INTERNATIONAL MONETARY FUND

Measures of Macroeconomic Shocks and Uncertainty for Asia-Pacific Economies

Prepared by Iris Claus and Leo Krippner

WP/26/174

2026
AUG

WORKING PAPER

 WP/26/174

IMF Working Paper
Institute for Capacity Development

Measures of Macroeconomic Shocks and Uncertainty for Asia-Pacific Economies*

Prepared by Iris Claus† and Leo Krippner‡

Authorized for distribution by Paul Cashin
August 2026

ABSTRACT: We develop standardized time-series measures of shocks and uncertainty for output growth and inflation across 14 Asia-Pacific economies using data from Consensus Economics survey of professional economic forecasters. They are based on changes in mean forecasts and standard deviations, adjusted for intra-year patterns that arise from the annual-average percent changes basis of the data. As an example of how our shock measures may be used, we apply the vector autoregression method of sign restrictions to decompose output growth and inflation shocks into fundamental demand and supply shocks. Those shocks and our uncertainty measures align well with major economic events such as the Asian Financial Crisis, the COVID-19 pandemic, and geopolitical conflicts. We show that our uncertainty measures most closely reflect the concept of macroeconomic uncertainty, with apparent differences to established economic policy uncertainty measures across all comparable economies, but a close relationship with two well-known macroeconomic uncertainty measures produced only for the United States. Hence, extending the measures of macroeconomic uncertainty to all 14 economies in our analysis, along with the shocks for those economies, creates a valuable dataset for economic policy setting and empirical research.

RECOMMENDED CITATION: Claus, I. and Krippner L. 2026. “Measures of Macroeconomic Shocks and Uncertainty for Asia-Pacific Economies.” IMF Working Paper WP/26/174, International Monetary Fund, Washington, DC. https://doi.org/10.5089/9798229055550.001

JEL Classification Numbers:	C1 econometric and statistical methods and methodology: general E3 prices, business fluctuations, and cycles
Keywords:	macroeconomic shocks; uncertainty; Asia-Pacific; Consensus Economics survey of professional economic forecasters; demand and supply shocks; vector autoregression; sign restrictions
Authors’ email addresses:	iclaus@IMF.org leokrippner@smu.edu.sg

* The authors would like to thank Diego Arturo, Rodriguez Guzman, Serdar Kabaca, Yoko Konishi, Andrea Pescatori, Mikhail Pranovich, Pau Rabanal, Chris Redl, Kwanho Shin, Oum Sothea, Tugrul Vehbi, Wing Thye Woo, and Yizhi Xu for valuable comments and suggestions.

† IMF – Singapore Regional Training Institute.

‡ Sim Kee Boon Institute of Singapore Management University, University of Waikato, and Centre for Applied Macroeconomic Analysis of the Australian National University.

WORKING PAPERS

Measures of Macroeconomic Shocks and Uncertainty for Asia-Pacific Economies

Prepared by Iris Claus and Leo Krippner

1 Introduction

In this article, we create measures of shocks and uncertainty for output growth and inflation for two northern Pacific-rim countries, i.e. the United States and Canada, and 12 Asia-Pacific economies, i.e. Australia, Japan, China (i.e. Chinese mainland), Hong Kong Special Administrative Region (SAR), India, Indonesia, Malaysia, New Zealand, Singapore, Korea, Taiwan Province of China, and Thailand from Consensus Economics survey of professional economic forecasters. To illustrate how our dataset may be used, we then present results from two empirical exercises, i.e. inferring fundamental demand and supply shocks from our shock series for Gross Domestic Product growth and Consumer Price Index inflation (hereafter GDP and CPI), and comparing our GDP and CPI uncertainty measures with well-known measures for economic policy uncertainty and macroeconomic and financial uncertainty.

One motivation for creating our GDP and CPI shock and uncertainty dataset is to help with monitoring and forecasting of macroeconomic developments, particularly by institutions involved with public policy. Examples are central banks and finance ministries which, respectively, use regular in-house forecasting rounds to inform monetary policy and fiscal policy decisions. Comprehensive, standardized, and regularly updated quantitative data on GDP and CPI shocks provides timely indications of changes to the central outlook for their own economy and trading partners, while GDP and CPI uncertainty indicates potential scenarios and additional channels (discussed in the following paragraph) that can directly affect GDP, all of which are important for considering policy responses and/or contingency planning.

A second motivation is to extend datasets available for applied research, and particularly for Asia-Pacific economies where relevant data may not already be available. Regarding uncertainty, there is a large and well-established literature that it influences economic activity through a range of channels, such as the value of waiting to invest and precautionary saving.[1] But empirical research on the potential effects of uncertainty first requires uncertainty to be quantified with appropriate proxies.

Cascaldi-Garcia et al. (2023) provides a comprehensive review of four categories of uncertainty measurement that have been developed in the literature, albeit mainly from a US perspective, i.e. newspaper-based, survey-based, econometric-based, and market-based measures. One well-known example within the first category are the Baker, Bloom, and Davis (BBD, 2016) economic policy uncertainty measures, and series of that type are available for nine Asia-Pacific economies, i.e. Australia, China, Hong Kong SAR, India, Japan, New Zealand, Pakistan, Singapore, and Korea, and of course for the United States and Canada. Two well-known examples in the econometric category, but only for the United States, are the macroeconomic and financial uncertainty measures from Jurado, Ludvigson, and Ng (JLN, 2015) and Ludvigson, Ma, and Ng (LMN, 2021).[2]

[1] Fernández-Villaverde and Guerrón-Quintana (2020) provides a thorough overview, and table 1 in Harris, Weenink, Grubb, and Turnip (2025) contains a useful summary of uncertainty channels to the real economy.

[2] The BBD data for the United States and for other countries covered by the BBD authors are available at www.policyuncertainty.com. The site also includes newspaper-based uncertainty estimates that follow the BBD method, but are produced by other authors. The JLN and LMN macroeconomic and financial measures of uncertainty can be downloaded from www.sydneyludvigson.com/macro-and-financial-uncertainty-indexes.

The uncertainty measures we develop are survey-based and, aside from Pakistan where CE data are not available, we can obtain uncertainty measures for all of the economies just mentioned and further expand the dataset to include Malaysia, Indonesia, Taiwan Province of China, and Thailand. Indeed, one reason for basing our uncertainty and shock measures on CE data is that further expansion can be produced at a relatively low cost, because the CE dataset we use already has a long history available including for other economies and updates are monthly (although access is by subscription). In this regard, we are not the first to use CE data to gauge uncertainty; Denis and Kannan (2013), Greig, Rice, Vehbi, and Wong (2018), Ozturk and Sheng (2018), and Harris, Weenink, Grubb, and Turnip (2025) are examples we are aware of in the literature.[3] However, as we return to further below, a contribution we make in this article is to explicitly address the intra-year nature of the CE time series data, which leads us to a distinctly different method from that of the other authors.

Regarding economic shocks, they are a central aspect of macroeconomics and so the related literature is far too large to attempt a summary. We refer readers to Ramey (2016) for the key themes, and note that research on the topic remains active, e.g. a recent example is Dery and Serletis (2026). For the purposes of our article, we discuss the main concepts from the perspective of vector autoregression models (VARs), because they are a familiar and widely used tool for empirically identifying macroeconomic shocks, and also because we can usefully discuss parallels and contrasts with our demand and supply shock application.

Hence, an estimated reduced-form VAR provides the anticipated components of the included macroeconomic variables, and the VAR residuals provide measures of the unanticipated components, i.e. shocks to the variables. Structural VAR analysis informed by economic principles then imposes restrictions on how the VAR residuals should in principle comove in response to fundamental shocks within the economy. For example, a positive demand shock should produce simultaneous positive shocks to GDP and CPI, and a positive supply shock should produce a simultaneous positive shock to GDP and a negative shock to CPI. In reverse then, a time series of GDP and CPI shocks should provide the basis for identifying a time series of demand and supply shocks. This may be achieved with the method of sign restrictions, i.e. rotating output and inflation shocks from the reduced-form VAR to create and retain candidate structural VARs that meet the given demand and supply principles; see Fry and Pagan (2011) for a detailed discussion. Another method of identification is the more-recent development of proxy VARs, where directly observed data for economic events related to the VAR variables are used to help identify fundamental shocks in the VAR, e.g. using known oil shock events to help identify supply shocks. Bruns and Lütkepohl (2026) provides a thorough review of that literature.

Once a VAR is estimated and identified, the impulse response functions associated with the fundamental shocks in the VAR indicate how those shocks propagate into

[3] We also follow the terminology of "uncertainty" used in those references, although those and our measures are actually forecast disagreement, i.e. dispersion in point predictions. More specialized literature distinguishes disagreement from survey-based uncertainty based on subjective probability distributions from each forecaster; e.g. see D'Amico and Orphanides (2008) and Ploenzke, Rich, and Tracy (2015). Where both are available, e.g. for the United States and European Surveys of Professional Forecasters, the two measures can differ materially, e.g. see Abel, Rich, Song, and Tracy (2016), and D'Amico and Orphanides (2008) notes a correlation of 0.4. CE survey results do not include probability distributions.

future outcomes of the VAR variables. Identifying shocks in real time helps to inform forecasts and potential policy choices, e.g. using monetary policy to offset the effects of demand shocks on output and inflation, but "looking through" supply shocks.

Our approach obtains GDP and CPI shocks directly, based on changes to forecasts between monthly updates, which avoids potential issues from publication lags and omitted variables that can arise with VAR estimation and identification. CE data have essentially no publication lags because they are forecasts collected simultaneously on the survey date for each month and released three days later. And there are no omitted variables in principle, because the GDP and CPI shocks are ex-ante data, i.e. based on forecasts that should consider all influences, rather than the ex-post data used to estimate VARs. The VAR sign restrictions method can then be used with our GDP and CPI shocks to infer demand and supply shocks, as in our first empirical exercise. Regarding the proxy VAR literature, the parallel of our approach is using observed data directly as measures of shocks. However, we have a regular time series of shocks, rather than occasional events.[4]

To create our shock and uncertainty measures, we use the summary results calculated and reported by CE each month for their surveys of GDP growth and CPI inflation forecasts (hereafter abbreviated to GDP and CPI) from professional economic forecasters. Specifically, our shock estimates are based on changes in the reported means of the survey results for GDP and CPI, and our uncertainty estimates are based on the reported standard deviations. We provide a complete description of the Consensus Economics dataset in section 2, but the most important aspects for the discussion here are that the length of the data is 30-35 years, and the forecasts collected and reported each month are annual-average percent changes (AAPCs) for the current calendar year and the following calendar year.

The key aspect in creating our shock and uncertainty dataset from the CE data is to estimate and adjust for what we call a "Year Advancement Pattern" (hereafter YAP), which is a particular type of seasonality inherent from progressive monthly updates of calendar-year AAPC forecasts. For example, mean changes and standard deviations for the current year should in principle decline with each monthly update because fewer remaining months in the year remain to be forecasted. Data for the following year should also have a regular pattern as the forecasting horizon shortens. Using an appropriate YAP estimate allows the data for the current and following years to be standardized to the same annual basis, and we can then combine the results for both years into a single measure.

As detailed in section 3, for the CE standard deviation series, we estimate and evaluate two candidate YAPs, i.e. a fully flexible dummy variable for each month that we label MYAP, and a YAP constrained to be linear across each year that we label LYAP. We find for CE standard deviation data across all economies that the MYAP is most appropriate for the current year, and the LYAP for the following year. Hence, for each economy, we apply those respective adjustments to each year's standard deviation series and then combine the results into our measures of uncertainty for that economy,

[4]We also note that our measurement of GDP and CPI shocks aligns well with the measurement of monetary policy surprises via unanticipated target and path surprises for policy interest rates, e.g. see Gürkaynak, Sack, and Swanson (2005), Gürkaynak, Sack, and Swanson (2007), and Swanson (2021), with the latter additionally including large-scale asset purchases. We do not include a monetary policy variable in our analysis, for reasons we discuss at the end of our conclusion.

i.e. a GDP and a CPI uncertainty series. Our checks confirm that no monthly pattern remains in the uncertainty measures we generate for each economy. Prior uses of CE data to gauge uncertainty, as we referred to earlier, are all based on pro-rated weighted-averages of unadjusted current- and following-year data. We show in section 3.3 that this method in principle retains a YAP, which makes data produced with that method less suitable as a quantitative measure of uncertainty, essentially because its level has a dependence on the month of the year.

For CE mean changes, as detailed in section 4, we find that the LYAP is most appropriate for both the current and following years. Hence, for each economy, we apply LYAP adjustments to the mean-change for both years, and then combine those results into our GDP and CPI shock measures for that economy.

For the empirical application in section 5.1, we use our estimates of GDP and CPI shocks for each economy to produce estimates of supply and demand shocks, based on the sign restrictions method discussed earlier. In section 5.2 we compare our uncertainty measures to the BBD economic policy uncertainty measures for the economies where both are available and our US measure to the JLN and LMN macroeconomic and financial uncertainty measures.

The article proceeds as already indicated above, i.e. section 2 introduces the CE dataset, section 3 develops our method for creating uncertainty measures, discusses the results, and compares our method to other CE-based measures, section 4 develops our method for creating shock measures and discusses the results, and section 5 contains our empirical exercise of extracting demand and supply shocks and the comparison of our and the BBD, JLN, and LMN measures of uncertainty. We conclude in section 6 with a brief summary and a discussion of potential extensions of our work.

2 Overview of Consensus Economics dataset

The shock and uncertainty dataset we develop is based on the monthly surveys of professional forecasters undertaken by Consensus Economics, published as Consensus Forecasts for the United States and Canada, and as Asia Pacific Consensus Forecasts for the other economies included in our analysis, with Japan appearing in both. The surveys are conducted on a given day early in the month, typically Monday of the first week to 1994 and Monday of the second week thereafter, and forecasts are collected for a range of macroeconomic variables.

Table 1 provides summary information for the economies we include in our dataset. The United States, Canada, and Japan have the longest time series, from October 1989. We omit the first four observations, October 1989 to January 1990, because the January observation appears to relate to the previous calendar year instead of switching to current calendar year, i.e. 1989 to 1990. All other January observations immediately switch to the current calendar year, as we discuss further below. Aside from Australian data, which start in November 1990, the data for the other economies are from December 1994. CE data for the Philippines are only available from April 2009, and for Vietnam from February 2020. We consider those time series too short at this stage for the type of analysis we apply, and so we have not included them in our dataset.

The forecasts collected for both GDP and CPI are annual-average percent changes

(AAPCs) for the current calendar year and the following calendar year, which we label in table 1 as GDP0, GDP1, CPI0, and CPI1. The exception is India, where forecasts are for the current and following fiscal years to the end of March. The two columns below each of the four categories give the minimum and maximum number of responses to each category.[5] The counts vary across economies, reflecting differences in the pool of forecasting institutions available to survey, and within each economy the maximums and minimums indicate that the number of respondents varies across time. The minimum is 9, for India and Indonesia, and the minimum number of responses for all other economies are between 10 and 20.

Table 1: Summary information for CE data on the 14 economies

			Number of respondents							
			GDP0		GDP1		CPI0		CPI1	
Item	Start	T	Min	Max	Min	Max	Min	Max	Min	Max
United States	2/90	435	19	33	19	33	19	33	19	33
Japan	2/90	435	18	27	12	27	15	26	12	25
Canada	2/90	435	11	20	11	20	11	20	11	20
Australia	11/90	426	13	24	13	24	13	24	13	24
New Zealand	12/94	377	11	17	11	17	10	17	10	17
Malaysia	12/94	377	11	21	11	21	11	21	11	21
Singapore	12/94	377	12	20	12	20	12	20	11	20
China	12/94	377	13	29	12	28	13	29	12	28
Hong Kong SAR	12/94	377	12	22	12	21	12	22	8	21
India	12/94	377	9	23	9	23	9	23	9	23
Indonesia	12/94	377	9	24	9	24	9	24	9	24
Korea	12/94	377	13	22	13	21	13	22	13	21
Taiwan Province of China	12/94	377	12	21	12	20	11	21	11	20
Thailand	12/94	377	11	27	11	27	11	27	11	27

Source: Authors' calculations

Consensus Economics then calculate summary statistics from the collected survey results. We use the CE-calculated means and standard deviations for GDP and CPI, for the current and following year, as the data for our analysis that we detail in sections 3 and 4. Hence, there are eight data series in total, i.e. four mean results that we respectively label as GDP0mn, GDP1mn, CPI0mn, and CPI1mn, and four standard deviation results that we respectively label as GDP0sd, GDP1sd, CPI0sd, and CPI1sd.

While the minimum number of responses for each economy, and even some of the maximums, are relatively small for typical mean and standard deviation calculations, we nevertheless consider the data to be usable for our purposes. One reason is that the forecasts are made by professionals, who tend to be better trained, more experienced, technically proficient, and well-informed on economic developments through

[5] The counts are only provided by CE from February 2005 for the United States, Canada, and Japan, and from April 2007 for the remaining economies. The counts for earlier years could be obtained from the unit record data, but we have not yet processed the dataset to that level.

their networks and data sources, relative to other groups in the economy, particularly compared to surveyed consumers/households. The averaging of even a small number of forecasts additionally confers the "wisdom of crowds". Empirically, Verbrugge and Zaman (2021) and Morley, Tian, and Wang (2026) are examples that establish the tangible benefits of professional forecasts relative to other groups.

Of course, macroeconomic forecasting has, to put it euphemistically, a high noise-to-signal ratio, so forecasts from professionals are really just "less worse" than competing forecasts. However, the second and perhaps more important reason we consider the CE data usable for our purposes is that we do not use the forecasts in their own right, but their changes to gauge shocks and their dispersion to gauge uncertainty. That is, from our perspective, the information value in the CE data depends not on potential accuracy, but on the much more reasonable proposition that individual professionals are making well-considered forecasts based on the information, judgement, opinions, experience that each has available to them. The "less worse" performance of the forecasts in aggregate, noted above, is evidence consistent with that proposition.

As an example of the eight data series for each economy, figure 1 plots them for the United States. We typically use the United States to illustrate the data and associated intermediate results as we proceed with our analysis to create our dataset; including those results for all economies would make the paper prohibitively lengthy, but we will provide any of those results on request. To illustrate the nature of AAPC forecasts for calendar years, we have also included in panels 1 and 3 of figure 1 realized calendar-year AAPCs calculated from GDP and CPI data. Hence, for example, GDP0mn at any point during a given year is the forecast for the AAPC GDP ultimately realized for that year. Table 2 lists the eight events that we have indicated in each of the panels of figure 1. Note that the CE survey date following an event is often in the month after the event.

Panels 1 and 3 of figure 1 illustrate two aspects associated with the GDP0mn and CPI0mn data, which are most evident for the 2020 year given the large forecast revisions at that time due to the COVID-19 pandemic. First, the GDP0mn and CPI0mn data have distinct steps as the survey month moves into a new year. Second, as months progress within a given year, the GDP0mn and CPI0mn data converge to their respective realized AAPCs for that year. This convergence reflects that AAPC forecasts for the current calendar year increasingly incorporate published data, leaving progressively less contribution from genuine forecasts for the remaining months of the year.

For the same reason, the GDP0sd and CPI0sd data show a clear pattern of decline. That is, the year-to-date realized data are fixed (to a close approximation, i.e. subject only to typically minor revisions), so forecasters should have no material disagreement about that component of the current year. Only potential outcomes for the remainder of the year should be subject to disagreement, so dispersion in the AAPC forecasts should decrease as the year advances.[6]

It is these "Year Advancement Patterns", or YAPs, that we seek to correct when developing our shock and uncertainty measures in sections 3 and 4. While they are very apparent in the current year, there may also be YAPs for the following year data,

[6]The decline is not necessarily to zero due to publication lags in the data, e.g. in the December survey, the US CPI is missing realized data for November, and December data would not yet be collected. For GDP in particular, there may also be differences in view on potential revisions of data already released.

i.e. GDP1mn, CPI1mn, GDP1sd, and CPI1sd, due to aspects such as the shortening forecast horizon and/or how following year forecasts are revised based on data published in the current year.

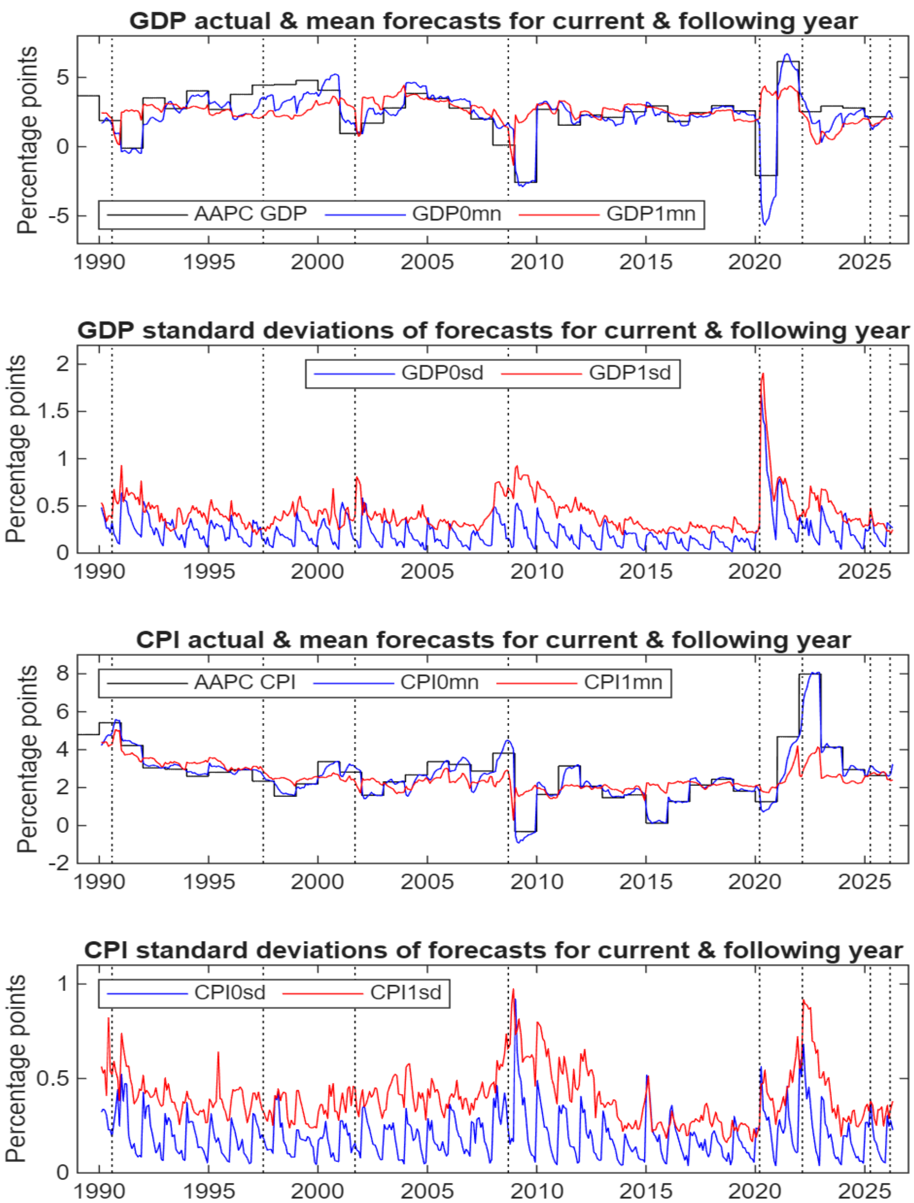

Figure 1: CE-calculated means and standard deviations of annual-average percent change (AAPC) forecasts for the United States. The ex-post realized AAPC GDP and CPI series are plotted for comparison. The eight events indicated with vertical dotted lines are summarized in table 2.

Source: Consensus Economics, US Bureau of Economic Analysis, US Bureau of Labor Statistics, and authors' calculations

We can formalize the concepts underlying the YAP to some extent, particularly for the current year, and we have included that discussion in Appendix A. But the formalization shows the limits to establishing theoretical YAPs from first principles, such as using expectations based on a diffusion process, particularly for the following-year data.

We therefore proceed on a purely empirical basis, while respecting three aspects that the YAP formalization highlights. The first is confirmation that mean changes are valid for calculating current-year GDP and CPI shocks, because the monthly difference eliminates the contribution from already-published data that is common to both months. Second, YAPs for absolute mean changes and for standard deviations should be monotonic over the year, which arises essentially from the dependence on a monotonic advancement of months remaining in the year. Third, YAPs for the following-year mean changes should have strongly positive slopes, because forecast revisions made in later months of the year increasingly fall into the following year.

Table 2: Summary of events indicated in time series figures

	Onset date	Event on selected date
1. Iraq annexes Kuwait	2-Aug-1990	Iraq invades Kuwait, and annexation completed on 4 August
2. Asian Financial Crisis	2-Jul-1997	Thai government abandons currency peg to the US dollar
3. Al-Qaeda attacks US	11-Sep-2001	Al-Qaeda destroys World Trade Centre Twin Towers, killing 2,997 people
4. Global Financial Crisis	15-Sep-2008	Lehman Brothers investment bank files for bankruptcy
5. COVID-19 pandemic	11-Mar-2020	World Health Organization declares a global pandemic
6. Russia/Ukraine war	24-Feb-2022	Russia launches full-scale invasion of Ukraine
7. Global trade disruptions	2-Apr-2025	US President Trump announces "Liberation Day" tariffs
8. US-Israel/Iran conflict	28-Feb-2026	The US and Israel launch surprise air strikes on multiple sites in Iran

Source: Authors' compilation

Estimating and applying the YAP is most straightforward for creating uncertainty measures from the CE standard deviation data, so we proceed with that first in the following section. Section 4 contains the analysis for creating shock measures from the CE mean data.

3 Uncertainty from standard deviation data

This section proceeds in two parts. In section 3.1 we detail our method for estimating the YAP, then YAP-adjusting the CE standard deviation series to create our uncertainty series. As previously mentioned, we use the United States to illustrate our process. The uncertainty results for all economies are contained in section 3.2, along with a brief commentary on the most material events in the history of those time series. In section 3.3, we outline the differences in our approach relative to prior uses of CE data to measure uncertainty.

3.1 YAP estimations

In this section, we use $\sigma_0(t)$ as our generic notation for the CE standard deviation data in the current year, which represents either GDP0sd or CPI0sd. Similarly, $\sigma_1(t)$ further below is for the following year CE standard deviation data, i.e. either GDP1sd or CPI1sd.

For the current year, we first apply the transformation $\log[\sigma_0(t)]$ to the data, which effectively puts the analysis on a multiplicative basis. That is, our adjustments for the estimated YAPs are intended to remove the declining ratios as time advances from the full-year January levels, hence normalizing the entire series to a common annual basis.[7]

The first YAP we test for the $\log[\sigma_0(t)]$ series is the simplest and most flexible, i.e. a separate dummy variable is used in each month, which we denote the Monthly YAP, or MYAP. Obviously this is a standard seasonal dummy regression but to avoid potential confusion in subsequent discussion, we reserve the terminology "seasonal" for discussing genuine variations in particular months after the YAP is allowed for.

The MYAP regression is:

$$\begin{bmatrix} \log[\sigma_0(1)] \\ \vdots \\ \log[\sigma_0(T)] \end{bmatrix} = \begin{bmatrix} I_{12} \\ \vdots \\ I_{12} \end{bmatrix} \begin{bmatrix} \beta_1 \\ \vdots \\ \beta_{12} \end{bmatrix} + \begin{bmatrix} \varepsilon_{M,1} \\ \vdots \\ \varepsilon_{M,T} \end{bmatrix} \tag{1}$$

which may be abbreviated to $Y = X_M\beta + \varepsilon_M$, where Y is the $T \times 1$ vector of data $\log[\sigma_0(t)]$, $\beta = [\beta_1, \ldots, \beta_{12}]'$ is the 12×1 vector of coefficient estimates, and ε_M is the $T \times 1$ vector of residuals. The $T \times 12$ matrix $X_M = [I_{12}, \ldots, I_{12}]'$ is notation for the repeated 12×12 identity matrix I_{12}, which is the dummy of 1 for each respective month and zeros otherwise. The rows in the first and last I_{12} matrices are assumed to be appropriately truncated so the dummy matrix aligns with the first and last month of Y in the sample. The OLS estimate of β is then $\beta = (X_M'X_M)^{-1} X_M'Y$.

The left-hand panels of figure 2 plot the MYAP β estimates for GDP and CPI with β_1 normalized to zero, i.e. $\beta - \beta_1 = [0, \beta_2 - \beta_1, \ldots, \beta_{12} - \beta_1]'$. Panels 1 and 3 of figure 3 plot the GDP and CPI data $\log[\sigma_0(t)]$ with no YAP adjustment, which we label as CNST, and the MYAP-adjusted series $\log[\tilde{\sigma}_0(t)] = \beta_1 + \varepsilon_{M,t}$. All elements in the series $\log[\tilde{\sigma}_0(t)]$ are now on the same full-year basis as the values $\log[\sigma_0(t)]$ for just the January months.

The left-hand panels of figure 2 also show the results from the Linear YAP (LYAP) we test, i.e. $Y = \alpha_L + X_L\delta + \varepsilon_L$, where α_L is a $T \times 1$ constant vector, and X_L is a $T \times 1$ vector that repeats the 12×1 vector $[12, \ldots, 1]'$ associated with January to December (again with appropriate truncations in the first and last vectors to match the first and last month of Y). The LYAP is the simplest parametric form that is

[7] Using the logarithm also means there is no difference between analysis using standard deviations or variances, because $\log\left[\{\sigma_0(t)\}^2\right] = 2\log[\sigma_0(t)]$.

guaranteed to be monotonic which, as discussed in section 2 and detailed in Appendix A, is a property that a YAP should have in principle.[8] Panels 1 and 3 of figure 3 plot the LYAP-adjusted series for GDP and CPI, $\log\left[\tilde{\sigma}_0(t)\right] = \beta_1 + \varepsilon_{L,t}$.

For comparison to the other two results, figure 2 also shows the results for a regression on a constant, i.e. $Y = \alpha_C + \varepsilon_C$, where α_C is a $T \times 1$ constant vector, which we label CNST. The CNST series for GDP and CPI in figure 3 is $\log\left[\tilde{\sigma}_0(t)\right] = \beta_1 + \varepsilon_{C,t}$, which is the $\log\left[\sigma_0(t)\right]$ data itself apart from a constant, i.e. $\log\left[\tilde{\sigma}_0(t)\right] = \beta_1 + \varepsilon_{C,t} = [\beta_1 - \alpha_C] + \alpha_C + \varepsilon_{C,t} = [\beta_1 - \alpha_C] + Y_t = [\beta_1 - \alpha_C] + \log\left[\sigma_0(t)\right]$.

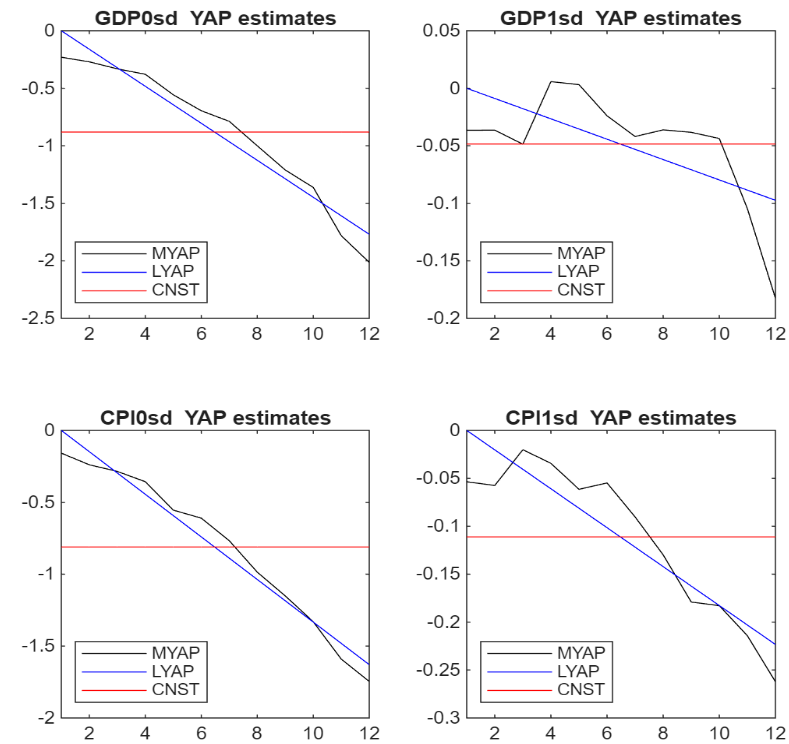

Figure 2: Year Advancement Pattern (YAP) estimates for the CE United States standard deviation data.

Source: Authors' estimations

Table 3 contains the statistical results underlying figures 2 and 3. The rows GDP0sd and CPI0sd correspond to the regression results plotted in the left-side panels of figure 2, i.e. for the current-year standard deviations of GDP and CPI. In short, they strongly suggest using the MYAP for current year standard deviations, for three reasons. First, column CNST/MYAP shows that the MYAP estimates for GDP0sd and CPI0sd are strongly preferred relative to the respective CNST estimates, i.e. the 0.0% p-values strongly reject the constrained regression estimates CNST versus the unconstrained MYAP estimates. Note that these tests are based on the difference in squared residuals

[8]In Appendix A.5, we equivalently develop the LYAP as a restricted regression with fixed monthly weights. That form allows for testing YAPs based on alternative weighting schemes that we may progress in future work, e.g. a montonic function based on the diffusion principles we discuss in Appendix A.4.

allowing for the Newey-West Heteroskedastic and AutoCorrelated correction (NW-HAC) with a window of 12 (i.e. the annual length of the MYAP regression data). An F-test would typically be used for testing constraints on a regression, but F-tests assume independent and normally distributed residuals, and figures 1 and 3 show that $\log[\tilde{\sigma}_0(t)]$ and hence $\varepsilon_{M,t}$ clearly display heteroskedasticity and serial correlation.[9] Column CNST/LYAP also rejects the CNST versus LYAP estimates, which confirms again that there are significant changes in the current-year standard deviations as the year progresses.

The second reason for preferring the MYAP is that column LYAP/MYAP rejects the LYAP versus the MYAP (p-value of 0.1% for GDP0sd, and 1.2% for CPI0sd). As apparent in figure 2, the reason is essentially that the LYAP is too simplistic to capture the nuanced shape of the MYAP, with material differences relative to the MYAP apparent around the start and end months of the year. Third, aside from the statistical results, the MYAPs for the left-side figures in figure 2 both show monotonic declines. That property is fully consistent with our intention of removing from the standard deviation data an estimated component due only to the surveys progressing through the year.

Table 3: United States standard deviation regression results

		p-values in % (1)			Jan.	Dec. to Jan.	
		CNST/	LYAP/	CNST/	geom.	ratio (3)	
Item	T	MYAP	MYAP	LYAP	mn (2)	MYAP	LYAP
GDP0sd	435	0.0	0.1	0.0	0.33	0.17	0.17
GDP1sd	435	7.6	13.6	17.4	0.40	0.86	0.91
GDP01sd	435	26.5	27.7	41.3	0.52	0.99	1.02
CPI0sd	435	0.0	1.2	0.0	0.32	0.20	0.20
CPI1sd	435	0.1	13.4	0.2	0.40	0.81	0.80
CPI01sd	435	22.3	24.5	38.6	0.50	1.03	1.02

Notes: (1) Constrained versus unconstrained regressions. (2) Geometric mean for January data. (3) Ratio of December to January geometric means.
Source: Authors' estimations

The last three columns in table 3 illustrate the materiality of the YAPs in practice and why it needs to be tested and allowed for if necessary. Hence, the January geometric mean is $\exp(\beta_1) = 0.33$ for GDP0sd and $\exp(\beta_1) = 0.32$ for CPI0sd, where β_1 are the MYAP coefficient estimates for January, i.e. the means of $\log(\text{GDP0sd})$ and $\log(\text{CPI0sd})$ for just the January months. The adjacent MYAP column shows the point estimates of the relative magnitudes of the geometric averages between the January and December standard deviations, obtained as $\exp(\beta_{12} - \beta_1) = 0.17$ for GDP0sd and $\exp(\beta_{12} - \beta_1) = 0.20$ for CPI0sd, where β_{12} are the MYAP coefficient estimates for December.[10]

[9] Nevertheless, in nearly all cases, the F-tests give the same rejections or non-rejections as the NW-HAC tests.

[10] Associated p-values for the 0.17 and 0.20 point estimates are redundant, given the NW-HAC test has already established statistical significances relative to CNST.

To interpret the results above, 0.33 and 0.32 for January are standard deviations for the full current year, and the December ratios indicate that the December standard deviations are on average 0.17 and 0.20 of the January values, i.e. hence outright geometric averages of $0.17 \times 0.33 = 0.056$ and $0.20 \times 0.32 = 0.064$. Adjusting for the MYAP puts the December standard deviations on the same annual basis as January so the values are comparable. The adjustments using $\beta_2, \ldots, \beta_{11}$ do the same for the February to November standard deviations. The MYAP-adjusted results are plotted in panels 1 and 3 of figure 3, showing that the values across each year are now all comparable. For completeness, the last column gives the December to January ratios for the LYAP, which are also 0.17 and 0.20.

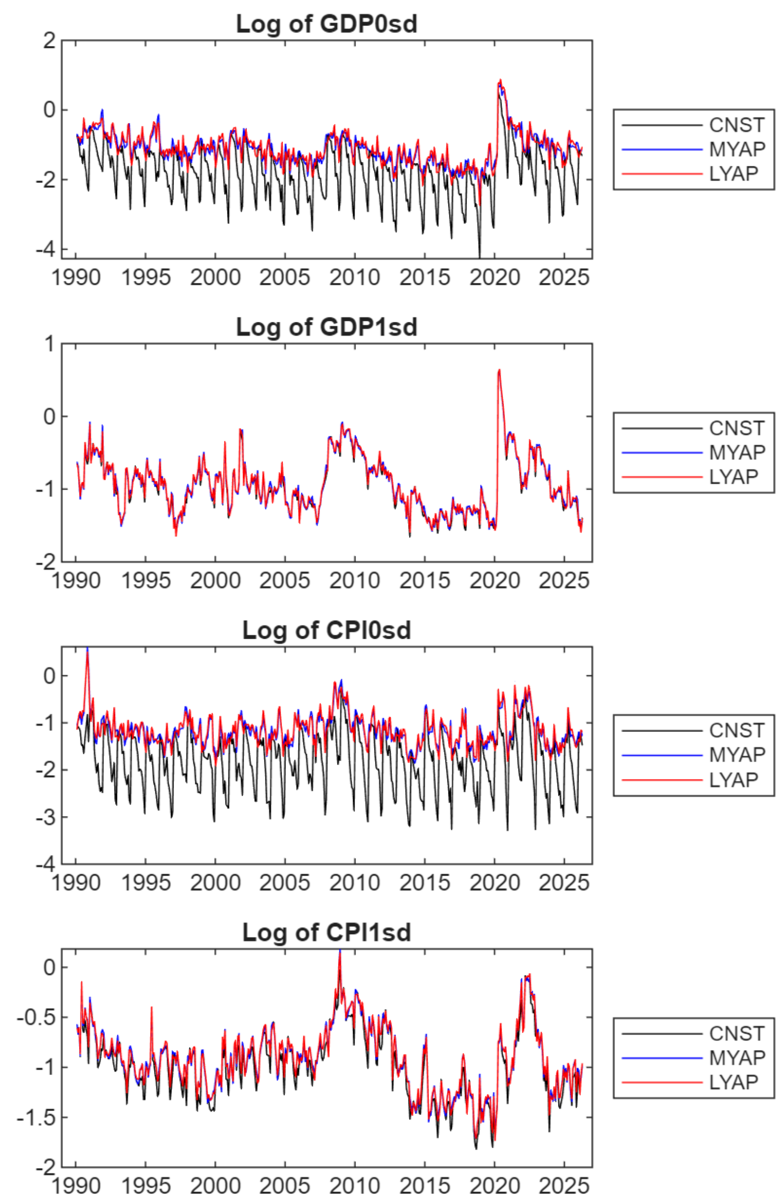

Figure 3: YAP-adjusted series for log of standard deviations.

Source: Authors' estimations

Turning to our analysis for the following-year standard deviation data, the rows GDP1sd and CPI1sd in table 3 summarize the key statistics associated with the regression results plotted in the right-side panels of figure 2 (and note that these panels have much smaller y-axis scales than for current-year results on the left-side). In short, these results suggest choosing the LYAP for the following-year standard deviations. Our choice follows from figure 3 showing that LYAP allows for the mild decline over the year relative to the CNST (and table 3 shows the difference is statistically significant for CPI1sd), and our choice also avoids the non-monotonicity relative to the MYAP (and the difference is insignificant for both GDP1sd and CPI1sd).

The January geometric mean column shows that standard deviations for the following year (i.e. 0.40 for both GDP and CPI) are larger than for the current-year values previously discussed (i.e. 0.33 and 0.32). The last two columns show that the relative magnitudes between January and December decline, but mildly with respect to that ratio for the current year.

Once the MYAP-normalized current year series, $\log\left[\tilde{\sigma}_0\left(t\right)\right]$, and the LYAP-normalized following year series, $\log\left[\tilde{\sigma}_1\left(t\right)\right]$, are obtained, we calculate the respective annualized levels as $\tilde{\sigma}_0\left(t\right) = \exp\left(\log\left[\tilde{\sigma}_0\left(t\right)\right]\right)$ and $\tilde{\sigma}_0\left(t\right) = \exp\left(\log\left[\tilde{\sigma}_0\left(t\right)\right]\right)$. Then we use $\tilde{\sigma}_{01} = \sqrt{\left[\tilde{\sigma}_0\left(t\right)\right]^2 + \left[\tilde{\sigma}_1\left(t\right)\right]^2}$ to combine those results into an uncertainty measure for GDP and for CPI, which we label respectively GDP01sd and CPI01sd. Rows 3 and 6 in table 3 show that there is no remaining MYAP or LYAP in the GDP01sd and CPI01sd series, i.e. the tests versus a CNST are statistically insignificant and the January to December ratios are approximately 1.

We have repeated the analysis above on all of the remaining 13 economies and, in short, we find that our choice of MYAP adjustment for the current-year standard deviations and LYAP adjustment for the following-year standard deviations applies very well to all other economies. It would be infeasible to present the figures and tables for all 14 economies in this article, but any or all are available on request. Here we provide an overview of the results with table 4 containing the same categories as the United States table, but each element contains just a single summary entry across all economies.

Hence, instead of p-values for the regression comparisons, we report the number of economies that have rejections at the 5% level. Consistent with United States results for GDP0sd and CPI0sd, all economies reject CNST versus MYAP or LYAP, and 12 economies reject LYAP versus MYAP. Additionally, our inspection of the current year regression results for each economy show the MYAP profiles to be either monotonic or close to monotonic for all economies (akin to those on the left-side of figure 2). Therefore, the MYAP is most applicable across all economies for adjusting the current-year standard deviation results.

For the following year, CPI1sd shows some rejections of CNST (there are some 10% rejections for GDP1sd, not shown in table 4), no economy rejects LYAP versus MYAP, and the regression figures often show material non-monotonicity for MYAP (akin to those on the right side of figure 2). Therefore, LYAP is most applicable across all economies when adjusting the following-year standard deviation results.

The remaining three columns are respectively, from the set of all economies, the maximum geometric mean of January values, the maximum ratio of December to January data, and the maximum December to January ratio from the LYAP regressions.

For example, the entry of 0.56 on the right-most column of the GDP0sd row indicates that for the LYAP estimates, no economy had a December-to-January ratio higher than 0.56. For CPI0sd, no ratio was higher than 0.57. These results show that the current-year YAP is always economically material, i.e. across all economies the December standard deviations at least need to be approximately doubled to standardize them to the same annual basis as the January standard deviations. The combined results GDP01sd and CPI01sd have ratios close to one, consistent with no remaining YAP.

Table 4: Standard deviation results summary for all economies

Item	Min. T	Number of of 5% p-values (1) CNST/ MYAP	LYAP/ MYAP	CNST/ LYAP	Max. Jan. geom. mn (2)	Maximum Jan. to Dec. ratio (3) MYAP	LYAP
GDP0sd	377	14	12	14	0.60	0.56	0.64
GDP1sd	377	0	0	0	0.67	1.05	1.07
GDP01sd	377	0	0	0	0.95	1.04	1.04
CPI0sd	377	14	12	14	0.70	0.57	0.59
CPI1sd	377	3	0	2	0.78	1.11	1.04
CPI01sd	377	0	0	0	1.10	1.05	1.04

Notes: (1) Constrained vs unconstrained regressions. (2) Geometric mean for January data. (3) Ratio of December to January geometric means.
Source: Authors' estimations

3.2 Uncertainty results

Figure 4 plots our GDP01sd and CPI01sd results for all economies. For the United States, there were notable local peaks of GDP uncertainty associated with the 9/11 Al-Qaeda attacks and the Global Financial Crisis (GFC), and the all-time peak was associated with the COVID-19 pandemic (COVID). CPI uncertainty was highest for the Iraq annexation of Kuwait, the GFC, and the Russia/Ukraine war.

COVID was associated with all-time peaks of GDP uncertainty for most economies in our dataset. The Asian Financial Crisis and its aftermath was also a period of high GDP uncertainty for many Asian economies, notably Thailand, Korea, Malaysia, and Indonesia. Those economies also had high CPI uncertainty, with an off-scale peak of 27 percentage points for Indonesia in September 1998.

For the global trade policy disruptions beginning on 2-Apr-2025, there are upticks in uncertainty apparent for many economies, most notably for Taiwan Province of China. The relatively mild impact on Asian economies in general may seem surprising, given that the BBD series and trade policy uncertainty increased substantially around this event.[11] The small and temporary increases may be due to the most onerous of the potential changes in the trade environment being relaxed within the month of April before the May 2025 CE survey.

[11] See section 5.2 for the BBD series and a discussion of the differences with our series. For trade policy uncertainty analogous word-count measures are available for the United States, China, Japan, and Korea at www.policyuncertainty.com/trade_uncertainty.html.

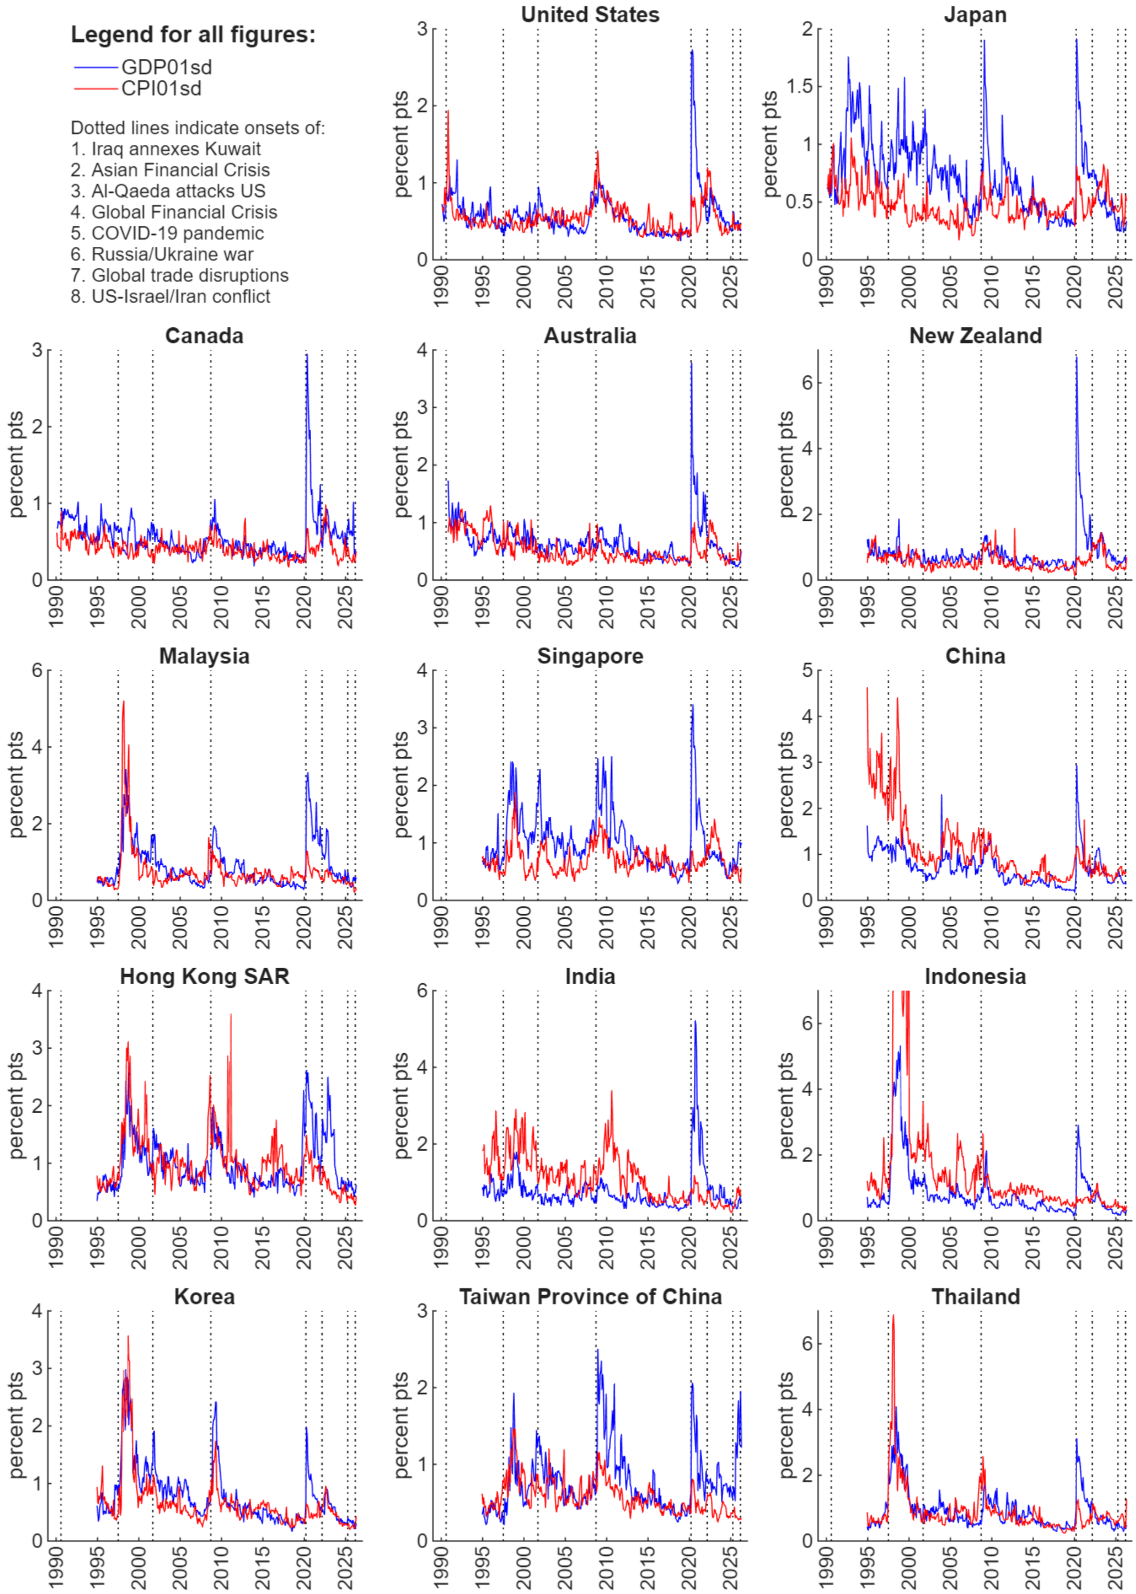

Figure 4: Uncertainty estimates of GDP and CPI. Indonesia CPI01sd has off-scale values from Feb-1998 to Dec-1999, with a peak of 27 percentage points for Sept-1998.

Source: Authors' estimations

The US-Israel/Iran conflict has prompted only moderate increases in GDP and CPI uncertainty for some economies, even while shocks to CPI and supply have generally been large (as will be discussed in sections 3 and 4). Implicitly then, it seems that forecasters have tended to agree on the likely macroeconomic effects from the conflict.

3.3 Comparison to other approaches

Denis and Kannan (2013), Greig, Rice, Vehbi, and Wong (2018), and Ozturk and Sheng (2018) are the examples we are aware of that use CE data to create measures of uncertainty for GDP,[12] with all based on the pro-rated weighted-average of each individual CE respondents' current- and next-year forecasts. Specifically, for month m (where m runs from 1 to 12) of year t, the 1-year ahead forecast for respondent k, denoted $F1_{k,m}$, is calculated as:

$$F1_{k,m} = \frac{13-m}{12} FC_{k,m} + \frac{m-1}{12} FN_{k,m} \tag{2}$$

where $FC_{k,m}$ and $FN_{k,m}$ are respectively the current- and next-year forecasts. The measure of uncertainty is then the standard deviation of the set of results for all $F1_{k,m}$, which we denote for convenience as stdev$(F1_{k,m})$.

The main issue with this approach, related to our discussion of the CE data in section 2 and detailed in Appendix A, is that $FC_{k,m}$ is an AAPC forecast, so it will increasingly incorporate published data as the year progresses, leaving progressively less contribution from genuine forecasts for the remaining months of the year. Because there is essentially no uncertainty for published data, there will be a YAP with monotonic decline for the variance of $FC_{k,m}$ by month. Similarly, a YAP will also be inherent in the variance of $FN_{k,m}$ due to the declining forecast horizon, albeit less pronounced than for $FC_{k,m}$. The calculation stdev$(F1_{k,m})$ will therefore include YAP contributions from the variances of $FC_{k,m}$ and $FN_{k,m}$ (there is also a term from the covariance of $FC_{k,m}$ with $FN_{k,m}$), even before the weights $(13-m)/12$ and $(m-1)/12$ are applied (and the latter may further influence the stdev$(F1_{k,m})$ YAP).

We provide a detailed exposition of these aspects in section A.6 of Appendix A, with the benefit of our formal discussion in the earlier sections of Appendix A on the mathematics underlying the YAPs in our analysis. However, we have not yet investigated the potential issues with using stdev$(F1_{k,m})$ as a measure of uncertainty for monitoring and research, which will require the CE unit record data.

4 Shocks from changes in CE mean data

This section proceeds in two parts. In section 4.1 we detail our method for estimating the YAP from the CE mean-change data, and then YAP-adjusting the data to create our shock series. We again use the United States to illustrate our process. The shock

[12] Ballantyne, Gillitzer, Jacobs, and Rankin (2016) uses CE data to create a measure of forecast disagreement for Australian CPI inflation and, noting the mechanical decline in later months of each year due to fewer unknown inflation outcomes, dummy variables are used to remove the predictable variation in disagreement. That approach coincides with the MYAP we have chosen from the three alternatives tested.

results for all economies are contained in section 4.2, along with a brief commentary on the most material events in the history of those time series.

4.1 Empirical estimation

In this section, we use $\Delta\mu_0(t)$ as our generic notation for the change in CE mean data in the current year, and we reserve the names GDP0mn and CPI0mn to denote the changes in the means of GDP and CPI (i.e. GDP0mn and CPI0mn do not refer to the mean levels). Similarly, $\Delta\mu_1(t)$ further below is for the change in CE mean data for the following year, i.e. either GDP1mn or CPI1mn.

We first calculate the changes in the means, i.e. $\Delta\mu_0(t) = \mu_0(t) - \mu_0(t-1)$ and $\Delta\mu_1(t) = \mu_1(t) - \mu_1(t-1)$ from the CE mean data. Note that current-year changes for January need to be calculated as $\Delta\mu_0(t \in \text{Jan}) = \mu_0(t \in \text{Jan}) - \mu_1(t \in \text{Dec})$, to account for the forecasted calendar years advancing between the December and January surveys. The data does not allow the analogous calculation for the following-year series, i.e. there are no forecasts for the third year so $\Delta\mu_1(t \in \text{Jan}) = \mu_1(t \in \text{Jan}) - \mu_2(t \in \text{Dec})$ cannot be calculated. Therefore, January of the following year has no mean-change data and so no results are given in our analysis.

Changes in the mean can be positive or negative, so for our YAP analysis we apply the transformations $\log[|\Delta\mu_0(t)|]$ and $\log[|\Delta\mu_1(t)|]$ to the data, where $|\cdot|$ is the absolute value (as discussed further below, we retain the signs of the mean changes and reapply them later to obtain our shock time series). Hence, our YAP analysis for mean changes is again on a multiplicative basis, but with respect to the magnitude of mean changes rather than levels as used for standard deviations in the previous section.[13]

We then apply the MYAP, LYAP, and CNST estimations, as discussed in the previous section, to $\log[|\Delta\mu_0(t)|]$ and $\log[|\Delta\mu_1(t)|]$. Note that because the data $\Delta\mu_1(t \in \text{Jan})$ cannot be obtained from the CE mean data, as discussed above, the following-year YAP estimations are based on the monthly dummy variables without the January observation. Hence, the following-year YAP plots below do not have January estimates, and the YAP-adjusted time series of mean changes for the following year do not have January values.

In short, the results show that LYAP adjustments are the best choice for both the current and following year. To make that case, we discuss the results for the United States summarized in figures 5 and 6, and table 5 below.

First, all panels in figure 5 show that MYAP has large deviations from monotonicity. This indicates that adjusting the series using the MYAP would be unsuitable, because it would remove genuine seasonality that should be retained. That is, the publication schedule of realized data often results in shocks to the AAPC forecasts, due to realized forecast errors and any reconsiderations of forecasts based on the newly-published data. That is particularly the case for GDP forecasts which, being published on a quarterly basis, should typically have the larger forecast revisions three months apart. This is evident in the top-left panel of figure 5, which clearly shows large GDP forecast revisions in May and August. Second, it is unsuitable to use no adjustment, i.e. the

[13] Taking the logarithm also means there is no difference between analysis using absolute values or squared values (where the latter would also produce a positive series that can be transformed with the logarithm), because $\log\left[\{\Delta\mu_0(t)\}^2\right] = 2\log[|\Delta\mu_0(t)|]$.

CNST case, because there is clearly a material, and typically statistically significant, pattern in the magnitudes of mean changes as the year advances.

Regarding the LYAP profiles, for both GDP and CPI the current year has the expected downward slope, given there are progressively fewer genuine forecast months as the year advances, so potential revisions for AAPC forecasts become smaller. The following year for both has a strong upward slope, which indicates larger revisions in magnitude as the year progresses. We simply use those empirical following-year LYAP estimates for our adjustment, but the positive slopes may be readily explained in principle; i.e. for the months late in the year, forecast changes to GDP and/or CPI outlook would be incorporated as smaller changes in the remainder of the current year, with larger changes for the remainder incorporated into the following year. Related, a simulation exercise we have included in appendix A.4 also establishes a positive slope for the following year.

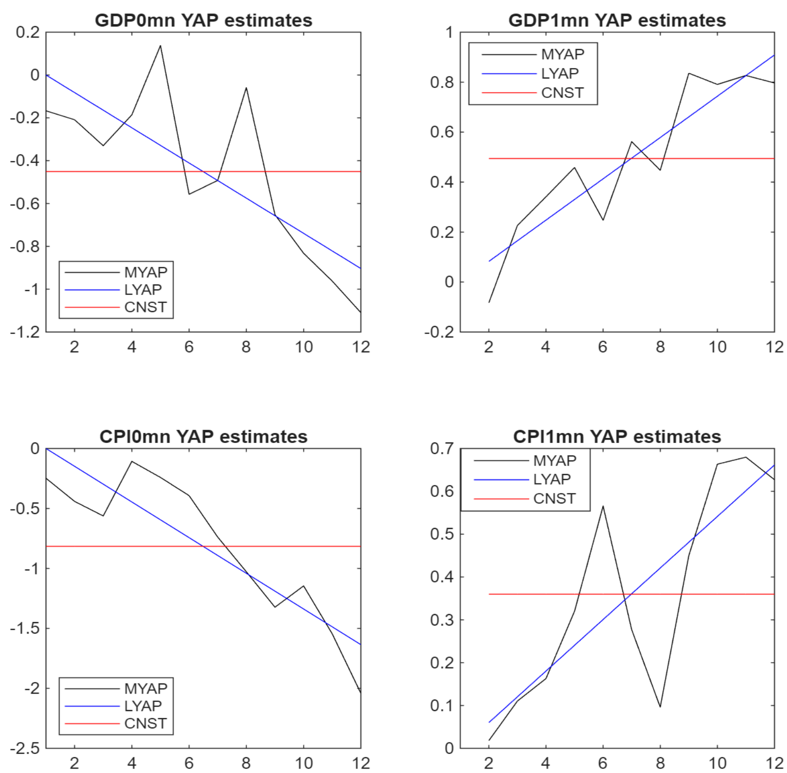

Figure 5: Year Advancement Patterns (YAPs) estimated for United States magnitudes of mean-change data.

Source: Authors' estimations

The LYAP adjustments are made to the current- and following-year log magnitude mean changes, i.e. $\log[|\Delta\mu_0(t)|]$ and $\log[|\Delta\mu_1(t)|]$, to produce the standardized series of log magnitude mean changes, which we denote as $\log[|\Delta\tilde{\mu}_0(t)|]$ and $\log[|\Delta\tilde{\mu}_1(t)|]$. We have not plotted these results, as we did for standard deviations, because the differences between the data and results are not visually apparent due to the higher variability and range of these series. However, figure 5 and the statistical results in table 5 show that the MYAP and LYAP adjustments are strongly significant relative to the unadjusted series CNST for the current- and following-year (i.e. rows GDP0mn, GDP1mn, CPI0mn, and CPI1mn). The December to January ratios in the last two columns show small ratios for the current year and large ratios for the following year.

Table 5: United States mean-change regression results

		p-values in % (1)			Jan.	Dec. to Jan.	
		CNST/	LYAP/	CNST/	geom.	ratio (3)	
Item	T	MYAP	MYAP	LYAP	mn (2)	MYAP	LYAP
GDP0mn	434	0.1	1.1	0.5	0.12	0.39	0.41
GDP1mn	398	2.4	18.8	3.9	0.03	2.41	2.48
GDP01mn	434	2.0	2.0	49.2	0.10	1.07	1.01
CPI0mn	434	0.0	0.9	0.0	0.09	0.17	0.19
CPI1mn	398	2.5	8.4	6.6	0.03	1.84	1.94
CPI01mn	434	1.9	2.4	32.2	0.07	1.05	1.18

Notes: (1) Constrained vs unconstrained regressions. (2) Geometric mean for January data. (3) Ratio of December to January geometric means.
Source: Authors' estimations

We then use $|\Delta\tilde{\mu}_0(t)| = \exp(\log[|\Delta\tilde{\mu}_0(t)|])$ and $|\Delta\tilde{\mu}_1(t)| = \exp(\log[|\Delta\tilde{\mu}_1(t)|])$ to obtain standardized magnitudes of mean changes, reapply the original signs of the $\Delta\mu_0(t)$ and $\Delta\tilde{\mu}_1(t)$ data to $|\Delta\tilde{\mu}_0(t)|$ and $|\Delta\tilde{\mu}_0(t)|$, i.e. $\Delta\tilde{\mu}_0(t) = \operatorname{sgn}[\Delta\mu_0(t)]\cdot|\Delta\tilde{\mu}_0(t)|$ and $\Delta\tilde{\mu}_1(t) = \operatorname{sgn}[\Delta\mu_1(t)]\cdot|\Delta\tilde{\mu}_1(t)|$, where $\operatorname{sgn}[\cdot]$ is the sign function, and sum the results for both years to obtain our estimate of the shock series, i.e. $\Delta\tilde{\mu}_{01}(t) = \Delta\tilde{\mu}_0(t) + \Delta\tilde{\mu}_1(t)$.

The results for the United States are plotted in the top-right panel of figure 6 in the following section, and the results in table 5 show no LYAP remains in the combined series, i.e. GDP01mn CPI01mn.

Table 6 contains the summary across all 14 economies, any or all which are available on request. Most economies show results similar to those discussed for the United States above, although the statistical significance given in column CNST/LYAP shows that only several economies reject the CNST at the 5% level for both the current and following years. For consistency across all economies, we choose to apply the LYAP in all cases.

The combined results show that all economies have no LYAP remaining to the 5% level, with the exception of CPI01mn for New Zealand, which also has an unusually high January to December ratio of 1.79. Malaysia has the highest ratio for GDP01mn, at 1.66. Our inspection of these results show that they arise from high genuine seasonality in the current- and following-year data, i.e. large forecast revisions late in the year, but the LYAP-adjusted series are still preferable to no adjustment.[14] Finally, MYAP is unsuitable for the same reasons discussed in the United States case, i.e. our inspections show that the MYAP results are not monotonic due to genuine seasonality in the forecast data.

[14] As discussed in section 2, we will provide these results or those for any other economy on request.

Table 6: Summary of mean-change results for all economies

		Number of of 5% p-values (1)			Max.	Maximum Dec. to Jan. ratio (3)	
	Min.	CNST/	LYAP/	CNST/	Jan. geom.		
Item	T	MYAP	MYAP	LYAP	mn (2)	MYAP	LYAP
GDP0mn	376	13	13	3	0.12	2.60	1.64
GDP1mn	344	13	8	5	0.07	3.48	3.57
GDP01mn	376	13	13	0	0.12	3.57	1.66
CPI0mn	376	14	12	6	0.15	1.25	0.97
CPI1mn	344	13	7	3	0.07	2.93	2.47
CPI01mn	376	14	14	1	0.21	3.23	1.79

Notes: (1) Constrained vs unconstrained regressions. (2) Geometric mean for January data. (3) Ratio of December to January geometric means.
Source: Authors' estimations

4.2 Results

Figures 6 and 7 plot our GDP01mn and CPI01mn results for all economies. Our first observation from these results is that the terminology "shocks" we use should not be interpreted in the econometric sense as serially independent, because there is apparent serial correlation in each of the series. We have confirmed that it is strongly significant statistically in all cases by simply applying a single-lag autoregression to each series. We have similarly confirmed that the serial correlation is from the underlying data. These results themselves suggest interesting avenues for future research, including alternative approaches to using CE data (e.g. using residuals from univariate or vector autoregressions of the CE data). For this article, we continue with our terminology of "shocks" being the changes of monthly means of CE GDP growth and CPI inflation forecasts, respectively, standardized to a two-year horizon as described in the previous section.

For the United States, the largest GDP shock was the negative value associated with COVID, and the largest CPI shock was the positive value associated with the Russia/Ukraine war. Across all economies, COVID was associated with the largest negative GDP shocks for half of the economies in our dataset, with Japan almost in that group. The Asian Financial Crisis accounts for the largest negative GDP shocks in the remaining economies, except for Taiwan Province of China where the Al-Qaeda attacks on the United States made a larger impact. The largest positive shocks to CPI have more variation with respect to particular events, but the post-COVID inflationary period including the Russia/Ukraine war is associated with large positive CPI shocks in most economies.

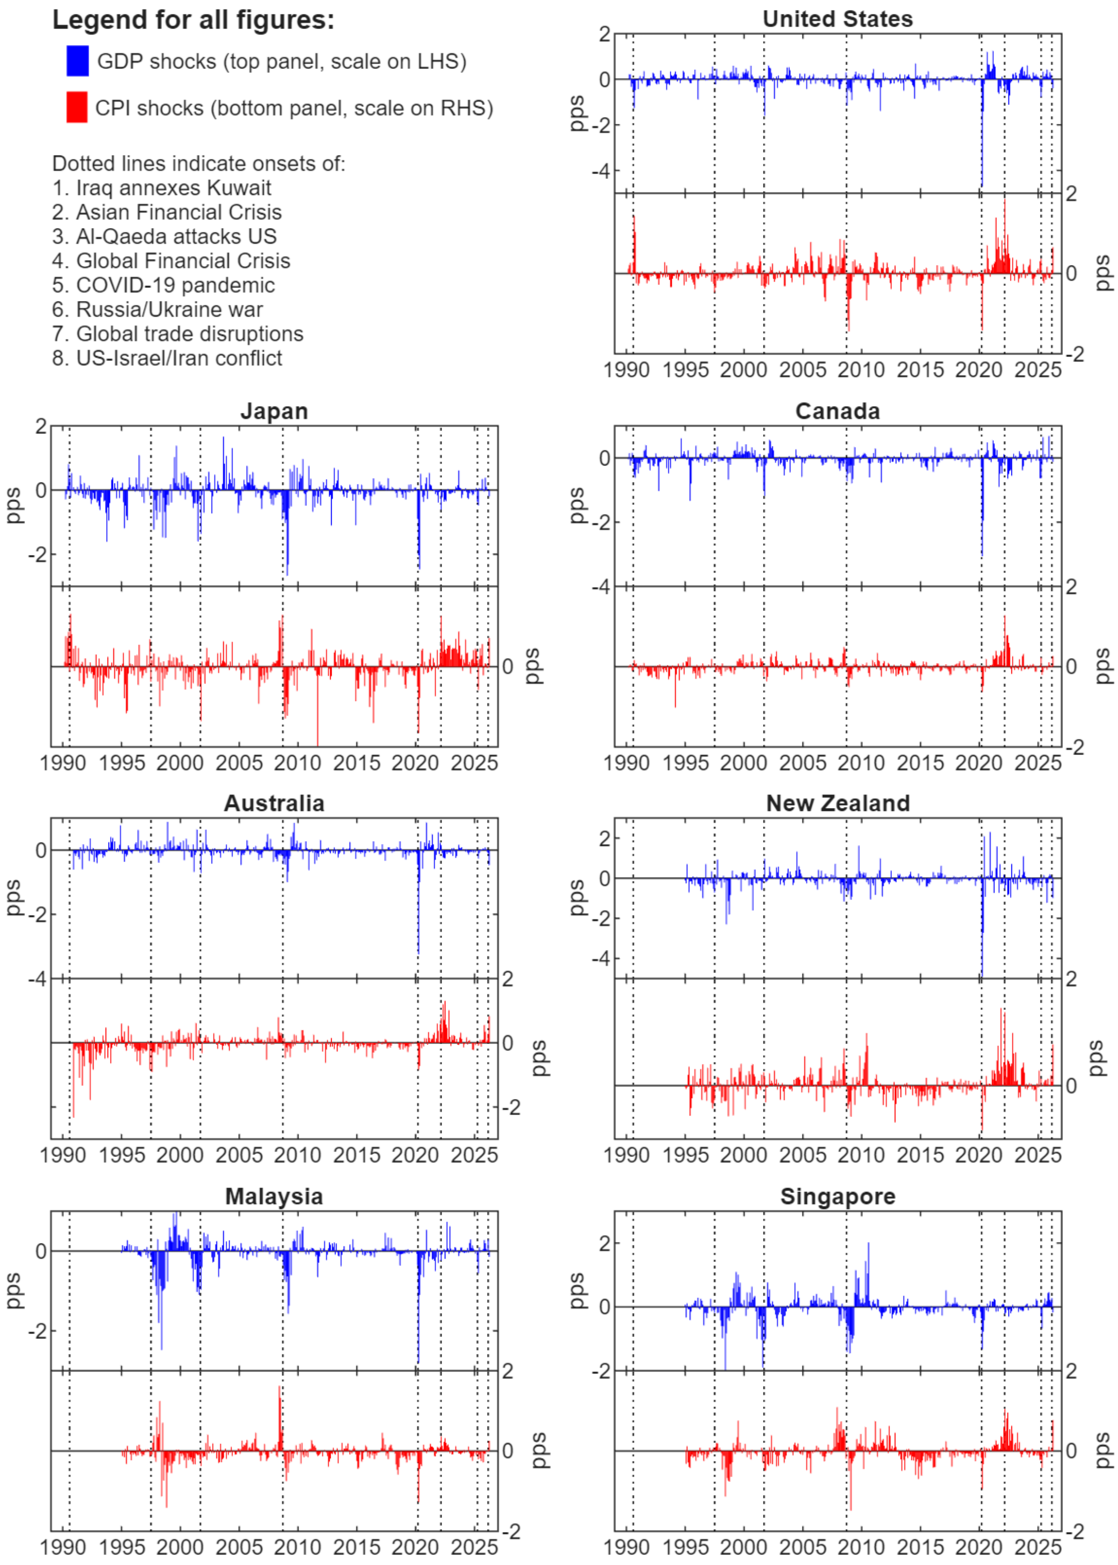

Figure 6: Shock estimates of GDP and CPI for the first seven economies in our dataset.

Source: Authors' estimations

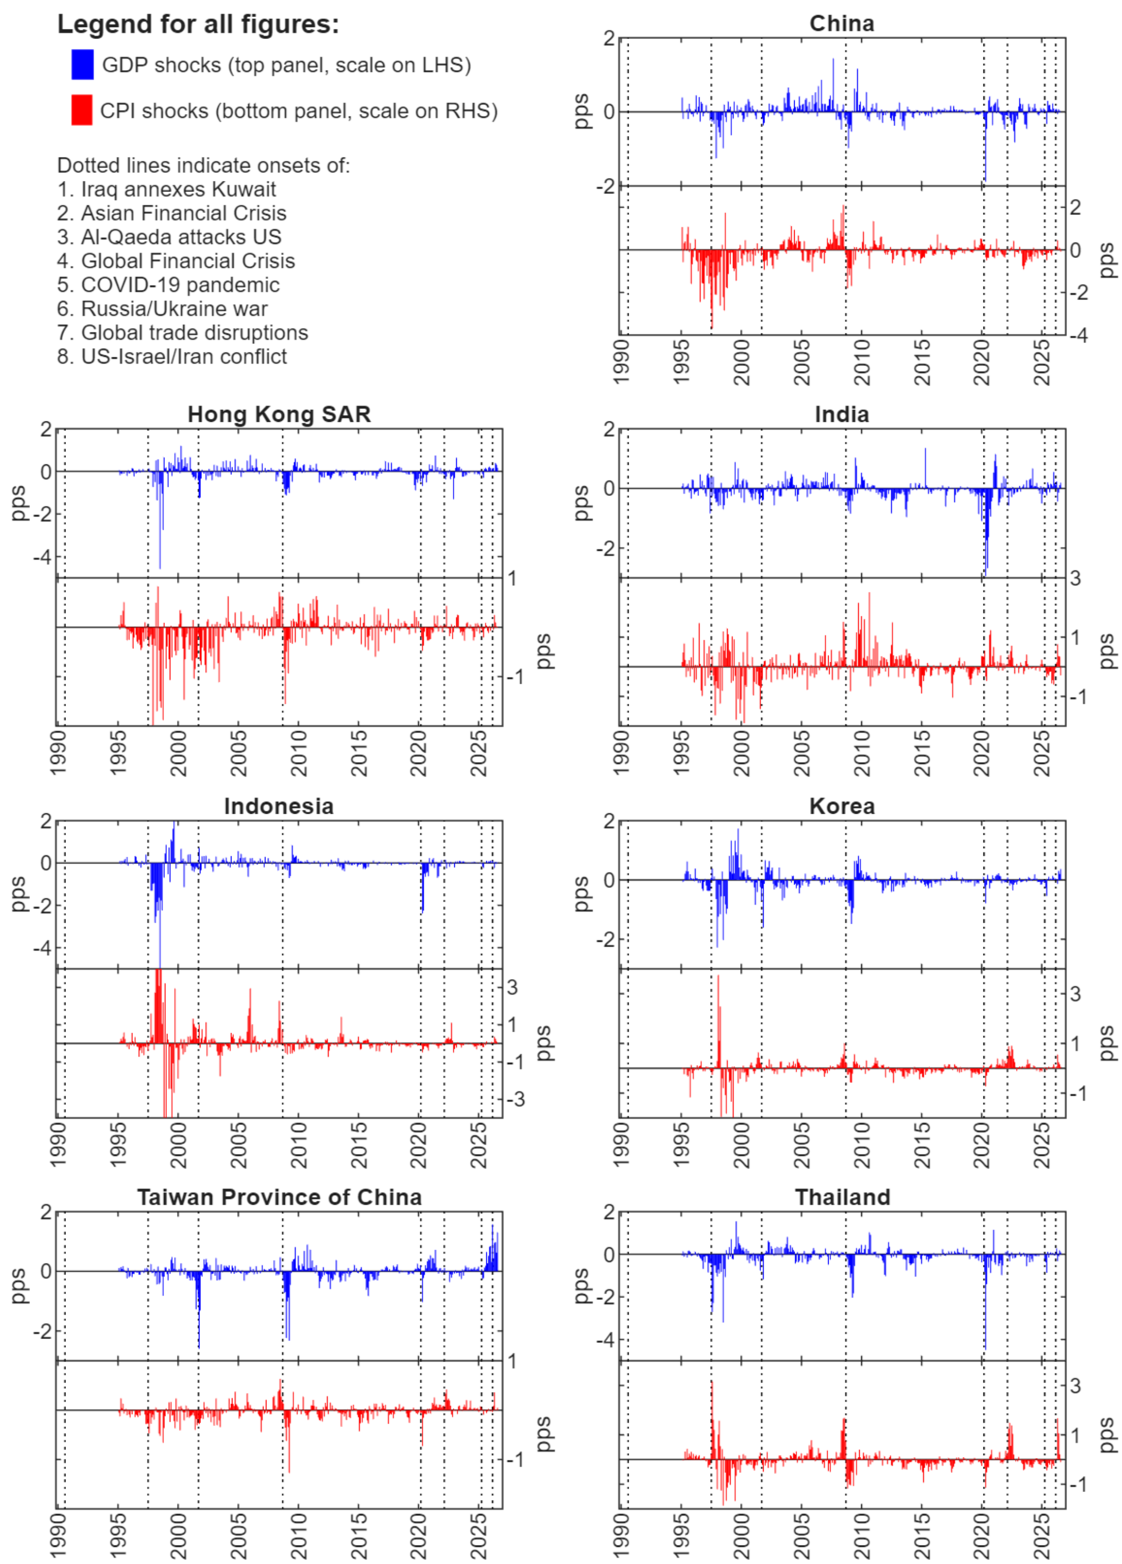

Figure 7: Shock estimates of GDP and CPI for the second seven economies in our dataset. Indonesia has off-scale values with a minimum of –7.6 for GDP, and a range of –16.1 to +19.4 for CPI.

Source: Authors' estimations

5 Empirical applications

This section proceeds in two parts. In section 5.1 we discuss how we transform our GDP and CPI shocks into fundamental demand and supply shocks, and then we present and discuss those results. In section 5.2, we compare our uncertainty measures to the BBD economic policy uncertainty measures for the economies where both are available and our US measure to the JLN and LMN macroeconomic and financial uncertainty measures.

5.1 Sign restrictions for demand and supply shocks

To outline our approach for inferring demand and supply shocks from our GDP and CPI shocks, i.e. our GDP01mn and CPI01mn series, we begin with our estimated expression for the United States, explain its concepts and derivation, and discuss the results we obtain from the expression. We then discuss why our demand and supply shock estimates should be interpreted with respect to several important caveats, which are associated with their derivation from forecast changes, in real time, for just two macroeconomic variables.

Hence, our expression for the United States that relates GDP and CPI shocks to demand and supply shocks is:

$$\begin{bmatrix} \mathrm{GDP}\,(t) \\ \mathrm{CPI}\,(t) \end{bmatrix} = \begin{bmatrix} +0.27 & +0.27 \\ -0.18 & +0.23 \end{bmatrix} \begin{bmatrix} \varepsilon_{S,t} \\ \varepsilon_{D,t} \end{bmatrix} \; ; \; \begin{bmatrix} \varepsilon_{S,t} \\ \varepsilon_{D,t} \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right) \tag{3}$$

where GDP(t) is GDP01mn(t), CPI(t) is CPI01mn(t), the 2×2 matrix is the response matrix, $\varepsilon_{S,t}$ is the supply shock, and $\varepsilon_{D,t}$ is the demand shock. The distribution of $[\varepsilon_{S,t}, \varepsilon_{D,t}]'$ indicates that the shocks are independent with unit variances.

The purpose of equation 3 is to express non-independent shocks to GDP and CPI (they have a correlation of 0.1155, with p-value 1.6% for the United States in our dataset) as responses to fundamental demand and supply shocks that are themselves independent. Hence, the quantitative interpretation of equation 3 is that a positive unit supply shock is associated with 0.27 percentage points (pps) increase in forecast GDP01mn and –0.18 pps decline in forecast CPI01mn, and a positive unit demand shock is associated with, respectively, 0.27 and 0.23 pps increases in forecast GDP01mn and CPI01mn. The signs within the response matrix are typical in the VAR literature, e.g. see Fry and Pagan (2011), with their basis in standard aggregate supply and demand principles for the macroeconomy.

The values in our response matrix are implicitly on the basis of a two-year horizon given our method of producing GDP and CPI shocks, i.e. the annualized value for the current year plus the full annual value for the following year. In other words, our fundamental shocks at time t may effectively be interpreted as imparting expected average changes in GDP and CPI outcomes over the next two years. The analogy with a VAR, although not necessarily directly comparable, would be averaging the GDP and CPI impulse responses from the VAR over the next two years.

To generate equation 3, we use the sign restriction method detailed in Fry and Pagan (2011). We have relegated the details to Appendix B, and here we outline the essential three key steps for the sign restriction method in the context of our application. The first step is to derive a base model with the form of equation 3.

The response matrix A_0 for the base model can be arbitrary; the only requirement is that the shocks $\varepsilon_{0,t} = [\varepsilon_{0,S,t}, \varepsilon_{0,D,t}]'$ for the base model need to be independent with unit variances. We derive our base model using an eigensystem decomposition of the covariance matrix calculated from our $[\mathrm{GDP}(t), \mathrm{CPI}(t)]'$ data. The second step is to use unitary matrices Q_j to simultaneously transform the base model response matrix A_0 as $A_j = A_0 Q_j$ and the shocks as $\varepsilon_{j,t} = Q_j' \varepsilon_{0,t}$, seeking new models where the signs of the transformed response matrix A_j match the signs imposed from economic principles.[15] For our application, Q_j is the 2×2 rotation matrix $Q(\theta_j)$ with the rotation angle θ_j, and the signs we seek to match are those in the response matrix of equation 3. The third step is choosing how to summarize the set of new models that meet the sign requirements.[16] For our application, we choose the model where positive supply shocks have the same effect on GDP as positive demand shocks. Indeed, with that normalization, as detailed in Appendix B, we can analytically solve for a unique value of θ_j, rather than generating a set of models using a range of θ_j values. Denoting the unique angle as $\tilde{\theta}$, the response matrix in equation 3 is then $\tilde{A} = A_0 Q\left(\tilde{\theta}\right)$. Our shock estimates associated with $\tilde{A}$ are obtained from our $[\mathrm{GDP}(t), \mathrm{CPI}(t)]'$ data and the inverse of $\tilde{A}$, i.e.

$$\begin{bmatrix} \varepsilon_{S,t} \\ \varepsilon_{D,t} \end{bmatrix} = \begin{bmatrix} +2.1 & -2.4 \\ +1.6 & +2.4 \end{bmatrix} \begin{bmatrix} \mathrm{GDP}(t) \\ \mathrm{CPI}(t) \end{bmatrix} \tag{4}$$

The top-right panel of figure 8 plots the series of inferred supply and demand shocks for the United States based on equations 3 and 4, along with indicators for the dates of the eight selected events we have detailed in table 2. Table 7 contains the United States GDP and CPI shocks associated with each of these events, and the inferred supply and demand shocks.

Table 7: United States shocks for selected events

Events (1) \ shocks	GDP	CPI	Supply	Demand
1. Iraq annexes Kuwait	–1.3	1.4	–6.0	1.4
2. Asian Financial Crisis	0.1	–0.5	1.4	–0.9
3. Al-Qaeda attacks US	–1.6	–0.4	–2.4	–3.5
4. Global Financial Crisis	–1.2	–0.6	–1.0	–3.2
5. COVID-19 pandemic	–4.7	–1.4	–6.3	–11.1
6. Russia/Ukraine war	–0.5	1.9	–5.4	3.7
7. Global trade disruptions	–0.8	0.4	–2.5	–0.4
8. US-Israel/Iran conflict	–0.4	0.7	–2.4	0.9

Note: (1) The event dates are in table 1 and are indicated in the figures.
Source: Authors' estimations

In nearly all cases, the supply and demand shocks are consistent with the nature of the events from the perspective of the United States. That is, the shocks by event

[15] The distribution requirements are automatically maintained because $[\mathrm{GDP}(t), \mathrm{CPI}(t)]' = A_j \varepsilon_{j,t} = A_0 Q_j Q_j' \varepsilon_{0,t} = A_0 \varepsilon_{0,t}$.

[16] Alternatives are simply retaining the whole set, given all are valid models that conform to the chosen sign restrictions, or using external information to select which models have shock results that best correspond to known shock events.

number in table 7 with brief context are: (1) large negative supply shock, consistent with the resulting spike in oil prices; (2) immaterial shocks, consistent with the main effects being in South-East Asia; (3) moderate negative supply and demand shock, consistent with sharp falls in consumer and business sentiment; (4) a moderate negative demand shock at the onset (with further negative demand shocks in following months), consistent with severe credit disruptions; (5) huge negative supply and demand shock, consistent with stay-at-home directives affecting production and consumption; (6) a large negative supply shock, consistent with sharp rises in oil prices; (7) a moderate negative supply shock, consistent with reduced volume and increased price of imports; (8) a moderate negative supply shock. One potential anomalous result is the positive demand shock associated with the Russia/Ukraine war, which we discuss further below.

There are several important caveats that should be kept in mind when interpreting our results. First, as mentioned earlier, they are for a single representative model. Other rotations would produce variations of our demand and supply shock estimates while remaining consistent with the required signs.

Second, because our results are inferred from forecast revisions, they should not be viewed as supply and demand shocks that might be estimated ex-post from a structural model with a full historical dataset, but as supply and demand shocks *perceived* by the set of forecasters in real time (i.e. between CE survey dates). Our shock series therefore reflect any dynamics associated with real-time forecasting, such as initial under- or over-reactions by forecasters to shock events, and corrections in subsequent surveys. The consistency of our estimated shock series with respect to known shock events suggests they provide a useful indication of fundamental shocks in real time, and they should be helpful for estimating a shock series within a structural model. Another useful exercise would be to compare structural model results with our shock series.

Third, because we only have two macroeconomic variables in our dataset, i.e. our GDP01mn and CPI01mn series, we can only infer two shocks, i.e. demand and supply shocks in our application. Hence, those shocks will inevitably reflect other fundamental shocks that would in principle be identifiable using a wider dataset. For example, without monetary and/or fiscal variables in our dataset, fundamental shocks from those aspects will be subsumed into our inferred shocks. That is, unanticipated easings (tightenings) in monetary and/or fiscal policy tend to raise (lower) GDP and CPI, which are the same responses as positive (negative) demand shocks. In that regard, our demand shocks will include demand induced by monetary policy and/or fiscal policy shocks. The COVID episode is a case in point: our estimates show a large initial negative demand shock on the event, subsequently followed by a series of positive demand shocks. However, the later likely reflects the prompt and large stimulus from monetary and fiscal policy following the onset of COVID. Moreover, our estimated positive demand shock following the onset of the Russia/Ukraine war is unlikely to be related to that event (the associated oil price shock is a classic supply shock according to economic principles), but due to other shocks that we cannot explicitly account for with only the two macroeconomic variables of GDP and CPI in our dataset.

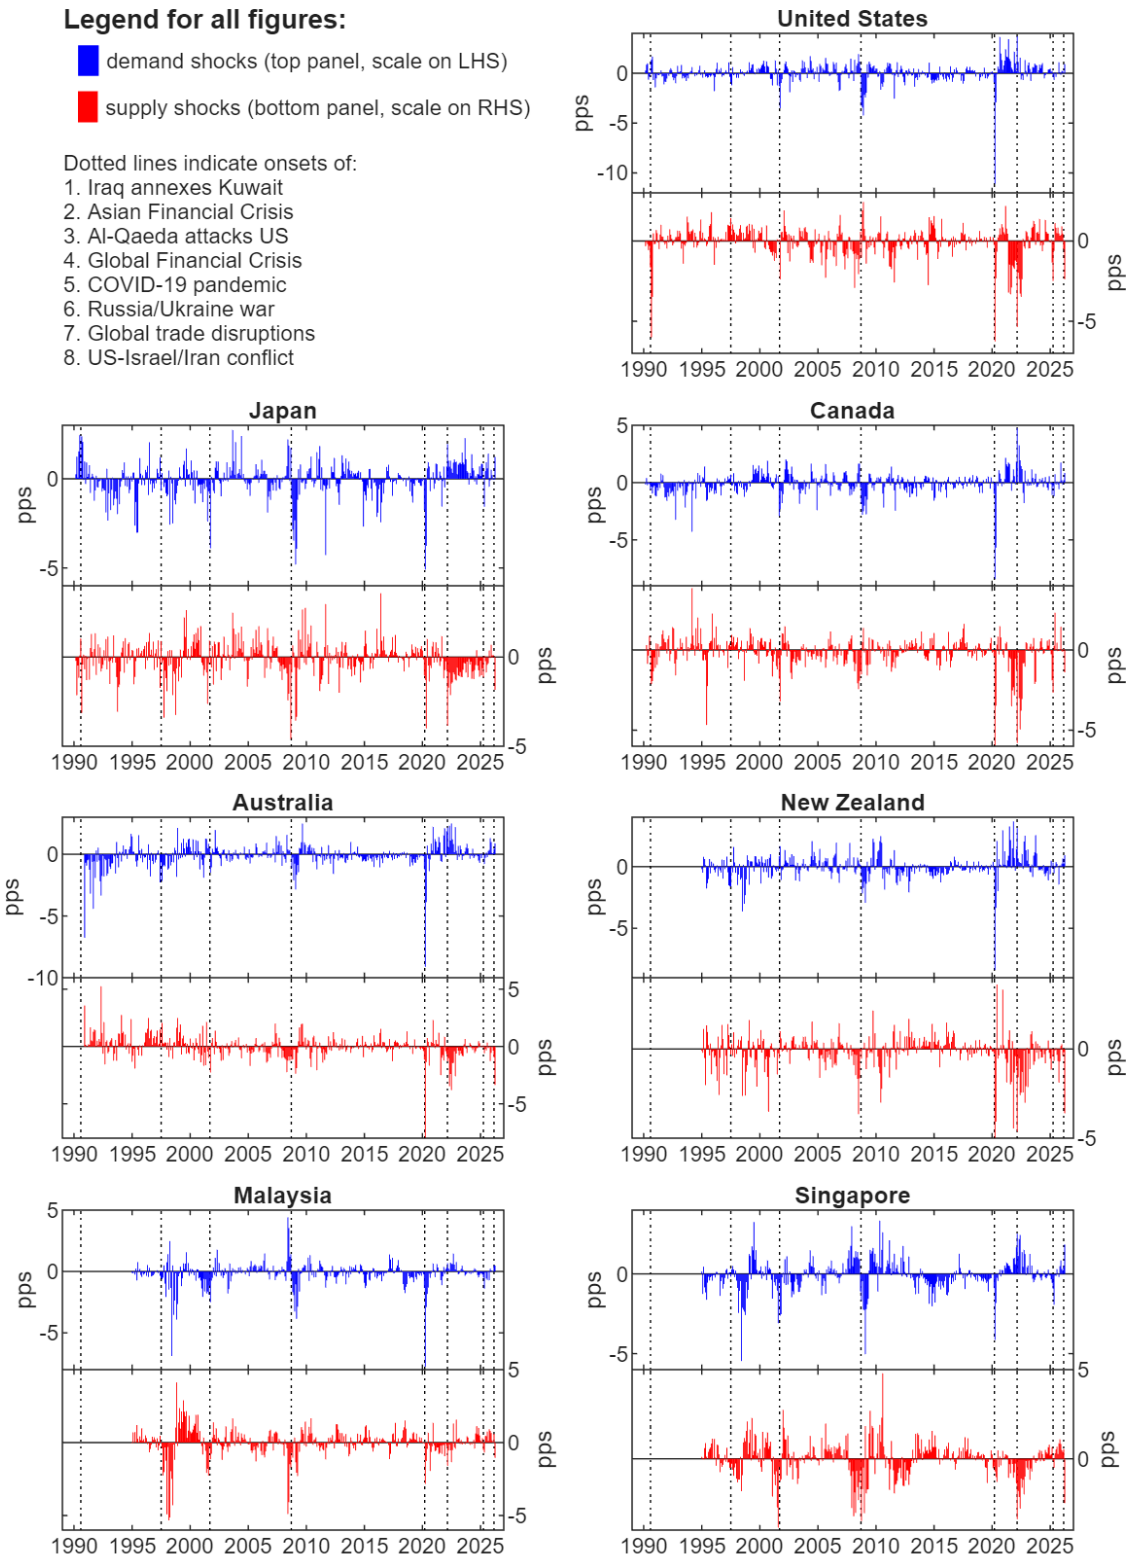

Figure 8: Estimates of demand and supply shocks for the first seven economies in our dataset.

Source: Authors' estimations

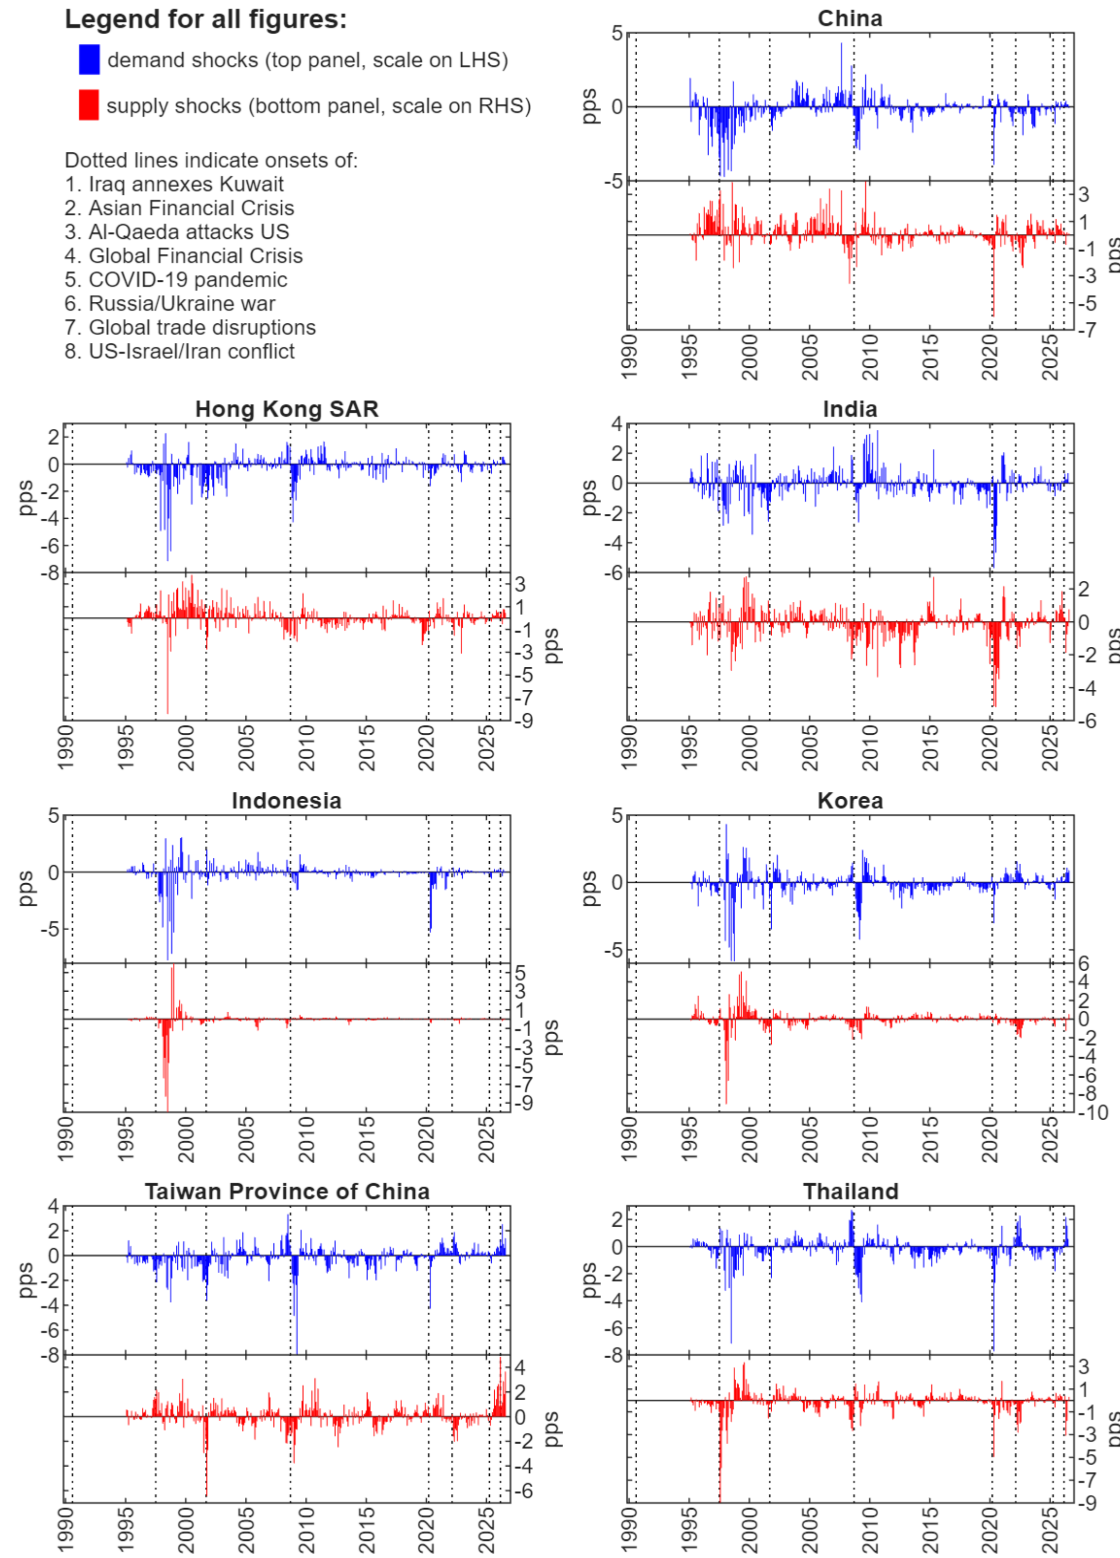

Figure 9: Estimates of demand and supply shocks for the second seven economies in our dataset.

Source: Authors' estimations

The remainder of the panels in figure 8 plot the results for six more economies, and figure 9 plots the results for the remaining seven economies we cover in our dataset. Many of the shocks already noted for the United States appear to some extent in the other figures. In particular, COVID was generally viewed as a large shock to supply and demand, but less so in Hong Kong SAR and Korea compared to other economies. In Asian economies professional forecasters, not surprisingly, regarded the Asian Financial Crisis as a large negative supply and demand shock albeit with lagged effects in some economies except in China (i.e. Chinese mainland) and Taiwan Province of China where it was perceived as mainly a demand shock.

5.2 Comparing uncertainty measures

In this section we compare our measures of uncertainty with the BBD dataset for the Pacific-rim and Asia-Pacific economies, which we downloaded from the website www.policyuncertainty.com. This site contains the uncertainty measures produced by BBD along with the uncertainty measures produced by other authors following the BBD method.

The list of economies in the BBD dataset that overlap with the economies in our dataset are the United States, Japan, Canada, Australia, New Zealand, Singapore, China, Hong Kong SAR, India, and Korea. We have also included Pakistan from the BBD dataset, which is the one Asian economy not in our dataset.

For comparison to the BBD data, we calculate a single series GDP01/CPI01sd from our GDP and CPI uncertainty data GDP01sd and CPI01sd, i.e. GDP01/CPI01sd $= \sqrt{[\text{GDP01sd}]^2 + [\text{CPI01sd}]^2}$. We hereafter refer to the GDP01/CPI01sd series as CE uncertainty (CEU).

Figure 10 plots, for each economy mentioned above, the BBD uncertainty series and our GDP01/CPI01sd series, both standardized as z-scores, i.e. $[x - \text{mean}(x)]/\text{stdev}(x)$ for the time series x, so they are on the same scale.

The notable aspect from each of the economy panels in figure 10 is that the BBD data and our series have quite different profiles, which is most apparent from the differing peaks in the series. That is, with the exception of New Zealand, the BBD series have maximums near the end of the sample, associated with the onset of global trade policy uncertainty on 2-Apr-2025. Most of our series have maximums after the onset of the COVID-19 pandemic, with the Hong Kong SAR and Korea maximums after the onset of the Asian Financial Crisis, and China prior to then.

The differences between the BBD and our series is also evident statistically. Hence, in table 8 we have provided two correlation statistics and their associated p-values for the null hypothesis of zero versus an alternative of positive correlation. The Pearson entry is the standard linear correlation coefficient. The Kendall Tau value is a non-parametric rank correlation statistic based on concordant and discordant pairs, which is more appropriate for non-normally distributed datasets like those in figure 10 (and also for small samples and/or with many ties, although neither aspect applies in our application). The Kendall results in all cases parallel those for the Spearman rank correlation coefficient, so we have not included results for the latter.

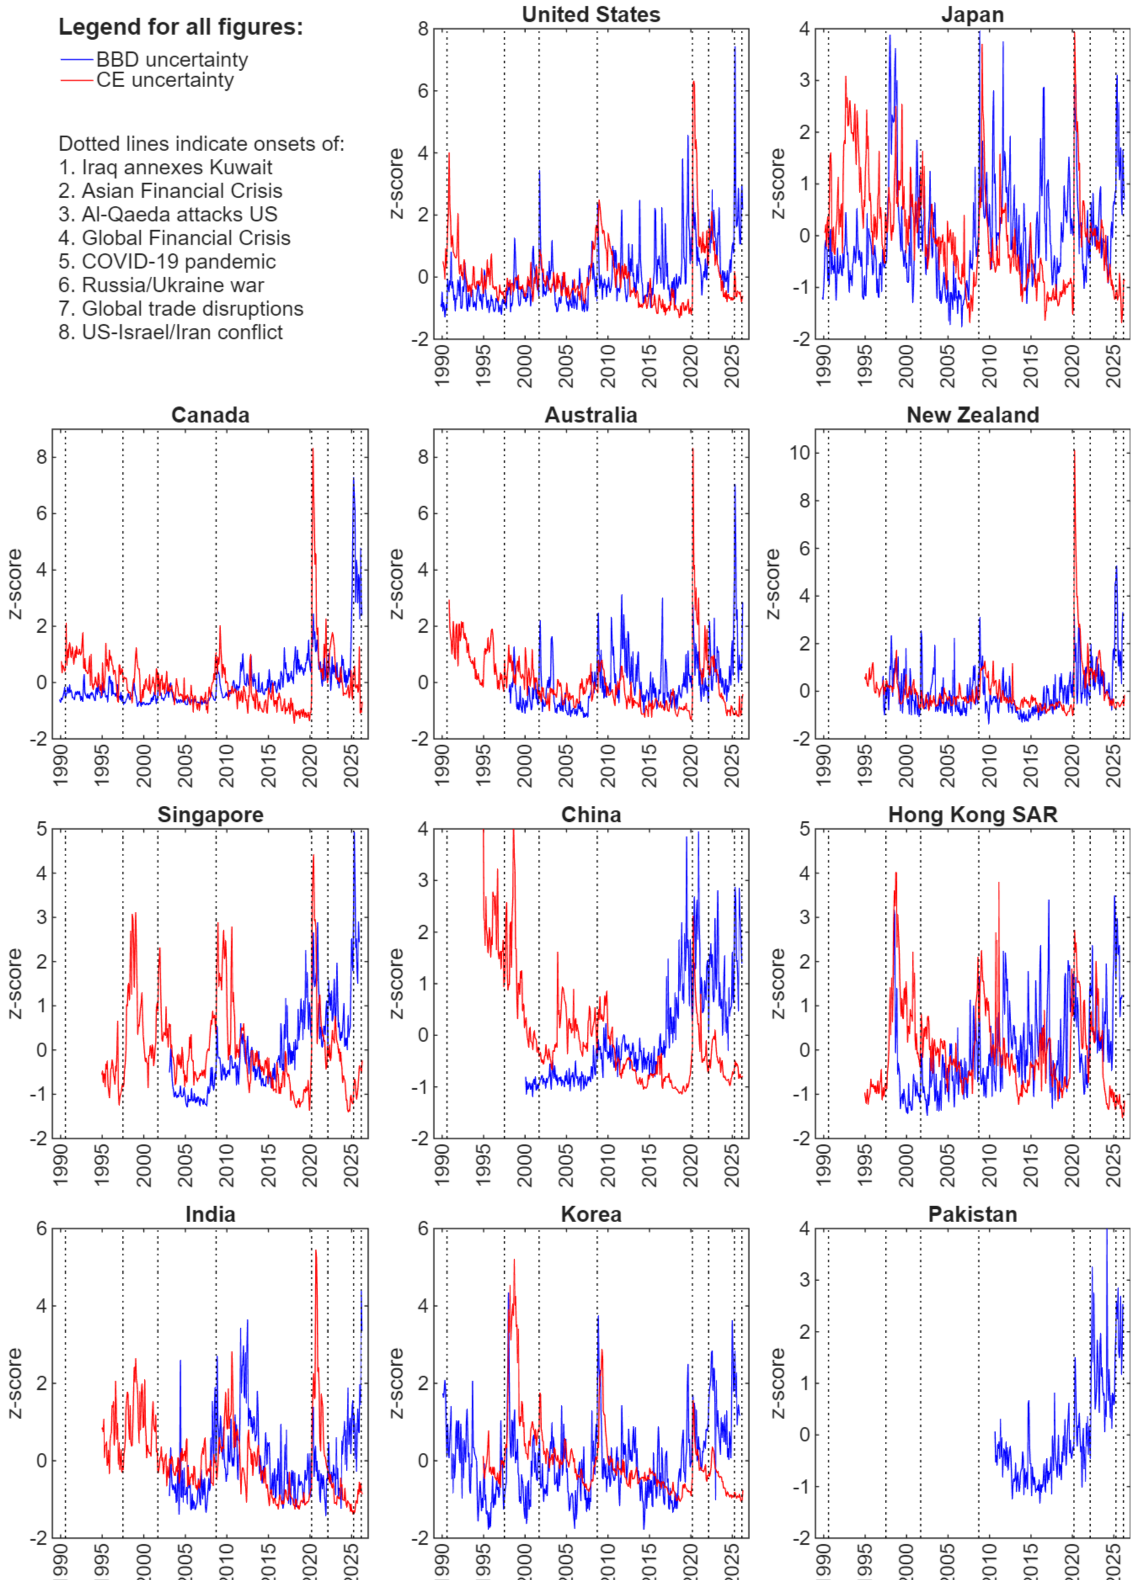

Figure 10: Comparison of BBD data and CE uncertainty estimates.

Source: www.policyuncertainty.com and authors' estimations

The Pearson p-values indicate six significant positive correlations, although the associated coefficient themselves are only around 0.1, apart from 0.33 for New Zealand. The Kendall p-values indicate that New Zealand's correlation of 0.22 is strongly significant, with the 0.05 correlation for India just within the 10% level of significance. The remaining cases reject a positive correlation.

Another statistical perspective showing the difference between the BBD series and ours is contained in the right side of table 8, i.e. testing for a time trend. For the 11 BBD series, including Pakistan, all economies except New Zealand show statistically significant positive time trends (and all to 1% significance, except India to 10%). The "slope" entries are the estimated time-trend contributions for each economy over the respective samples available. Specifically, it is the estimated monthly coefficient for the time trend multiplied by the number of points used in the regression. This calculation therefore gives the time-trend contribution to the difference between the start and end of the available sample. For our 10 series, only New Zealand has a significant (to the 10% level) slope, which is positive.

Table 8: Statistical comparison of CE and BBD uncertainty

	BBD/CEU correlations				Estimated time trends (1)			
	Pearson		Kendall		BBD		CEU	
Item	coeff.	p-v%	Tau	p-v%	slope	p-v%	slope	p-v%
United States	0.13	0.4	0.02	32.0	1.9	0.0	–0.2	64.7
Japan	0.09	3.7	0.03	21.3	0.8	0.4	–1.8	100.0
Canada	0.11	1.4	–0.04	87.9	2.0	0.0	–0.5	87.4
Australia	0.08	7.1	–0.01	55.9	1.3	0.0	–0.3	78.5
New Zealand	0.20	0.0	0.17	0.0	0.8	3.4	0.4	15.1
Singapore	0.03	33.0	–0.03	79.1	2.7	0.0	–0.6	91.7
China	–0.21	100.0	–0.28	100.0	2.8	0.0	–0.8	100.0
Hong Kong SAR	–0.06	85.1	–0.11	99.9	1.7	0.0	–1.2	99.6
India	0.01	42.7	0.05	10.6	0.5	7.9	–0.7	96.5
Korea	0.11	1.6	0.00	50.0	1.4	0.0	–1.7	100.0
Pakistan	n/a	n/a	n/a	n/a	0.7	0.0	n/a	n/a

Note: (1) Slope is explained in the paragraph immediately above this table, and the slope p-values (p-v%) use Newey-West HAC estimates.
Source: Authors' estimations

Figure 11 plots, again as z-scores, our CEU series for the United States and the JLN macroeconomic and financial measures of uncertainty. These are available from www.sydneyludvigson.com/macro-and-financial-uncertainty-indexes.[17] The third measure from that website is the real macroeconomic uncertainty series from LMN, which uses just the real activity variables from the macroeconomic series from JLN. For clarity we have not included the LMN real macroeconomic uncertainty measure in figure 11 because as a z-score it is very similar to the macroeconomic measure over our sample period.

[17] We have used the JLN and LMN results for the 3-month horizon results, but correlations between the 1-, 3-, and 12-month series are close to 1.

The profile of the JLN macroeconomic series is very similar to our CEU, particularly in the peaks around the Al-Qaeda/US, GFC, and COVID events, and somewhat around the Iraq/Kuwait event. The financial series is most similar around the GFC and COVID peaks.

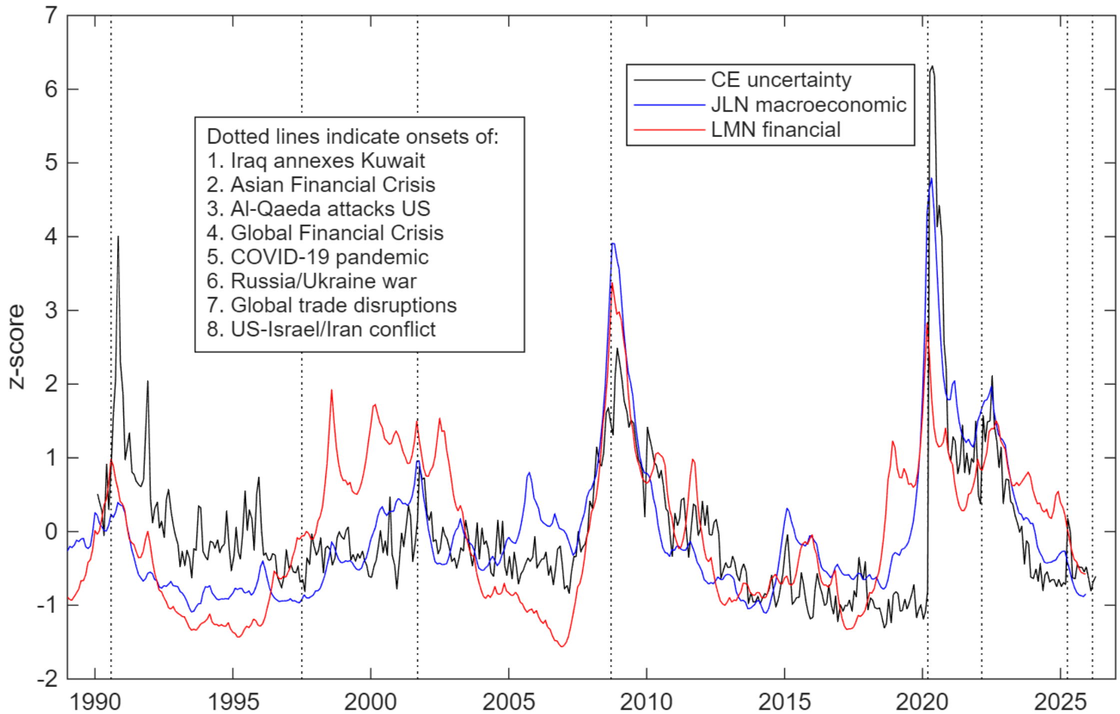

Figure 11: Comparison of JLN and LMN uncertainty estimates with our CEU series.

Source: www.sydneyludvigson.com/macro-and-financial-uncertainty-indexes and authors' estimations

Table 9 contains the correlations for the three JLN/LMN series, our United States CEU series, and the individual GDP01sd and CPI01sd series plotted in the top-right panel of figure 10. The standard correlations (Pearson) are high, and the rank correlation (Kendall) results are moderate. All are extremely significant (a maximum of 2e-6, i.e. 2e-4%), so the p-values are not reported in the table.

The highest correlations are between the JLN macroeconomic uncertainty and our CEU, consistent with the CEU being composed from surveys that are themselves forecasts of the key macroeconomic variables for the economy, i.e. GDP and CPI. For the JLN real uncertainty measure, the correlation is higher with our GDP01sd series than for CPI01sd, consistent with GDP01sd representing uncertainty about the real variable GDP. The correlations of LMN financial uncertainty with our series are generally lowest, reflecting the difference in the concepts being measured. For comparison, the correlation of JLN macroeconomic uncertainty with LMN financial uncertainty is 0.69 Pearson and 0.45 Kendall.

Table 9: Correlations of JLN and LMN with CE uncertainty

Correlations (1)	Pearson			Kendall		
JLM&LMN\CE	GDP	CPI	CEU	GDP	CPI	CEU
Macroeconomic	0.64	0.52	0.69	0.30	0.31	0.36
Financial	0.49	0.37	0.51	0.33	0.17	0.28
Real	0.64	0.35	0.63	0.20	0.15	0.20

Note: (1) All correlations are well below 1% significance.

Source: Authors' estimations

Note that we have also undertaken the same time trend estimations explained for our BBD analysis. In all cases they show statistically insignificant and economically immaterial trends for the CE series, and significant but immaterial trends for the JLN series (i.e. a maximum, for the real series, of 0.16 z-score units over the 35-year period).

The first key theme from the results and discussion above are that the BBD economic policy uncertainty measures represent a different perspective on uncertainty than our macroeconomic uncertainty measures. As such, our series should not be used as a policy uncertainty measure. Second, the BBD series for the United States is also quite different from the JLN and LMN macroeconomic and financial uncertainty measures, while our United States uncertainty measures compare well to the JLN and LMN macroeconomic series, i.e. both our overall CEU series with the JLN measure, and our GDP series to the LMN real measure. The implication is that the uncertainty measures for the economies we cover should best be interpreted as measures of macroeconomic uncertainty, which is also most consistent with the data they are derived from. As such, our series offer measures of macroeconomic uncertainty for a wider set of economies than JLN and LMN, and also a complementary perspective to the BBD measures of policy uncertainty. Both are useful aspects for economic policy setting and applied empirical research.

6 Conclusion

In this article, we have used the Consensus Economics survey of professional economic forecasters to create a dataset of shocks and uncertainty, for GDP growth and CPI inflation, for 12 Asia-Pacific and two Pacific-rim economies. The contribution of our methodology, which we detail for the United States and then apply to all economies, is to explicitly correct for intra-year patterns in the month-by-month changes to mean forecasts, used to derive shocks, and in the standard deviations of forecasts, used to derive uncertainty. The patterns are inherent from the calendar-year annual-average percent change basis of the Consensus Economics forecasts for GDP and CPI.

As an illustration of using the GDP and CPI shocks from our dataset, we have applied the VAR sign restrictions method to infer the perceptions from professional forecasters of fundamental demand and supply shocks for each of the economies we cover. The results are consistent with the nature of eight major economic events over our sample period, e.g. the 1997 Asian Financial Crisis was regarded as a large negative supply and demand shock within many Asian economies, and the 2020 onset of COVID was regarded similarly across all economies in our dataset. The US-Israel/Iran conflict has been regarded as a moderate negative supply shock.

To compare our GDP and CPI uncertainty measures to the widely used measures developed by Baker, Bloom, and Davis (2016) we have combined them into a single measure. The differences between the BBD newspaper-based measures and our survey-based measures, for each of the 10 economies with data in common, are very apparent both visually and statistically. Notably, the Baker et al. (2016) measures often rank uncertainty associated with the 2025 global trade disruptions well above COVID, and the measures generally have significant positive time trends. Neither aspect features in any of our results. We have further shown that our United States uncertainty measures relate well to the econometric-based macroeconomic uncertainty measures developed by Jurado, Ludvigson, and Ng (2015) and Ludvigson, Ma, and Ng (2021), suggesting that our measures for other economies provide a useful gauge of that uncertainty concept.

There are several potential extensions to our analysis in this article. The first is extending the range of economies in our shock and uncertainty dataset to include the Western, Latin American, and Eastern European countries for which data are already available and updated monthly by Consensus Economics. Indeed, the breadth and length of the Consensus Economics dataset, along with its consistent compilation, was one of our motivations for using the CE data.

Second, our analysis could potentially be improved with alternative adjustments for the intra-year pattern of the data and/or using wider sets of data available from Consensus Economics. Regarding the former, further developing the diffusion process mentioned in section 2 and discussed further in appendix A.4 is one consideration. Regarding wider datasets, unit record data for survey respondents could be used rather than the means and standard deviations for GDP and CPI already calculated by Consensus Economics.

Third is to add a monetary policy perspective to our dataset. In principle, that could be based on the 3-month and 10-year interest rate forecasts surveyed by Consensus Economics. However, one aspect to contend with is the lack of sufficient data in those categories for many Asian economies. Another is that, even where complete interest rate data are available, the lower-bound constraint and unconventional methods of stimulus associated with the Global Financial Crisis and COVID would need to be appropriately allowed for.

Once we have a suitable monetary policy variable, the fundamental shocks from our approach, based on ex-ante data, may be compared and contrasted with those obtained from standard macroeconomic VARs, based on ex-post data. Our data may also prove useful as proxy variables to assist in identifying fundamental macroeconomic shocks.

References

Abel, J., R. Rich, J. Song, and J. Tracy (2016). The measurement and behavior of uncertainty: evidence from the ECB survey of professional forecasters. *Journal of Applied Econometrics 31*(3), 533–550. doi.org/10.1002/jae.2430.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics 131*(4), 1593–1636. doi.org/10.1093/qje/qjw024.

Ballantyne, A., C. Gillitzer, D. Jacobs, and E. Rankin (2016). Disagreement about inflation expectations. *Reserve Bank of Australia Research Discussion Paper RDP 2016-02*.

Bruns, M. and H. Lütkepohl (2026). Review of proxy vector autoregressive analysis. *Reviews of Economic Literature 1*, 1–37. doi.org/10.21627/crv7mx92.

Cascaldi-Garcia, D., C. Sarisoy, J. M. Londono, B. Sun, D. D. Datta, T. Ferreira, O. Grishchenko, M. R. Jahan-Parvar, F. Loria, S. Ma, M. Rodriguez, I. Zer, and J. Rogers (2023). What is certain about uncertainty? *Journal of Economic Literature 61*(2), 624–654. doi.org/10.1257/jel.20211645.

D'Amico, S. and A. Orphanides (2008). Uncertainty and disagreement in economic forecasting. *Federal Reserve Board Finance and Economics Discussion Series 2008-56*.

Denis, S. and P. Kannan (2013). The impact of uncertainty shocks on the UK economy. *International Monetary Fund Working Paper WP/13/66*. doi.org/10.5089/9781475552690.001.

Dery, C. and A. Serletis (2026). Macroeconomic shocks and business cycle dynamics: a comparative analysis. *Oxford Economic Papers 78*(1), 233–258. doi.org/10.1093/oep/gpaf029.

Fernández-Villaverde, J. and P. A. Guerrón-Quintana (2020). Uncertainty shocks and business cycle research. *Review of Economic Dynamics 37*(S1), S118–S146. doi.org/10.1016/j.red.2020.06.005.

Fry, R. and A. Pagan (2011). Sign restrictions in structural vector autoregressions: a critical review. *Journal of Economic Literature 49*(4), 938–960. doi.org/10.1257/jel.49.4.938.

Greig, L., A. Rice, T. Vehbi, and B. Wong (2018). Measuring uncertainty and its impact on a small open economy. *The Australian Economic Review 51*(1), 87–98. doi.org/10.1111/1467-8462.12255.

Gürkaynak, R. S., B. Sack, and E. T. Swanson (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking 1*(1), 55–93.

Gürkaynak, R. S., B. Sack, and E. T. Swanson (2007). Market-based measures of monetary policy expectations. *Journal of Business and Economic Statistics 25*(2), 201–212. doi.org/10.1198/073500106000000387.

Harris, B., S. Weenink, C. Grubb, and G. Turnip (2025). How does macroeconomic uncertainty impact the New Zealand economy. *Reserve Bank of New Zealand Analytical Note Series AN2025-06*.

Jurado, K., S. C. Ludvigson, and S. Ng (2015). Measuring uncertainty. *American Economic Review 105*(3), 1177–1216. doi.org/10.1257/aer.20131193.

Ludvigson, S. C., S. Ma, and S. Ng (2021). Uncertainty and business cycles: exogenous impulse or endogenous response. *American Economic Journal: Macroeconomics 13*(4), 369–410. doi.org/10.1257/mac.20190171.

Morley, J., J. Tian, and B. Z. Wang (2026). Disagreement over the nature of macroeconomic shocks. *Centre for Applied Macroeconomic Analysis Working Paper 21/2026*.

Ozturk, E. O. and X. S. Sheng (2018). Measuring global and country-specific uncertainty. *Journal of International Money and Finance 88*, 276–295. doi.org/10.1016/j.jimonfin.2017.07.014.

Ploenzke, M., R. Rich, and J. Tracy (2015). What does disagreement tell us about uncertainty? Liberty Street Economics. https://libertystreeteconomics.newyorkfed.org/2015/01/what-does-disagreement-tell-us-about-uncertainty/.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2, pp. 71–162. Elsevier. doi.org/10.1016/bs.hesmac.2016.03.003.

Swanson, E. T. (2021). Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics 118*, 32–53. doi.org/10.1016/j.jmoneco.2020.09.003.

Verbrugge, R. J. and S. Zaman (2021). Whose inflation expectations best predict inflation? *Economic Commentary, Federal Reserve Bank of Cleveland 2021-19*, 1–7. doi.org/10.26509/frbc-ec-202119.

A Year Advancement Pattern (YAP) mathematics

We use this appendix to establish the principles of Year Advancement Patterns (YAPs) for the CE data for the monthly consumer price index. The same calculation applies for GDP but on a quarterly basis. Section A.1 begins with the AAPC expressions underlying individual forecasts for the current and following years. Sections A.2 uses the AAPC expressions within the usual standard deviation calculation, as would be applied by CE, to show why a YAP should be inherent within the CE standard deviation data. Section A.3 calculates the mean, shows why the mean change contains the relevant forecast information from the CE data, and also why a YAP will be inherent in the mean-change data. In section A.4, we use a diffusion representation of monthly forecasts with the standard deviation and mean-change expressions to illustrate the monotonic nature of the YAP which, as discussed in section 2, we respect with our empirical YAP estimations. Finally, in section A.5 we discuss why prior uses of CE data to calculate uncertainty measures retain a YAP.

A.1 AAPC definitions

We denote the set of forecast responses obtained by CE, for each given survey time t, and for the current and following year, as $\{S_{0,k}\,(t)\}_{k=1}^{K}$ and $\{S_{1,k}\,(t)\}_{k=1}^{K}$, where K is the number of respondents in each set.

The AAPC definition that underlies each of the surveyed forecasts for the current year may be represented as:[18]

$$S_{0,k}(t) = \frac{\sum_{n=1}^{N} Z_0(t,n) + \sum_{n=N+1}^{12} \mathbb{E}_t\left[Z_{0,k}(t,n)\right]}{\sum_{n=1}^{12} Z_{-1}(t,n)} - 1 \tag{5}$$

where t is the calendar date of the survey, $S_{0,k}(t)$ is the survey expectation data for the current year (which is the year of the date t), N is the number of months in the current year for which consumer price index level data $Z_{0,k}(t,n)$ has already been published, $\mathbb{E}_t\left[Z_{0,k}(t,n)\right]$ are the expected index levels for the remainder of the current year, and $Z_{-1,k}(t,n)$ are the index levels for the previous year.

For clarity, here and in subsequent expressions, our notation implicitly assumes that all data for the previous year are published, and published data in the previous and current year are not subject to revision. Hence, what otherwise could potentially be individual values for each forecast response, i.e. $Z_{-1,k}(t,n)$ and $Z_{0,k}(t,n)$, are already replaced with single realized values $Z_{-1}(t,n)$ and $Z_0(t,n)$ common to all forecasters.

In practice, for the earlier months of the current year, one or more months at the end of the previous year would be "yet-unknown", i.e. still expected values because index levels are published with a lag. Similarly, the first summation in the numerator of equation 5 includes yet-unknown index levels for month N, and potentially some earlier months. The materiality of the approximation implicit in our notation would increase with longer publication lags and/or higher potential for revisions, and these aspects will differ among different economies and between GDP and CPI.

The AAPC for the following year is:

$$S_{1,k}(t) = \frac{\sum_{n=1}^{12} \mathbb{E}_t\left[Z_1(t,n)\right]}{\sum_{n=1}^{N} Z_0(t,n) + \sum_{n=N+1}^{12} \mathbb{E}_t\left[Z_{0,k}(t,n)\right]} - 1 \tag{6}$$

While seemingly analogous to $S_{0,k}(t)$, the presence of forecast components on the numerator and denominator of $S_{1,k}(t)$ introduces more potential complexity in the dynamics of $S_{1,k}(t)$, i.e. it can vary with changes in the forecast components of $S_{0,k}(t)$ and $\mathbb{E}_t\left[Z_1(t,n)\right]$.

A.2 YAP from standard deviation data

CE calculate the standard deviations of the sets of forecast responses $\{S_{0,k}(t)\}_{k=1}^{K}$ and $\{S_{1,k}(t)\}_{k=1}^{K}$. We denote those calculations as $\sigma_0(t) = \text{stdev}\left[\{S_{0,k}(t)\}_{k=1}^{K}\right]$ and $\sigma_1(t) =$

[18]We express the definition here as the annual sum percent change, which is easier for notation in the expressions we subsequently use. The annual sum percent change is obviously mathmatically equivalent to the AAPC, simply by scaling the numerator and denominator by 1/12 to get the average index levels for each year.

stdev$\left[\{S_{1,k}(t)\}_{k=1}^{K}\right]$, where for clarity in notation we use "stdev" to denote the usual standard deviation operator, i.e. stdev$\left[\{x\}_{k=1}^{K}\right] = \text{sqrt}\left[\frac{1}{K}\sum_{k=1}^{K}(x_k - \text{mean}[x])^2\right]$.

Using equation 5, $\sigma_0(t)$ is:

$$\sigma_0(t) = \frac{\text{stdev}\left(\displaystyle\sum_{n=N+1}^{12} \mathbb{E}_t[Z_{0,k}(t,n)]\right)}{\displaystyle\sum_{n=1}^{12} Z_{-1}(t,n)} \tag{7}$$

which is just the standard deviation of the forecast elements divided by $\sum_{n=1}^{12} Z_{-1}(t,n)$, because the latter and $\sum_{n=1}^{N} Z_0(t,n)$ are known scalar quantities (hence non-random). The YAP for $\sigma_0(t)$ is the mechanical tendency for it to decrease as the year advances simply because the standard deviation data for a given month reflects only the number of remaining months of the year that are still subject to forecasts.

The standard deviation expression for the following year, $\sigma_1(t) = \text{stdev}\left[\{S_{1,k}(t)\}_{k=1}^{K}\right]$, is simply the standard deviation operator applied to equation 6 as a whole. Unlike for $\sigma_0(t)$, the expression for $\sigma_1(t)$ does not simplify further due to forecast elements (i.e. random variables) being present in both the numerator and denominator. Specifically, for a set of positive random variables x and y, stdev$(x/y) \neq$ stdev(x)/stdev(y).

A.3 YAP from changes in mean data

CE calculate the means of the sets of forecast responses $\{S_{0,k}(t)\}_{k=1}^{K}$ and $\{S_{1,k}(t)\}_{k=1}^{K}$. We denote those calculations as $\mu_0(t) = \text{mean}\left[\{S_{0,k}(t)\}_{k=1}^{K}\right]$ and $\mu_1(t) = \text{mean}\left[\{S_{1,k}(t)\}_{k=1}^{K}\right]$, where for clarity in notation we use "mean" to denote the usual mean operator, i.e. mean$\left[\{x\}_{k=1}^{K}\right] = \frac{1}{K}\sum_{k=1}^{K} x_k$.

Using equation 5, $\mu_0(t)$ is:

$$\mu_0(t) = \frac{\displaystyle\sum_{n=1}^{N} Z_0(t,n) + \text{mean}\left(\displaystyle\sum_{n=N+1}^{12} \mathbb{E}_t[Z_{0,k}(t,n)]\right)}{\displaystyle\sum_{n=1}^{12} Z_{-1}(t,n)} - 1 \tag{8}$$

Hence, $\mu_0(t)$ condenses to a constant plus the mean of the forecast elements divided by $\sum_{n=1}^{12} Z_{-1}(t,n)$, because the latter and $\sum_{n=1}^{N} Z_0(t,n)$ are known scalar quantities.

Taking the first difference of the mean data series, i.e. $\Delta\mu_0(t) = \mu_0(t) - \mu_0(t-1)$, removes the levels $Z_0(t,n)$ that $\mu_0(t)$ and $\mu_0(t-1)$ have in common, leaving just the elements associated with forecasts. Specifically, multiplying $\mu_0(t) - \mu_0(t-1)$ by the constant $\sum_{n=1}^{12} Z_{-1}(t,n)$ to leave just the numerator elements for clarity, gives:

$$\begin{aligned}
& [\mu_0(t) - \mu_0(t-1)] \sum_{n=1}^{12} Z_{-1}(t,n) \\
= \; & \sum_{n=1}^{N} Z_0(t,n) + \text{mean}\left(\sum_{n=N+1}^{12} \mathbb{E}_t [Z_{0,k}(t,n)] \right) \\
& - \sum_{n=1}^{N-1} Z_0(t,n) - \text{mean}\left(\sum_{n=N}^{12} \mathbb{E}_{t-1} [Z_{0,k}(t-1,n)] \right) \\
= \; & Z_0(t,N) - \text{mean}\{\mathbb{E}_{t-1}[Z_{0,k}(t-1,N)]\} \\
& + \text{mean}\left(\sum_{n=N+1}^{12} \mathbb{E}_t [Z_{0,k}(t,n)] \right) - \text{mean}\left(\sum_{n=N+1}^{12} \mathbb{E}_t [Z_{0,k}(t-1,n)] \right)
\end{aligned} \tag{9}$$

which we have decomposed into two terms. The first term represents a forecast error, i.e. the realized index level for month N less the mean forecast for month N made in the previous month. The second term is a change in the mean forecasts for months $N+1$, i.e. the forecasts for the remainder of the current year updated to time t versus the forecast for those same months made in the previous month, time $t-1$.

Like for $\sigma_0(t)$, the YAP for $\Delta\mu_0(t)$ arises from only $1 + 12 - N$ months of the current year being subject to change, in this case one realized month versus its previous forecast, and changes in forecasts for the remaining $12 - N$ months.

The mean for the following year, $\mu_1(t) = \text{mean}\left[\{S_{1,k}(t)\}_{k=1}^{K}\right]$, is simply the mean operator applied to equation 6 as a whole. Like for the following-year standard deviation, the expression cannot be simplified further mathematically because forecast elements are present in both the numerator and denominator. In this case, for a set of positive random variables x and y, $\text{mean}(x/y) \neq \text{mean}(x)/\text{mean}(y)$.

The mean-change expression for the following year is $\Delta\mu_1(t) = \mu_1(t) - \mu_1(t-1)$. This again cannot be simplified further, for the reason explained above and the additional aspect that the denominators are different due to the advancement of N in the summations.

A.4 Nature of the standard deviation and mean-change data

Sections A.2 and A.3 show why YAPs should be inherent in CE standard deviations and mean-changes. The purpose of this section is to show that the YAPs should be monotonic. We illustrate this with a Monte-Carlo exercise, i.e. by simulating forecasts from individual respondents, and calculating the summary results from those simulations. We do not use the simulation results themselves for our YAP estimations, only the principle that our YAP estimations should be based on monotonic functional forms. However, we briefly discuss in section A.5 how simulated functional forms could be used in future work.

For our simulation, we represent each respondent by a random walk with drift to simulate monthly log changes r_t, i.e. $r_t = \delta + r_{t-1} + \varepsilon_t$, $\varepsilon_t \sim N(0, \sigma^2)$, cumulate those values in the log index level $R_t = R_{t-1} + r_t$, and apply $Z_t = \exp(R_t)$ to obtain the index levels.[19] For the base model, we use $\delta = 0.02/12$, hence trend growth of approximately 2% per year, and $\sigma = 0.01/12$, hence approximately 1% annual standard deviation around trend growth. We first fix r_t to $\delta, \ldots, 12\delta$ as the known values for each month of the previous year. For month M in the current year, we fix the known values of r_t for previous months of the current year to be $13\delta, \ldots, 12\delta + (M-1)\,\delta$. The stochastic simulated values $r_t = \delta + r_{t-1} + \varepsilon_t$ begin at month M in the current year, continuing through to December of the following year. Figure A.1 shows examples of the simulations, i.e. two simulated paths as at January (i.e. $M = 1$, so only values for the previous year are known), and June (i.e. $M = 6$, so values are known for the previous year and up to May for the current year).

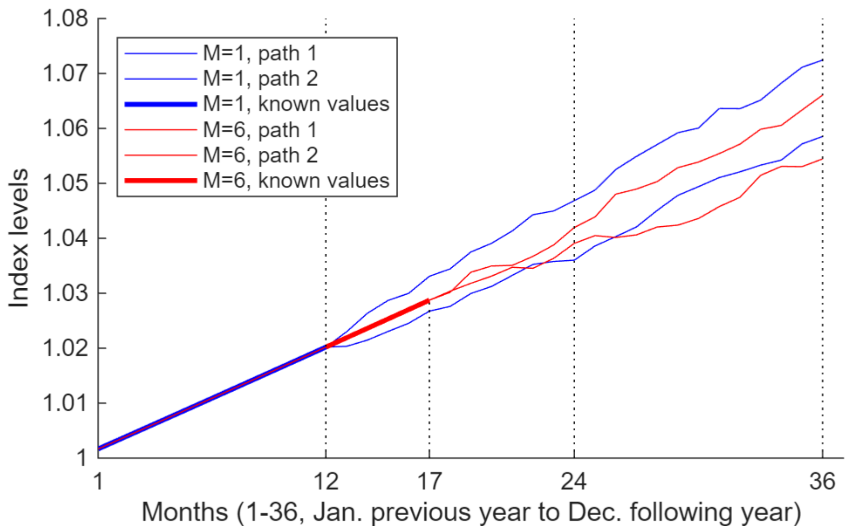

Figure A.1: Examples of two simulated paths as at January and June.

Source: Authors' estimations

Starting with $M = 1$, we undertake 100,000 simulations of the process described above, calculate the AAPCs for the current and following year, and for each year calculate the means and standard deviations of those 100,000 AAPC results.[20] We repeat this exercise for each month M from 1 to 12, and those results establish the standard deviation and mean-change YAPs from the base model.

Figure A.2 plots the YAP results for $\sigma_0(t)$ and $\sigma_1(t)$, which are both monotonic downward. The results for $|\Delta\mu_0(t)|$ and $|\Delta\mu_1(t)|$ in the base model simulation will be approximately zero, because the stochastic components around the steady trend growth rate are essentially averaged out. To simulate a change in the mean expectations of growth, we add a further trend growth component γ to the stochastic part of the simulations. That is, we use $r_t = \delta + \gamma + r_{t-1} + \varepsilon_t$ from month M onwards, with

[19] This relatively simple process, i.e. homoskedastic Gaussian, would likely admit closed-form solutions for $\sigma_0(t)$, $\sigma_1(t)$, $|\Delta\mu_0(t)|$, and $|\Delta\mu_1(t)|$, but the calculations and resulting expressions would nevertheless be onerous. For our purposes, simulations with representative parameters is sufficient (and it also allows for more complex processes that may not have closed-form solutions).

[20] Of course, 100,000 is many magnitudes greater than the typical number of respondents. The large number is required to remove sampling variability, leaving just the actual functional form that results from the simulation.

$\gamma = 0.01/12$, so approximately 1% additional growth per year relative to the base model.[21]

Figure A.3 plots the results for $|\Delta\mu_0(t)|$ and $|\Delta\mu_1(t)|$, which are monotonic downward for the former, and monotonic upward for the latter. The results for $\sigma_0(t)$ and $\sigma_1(t)$ will be identical to those in figure A.2, because the stochastic component around the trend, i.e. $r_t = r_{t-1} + \varepsilon_t$, remains the same.

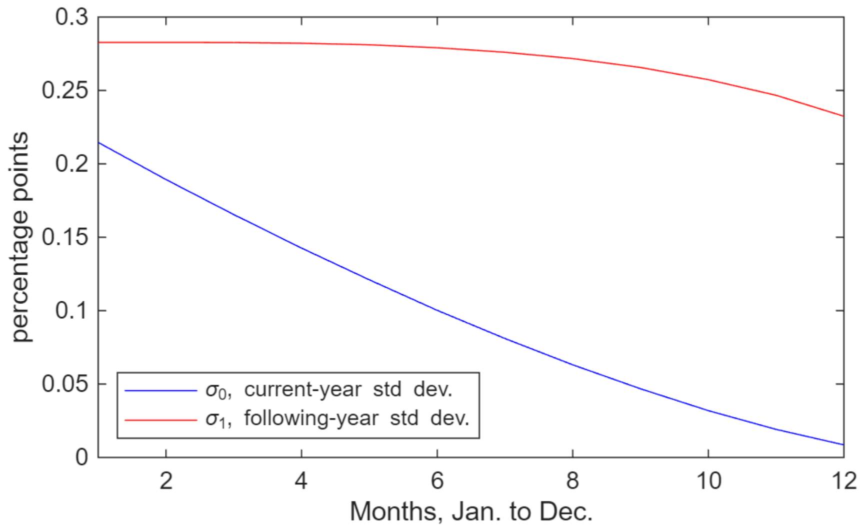

Figure A.2: YAPs for current- and following-year standard deviations obtained from the simulation.

Source: Authors' estimations

The nature of the results from our models above remain the same with alternative parameters, i.e. varying δ, σ, and γ. That is, $\sigma_0(t)$ and $\sigma_1(t)$ are always monotonic downward over the year, and $|\Delta\mu_0(t)|$ and $|\Delta\mu_1(t)|$ are always monotonic downward and monotonic upward, respectively. Importantly, the nature of the $|\Delta\mu_0(t)|$ and $|\Delta\mu_1(t)|$ results remain the same whether a positive or negative value of γ is used to change the trend growth rate.

[21] This simple step-change example does not really require simulations; it could be solved to give pro-rated changes across the year based on δ and γ.

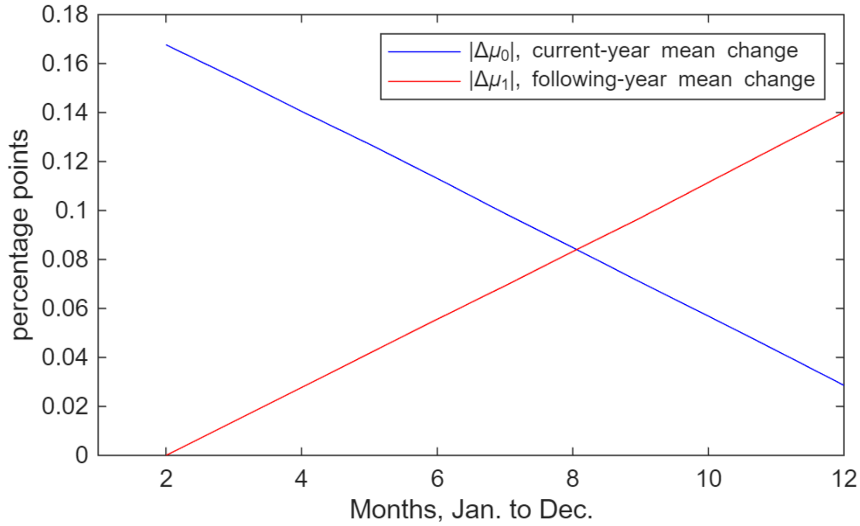

Figure A.3: YAPs for current- and following-year mean change obtained from the simulation.

Source: Authors' estimations

The set-up for our models above can readily be changed to allow for data publication lags, by simply beginning the stochastic components of each run one or several months earlier than the base model. Hence, we have tested our simulation with various publication lags, and we again obtain the nature of results we have already discussed above.

A.5 Alternative weightings for YAP regressions

The functional forms that arise from the simulations in the previous section offer the potential for producing YAPs directly, rather than estimating them as in sections 3 and 4, or at least providing an alternative monotonic functional forms that may be better than the LYAP we have used. Here we briefly provide some background to our considerations, but also why we decided not to progress it at this stage. We may revisit it in future work.

The MYAP regression from section 3 is:

$$Y = \begin{bmatrix} I_N \\ \vdots \\ I_N \end{bmatrix} \begin{bmatrix} \beta_1 \\ \vdots \\ \beta_{12} \end{bmatrix} + \begin{bmatrix} \varepsilon_{M,1} \\ \vdots \\ \varepsilon_{M,T} \end{bmatrix} \tag{10}$$

where $\beta = [\beta_1, \ldots, \beta_{12}]$ is unconstrained. It can be constrained to a LYAP by expressing the constraint within the standard notation $R\beta = q$, and specifying R and q so β is a linear function. Hence, using $R = I_{12}$, and $q = [\delta + 12\gamma, \ldots, \delta + \gamma]'$ gives:

$$\begin{aligned} I_{12} \begin{bmatrix} \beta_1 \\ \vdots \\ \beta_{12} \end{bmatrix} &= \begin{bmatrix} \delta + 12\gamma \\ \vdots \\ \delta + \gamma \end{bmatrix} \\ \begin{bmatrix} \beta_1 \\ \vdots \\ \beta_{12} \end{bmatrix} &= \begin{bmatrix} 1 & 12 \\ \vdots & \vdots \\ 1 & 1 \end{bmatrix} \begin{bmatrix} \delta \\ \gamma \end{bmatrix} \end{aligned} \tag{11}$$

and substituting the result into equation 10 gives:

$$Y = \left(\begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} \otimes \begin{bmatrix} 1 & 12 \\ \vdots & \vdots \\ 1 & 1 \end{bmatrix} \right) \begin{bmatrix} \delta \\ \gamma \end{bmatrix} + \begin{bmatrix} \varepsilon_{L,1} \\ \vdots \\ \varepsilon_{L,T} \end{bmatrix} \tag{12}$$

where the Kronecker product is notation to represent the 12×2 matrix being repeated to match the length of Y. This form of constraint is readily generalizable to arbitrary monthly weights, i.e.:

$$Y = \left(\begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} \otimes \begin{bmatrix} 1 & W_1 \\ \vdots & \vdots \\ 1 & W_{12} \end{bmatrix} \right) \begin{bmatrix} \delta \\ \gamma \end{bmatrix} + \begin{bmatrix} \varepsilon_{W,1} \\ \vdots \\ \varepsilon_{W,T} \end{bmatrix} \tag{13}$$

Hence, results from the simulations in section A.4 could be used with $W_n = \log[S(n)]$ to obtain a 12×1 functional form for estimating the YAP. But the main issues in this set-up are that the simulated weights are based on selected parameters and, because the regression is based on log values, e.g. $Y_t = \log(\text{GDP0sd})$, then the ultimate result is that δ scales and γ appears as a power of the weights, i.e. $\exp(\delta + \gamma W_n) = \exp(\delta) \cdot [\exp(W_n)]^{\gamma}$. All of these aspects amount to a fairly arbitrary regression, despite being based on a justified theoretical process.

A non-arbitrary alternative would be to directly estimate the parameters for the simulation process, i.e. σ and γ for the standard deviations, and potentially with a parameter L for the publication lag. Hence, the functional form would be $\log[S(n)] = \log[S(n, \gamma, \sigma, L)]$. However, $S(n, \gamma, \sigma, L)$ is a non-linear function, so estimating the parameters n, γ, σ, and L would require a non-linear regression.

Based on our considerations above, we decided at this stage to follow the pragmatic approach of using MYAP and LYAP. They are still an arbitrary choice, but have the advantage of being simple and robust relative to the alternatives.

A.6 Comparison to other approaches

As discussed in section 3.3, CE data has been used to calculate uncertainty measures from the standard deviations of the results $\{F1_{k,m}\}_{k=1}^{K}$, where each $F1_{k,m}$ is the pro-rated weighted-average of current- and following-year forecasts from individual survey respondents, i.e. repeated here for convenience:

$$F1_{k,m} = \frac{13 - m}{12} FC_{k,m} + \frac{m - 1}{12} FN_{k,m} \tag{14}$$

We denote this method of calculating uncertainty as $\sigma(F1) = \text{stdev}\Big(\{F1_{k,m}\}_{k=1}^{K}\Big) = \sqrt{\Omega(F1)}$, where the variance is $\Omega(F1) = \text{var}\Big(\{F1_{k,m}\}_{k=1}^{K}\Big)$. Also, for notational convenience below, we further express the weights in equation 14 as $w_0 = (13-m)/12$ and $w_1 = (m-1)/12$.

$FC_{k,m}$ is $S_{0,k}(t)$ in our notation from section A.1, where m is the month for our time t, and $S_{0,k}(t)$ contains a known component and a forecast component. Hence, we can express $FC_{k,m}$ as $FC_{k,m} = FC_{k,m}^{a} + FC_{k,m}^{b}$, where $FC_{k,m}^{a} = \sum_{n=1}^{N} Z_0(t,n) / \sum_{n=1}^{12} Z_{-1}(t,n)$, the known component of $FC_{k,m}$, and $FC_{k,m}^{b} = \sum_{n=1}^{12} \mathbb{E}_t[Z_{0,k}(t,n)] / \sum_{n=1}^{12} Z_{-1}(t,n)$, the forecast component of $FC_{k,m}$. The following-year value $FN_{k,m}$ is $S_{1,k}(t)$ in our notation from section A.1, and it is purely a forecast.

Using the expressions above, $F1_{k,m} = w_0\left(FC_{k,m}^{a} + FC_{k,m}^{b}\right) + w_1 FN_{k,m}$. The variance of $\{F1_{k,m}\}_{k=1}^{K}$ may then be expressed as:

$$
\begin{aligned}
\Omega(F1) &= \text{var}\left[w_0\left(FC_{k,m}^{a} + FC_{k,m}^{b}\right)\right] + \text{var}\left[w_1 FN_{k,m}\right] \\
&\quad +2\text{cov}\left[w_0\left(FC_{k,m}^{a} + FC_{k,m}^{b}\right), w_1 FN_{k,m}\right] \\
&= w_0^2 \text{var}\left[FC_{k,m}^{b}\right] + w_1^2 \text{var}\left[FN_{k,m}\right] + 2w_0 w_1 \text{cov}\left[FC_{k,m}^{b}, FN_{k,m}\right] \\
&= w_0^2 \cdot w_0^2 \Omega_0 + w_1^2 \Omega_1 + 2w_0 w_1 \cdot w_0 \Omega_{01} \qquad (15)
\end{aligned}
$$

where we have omitted the set notation $\{\cdot\}_{k=1}^{K}$ for clarity, the variance of the known component $FC_{k,m}^{a}$ has been set to a zero contribution, and the results $\text{var}\big[FC_{k,m}^{b}\big] = w_0^2 \Omega_0$ and $\text{cov}\big[FC_{k,m}^{b}, FN_{k,m}\big] = w_0 \Omega_{01}$ reflect that $FC_{k,m}^{b}$ represents the remaining fraction of the year w_0. The values Ω_0 and Ω_1 are simplified notation for the typical values of the full-year variances for the current and following year, and Ω_{01} is the typical covariance between the full current year and the following year. Of course, Ω_0, Ω_1, and Ω_{01} are all cross-sectional, i.e. at a particular point in time across the set of respondent forecasts at that time, and so they would vary over time.

The key point from the above exposition is that the contribution of $FC_{k,m}$ to $\Omega(F1)$ declines steadily over the year, i.e. with weight w_0^2 on Ω_0 and weight w_0 on Ω_{01}, even before w_0 is applied again for the weighted-average calculation. $w_0^2 \Omega_0$ and $w_0 \Omega_{01}$ represent the steadily shrinking genuine forecast component in the remaining months of the year, i.e. the YAP we have detailed in prior sections of this appendix.

The calculations of $\Omega(F1)$ and $\sigma(F1)$ will therefore have a YAP. However, the materiality of the YAP in $\sigma(F1)$ relative to our YAP-adjusted series $\tilde{\sigma}_{01}$ remains to be tested using the unit record data. The key difference between the approaches is that $\sigma(F1)$ includes a covariance term, while ours does not because it is based on the variances $\tilde{\sigma}_0^2$ and $\tilde{\sigma}_1^2$ derived from the CE standard deviations reported for the current and following year.

B Sign restrictions

For our sign restrictions used in our demand and supply shocks application, we follow Fry and Pagan (2011, hereafter FP). First we set up our base model for A_0 and $\varepsilon_{0,t}$

using an abbreviated version of equation 3 from the main text, i.e.:

$$X_t = A_0 \varepsilon_{0,t} \ ; \ \varepsilon_{0,t} \sim N(0, I_2) \tag{16}$$

where $X_t = [\mathrm{GDP}(t), \mathrm{CPI}(t)]$, A_0 is a 2×2 response matrix, and $\varepsilon_{0,t} = [\varepsilon_{1,t}, \varepsilon_{2,t}]'$ are uncorrelated innovations with unit variances, i.e. the covariance matrix $\mathbb{E}\left[\varepsilon_{0,t}\varepsilon_{0,t}'\right] = \Omega_{\varepsilon_0}$ is the 2×2 identity matrix I_2. From this expression, the covariance matrix for X_t is therefore $\Omega_X = \mathbb{E}[X_t X_t'] = A_0 \Omega_{\varepsilon_0} A_0' = A_0 A_0'$.

The empirical estimate of Ω_X from our GDP and CPI shock series is $\Omega_X = \frac{1}{T} X X'$, where X is the $2 \times T$ matrix of the time series $\{X_t\}_{t=1}^{T}$. We use an eigensystem decomposition of Ω_X to obtain a base model for A_0 and $\varepsilon_{0,t}$ that has the required properties in equation 16.[22] Hence, $\Omega_X = V D V^{-1} = V D^{\frac{1}{2}} D^{\frac{1}{2}} V^{-1}$, where V is the eigenvector matrix, D is the diagonal eigenvalue matrix, and $V' = V^{-1}$ because Ω_X is a symmetric matrix.

We can now set $A_0 = V D^{\frac{1}{2}}$ and $\varepsilon_0 = D^{-\frac{1}{2}} V^{-1} X = D^{-\frac{1}{2}} V' X$, which have the properties required for the base model in equation 16. That is, $A_0 \varepsilon_{0,t} = V D^{\frac{1}{2}} D^{-\frac{1}{2}} V^{-1} X = V V^{-1} X = X$, $A_0 A_0' = V D^{\frac{1}{2}} D^{\frac{1}{2}} V' = V D V^{-1} = \Omega_X$, and $\frac{1}{T}\varepsilon_0 \varepsilon_0' = \frac{1}{T} D^{-\frac{1}{2}} V' X \left(D^{-\frac{1}{2}} V' X\right)'$ $= D^{-\frac{1}{2}} V' \frac{1}{T} X X' V D^{-\frac{1}{2}} = D^{-\frac{1}{2}} V' V D V' V D^{-\frac{1}{2}} = D^{-\frac{1}{2}} D D^{-\frac{1}{2}} = I_2$.

The sign restriction method uses unitary matrices Q_j, hence $Q_j Q_j' = Q_j' Q_j = I$, to transform the base model into $X = A_j \varepsilon_j$, where $A_j = A_0 Q_j$ and $\varepsilon_j = Q_j' \varepsilon_0$, with the objective of seeking results where the signs of the matrix elements A_j align with the desired signs contained in the sign matrix. The transformed model $X = A_j \varepsilon_j$ always maintains the required properties in equation 16, because $A_j \varepsilon_j = A_0 Q_j Q_j' \varepsilon_0 = A_0 \varepsilon_0 = X$, $A_j A_j' = A_0 Q_j Q_j' A_0' = A_0 A_0'$, and $\frac{1}{T}\varepsilon_j \varepsilon_j' = Q_0' \varepsilon_0 \varepsilon_0' Q_0 = Q_0' I Q_0 = Q_0' Q_0 = I$.

In the context of our two-variable case, Q_j is the 2×2 rotation matrix, i.e.:

$$Q_j = \begin{bmatrix} \cos(\theta_j) & -\sin(\theta_j) \\ \sin(\theta_j) & \cos(\theta_j) \end{bmatrix} \tag{17}$$

where θ is the rotation angle in radians.[23]

In the two-variable case, seeking results where the signs of the elements from $A_j = A_0 Q_j'$ align with the sign matrix may be represented explicitly as:

$$A_j = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \cos(\theta_j) & -\sin(\theta_j) \\ \sin(\theta_j) & \cos(\theta_j) \end{bmatrix} \approx \begin{bmatrix} + & + \\ - & + \end{bmatrix} \tag{18}$$

where the 2×2 matrix with elements a, b, c, d is $A_0 = V D^{-\frac{1}{2}}$, and "$\approx$" is our notation to represent "aligns with", i.e. each element of $A_0 Q_j$ has the signs of the corresponding element in the sign matrix. The search would typically proceed with a systematic grid for θ, and there would be a set of angles producing $A_0 Q_j$ results that align with the sign matrix.

For our exercise, we choose to obtain a unique model by imposing a normalization condition that positive supply and demand shocks have the same effect on GDP, i.e.

[22] FP use a Cholesky decomposition to set up a base model, but the key aspect is to obtain $\Omega_{\varepsilon_0} = I_2$.

[23] It is straightforward to verify that Q_j is a unitary matrix, i.e. $Q_j Q_j = Q_j' Q_j = I_2$. That is, the $(1,1)$ and $(2,2)$ elements of $Q_j Q_j'$ are $[\cos(\theta_j)]^2 + [\sin(\theta_j)]^2 = 1$, and the $(1,2)$ and $(2,1)$ elements are $\cos(\theta_j)\sin(\theta_j) - \cos(\theta_j)\sin(\theta_j) = 0$.

two elements in the first row of A_0Q_j are equal, i.e. $[A_0Q_j]_{11} = [A_0Q_j]_{12}$. With that restriction, the simple form with the two variables allows us to solve for θ_j analytically, i.e.:

$$\begin{aligned}
a\cos(\theta_j) + b\sin(\theta_j) &= -a\sin(\theta_j) + b\cos(\theta_j) \\
(a+b)\sin(\theta_j) &= (b-a)\cos(\theta_j) \\
\frac{\sin(\theta_j)}{\cos(\theta_j)} &= \tan(\theta_j) = \frac{b-a}{a+b} \\
\theta_j &= \arctan\left(\frac{b-a}{a+b}\right) \qquad (19)
\end{aligned}$$

If θ_j produces a negative result for $A_0Q_j[1,1]$ and $A_0Q_j[1,2]$, then the angle $\theta_j + \pi$ will produce the positive result.

Our calculation of A_0 for the United States is:

$$A_0(\mathrm{US}) = \begin{bmatrix} 0.0606 & -0.3784 \\ -0.2875 & -0.0798 \end{bmatrix} \qquad (20)$$

which results in an angle $\theta_j + \pi = 0.6994\pi = 2.1973$, so our calculation of A_j from equation 19 is:

$$A_j(\mathrm{US}) = \begin{bmatrix} 0.2718 & 0.2718 \\ -0.1844 & 0.2342 \end{bmatrix} \qquad (21)$$

Our dataset at this stage does not include a monetary policy variable, for reasons we mention in the conclusion, and so our fundamental shock analysis does not allow monetary policy shocks to be identified. Therefore, shocks to GDP and CPI that may have arisen from monetary policy shocks are attributed to demand shocks. To show this, we use the set of sign restrictions from table 3 of FP for a VAR with GDP, CPI, and a monetary policy variable, i.e.:

$$\begin{bmatrix} \Delta\mathbb{E}_t[Y_{t+\tau}] \\ \Delta\mathbb{E}_t[\pi_{t+\tau}] \\ \Delta\mathbb{E}_t[M_{t+\tau}] \end{bmatrix} = \begin{bmatrix} + & + & - \\ - & + & - \\ - & + & + \end{bmatrix} \begin{bmatrix} \varepsilon_{S,t} \\ \varepsilon_{D,t} \\ \varepsilon_{M,t} \end{bmatrix} \qquad (22)$$

A positive (tightening) monetary policy shock $+\varepsilon_{M,t}$ would therefore produce the response $(-\Delta\mathbb{E}_t[Y_{t+\tau}], -\Delta\mathbb{E}_t[\pi_{t+\tau}], +\Delta\mathbb{E}_t[M_{t+\tau}])'$. But without a monetary policy variable, the response would show as just $(-\Delta\mathbb{E}_t[Y_{t+\tau}], -\Delta\mathbb{E}_t[\pi_{t+\tau}])'$, which is the response to a negative demand shock. Similarly, a positive (easing) fiscal policy shock, hence positive responses to output and inflation, would be attributed as a positive demand shock. Our inferred series of demand shocks therefore include genuine demand shocks along with policy and other shocks that stimulate demand.

PUBLICATIONS

Measures of Macroeconomic Shocks and Uncertainty for Asia-Pacific Economies
Working Paper No. WP/2026/174