Appendix of Supplemental Results

for the paper

"A Joint Impulse Response Function for Vector Autoregressive Models"

Thomas F. P. Wiesen
University of Maine School of Economics
5782 Winslow Hall
Orono, ME 04469, USA
thomas.wiesen@maine.edu

Paul M. Beaumont
Florida State University Department of Economics
Room 276 BEL, 113 Collegiate Loop
Tallahassee, FL 32306, USA
beaumont@fsu.edu

A.1 Simulated Impulse Responses Comparing the jIRF to the Sum of the oIRFs

In section 4 of the paper, we illustrate the properties of the jIRF and compare the jIRF to the sum of the gIRFs. For empirical researchers, it may be additionally interesting to compare the jIRF to the sum of the orthogonalized impulse response functions (oIRF). This section of the appendix makes those comparisons. The oIRF, which utilizes Cholesky decomposition, depends on the order of the variables of the VAR. Therefore, in this simulation, we compute all $K!$ permutations of the oIRF sums and compare those to the jIRF.

As in section 4, we simulate a 4-variable VAR(1) model using a symmetric Toeplitz matrix as the VAR's coefficient matrix whose diagonal elements are all equal to 0.55 and whose off-diagonal elements are all equal to 0.1 (see equation (21) in the paper). Our interest lies in measuring how variable 1 responds due to simultaneous shocks from variables 2 and 3. We compute all $K! = 4! = 24$ permutations of the oIRF. After the oIRFs are obtained, we permute the order of the variables back to their original ordering so that elements of the oIRFs can be directly compared. For each oIRF permutation, we then add the orthogonalized impulse response of variable 1 due to a shock from variable 2 *plus* the orthogonalized impulse response of variable 1 due to a shock from variable 3. Since there are $K! = 24$ of these oIRF sums, we show only the high and low range of the computed 24 oIRF sums illustrated by the gray bars in Figure A.1 below. Overlayed on these plots are the jIRFs, shown as a black solid line with circle markers, measuring how variable 1 responds due to joint, simultaneous one standard deviation shocks from variables 2 and 3. The jIRFs shown in the plots below are identical to the jIRFs shown in Figure 1 of the paper. Each of the eight graphs in Figure A.1 corresponds to a different shock correlation (covariance) matrix and corresponds to the same pattern as in Figure 1 of the paper. Specifically, the shock correlation between variables 2 and 3 (namely, ρ_{23}) is 0.8, 0.3, 0, and -0.25 in the first, second, third, and fourth rows, respectively. In the first column of plots, the shock correlation between variables 1 and 2 is $\rho_{12} = 0.5$, and the shock correlation between variables 1 and 3 is $\rho_{13} = -0.1$, so that the correlations ρ_{12} and ρ_{13} have different signs. In the second column of plots, the shock correlation between variables 1 and 2 is still $\rho_{12} = 0.5$, but the shock correlation between variables 1 and 3 is $\rho_{13} = 0.2$, so that the correlations ρ_{12} and ρ_{13} have the same sign. All the other shock correlations are constant across the plots. The specific shock correlation (covariance) matrix used in each plot is shown below each graph.

As evidenced by the fourth row of plots, the jIRF need not be bounded by the range of the oIRF sums. This is not surprising given that the jIRF and oIRF are quite different in construction and motivation. It is worth stressing that the gIRF, not the oIRF, is the tool most comparable to our proposed jIRF because the jIRF is a multi-shock "generalization" of the gIRF. Both the gIRF and jIRF are agnostic to structural identification, while remaining consistent with the data; namely, the gIRF and jIRF are unique, but the structural parameters need not be explicitly identified. On the other hand, the oIRF relies on a specific, although possibly incorrect, structural identification and depends on the recursive ordering of the variables. This recursive ordering imposes a structural restriction to the system that, when done properly, must be justified by economic theory. In practice, calculating all $K!$ permutations of the oIRFs is not recommended as many of these orderings will contain "incorrect" recursive structures. Here, we compute all $K!$ permutations of the oIRF sums simply for the purpose of comparison and to juxtapose the jIRF to existing econometric tools.

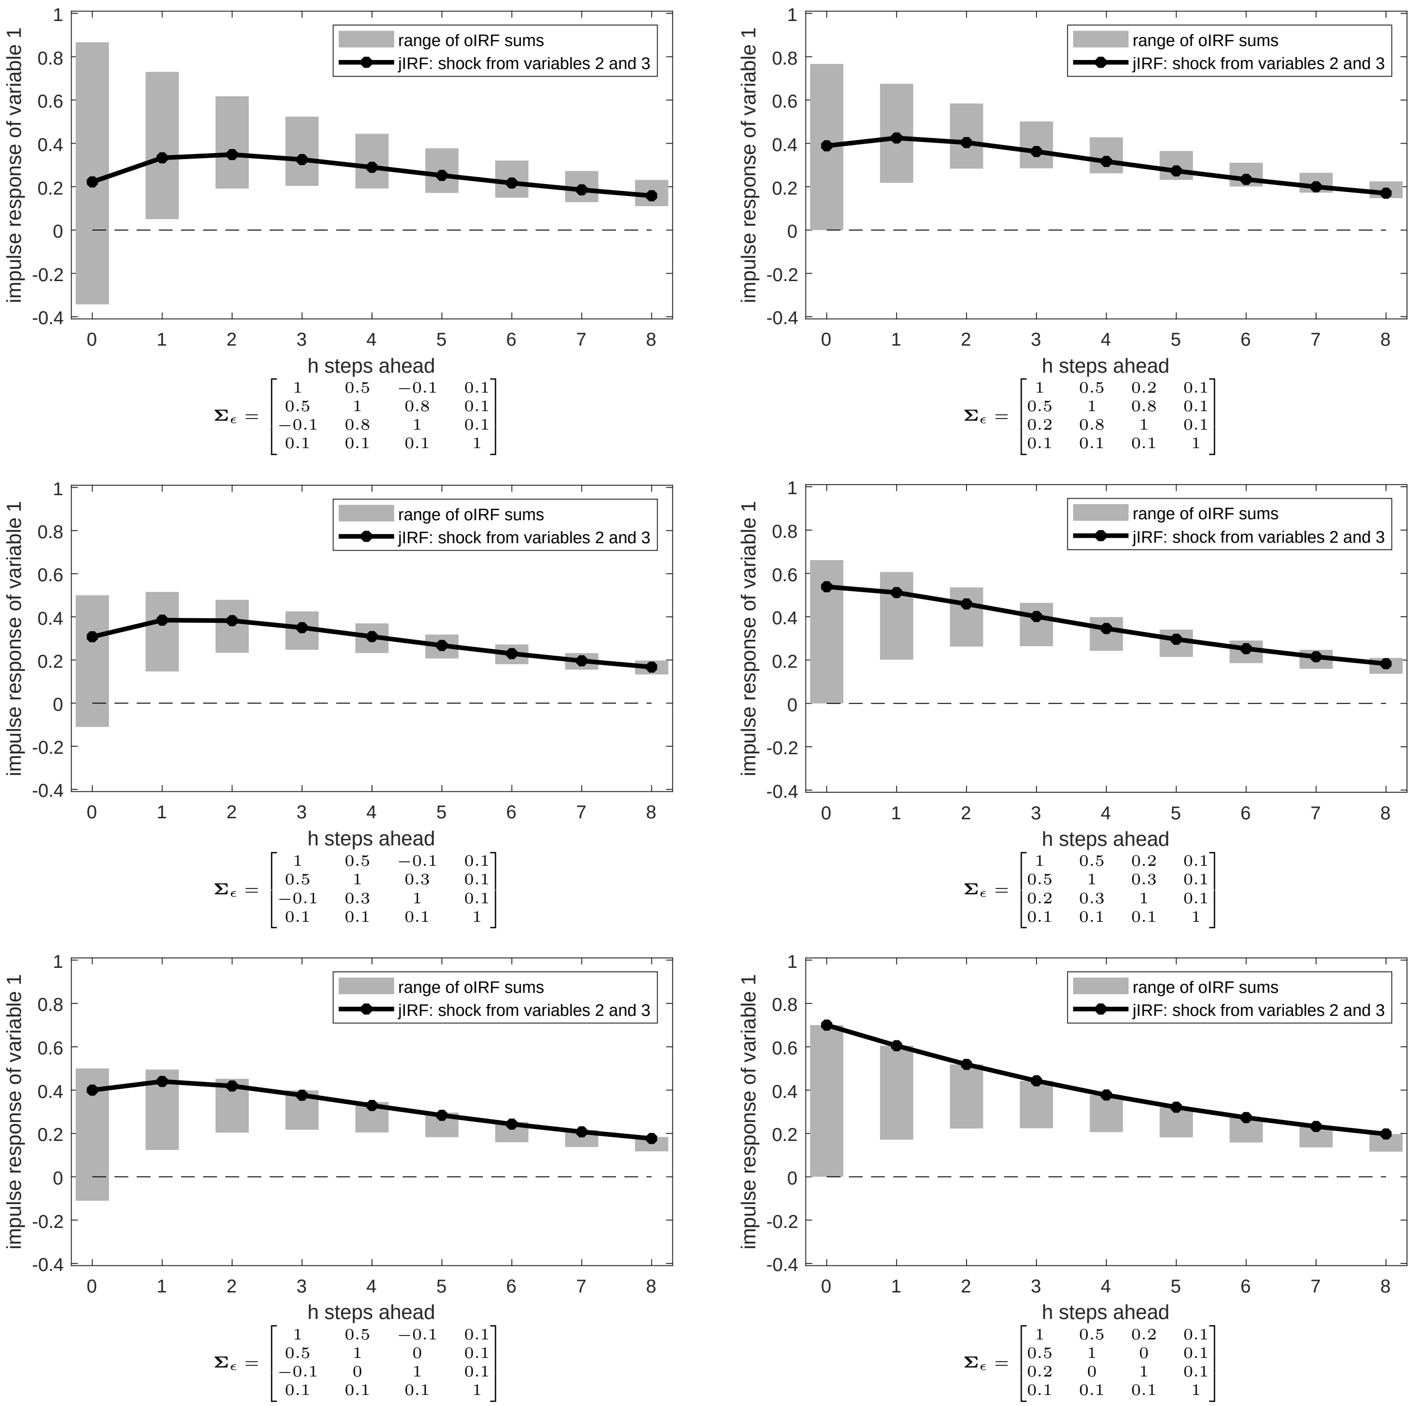

$$\Sigma_\epsilon = \begin{bmatrix} 1 & 0.5 & -0.1 & 0.1 \\ 0.5 & 1 & 0.8 & 0.1 \\ -0.1 & 0.8 & 1 & 0.1 \\ 0.1 & 0.1 & 0.1 & 1 \end{bmatrix} \qquad \Sigma_\epsilon = \begin{bmatrix} 1 & 0.5 & 0.2 & 0.1 \\ 0.5 & 1 & 0.8 & 0.1 \\ 0.2 & 0.8 & 1 & 0.1 \\ 0.1 & 0.1 & 0.1 & 1 \end{bmatrix}$$

$$\Sigma_\epsilon = \begin{bmatrix} 1 & 0.5 & -0.1 & 0.1 \\ 0.5 & 1 & 0.3 & 0.1 \\ -0.1 & 0.3 & 1 & 0.1 \\ 0.1 & 0.1 & 0.1 & 1 \end{bmatrix} \qquad \Sigma_\epsilon = \begin{bmatrix} 1 & 0.5 & 0.2 & 0.1 \\ 0.5 & 1 & 0.3 & 0.1 \\ 0.2 & 0.3 & 1 & 0.1 \\ 0.1 & 0.1 & 0.1 & 1 \end{bmatrix}$$

$$\Sigma_\epsilon = \begin{bmatrix} 1 & 0.5 & -0.1 & 0.1 \\ 0.5 & 1 & 0 & 0.1 \\ -0.1 & 0 & 1 & 0.1 \\ 0.1 & 0.1 & 0.1 & 1 \end{bmatrix} \qquad \Sigma_\epsilon = \begin{bmatrix} 1 & 0.5 & 0.2 & 0.1 \\ 0.5 & 1 & 0 & 0.1 \\ 0.2 & 0 & 1 & 0.1 \\ 0.1 & 0.1 & 0.1 & 1 \end{bmatrix}$$

Figure A.1: Comparison of the jIRF with the range of possible oIRF sums.

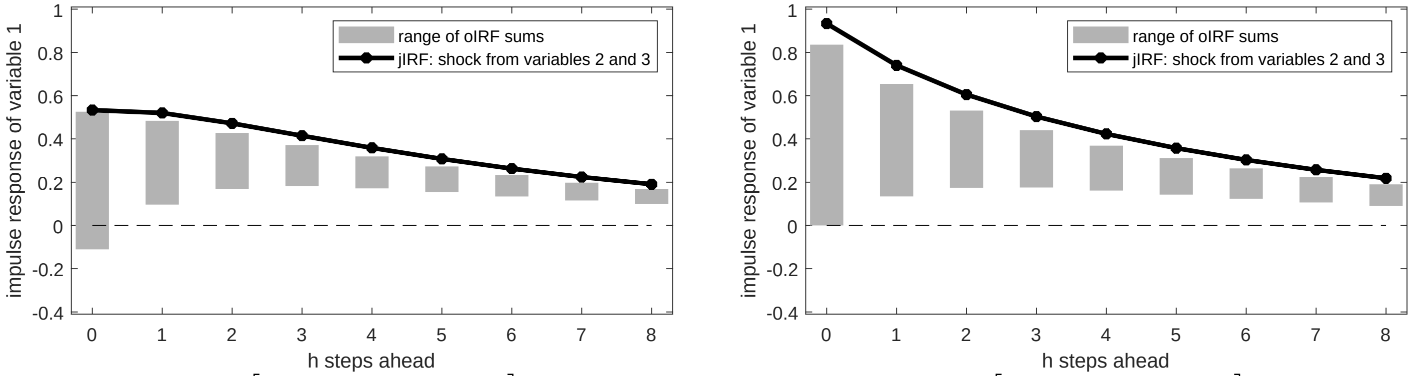

$$\mathbf{\Sigma}_\epsilon = \begin{bmatrix} 1 & 0.5 & -0.1 & 0.1 \\ 0.5 & 1 & -0.25 & 0.1 \\ -0.1 & -0.25 & 1 & 0.1 \\ 0.1 & 0.1 & 0.1 & 1 \end{bmatrix} \qquad \mathbf{\Sigma}_\epsilon = \begin{bmatrix} 1 & 0.5 & 0.2 & 0.1 \\ 0.5 & 1 & -0.25 & 0.1 \\ 0.2 & -0.25 & 1 & 0.1 \\ 0.1 & 0.1 & 0.1 & 1 \end{bmatrix}$$

Figure A.1: Comparison of the jIRF with the range of possible oIRF sums. (Continued)

A.2 Volatility Estimation from High, Low, Open, and Close Prices

This section summarizes how we computed daily stock market volatility using high, low, open, and close (HLOC) prices used in sections 6 and 7 of the paper. To compute the volatility estimates, we follow the notation of Yang and Zhang (2000) and use the TTR R code package by Joshua Ulrich.

Define the following notation for date t prices.

- O_t : opening price
- H_t : high price during the trading interval
- L_t : low price during the trading interval
- C_t : closing price
- $o_t = \ln O_t - \ln C_{t-1}$: the normalized open
- $u_t = \ln H_t - \ln O_t$: the normalized high
- $d_t = \ln L_t - \ln O_t$: the normalized low
- $c_t = \ln C_t - \ln O_t$: the normalized close

An early and commonly used estimator of daily volatility using HLOC data is the method of Parkinson (1980) which assumes an underlying price that follows a geometric Brownian motion process. The time t Parkinson estimator, denoted as $V_{p,t}$, is computed as

$$\begin{aligned} V_{p,t} &= \sqrt{\frac{N}{4n\ln(2)} \sum_{i=0}^{n-1} (u_{t-i} - d_{t-i})^2} \\ &= \sqrt{\frac{N}{4n\ln(2)} \sum_{i=0}^{n-1} \left(\ln\left(\frac{H_{t-i}}{L_{t-i}}\right)\right)^2}, \end{aligned} \tag{A.1}$$

where N is the number of trading days in the year and n is the number of trading days over which the volatility is estimated. In our application we use $n = 1$ and $N = 250$ to get annualized daily volatility, which we scale up by 100 so that the units are annualized percentages.

The Parkinson estimator uses only the daily high and low prices and assumes that: (1) there are no jumps between the previous day's closing price and the current day's opening price, and (2) there is no drift in the underlying geometric Brownian motion price process. If these conditions are violated, the Parkinson estimator will be biased and inefficient.

Rogers and Satchell (1991) propose an estimator that is independent of the drift rate in the underlying process and has a smaller estimator variance than $V_{p,t}$. The time t volatility estimator of Rogers and Satchell (1991), denoted as $V_{rs,t}$, can be computed as

$$\begin{aligned} V_{rs,t} &= \sqrt{\frac{N}{n} \sum_{i=0}^{n-1} \left(u_{t-i} \left(u_{t-i} - c_{t-i} \right) + d_{t-i} \left(d_{t-i} - c_{t-i} \right) \right)} \\ &= \sqrt{\frac{N}{n} \sum_{i=0}^{n-1} \left(\ln\left(\frac{H_{t-i}}{O_{t-i}}\right) \ln\left(\frac{H_{t-i}}{C_{t-i}}\right) + \ln\left(\frac{L_{t-i}}{O_{t-i}}\right) \ln\left(\frac{L_{t-i}}{C_{t-i}}\right) \right)}. \end{aligned} \tag{A.2}$$

Since $V_{rs,t}$ does not allow for opening jumps in prices, it may still underestimate the true daily volatility.

Yang & Zhang (2000) derive a minimum-variance unbiased estimator of volatility, denoted as $V_{yz,t}$ for time t, which is independent of both drift and opening jumps

$$V_{yz,t} = \sqrt{V_{o,t}^2 + kV_{c,t}^2 + (1-k)V_{rs,t}^2}, \tag{A.3}$$

where

$$V_{o,t} = \sqrt{\frac{N}{n-1} \sum_{i=0}^{n-1} (o_{t-i} - \bar{o})^2}, \tag{A.4}$$

$$\bar{o} = \frac{1}{n} \sum_{i=0}^{n-1} o_{t-i}, \tag{A.5}$$

$$V_{c,t} = \sqrt{\frac{N}{n-1} \sum_{i=0}^{n-1} (c_{t-i} - \bar{c})^2}, \tag{A.6}$$

$$\bar{c} = \frac{1}{n} \sum_{i=0}^{n-1} c_{t-i}, \tag{A.7}$$

and V_{rs_t} is the Rogers & Satchell estimator defined above. Here, the value $k = 0.07834101$ is chosen to minimize the variance and is dependent on n.

Note that the use of $o_t = \ln O_t - \ln C_{t-1}$ in the expression for $V_{o,t}$ requires that $n > 1$. In our application, we use $n = 2$ so that the first available volatility is for period $t = 3$. We again use $N = 250$ to get annualized daily volatility, which we scale up by 100 so that the units are annualized percentages.

In the application, we use the stock market volatility of the 11 systemically important financial institutions as listed by the Federal Reserve's Large Institution Supervision Coordinating Committee—Bank of America, Bank of New York Mellon, Barclays, Citigroup, Credit Suisse, Deutsche Bank, Goldman Sachs, JP Morgan Chase, Morgan Stanley, State Street, and Wells Fargo—as of August 2020.

Table A.1 shows the means and standard deviations of the daily volatility estimates for each bank using the Parkinson and Yang-Zhang methods. As expected, the Parkinson method underestimates both the mean and the standard deviation of the volatility estimates relative to the Yang-Zhang method.

As an example, Figure A.2 shows the daily time series of computed volatilities using the Parkinson and Yang-Zhang methods for Bank of America (BAC). Note that, particularly during the pandemic period, the volatility estimates from the Parkinson method considerably underestimate those using the Yang-Zhang method. The plots for the other banks are similar.

Table A.1: Means and standard deviations of the daily annualized percent volatility using the Parkinson and Yang-Zhang estimators.

		Parkinson (V_p)		Yang-Zhang (V_{yz})	
Financial Institution	Ticker	Mean	SD	Mean	SD
Bank of America	BAC	20.99	14.23	27.99	24.00
Bank of New York Mellon	BK	19.41	12.76	25.05	21.07
Barclays	BCS	16.31	12.20	29.79	27.81
Citigroup	C	22.05	16.56	29.30	27.23
Credit Suisse	CS	14.99	10.82	26.06	23.23
Deutsche Bank	DB	18.68	11.82	33.24	25.36
Goldman Sachs	GS	20.54	13.66	25.95	20.30
JP Morgan Chase	JPM	18.38	13.42	24.54	22.67
Morgan Stanley	MS	22.13	15.09	29.05	23.77
State Street	STT	22.86	15.79	29.40	24.59
Wells Fargo	WFC	21.02	15.17	27.14	24.15

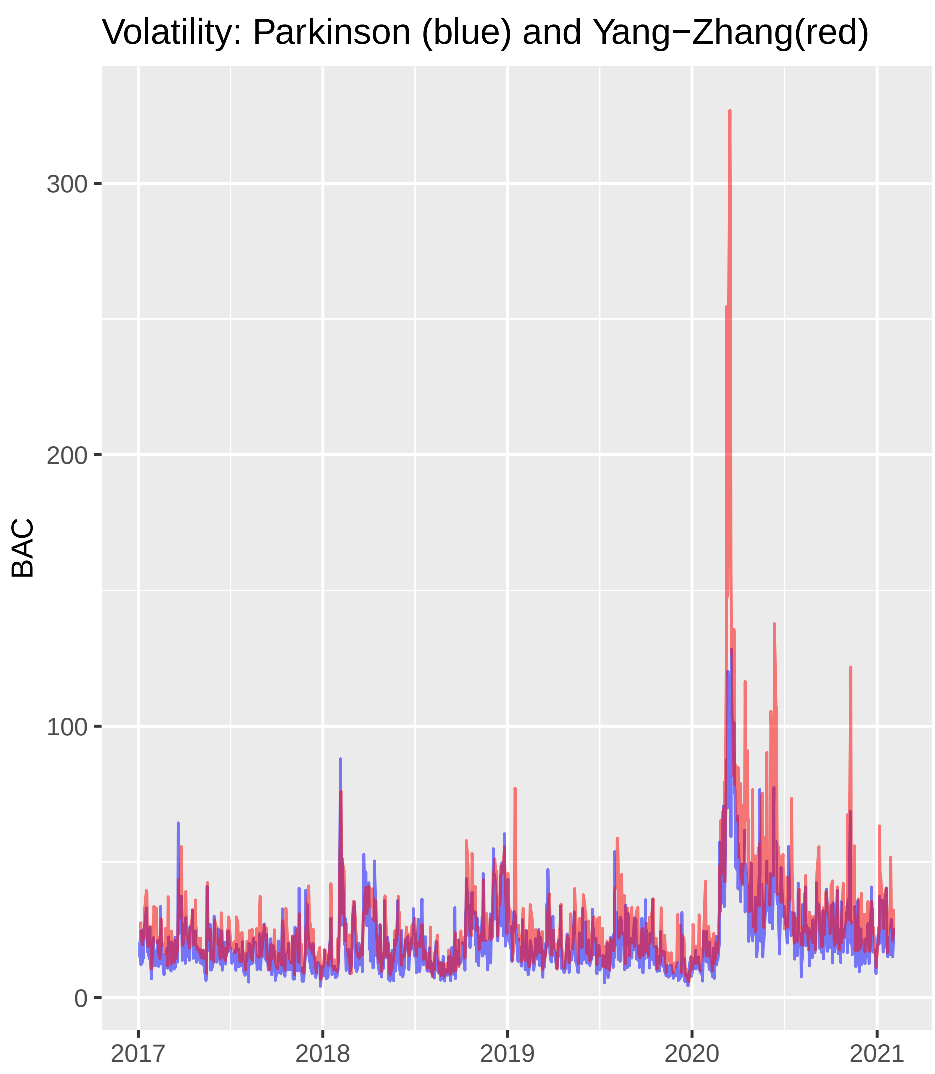

Figure A.2: Daily annualized percent volatility for Bank of America (BAC) using the Parkinson and the Yang-Zhang estimators.

A.3 Trans-Atlantic Volatility Transmissions: Additional Impulse Response Functions

In section 6 of the paper, we estimate a VAR(1) model in the pre-pandemic period (January 1, 2017 to January 31, 2020) and re-estimate the VAR in the pandemic period (February 3, 2020 to February 4, 2021) using daily equity market volatility of eleven systemically important financial institutions. These include eight American financial institutions (Bank of America, Bank of New York Mellon, Citigroup, Goldman Sachs, JP Morgan Chase, Morgan Stanley, State Street, and Wells Fargo) and three European financial institutions (Barclays, Credit Suisse, and Deutsche Bank).

The goal is to measure trans-Atlantic volatility transmissions from European to American financial institutions before and during the COVID-19 pandemic. To that end, we use the joint impulse response function (jIRF) to measure how the American banks respond due to joint, simultaneous shocks from the three European banks. For comparison purposes, the generalized impulse response functions (gIRF) are also computed, which measure how the American banks respond due to individual shocks from the three European banks. The impulse response functions for Bank of America, Citigroup, and JP Morgan Chase are shown in Figure 3 in the paper. This appendix section shows the impulse responses for Bank of New York Mellon, Goldman Sachs, Morgan Stanley, State Street, and Wells Fargo.

In Figure A.3 below, the left column of panels is for the pre-pandemic period, and the right column of panels is for the pandemic period. Row one of the panels shows the responses of Bank of New York Mellon (BK), row two shows the responses of Goldman Sachs (GS), row three shows the responses of Morgan Stanley (MS), row four shows the responses of State Street (STT), and row five shows the responses of Wells Fargo (WFC). In each of the panels of the figure below, there are five lines depicting the responses of the American financial institution due to various shocks from the European financial institutions. In each of the ten panels, the black line with circle markers is the jIRF measuring how the American financial institution's volatility responds due to joint, simultaneous shocks from the three European financial institutions, and the grey shaded region shows the 80% confidence interval for the jIRF. The blue dotted line is the gIRF measuring how the American financial institution's volatility responds due to a shock from Barclays (BCS) alone. The red dashed line is the gIRF measuring how the American financial institution's volatility responds due to a shock from Credit Suisse (CS) alone. The green dot-dash line is the gIRF measuring how the American financial institution's volatility responds due to a shock from Deutsche Bank (DB) alone. Finally, the solid brown line is the sum of the three gIRFs, and the shaded brown region shows the 80% confidence interval for the sum of the gIRFs. Notice the vertical scale in the pre-pandemic period panels is different from the vertical scale in the pandemic period panels.

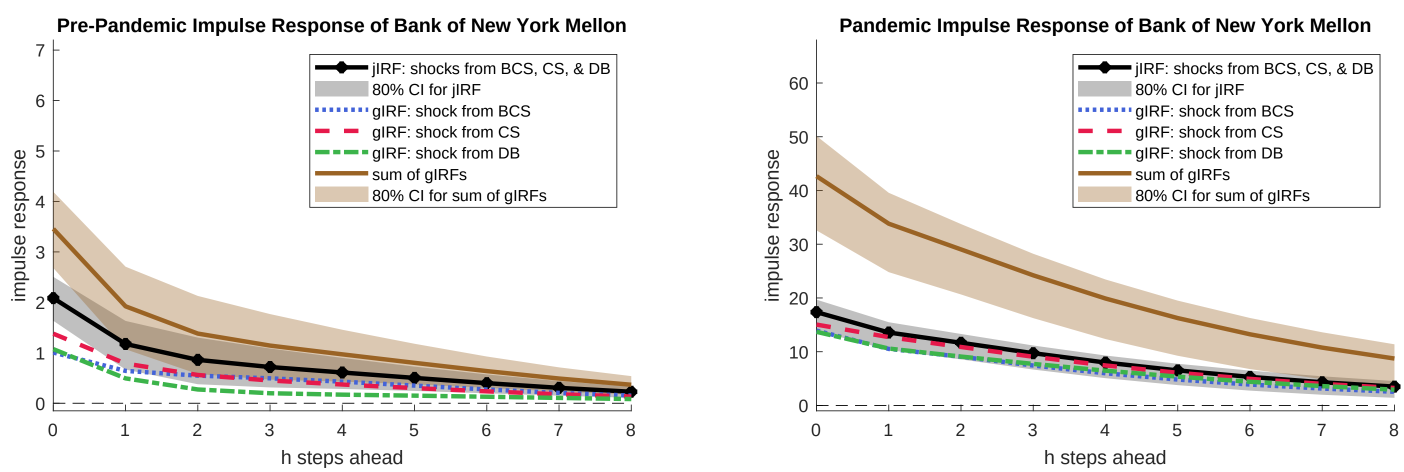

Figure A.3: Comparison of the jIRFs and gIRFs for American banks not shown in the paper.

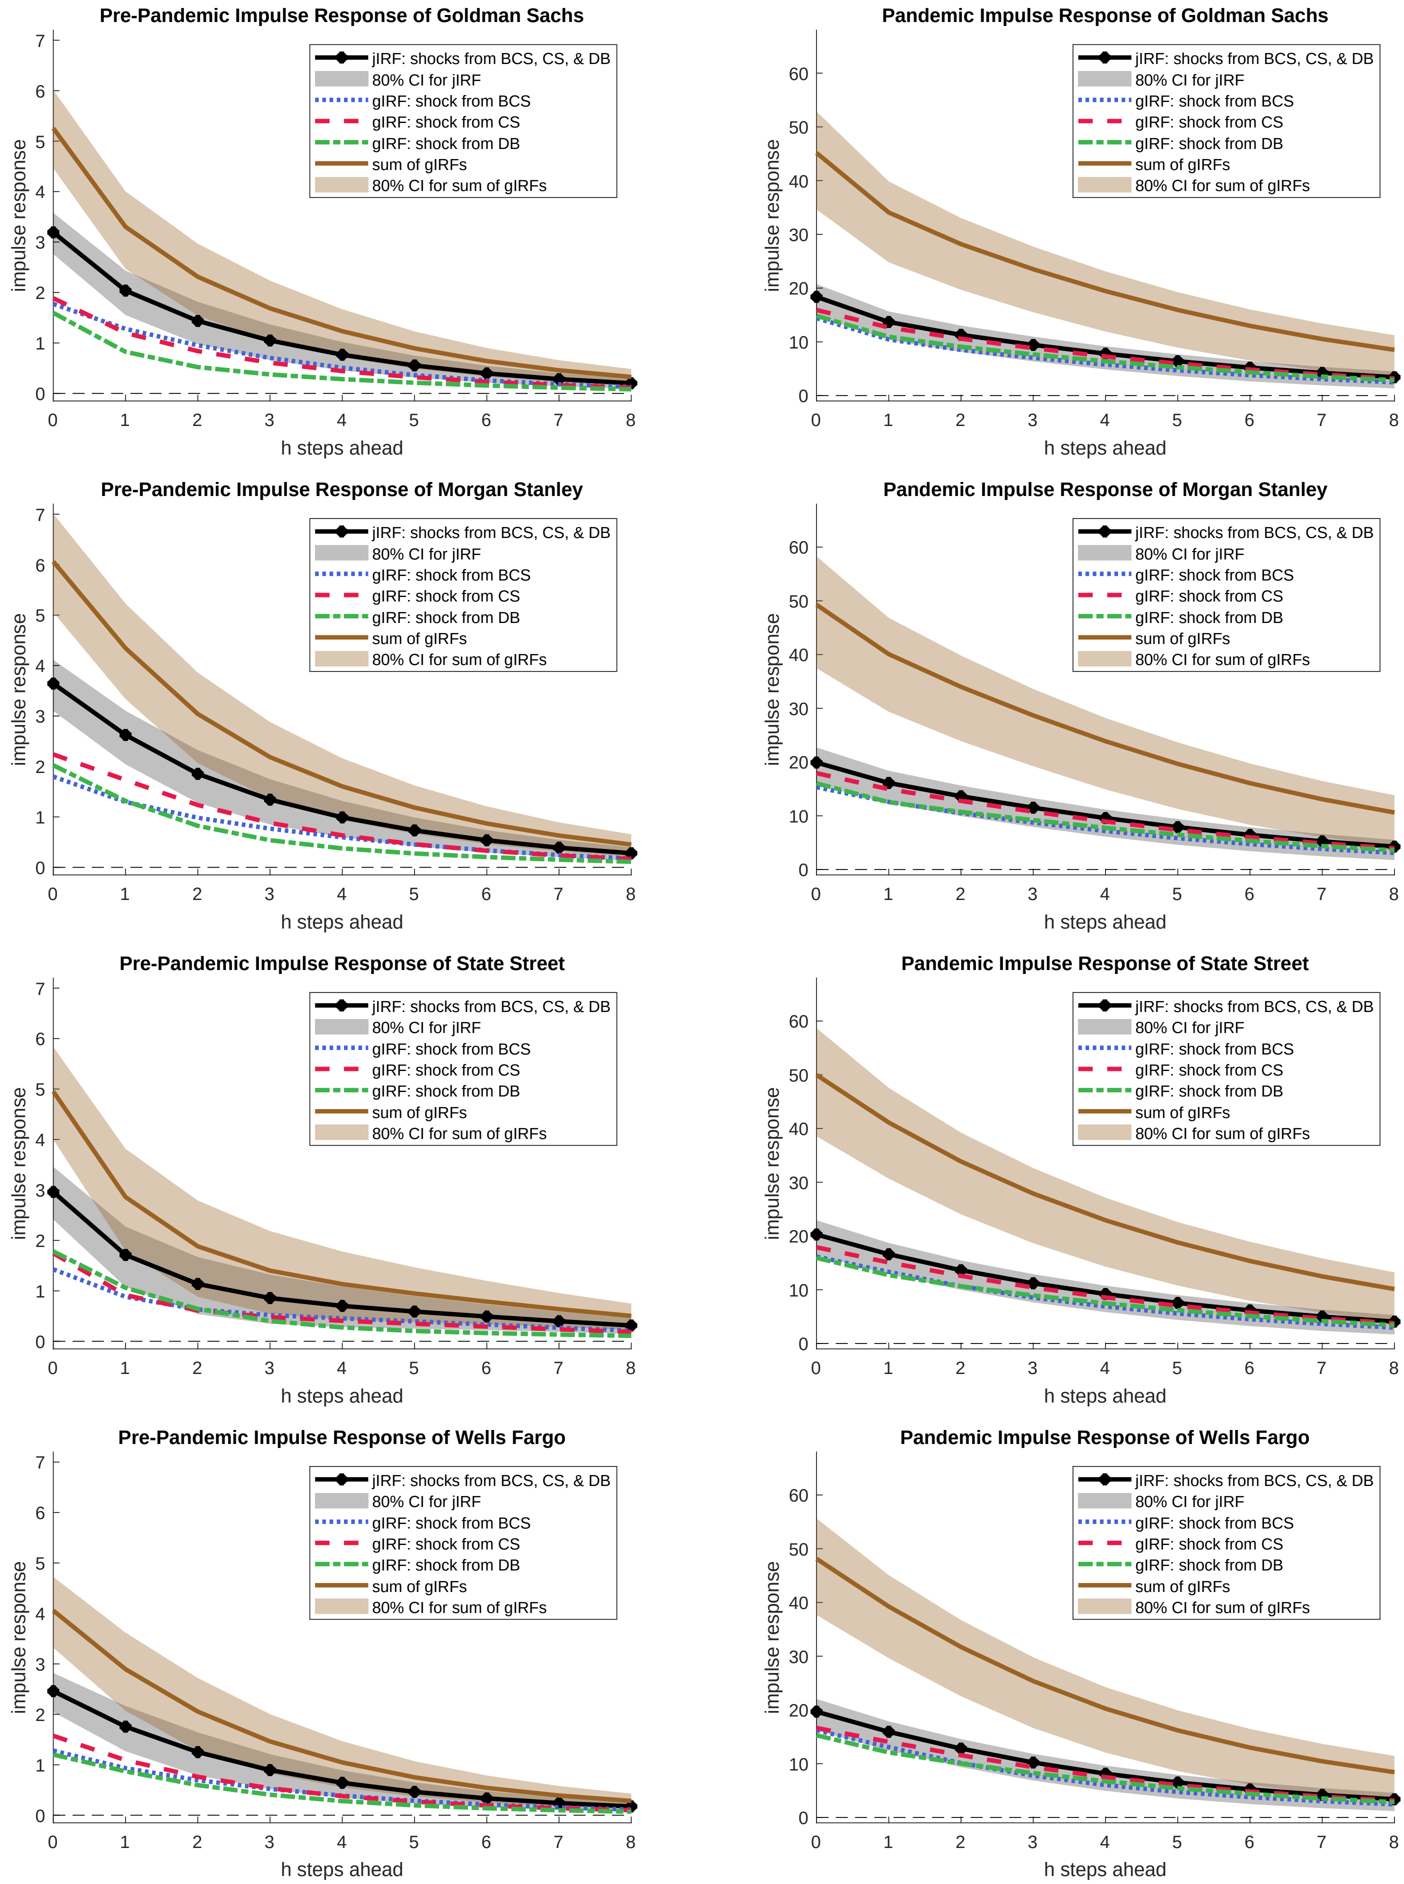

Figure A.3: Comparison of the jIRFs and gIRFs for American banks not shown in the paper. (Continued)

We are also interested in empirically confirming the difference between the sum of the gIRF and the jIRF in our application. Thus, we calculate the confidence intervals for the difference between the sum of the gIRFs (i.e., the gIRF for an American bank's volatility response due to a volatility shock from Barclays alone, plus the gIRF due to a shock from Credit Suisse alone, plus the gIRF due to a shock from Deutsche Bank alone) minus the jIRF due to joint, simultaneous shocks from the three European banks. The difference between the sum of the gIRFs and the jIRF for Bank of America, Citigroup, and JP Morgan Chase are shown in Figure 4 in the paper. The purple line in Figure A.4 shows the difference between the sum of the gIRFs and the jIRF for Bank of New York Mellon (first row), Goldman Sachs (second row), Morgan Stanley (third row), State Street (fourth row), and Wells Fargo (fifth row). The shaded purple regions show the 80% and 99% confidence intervals calculated using the bootstrap. The left column of panels is for the pre-pandemic period, and the right column of panels is for the pandemic period.

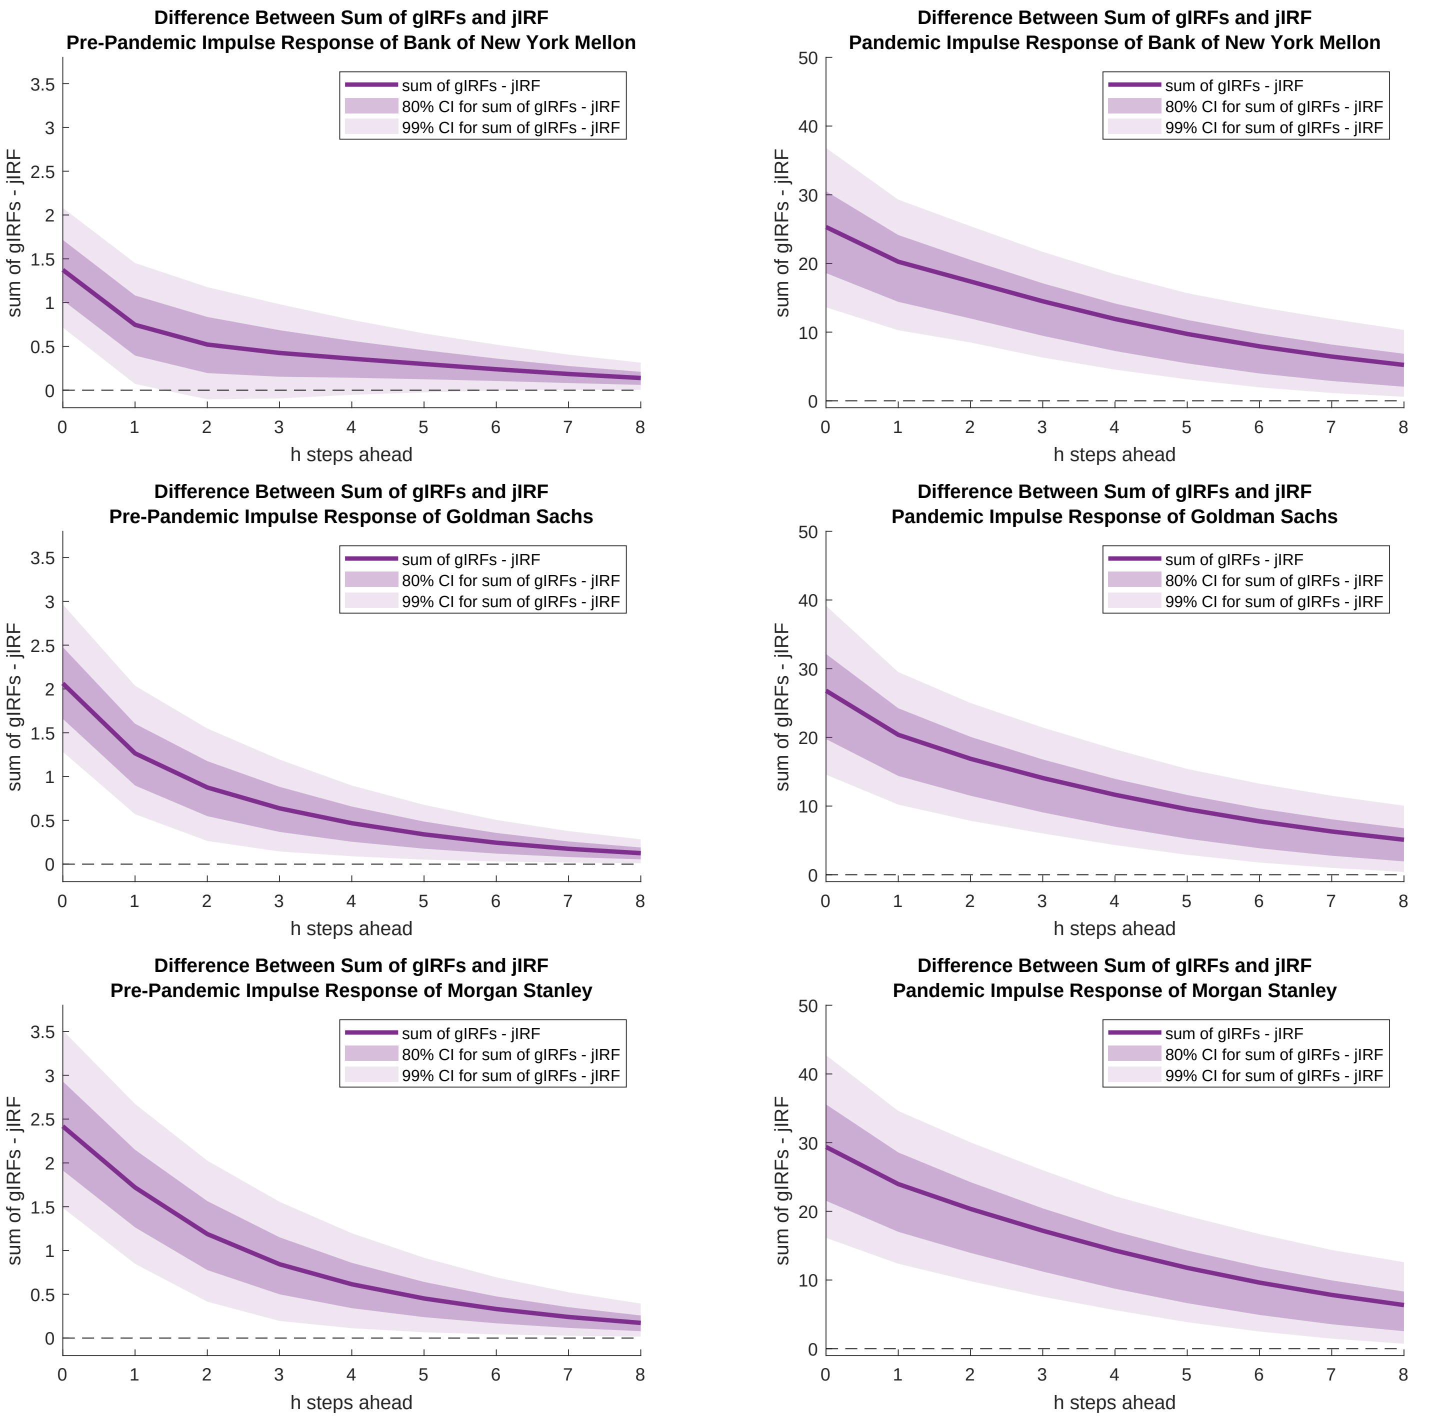

Figure A.4: Difference between the sum of the gIRFs and the jIRF for American banks not shown in the paper.

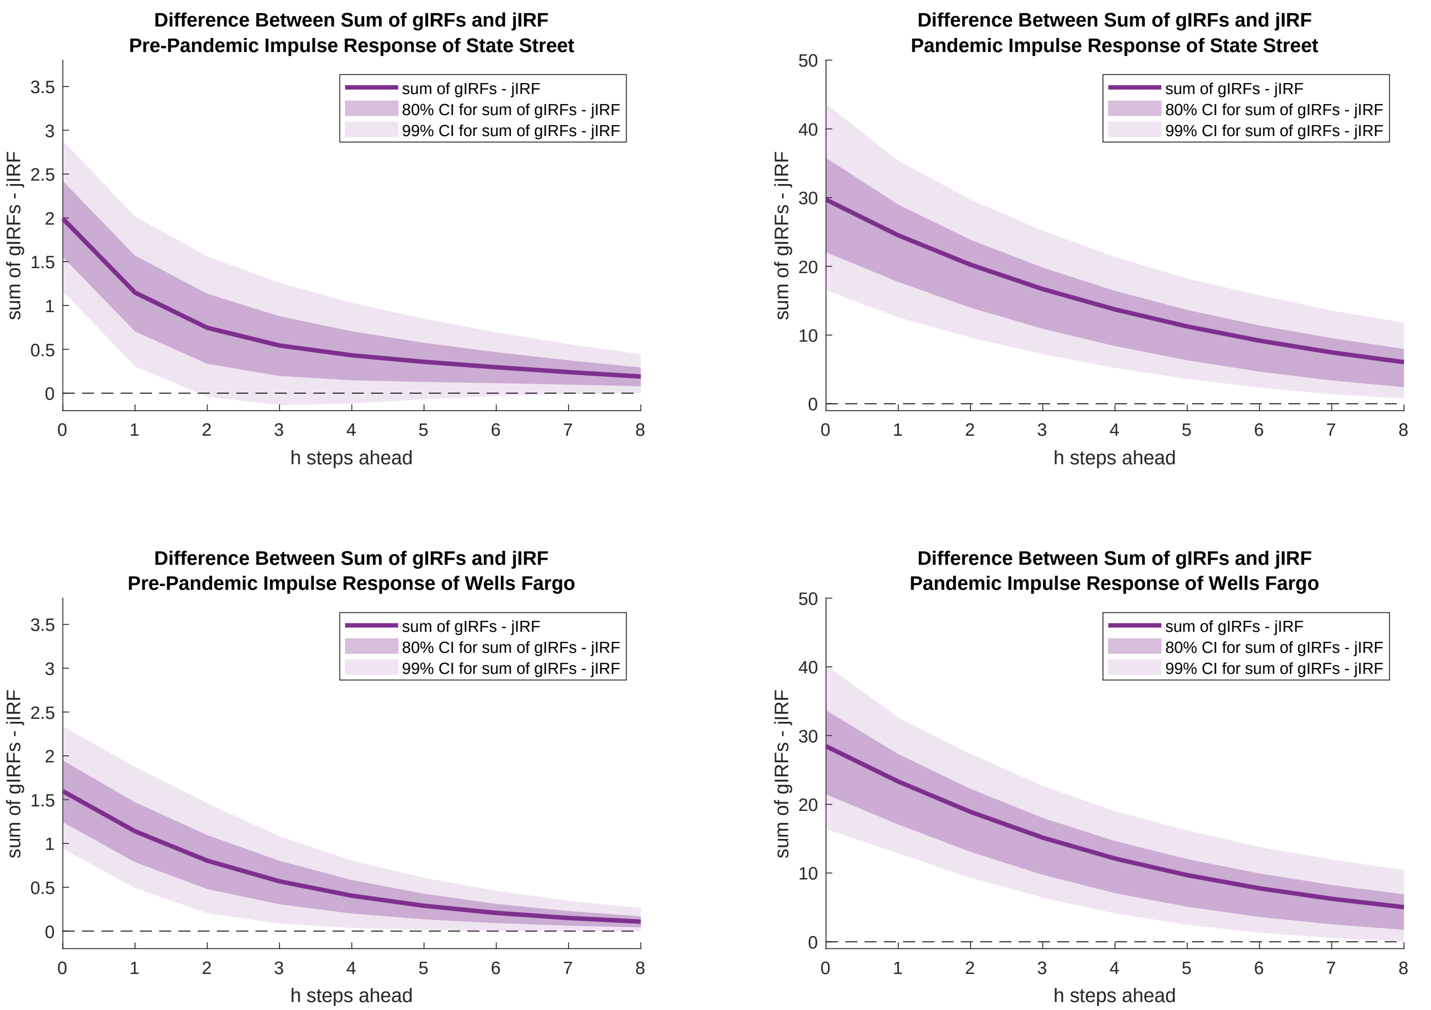

Figure A.4: Difference between the sum of the gIRFs and the jIRF for American banks not shown in the paper. (Continued)

A.4 Constructing a Summary Variable for the European Banks

To verify the efficacy of the jIRF approach, in section 7 of the paper we consider replacing the three individual European banks (Barclays, Credit Suisse, and Deutsche Bank) with three possible European summary variables: (1) the first principal component, (2) the market capitalization weighted average of the volatilities of the European banks, and (3) the volatility of the market capitalization weighted average of the HLOC prices of the European banks.

A.4.1 Principal Component Analysis of European Banks' Volatilities ("EPCA")

To obtain the first European summary variable, we initially compute the volatility of each of the three European banks and then perform principal component analysis using those three vectors of bank volatility. We use the first principal component, explaining the largest variance share, as a European volatility summary variable ("EPCA"). We center the data to have mean zero and scale the data to have unit standard deviation. This will scale the volatility measure but will not affect the results of the VAR analysis. For the Parkinson volatility method, the first component explains about 87 percent of the variance of the European banks, and for the Yang-Zhang volatility method, the first component explains about 84 percent of the variance of the European banks.

A.4.2 European Market Capitalization Weighted Average of Volatilities ("EINDa")

To obtain the second European summary variable, we employ daily market capitalization data for Barclays, Credit Suisse, and Deutsche Bank available from YCHARTS to form an index of volatility for the European banks. In particular, we construct a market capitalization weighted average of the computed volatilities for the European banks ("EINDa"). The time t market capitalization weighted average of volatility for the three European banks is computed as

$$EINDa_t = w_{BCS,t}V_{BCS,t} + w_{CS,t}V_{CS,t} + w_{DB,t}V_{DB,t}$$

where $V_{BCS,t}$, $V_{CS,t}$, and $V_{DB,t}$ are the daily volatilities of Barclays, Credit Suisse, and Deutsche Bank, respectively, using either the Parkinson or the Yang-Zhang estimator. The market capitalization weights at time t are defined as

$$w_{BCS,t} = \frac{MC_{BCS,t}}{MC_{BCS,t} + MC_{CS,t} + MC_{DB,t}}, \tag{A.8}$$

$$w_{CS,t} = \frac{MC_{CS,t}}{MC_{BCS,t} + MC_{CS,t} + MC_{DB,t}}, \tag{A.9}$$

and

$$w_{DB,t} = \frac{MC_{DB,t}}{MC_{BCS,t} + MC_{CS,t} + MC_{DB,t}}, \tag{A.10}$$

where $MC_{BCS,t}$, $MC_{CS,t}$, and $MC_{DB,t}$ are the time t market capitalizations of Barclays, Credit Suisse, and Deutsche Bank, respectively.

A.4.3 Volatility of the Market Capitalization Weighted HLOC Prices of the European Banks ("EINDb")

The third alternative for constructing a European summary variable again involves the market capitalization data. However, with this third method, we first compute the market capitalization weighted values for the HLOC prices of the three European banks and then compute the volatility using these weighted HLOC values. Whereas "EINDa" was a weighted average of three banks' volatilities, "EINDb" is the volatility of an index formed through the weighted averages of the three banks' HLOC equity prices.

A.4.4 Comparison of the Three European Summary Variables

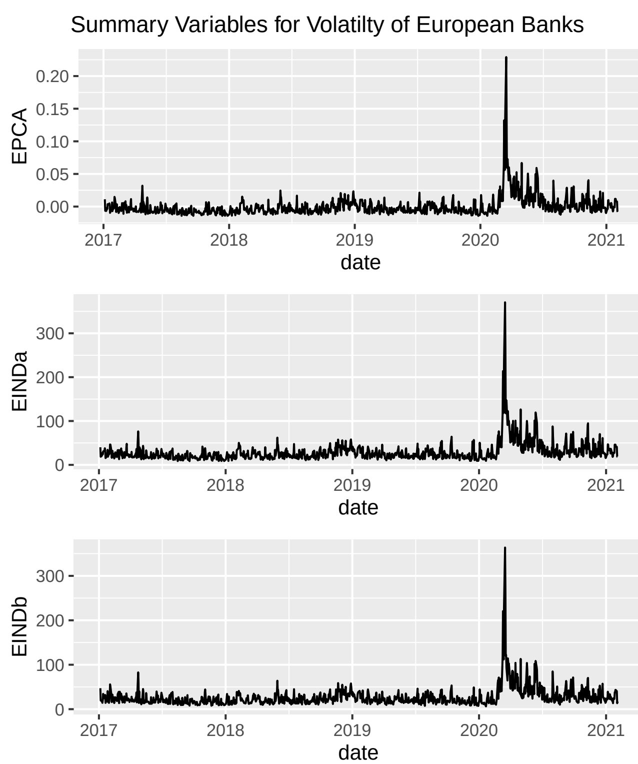

Figure A.5: Three summary variables for the European banks' volatility. EPCA: first principal component of volatilities; EINDa: market capitalization weighted average of volatilities; EINDb: volatility of market capitalization weighted average of HLOC prices.

Figure A.5 above shows the three alternative variables used to summarize the three European banks' volatility, "EPCA," "EINDa," and "EINDb." Except for the rescaling of EPCA, it is clear that the three different alternatives produce nearly identical results. Therefore, for succinctness, we only use the principal component ("EPCA") and the market capitalization weighted average of volatilities ("EINDa"), while omitting the volatility of the market capitalization weighted HLOC prices ("EINDb") in section 7 of the paper.

A.5 Comparing the jIRF to Single Variable Alternatives: Additional Impulse Response Functions

In section 7 of the paper, we use variables that summarize the volatility of the three European financial institutions. We then estimate VAR models that include one of the European summary variables and the volatilities of the eight American financial institutions, while excluding the volatilities of the three individual European financial institutions. We then obtain the gIRF measuring how the volatility of an American financial institution responds due to a shock from the European summary variable. These gIRFs are then compared to the jIRFs from section 6 of the paper and to section A.3 of this appendix.

The goal of this analysis is to see if the jIRF, which measures the total effect of joint, simultaneous shocks from the three European banks, is identical to a gIRF measuring the effect of a shock from the European summary variable. As explained in section A.4, we summarize the volatility of the three European banks in multiple ways; (1) using the first principal component obtained from principal component analysis, (2) using an index of market capitalization weighted averages of the three European bank volatilities. A third alternative—the volatility of the market capitalization weighted average of HLOC prices—was also considered, but we omit those IRFs because the results are nearly indistinguishable from the other results. The impulse responses for Bank of America, Citigroup, and JP Morgan Chase are shown in Figure 5 in the main text. This appendix shows the impulse responses for Bank of New York Mellon, Goldman Sachs, Morgan Stanley, State Street, and Wells Fargo.

In Figure A.6 below, we show the impulse responses of various American financial institutions due to volatility shocks from the European summary variables. The left column of panels is for the pre-pandemic period, and the right column of panels is for the pandemic period. The first row of panels is the impulse response for Bank of New York Mellon (BK), the second row is for Goldman Sachs (GS), the third row is for Morgan Stanley (MS), the fourth row is for State Street (STT), and the fifth row is for Wells Fargo (WFC). In each of the 10 panels, the orange line with upside down triangle markers is the gIRF measuring how the American bank's volatility responds due to a shock from the first principal component of European volatilities. To be clear, this gIRF was obtained from a VAR that included the eight American banks' volatilities and the first European principal component, while excluding the three separate volatilities for the European banks. In each panel, the pink line with right side up triangle markers is the gIRF measuring how the American bank's volatility responds due to a shock from the European index of market capitalization weighted volatilities. To be clear, this gIRF was obtained from a VAR that included the eight American banks' volatilities and the market capitalization weighted index of the European bank volatilities, while excluding the three separate volatilities for the European banks. For comparison, the black lines with circle markers in Figure A.6 are the jIRFs measuring how the American bank's volatility responds due to joint, simultaneous volatility shocks from the three European banks. These jIRFs are identical to the jIRFs shown in Figure A.3 of this appendix, which were obtained from a VAR that included the eight American banks and the three European banks. Notice the vertical scale in the pre-pandemic period panels is different from the scale in the pandemic period panels.

Notice that the gIRFs in Figure A.6 are consistently smaller than the jIRFs. This is because the European summary variables used to estimate the gIRFs are averages, and any average loses some of the total variance in the original variables. Thus, the gIRFs in Figure A.6 underestimate the effect of the shocks. In contrast, the jIRF allows us to measure the total impact of simultaneous shocks from the three European banks—there is no loss of variance with the jIRF. As argued in section 7 of the paper, the jIRF produces more appropriately sized responses compared to the gIRFs of single variable alternatives.

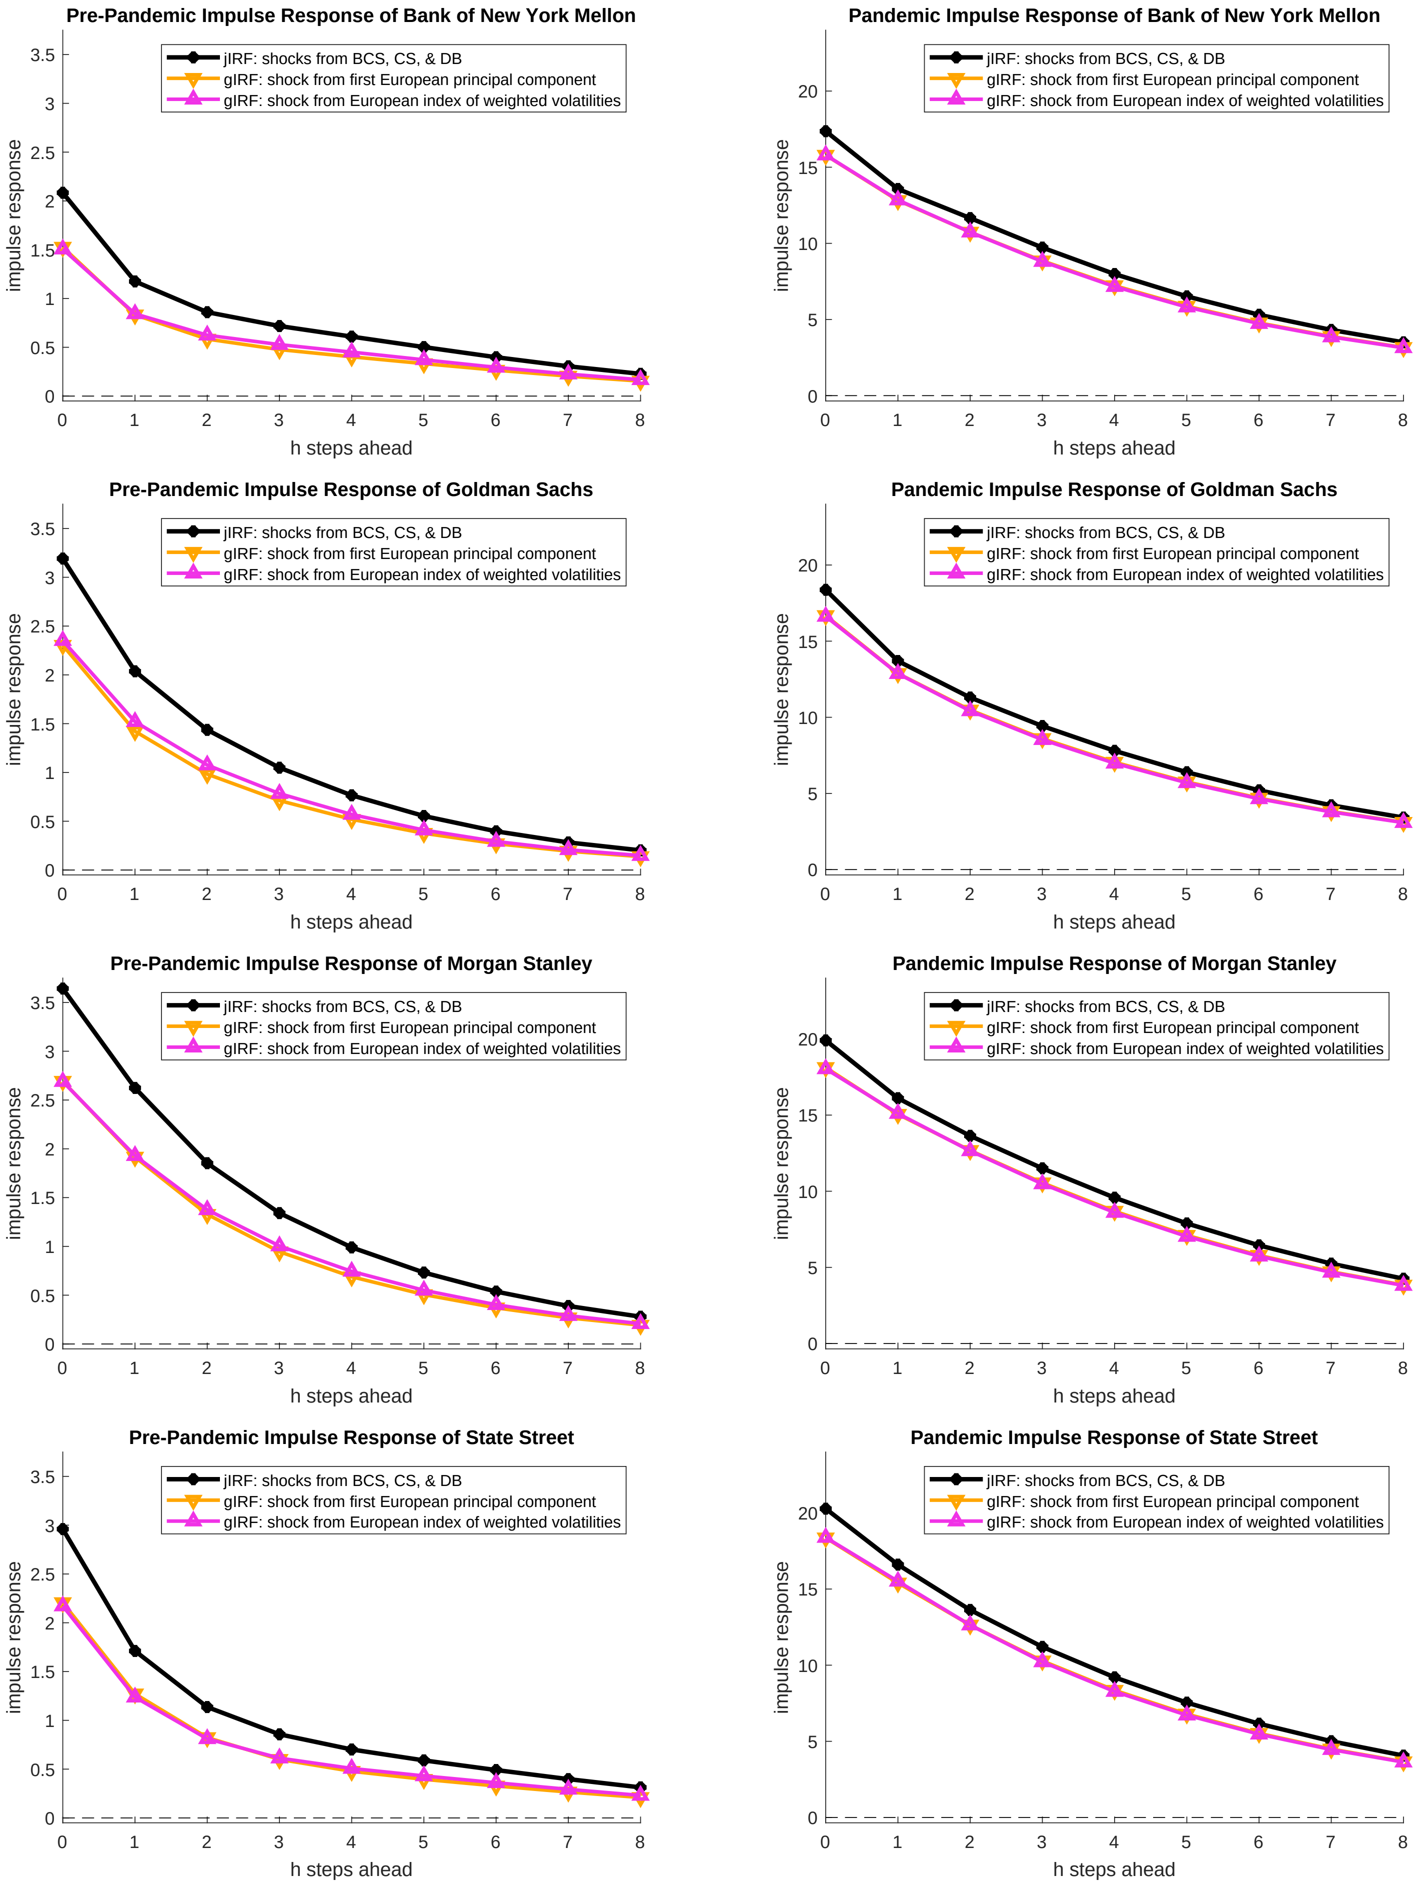

Figure A.6: The gIRFs using the European summary variables for American banks not shown in the paper. The jIRFs are included for comparison.

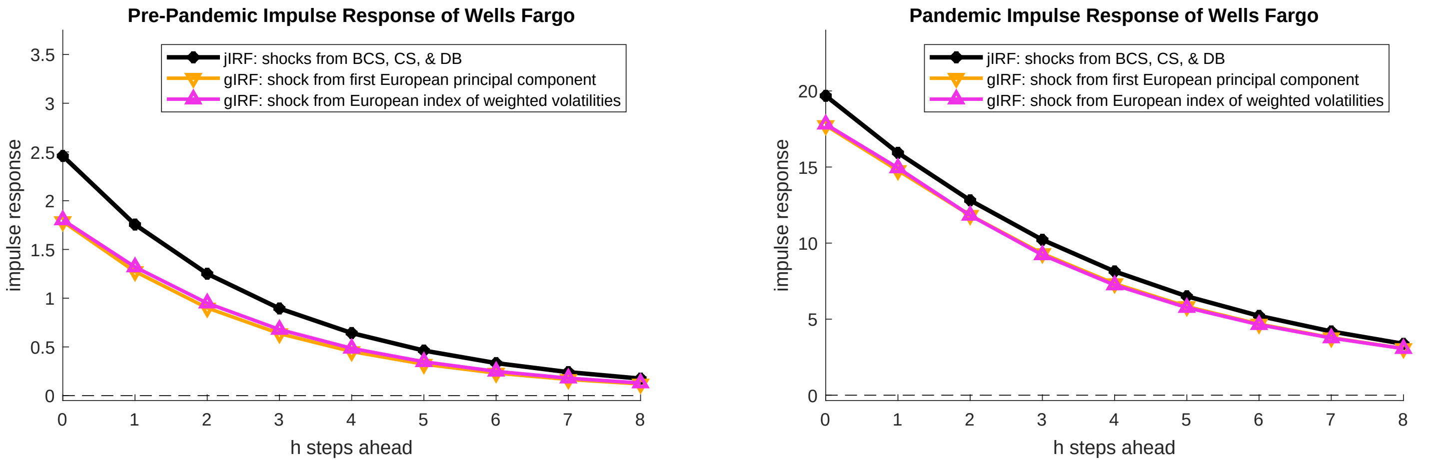

Figure A.6: The gIRFs using the European summary variables for American banks not shown in the paper. The jIRFs are included for comparison. (Continued)