirical evidence from the developed and developing world', *International Journal of Finance & Economics* **27**(4), 4224–4243.

Baker, S. R., Bloom, N. & Davis, S. J. (2016), 'Measuring economic policy uncertainty', *The Quarterly Journal of Economics* **131**(4), 1593–1636.

Ball, L. (1994), What determines the sacrifice ratio?, *in* 'Monetary policy', The University of Chicago Press, pp. 155–193.

Bowdler, C. (2009), 'Openness, exchange rate regimes and the Phillips curve', *Journal of International Money and Finance* **28**(1), 148–160.

Cecchetti, S. G. & Rich, R. W. (2001), 'Structural estimates of the US sacrifice ratio', *Journal of Business & Economic Statistics* **19**(4), 416–427.

Coibion, O., Gorodnichenko, Y. & Ulate, M. (2017), 'The cyclical sensitivity in estimates of potential output', *National Bureau of Economic Research* .

Cusbert, T. (2017), 'Estimating the NAIRU and the unemployment gap', *RBA Bulletin, June* pp. 13–22.

Daniels, J. P., Nourzad, F. & VanHoose, D. D. (2005), 'Openness, central bank independence, and the sacrifice ratio', *Journal of Money, Credit, and Banking* **37**(2), 371–379.

De Simone, F. N. & Clarke, S. (2007), 'Asymmetry in business fluctuations: International evidence on Friedman's plucking model', *Journal of International Money and Finance* **26**(1), 64–85.

Dolado, J., Pedrero, R. M.-D. & Ruge-Murcia, F. J. (2004), 'Nonlinear monetary policy rules: some new evidence for the US', *Studies in Nonlinear Dynamics & Econometrics* **8**(3).

Dupraz, S., Nakamura, E. & Steinsson, J. (2019), 'A plucking model of business cycles', *National Bureau of Economic Research* .

Eggoh, J. C. & Khan, M. (2014), 'On the nonlinear relationship between inflation and economic growth', *Research in Economics* **68**(2), 133–143.

Fernald, J. G. (2015), 'Productivity and potential output before, during, and after the Great Recession', *NBER Macroeconomics Annual* **29**(1), 1–51.

Fontanari, C., Palumbo, A. & Salvatori, C. (2020), 'Potential output in theory and practice: A revision and update of Okun's original method', *Structural Change and Economic Dynamics* **54**, 247–266.

Friedman, M. (1993), 'The "plucking model" of business fluctuations revisited', *Economic Inquiry* **31**(2), 171–177.

Furuoka, F. & Harvey, H. (2015), 'Estimation of new Keynesian Phillips curve in Malaysia', *Malaysian Journal of Business and Economics (MJBE)* .

Furuoka, F. et al. (2007), 'Does the "Phillips curve" really exist? New empirical evidence from Malaysia', *Economics Bulletin* **5**(16), 1–14.

Galbraith, J. K. (1997), 'Time to ditch the NAIRU', *Journal of Economic Perspectives* **11**(1), 93–108.

Gonçalves, S., Herrera, A. M., Kilian, L. & Pesavento, E. (2023), 'State-dependent local projections', *Federal Reserve Bank of Dallas Working Papers* **2302**.

Hall, R., Feldstein, M., Frankel, J., Gordon, R., Romer, C., Romer, D. & Zarnowitz, V. (2003), 'The NBER's business-cycle dating procedure', *Business Cycle Dating Committee, National Bureau of Economic Research* .

Hamilton, J. D. (1996), 'This is what happened to the oil price-macroeconomy relationship', *Journal of Monetary Economics* **38**(2), 215–220.

Hartley, J. S. (2021), 'Friedman's plucking model: New international evidence from Maddison Project data', *Economics Letters* **199**, 109724.

Hazell, J., Herreno, J., Nakamura, E. & Steinsson, J. (2022), 'The slope of the Phillips curve: evidence from US states', *The Quarterly Journal of Economics* **137**(3), 1299–1344.

Jackson, L. E., Kliesen, K. L. & Owyang, M. T. (2019), 'The nonlinear effects of uncertainty shocks', *Studies in Nonlinear Dynamics & Econometrics* **24**(4), 20190024.

Jordà, Ò. (2005), 'Estimation and inference of impulse responses by local projections', *American Economic Review* **95**(1), 161–182.

Khan, M. S. & Ssnhadji, A. S. (2001), 'Threshold effects in the relationship between inflation and growth', *IMF Staff papers* **48**(1), 1–21.

Kilian, L. & Kim, Y. J. (2011), 'How reliable are local projection estimators of impulse responses?', *Review of Economics and Statistics* **93**(4), 1460–1466.

Kim, C.-J. & Nelson, C. R. (1999), 'Friedman's plucking model of business fluctuations: tests and estimates of permanent and transitory components', *Journal of Money, Credit and Banking* pp. 317–334.

Kohlscheen, E., Moessner, R. & Rees, D. (2023), 'The shape of business cycles: a cross-country analysis of friedman's plucking theory', *BIS Working Papers* **1076**.

Kumar, A. & Orrenius, P. M. (2016), 'A closer look at the Phillips curve using state-level data', *Journal of Macroeconomics* **47**, 84–102.

Montiel Olea, J. L. & Plagborg-Møller, M. (2021), 'Local projection inference is simpler and more robust than you think', *Econometrica* **89**(4), 1789–1823.

Munir, Q., Mansur, K. & Furuoka, F. (2009), 'Inflation and economic growth in Malaysia: A threshold regression approach', *ASEAN Economic Bulletin* pp. 180–193.

Okun, A. M. (1962), 'Potential gnp: Its measurement and significance', *Proceedings of the Business and Economic Statistics Section of the American Statistical Association* pp. 98–104.

Plagborg-Møller, M. & Wolf, C. K. (2021), 'Local projections and VARs estimate the same impulse responses', *Econometrica* **89**(2), 955–980.

Rahman, S. & Serletis, A. (2010), 'The asymmetric effects of oil price and monetary policy shocks: A nonlinear VAR approach', *Energy Economics* **32**(6), 1460–1466.

Ruberl, H., Ball, M., Lucas, L. & Williamson, T. (2021), 'Estimating the NAIRU in Australia', **2021-01**.

Suah, J. L. (2023), 'An empirical take on the 'plucking' model', *BNM Working Papers* **WP3/2023**.

Tang, C. F. & Lean, H. H. (2007), 'Is the Phillips curve stable for Malaysia? new empirical evidence', *Malaysian Journal of Economic Studies* **44**(2), 95–105.

Taylor, M. P. & Davradakis, E. (2006), 'Interest rate setting and inflation targeting: Evidence of a nonlinear Taylor rule for the United Kingdom', *Studies in Nonlinear Dynamics & Econometrics* **10**(4).

Wellmann, S. (2023), 'The phillips curve in the euro area: New evidence using country-level data', *University of Tübingen Working Papers in Business and Economics* **156**.

Appendix

A Additional analysis on the plucking properties

Figure A.1

(a) Average Change in the Unemployment Rate During Contraction Against Subsequent Expansion: By Economy-Phase Pairs

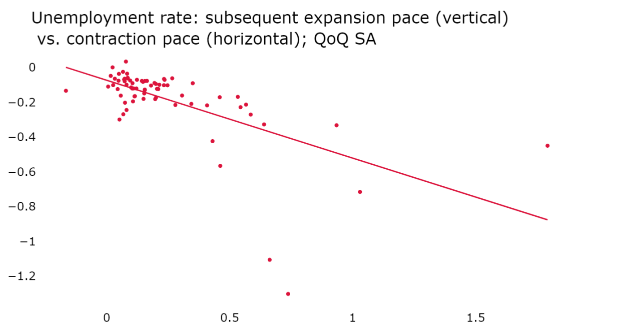

(b) Average Change in the Unemployment Rate During Expansion Against Subsequent Contraction: By Economy-Phase Pairs

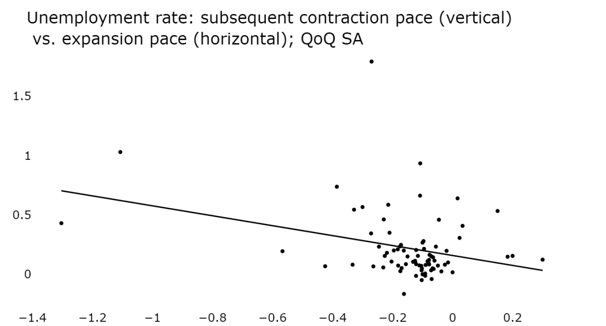

Subfigure (a) shows a scatterplot of the average quarterly change in the u-rate during a contraction (horizontal axis) against that of during the subsequent expansion (vertical axis). Subfigure (b) shows the reverse: the average quarterly change in the u-rate during an expansion (horizontal axis) against that of during the subsequent contraction (vertical axis). In both subfigures, each observation represent unique expansion-contraction, and contraction-expansion pairs. A version with outliers removed is available in figure A.3.

Figure A.2

(a) Amplitude of U-Rate During Subsequent Expansion Against Current Contraction; Outliers Removed

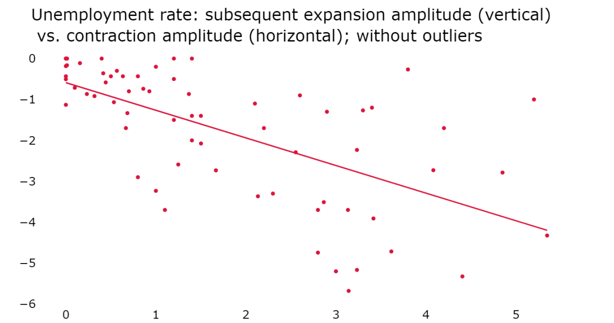

(b) Amplitude of U-Rate During Subsequent Contraction Against Current Expansion; Outliers Removed

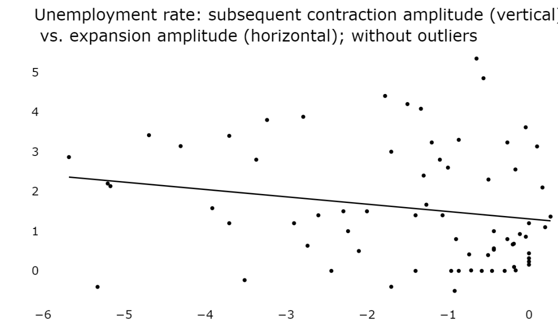

Subfigure (a) shows a scatterplot of the change in the u-rate during a contraction (horizontal axis) against that of during the subsequent expansion (vertical axis). Subfigure (b) shows the reverse: the change in the u-rate during an expansion (horizontal axis) against that of during the subsequent contraction (vertical axis). The highest and lowest 2.5% of observations along both axes are dropped. In both subfigures, each observation represents unique expansion-contraction (or contraction-expansion). In both subfigures, each observation represent unique expansion-contraction, and contraction-expansion pairs. A version with all observations included is available in section 1 in figure 1. The top 2.5%, and bottom 2.5% of observations along both axes are removed.

Figure A.3

(a) Average Change in the Unemployment Rate During Subsequent Expansion Against Current Contraction by Economy-and-Phase; Outliers Removed

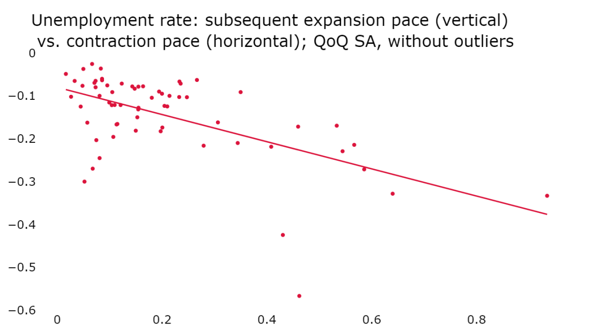

(b) Average Change in the Unemployment Rate During Subsequent Contraction Against Subsequent Expansion by Economy-and-Phase; Outliers Removed

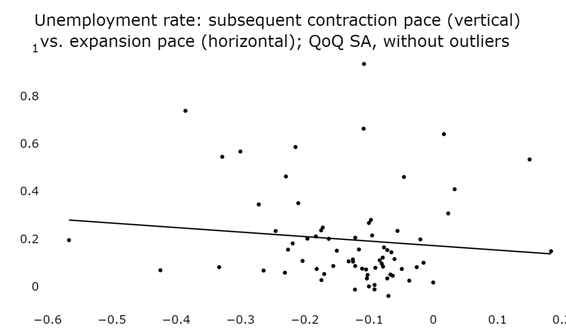

Subfigure (a) shows a scatterplot of the average quarterly change in the u-rate during a contraction (horizontal axis) against that of during the subsequent expansion (vertical axis). Subfigure (b) shows the reverse: the average quarterly change in the u-rate during an expansion (horizontal axis) against that of during the subsequent contraction (vertical axis). The highest and lowest 2.5% of observations along both axes are dropped. In both subfigures, each observation represents unique expansion-contraction (or contraction-expansion). A version with all observations included is available in figure A.1.

B Additional analysis on estimating the unemployment rate floors

Table B.1: Pre-COVID vintage (up to 2019 Q4): Tolerance thresholds (multiplier of the standard deviation of the u-rate)

Economy	X	**Percentage points**
United States	1.05	1.72
Germany	1.2	2.55
France	0.45	0.81
Italy	0.4	0.67
United Kingdom	0.3	0.71
Japan	0.25	0.31
Australia	0.8	1.39
Singapore	2.2	1.59
Korea	0.32	0.16
Hong Kong SAR	0.15	0.25
Mexico	0.5	0.42
Chile	2	3.31
Brazil	1.4	3.48
China	1	0.15
India	0.8	0.58
Malaysia	0.2	0.04
Thailand	0.1	0.07
Indonesia	0.7	0.76
Philippines	1.5	3.16

Figure B.1

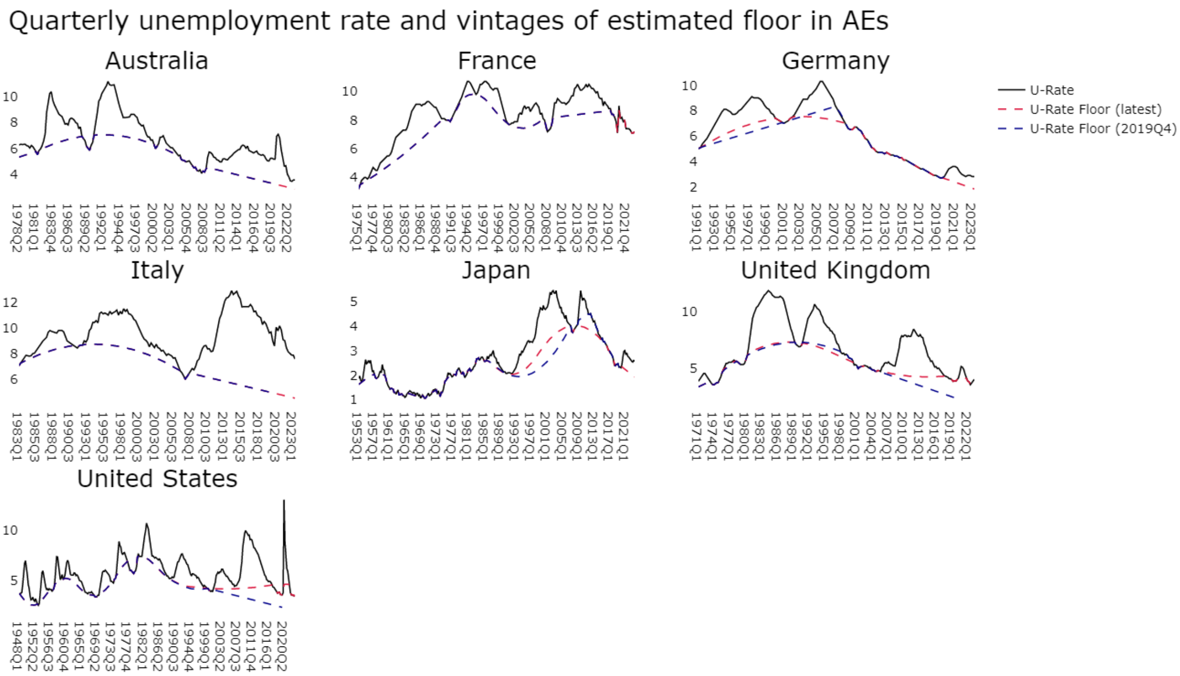

Figure B.2

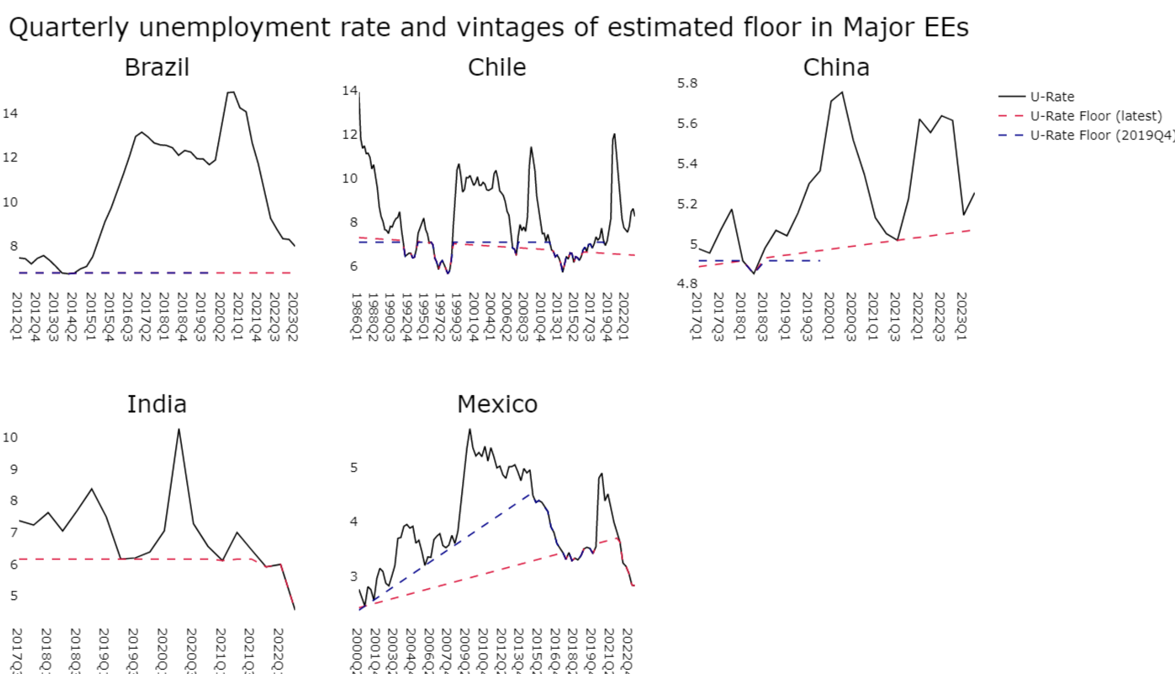

Figure B.3

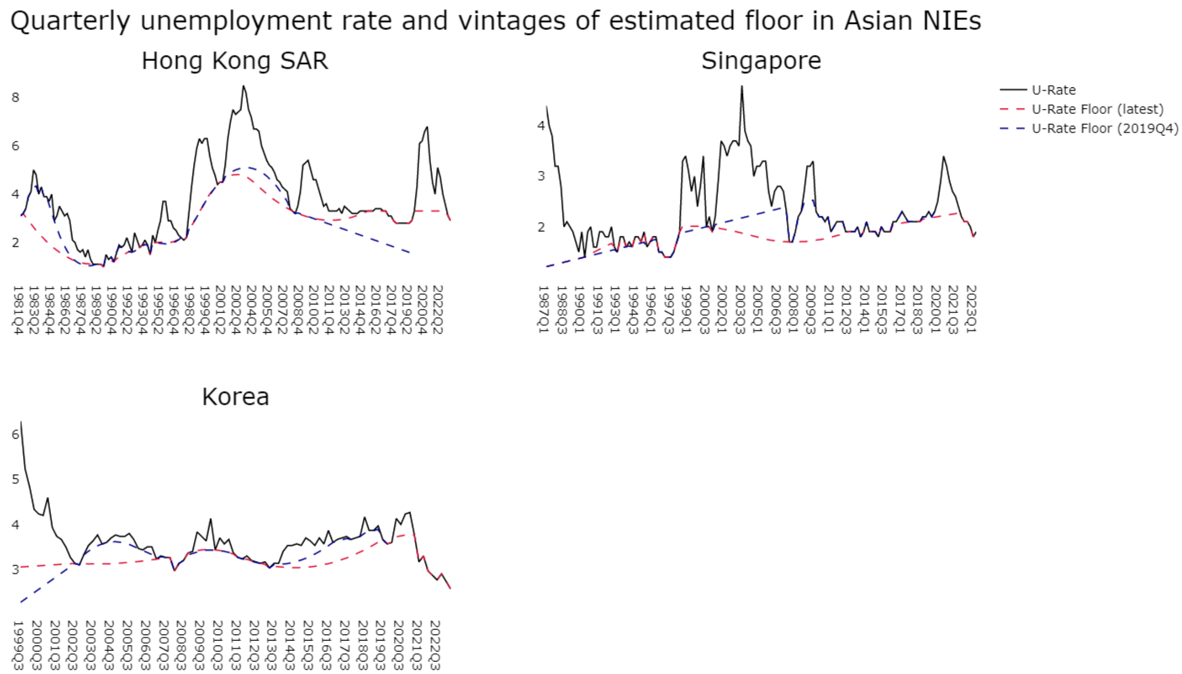

Figure B.4

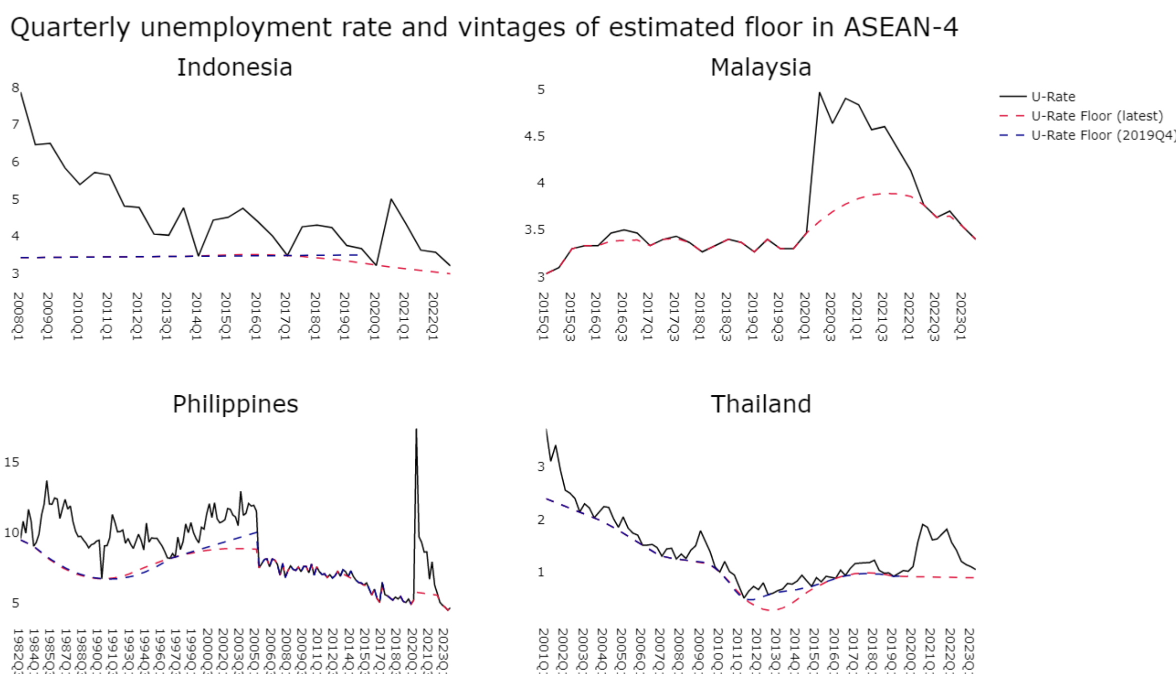

C Additional findings on the propagation of macroeconomic shocks

Figure C.1: All economies

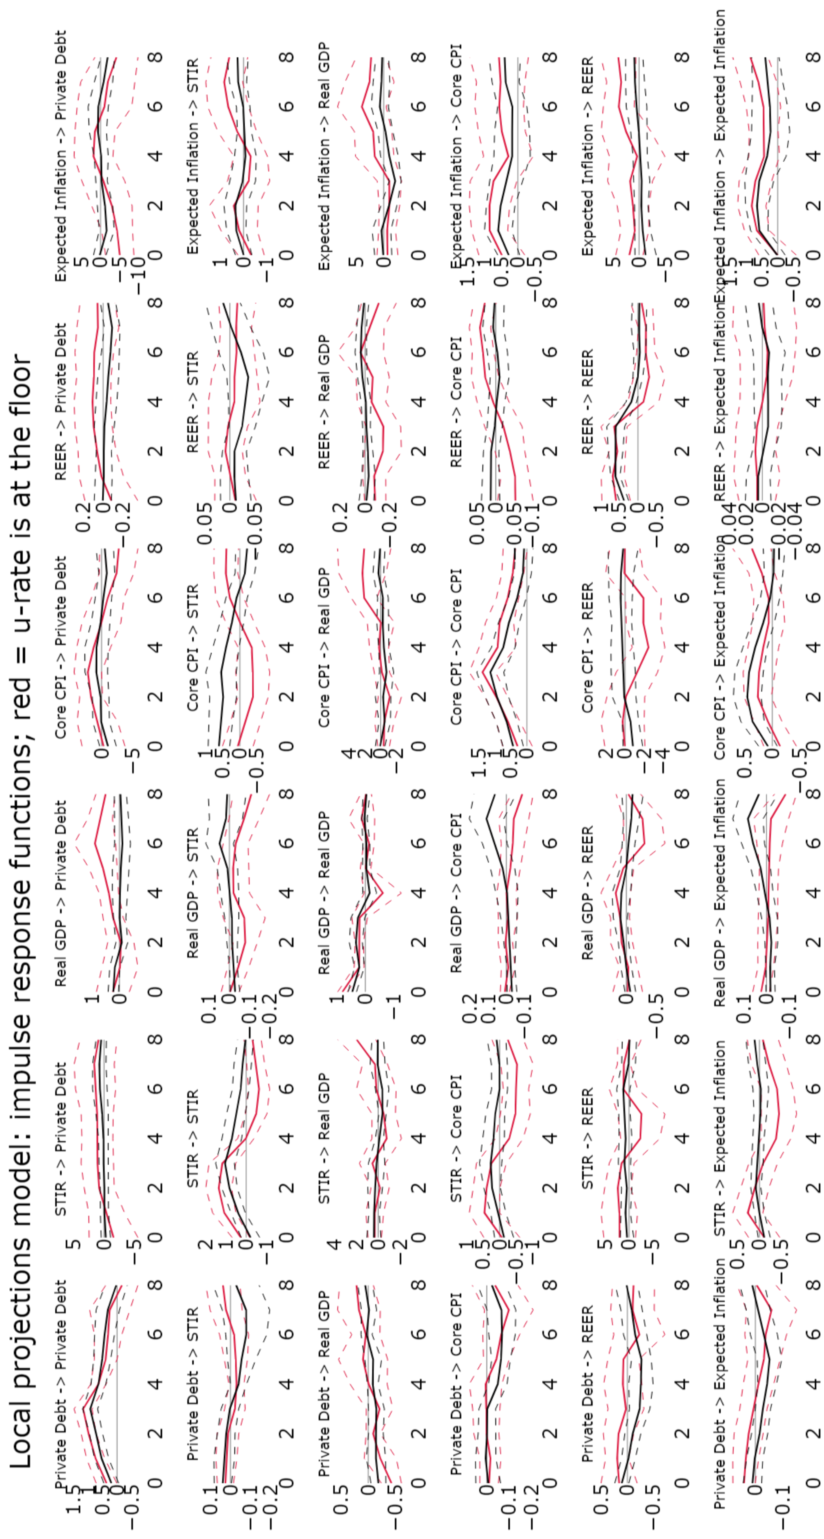

Subfigures show the estimated IRFs of the corresponding variables from various shocks labelled in the respective titles, which follow the convention "*shock variable* → *response variable*". The red lines indicate the IRFs for when the u-rate is at the floor, and black lines for when the u-rate is above the floor. The dotted lines correspond to the respective 95% confidence intervals, estimated with heteroscedastic-robust standard errors. Estimation, and visualisation, are implemented with the `localprojections` package in Python. The main text shows the same IRF estimates grouped by shock variable.

Figure C.2: Excluding Brazil, India and Malaysia

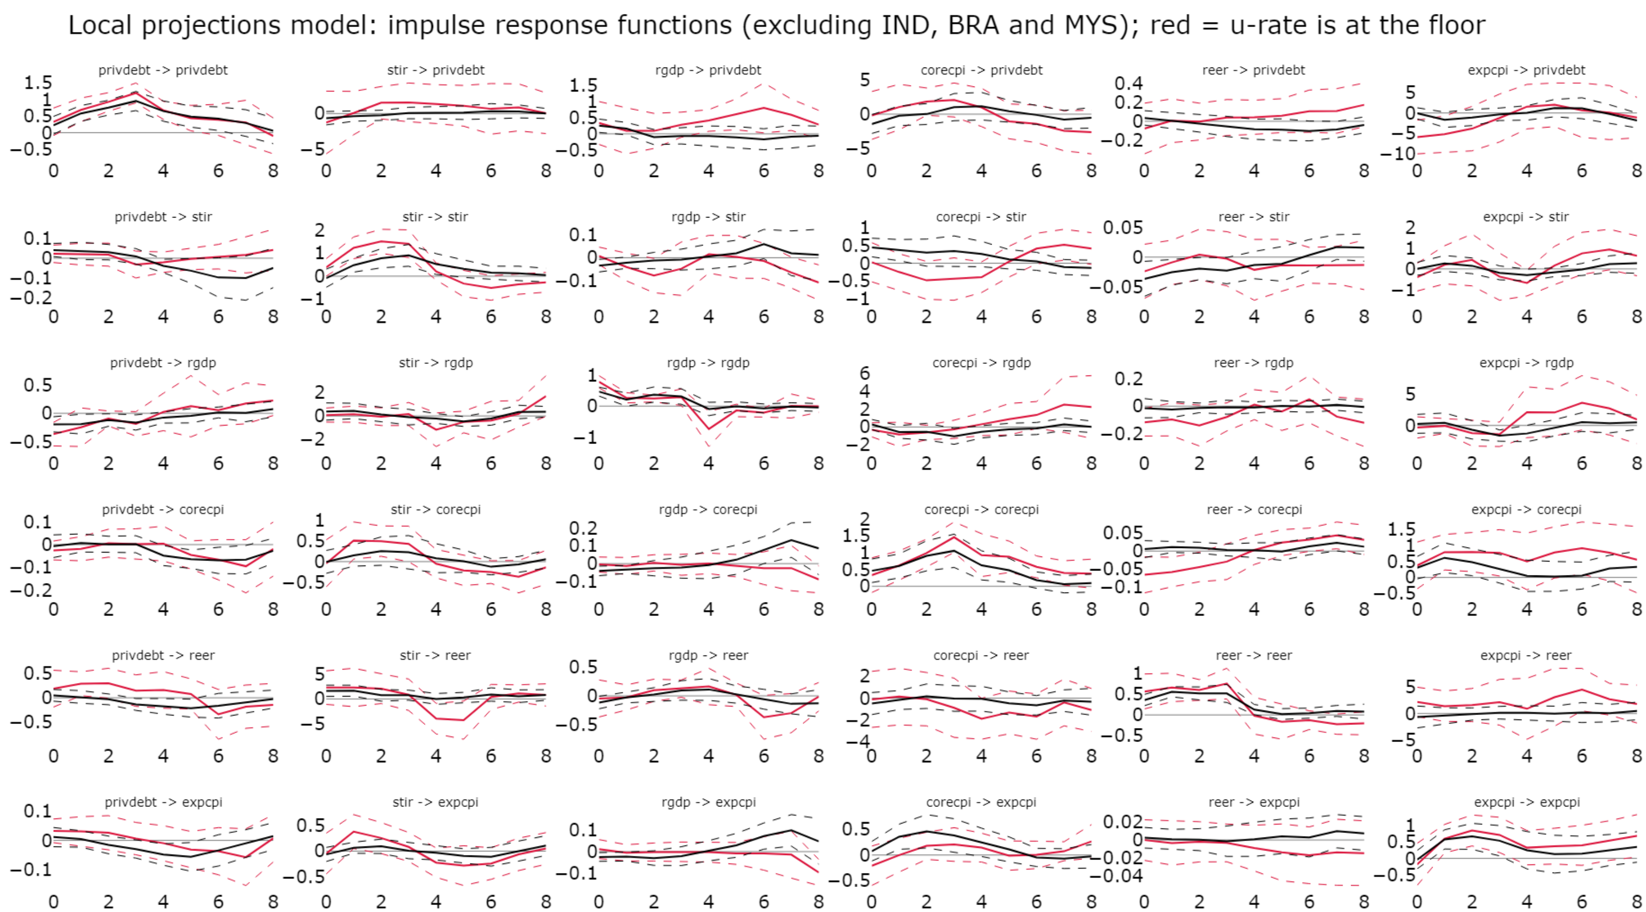

Estimates here further excluded Brazil, India and Malaysia. Estimated IRFs are similar to the full version, shown in figure C.1. Subfigures show the estimated IRFs of the corresponding variables from various shocks labelled in the respective titles, which follow the convention "*shock variable* → *response variable*". The red lines indicate the IRFs for when the u-rate is at the floor, and black lines for when the u-rate is above the floor. The dotted lines correspond to the respective 95% confidence intervals, estimated with heteroscedastic-robust standard errors. Estimation, and visualisation, are implemented with the `localprojections` package in Python.

Figure C.3: Advanced economies only

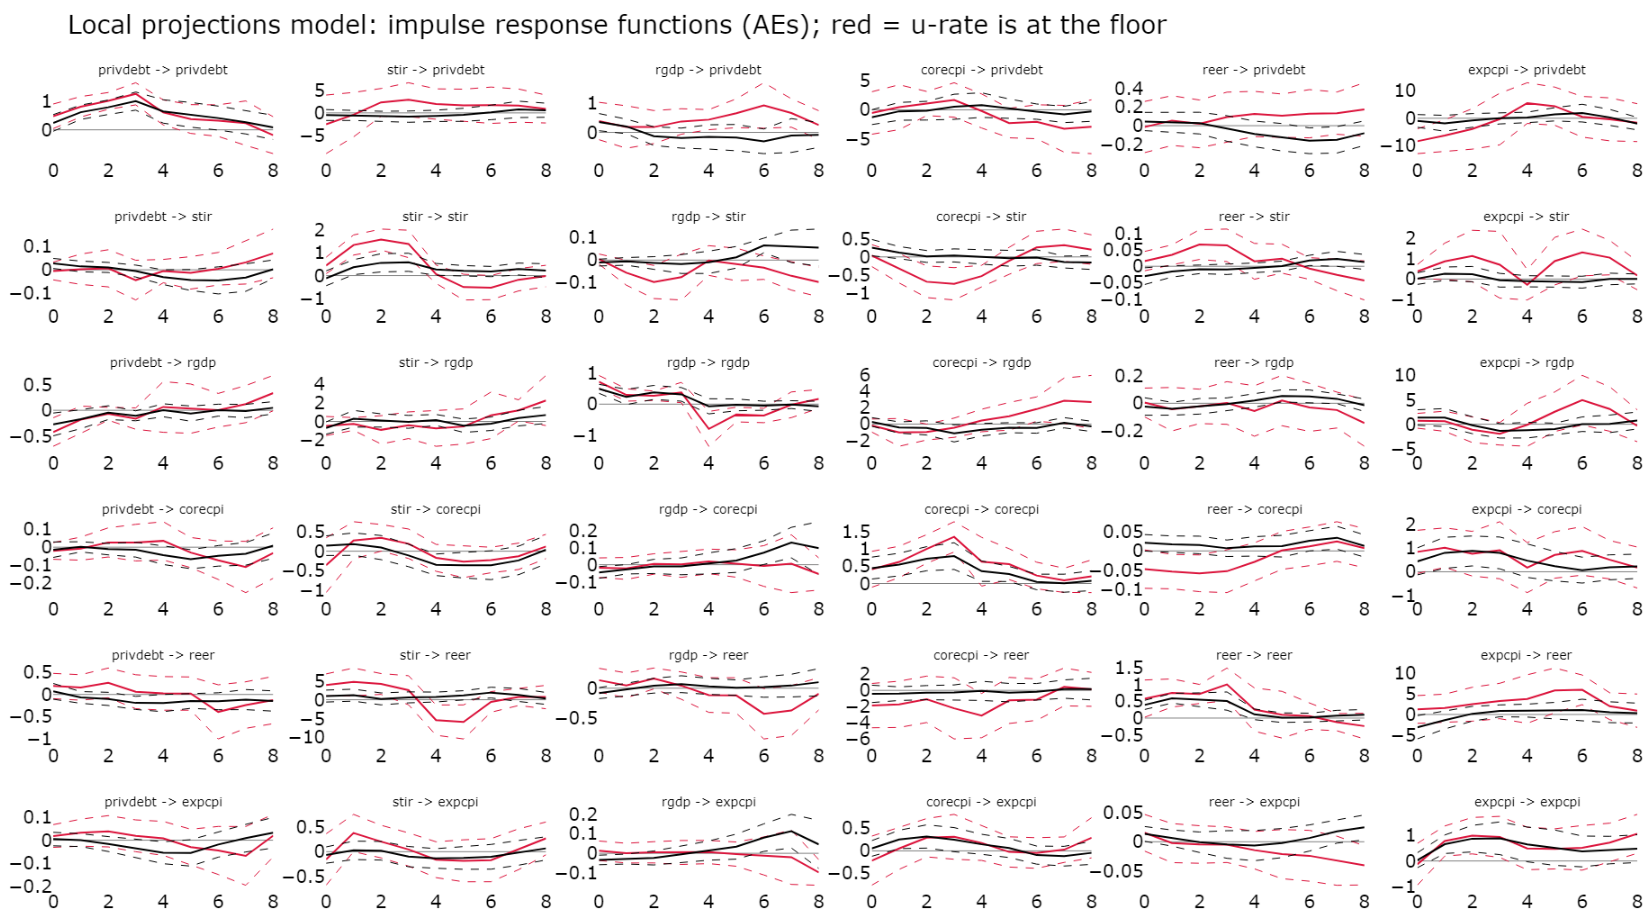

Estimates here contain only advanced economies. Estimated IRFs are similar to the full version, shown in figure C.1. Subfigures show the estimated IRFs of the corresponding variables from various shocks labelled in the respective titles, which follow the convention "*shock variable* → *response variable*". The red lines indicate the IRFs for when the u-rate is at the floor, and black lines for when the u-rate is above the floor. The dotted lines correspond to the respective 95% confidence intervals, estimated with heteroscedastic-robust standard errors. Estimation, and visualisation, are implemented with the `localprojections` package in Python.

Figure C.4: Emerging economies only

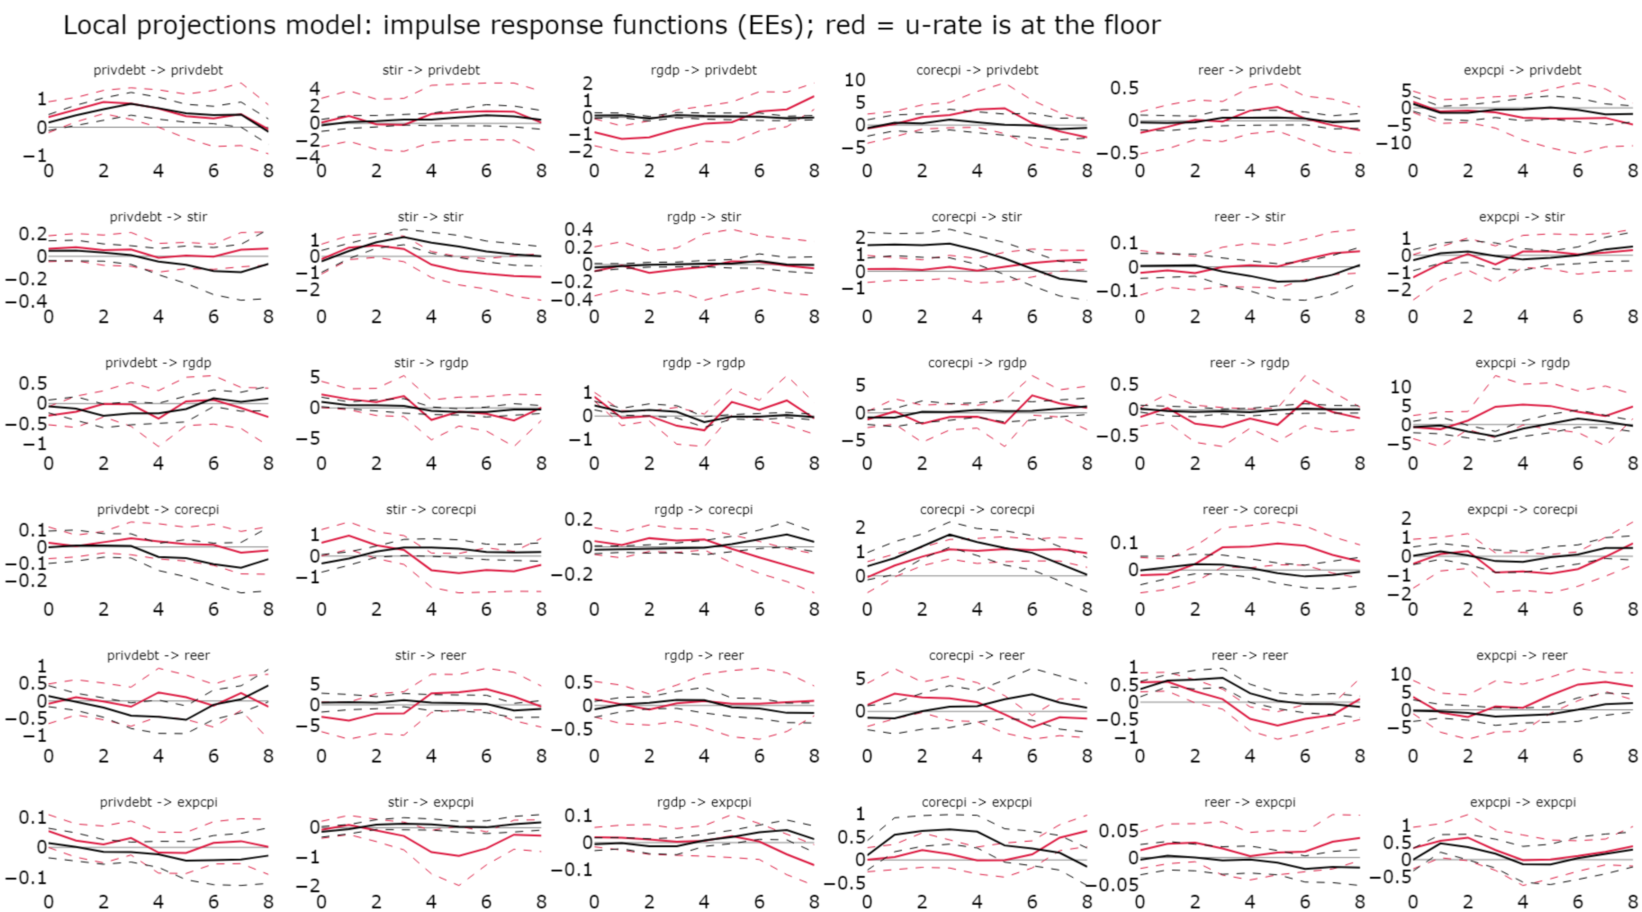

Estimates here contain only emerging economies. Estimated IRFs are similar to the full version, shown in figure C.1. Subfigures show the estimated IRFs of the corresponding variables from various shocks labelled in the respective titles, which follow the convention "*shock variable* → *response variable*". The red lines indicate the IRFs for when the u-rate is at the floor, and black lines for when the u-rate is above the floor. The dotted lines correspond to the respective 95% confidence intervals, estimated with heteroscedastic-robust standard errors. Estimation, and visualisation, are implemented with the `localprojections` package in Python.

Previous volumes in this series

1158 December	Financial development and the effectiveness of macroprudential and capital flow management measures	Yusuf Soner Başkaya, Ilhyock Shim, Philip Turner
1157 December	Fintech vs bank credit: How do they react to monetary policy?	Giulio Cornelli, Fiorella De Fiore, Leonardo Gambacorta and Cristina Manea
1156 December	Monetary policy frameworks away from the ELB	Fiorella De Fiore, Benoit Mojon, Daniel Rees, Damiano Sandri
1155 December	Monetary Tightening, Inflation Drivers and Financial Stress	Frederic Boissay, Fabrice Collard, Cristina Manea, Adam Shapiro
1154 December	Foreign investor feedback trading in an emerging financial market	Ingomar Krohn, Vladyslav Sushko, Witit Synsatayakul
1153 December	Foreign institutional investors, monetary policy, and reaching for yield	Ahmed Ahmed, Boris Hofmann, Martin Schmitz
1152 December	The Heterogeneous Impact of Inflation on Households' Balance Sheets	Clodomiro Ferreira, José Miguel Leiva, Galo Nuño, Álvaro Ortiz, Tomasa Rodrigo, and Sirenia Vazquez
1151 November	The financial origins of regional inequality	Anne Beck and Sebastian Doerr
1150 November	Markups and the asymmetric pass-through of cost push shocks	Enisse Kharroubi, Renée Spigt, Deniz Igan, Koji Takahashi, Egon Zakrajšek
1149 November	Housing affordability: a new data set	Nina Biljanovska, Chenxu Fu, Deniz Igan
1148 November	Firm heterogeneity, capital misallocation and optimal monetary policy	Beatriz González, Galo Nuño, Dominik Thaler, Silvia Albrizio
1147 November	Central Band Digital Currency and Privacy: A Randomized Survey Experiment	Syngjoo Choi, Bongseob Kim, Young-Sik Kim, Ohik Kwon
1146 November	On par: A Money View of stablecoins	Iñaki Aldasoro, Perry Mehrling, Daniel H. Neilson
1145 November	Dollar and Government Bond Liquidity: Evidence from Korea	Jieun Lee

All volumes are available on our website www.bis.org.