Filters & Impulse Response Functions

Macroeconomics (M8674), March 2024

Vivaldo Mendes, ISCTE

vivaldo.mendes@iscte-iul.pt

1. Filters

Types of Filters

- Main objective: to separate the **long-run trend** from the **short-run cyclical component** of a time series $y(t)$.
- There are various approaches to achieve this:
  - Linear filter
  - Linear filter with breaks
  - Nonlinear filters
- Nonlinear filters
  - Hodrick-Prescott filter (Hodrick & Prescott, 1997)
  - Band Pass filter (Baxter & King, 1999)
  - Hamilton filter (James Hamilton, 2017)
  - ...and some others

The Hodrick-Prescott filter

- The HP filter is the most used filter in macroeconomics
- It is given by the minimization problem

$$\min_{\tau_t} \sum_{t=1}^{T} \left\{ (y_t - \tau_t)^2 + \lambda [(\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1})]^2 \right\} \tag{1}$$

- where:
    - y_t is the observed time series
    - τ_t is the smooth trend that we want to obtain
    - λ is the parameter that we set to obtain the *desired* smoothness in the trend

The HP filter: Special Cases

The value given to parameter λ is a choice of ours:

$$\min_{\tau_t} \sum_{t=1}^{T} \left\{ (y_t - \tau_t)^2 + \lambda [(\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1})]^2 \right\}$$

- $\lambda = 0 \Rightarrow$ trivial solution because there are **no cycles**: $y_t = \tau_t, \forall t$
- $\lambda \to \infty \Rightarrow$ linear trend leads to **huge cycles** between y_t and τ_t
- $\lambda = 1600 \Rightarrow$ duration/amplitude of cycles acceptable for **quarterly data**
- $\lambda = 7 \Rightarrow$ duration/amplitude of cycles acceptable for **annual data**
- There is no "unquestionable" value for λ

The HP Filter: an Example

- Main objective: obtain cycles as % deviations from the trend
- This has an important implication:
  - Time series with a trend: *apply logs* to the data before extracting the trend and the cycles
  - Time series without a clear trend: *do not apply logs* to the data
- Quarterly data: "US_data.csv"
- A simple example:
  - Real GDP (GDP)
  - Consumer Price Index (CPI)
  - Unemployment Rate (UR)

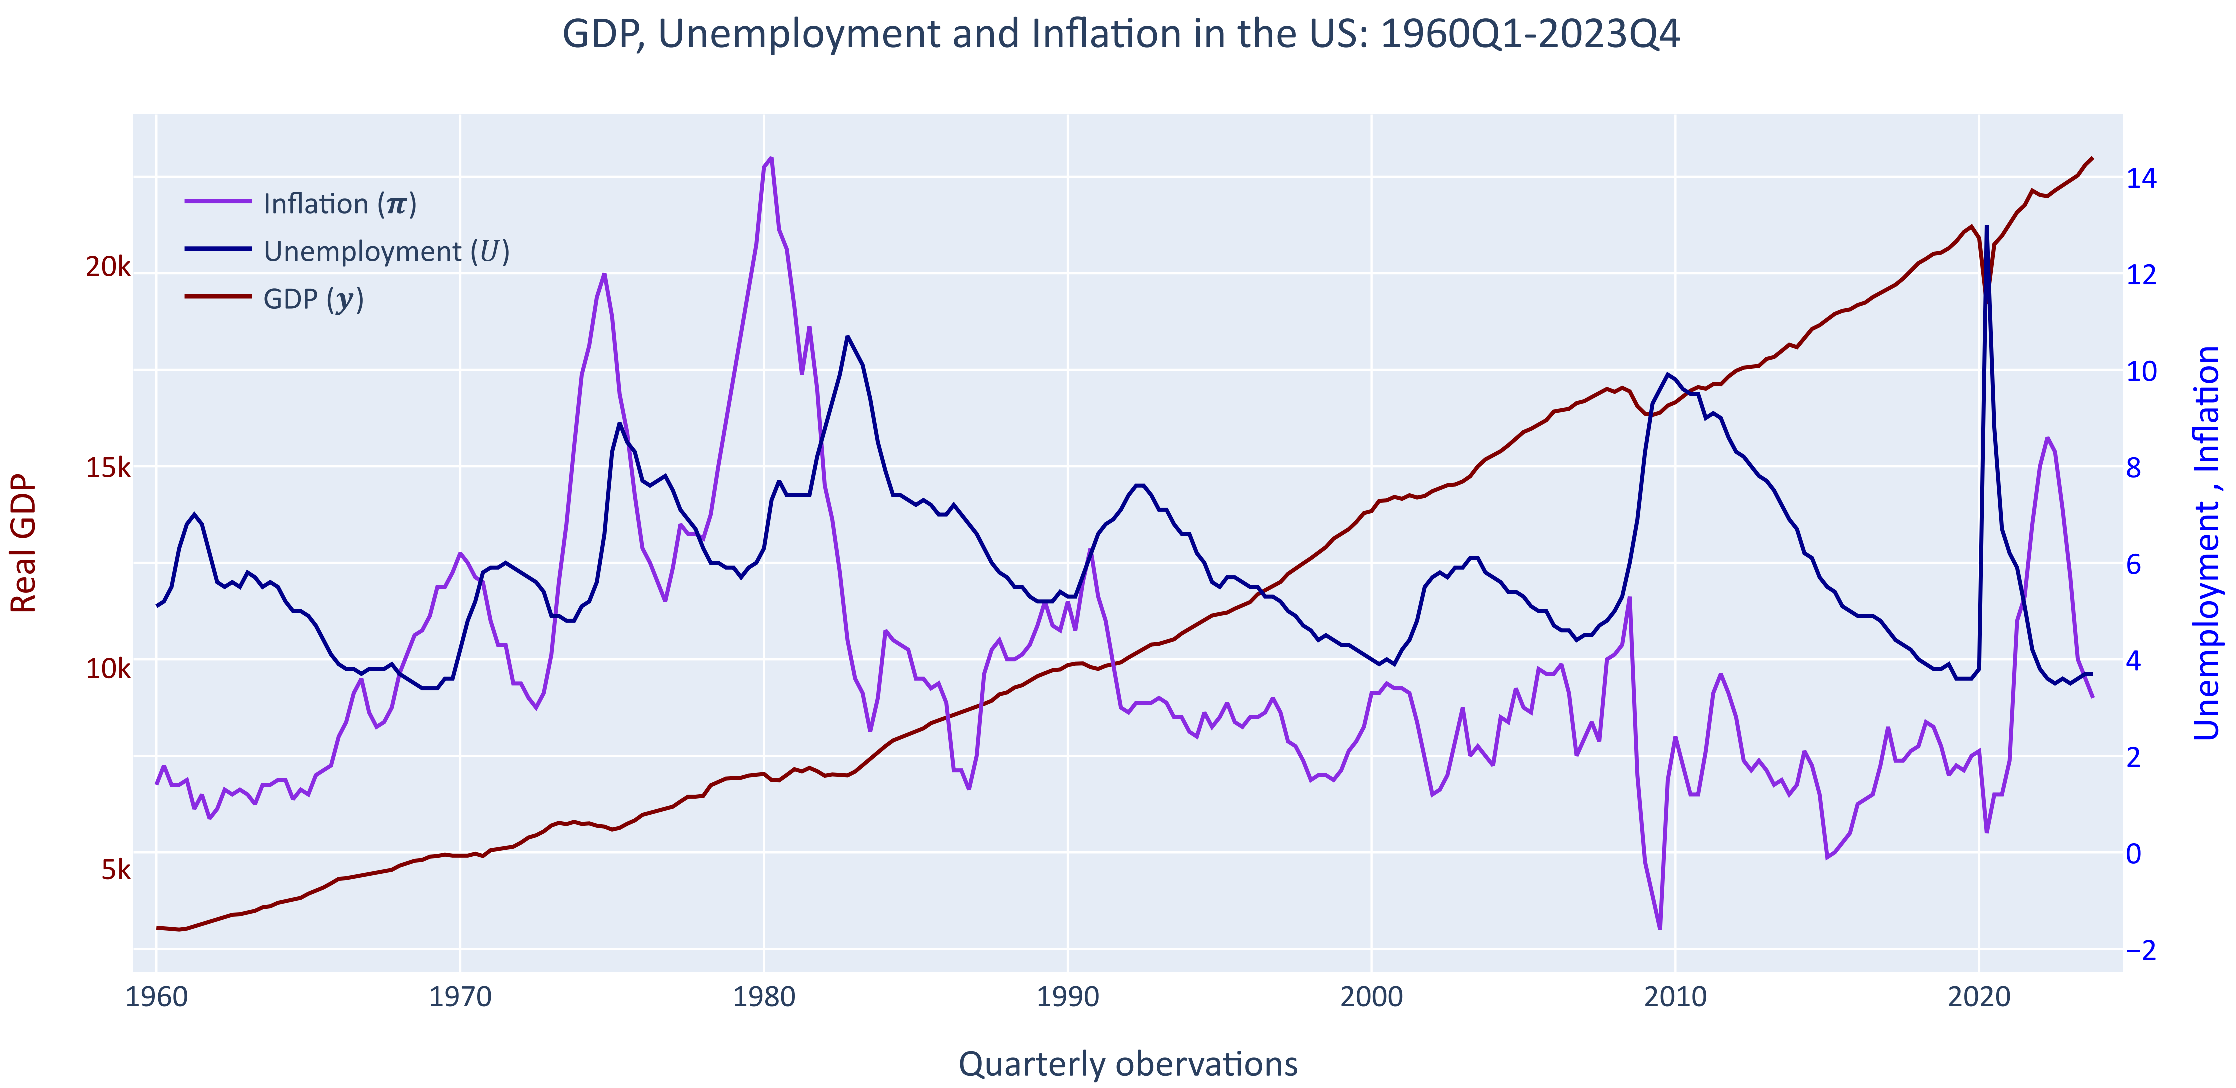

GDP, Unemployment and Inflation in the US: 1960Q1-2023Q4
Inflation (π)
Unemployment (U)
GDP (y)
Real GDP
5k
10k
15k
20k
Unemployment , Inflation
−2
0
2
4
6
8
10
12
14
1960
1970
1980
1990
2000
2010
2020
Quarterly obervations

Dealing with rows and columns in a Matrix

```
choco = 10×6 Matrix{Float64}:
  0.626776    0.262165    0.0625441  -0.672179   -1.18456   -0.557628
 -0.928241    0.669247   -1.52369     0.79812     1.6391    -0.290093
  0.714142    0.903336    0.209435   -0.330667    0.283155   1.29553
 -1.41671     0.315082    0.249256    0.136191   -0.889723   1.11908
 -0.300019   -0.12429    -0.328248    0.749681   -0.435038  -0.671975
 -1.08759    -1.78173     0.18719    -1.55239     0.721633   0.518601
 -0.737495   -0.50324    -0.0896977   0.705148    0.185982   0.502288
 -1.28098    -0.0335934  -0.858966   -0.412964    1.26839   -0.838636
  0.0447736  -0.743076   -1.04855     0.0814288  -1.62432    0.902365
  0.49268     0.728621    0.355343    0.517302    2.31929    0.589591

choco = randn(10,6)

choco_slice = 3×3 Matrix{Float64}:
 -0.672179  -1.18456   -0.557628
  0.79812    1.6391    -0.290093
 -0.330667   0.283155   1.29553

choco_slice = choco[1 : 3 , 4 : 6]
```

Choosing rows

separator

Choosing columns

Calculate the HP filter for a single variable

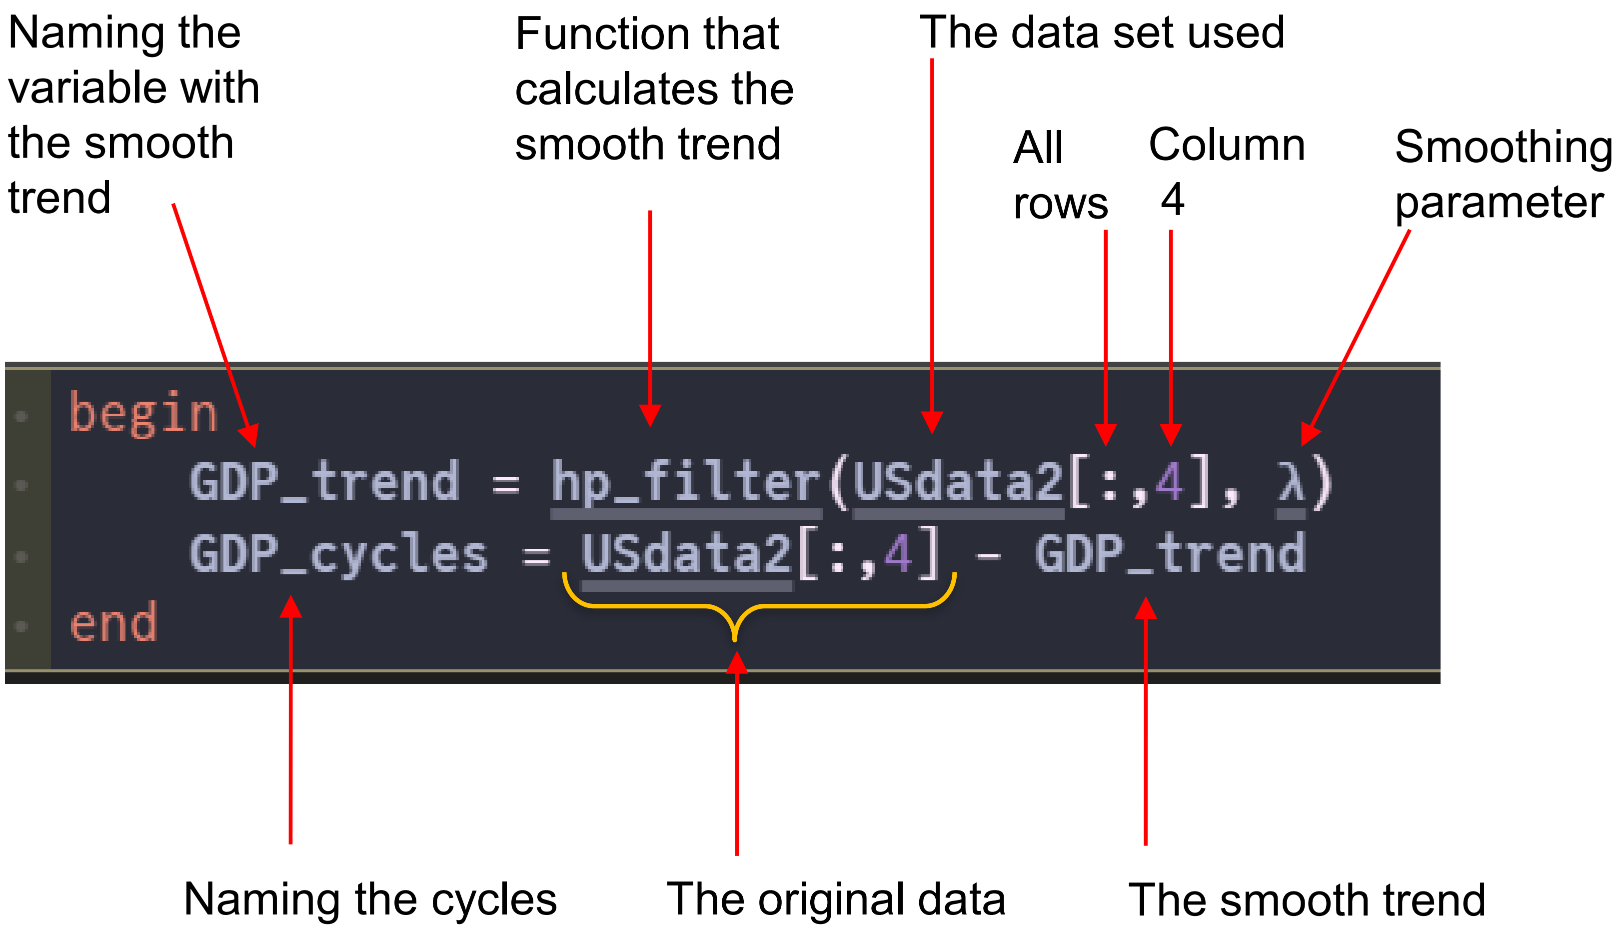

Calculate the HP filter for several variables

Number of columns to iterate: $n = 1,2,3,4$

Allocating space for the trend in each variable (4 altogether): a vector named as **hp_trend** with

287 rows

4 columns

All rows

The **for** loop

```
begin
    hp_trend = zeros(287,4)
    for n = 1:4
        hp_trend[:,n] = hp_filter(USdata2[:,n], λ)
    end
    hp_trend
end
```

Fill in with trend data

Function that calculates the smooth trend

The original data set

Number of columns that will be iterated, one-by-one

$n = 1,2,3,4$

Business cycles: Inflation and Unemployment

Inflation-gap	Cyclical unemployment
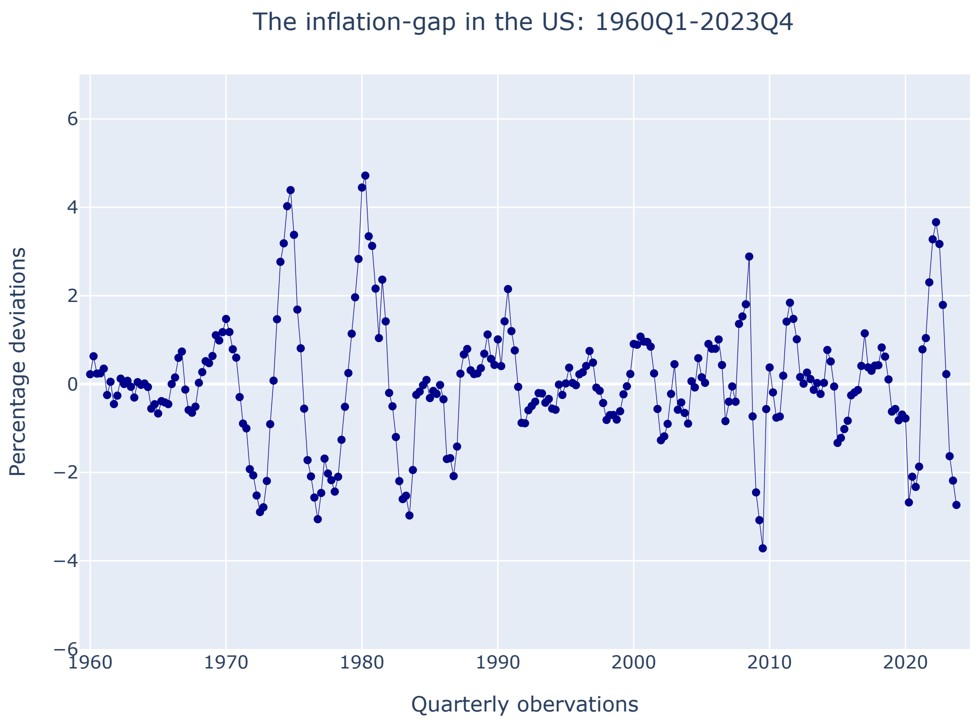	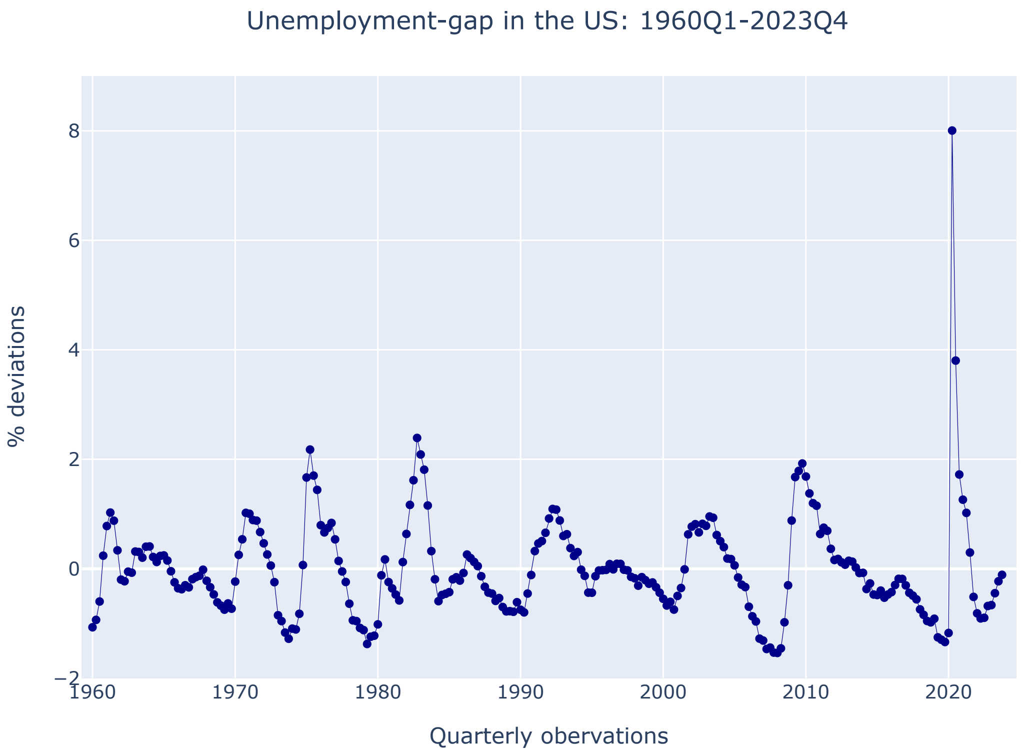

Business cycles: Output gap

Correctly done (using logs)	Incorrectly done (using levels)
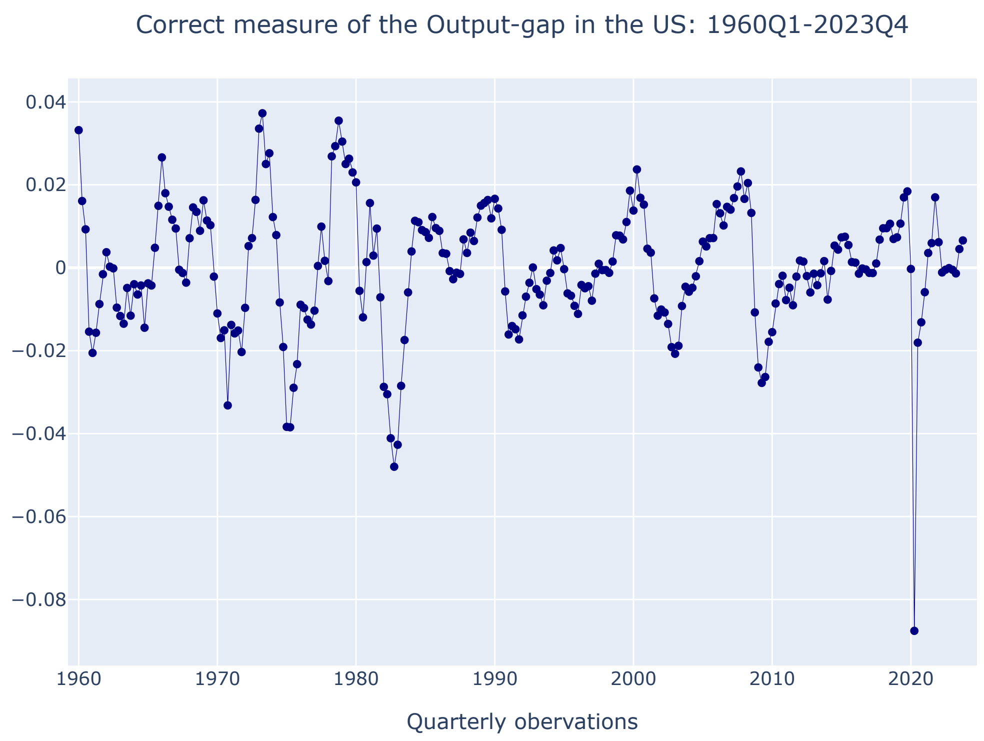	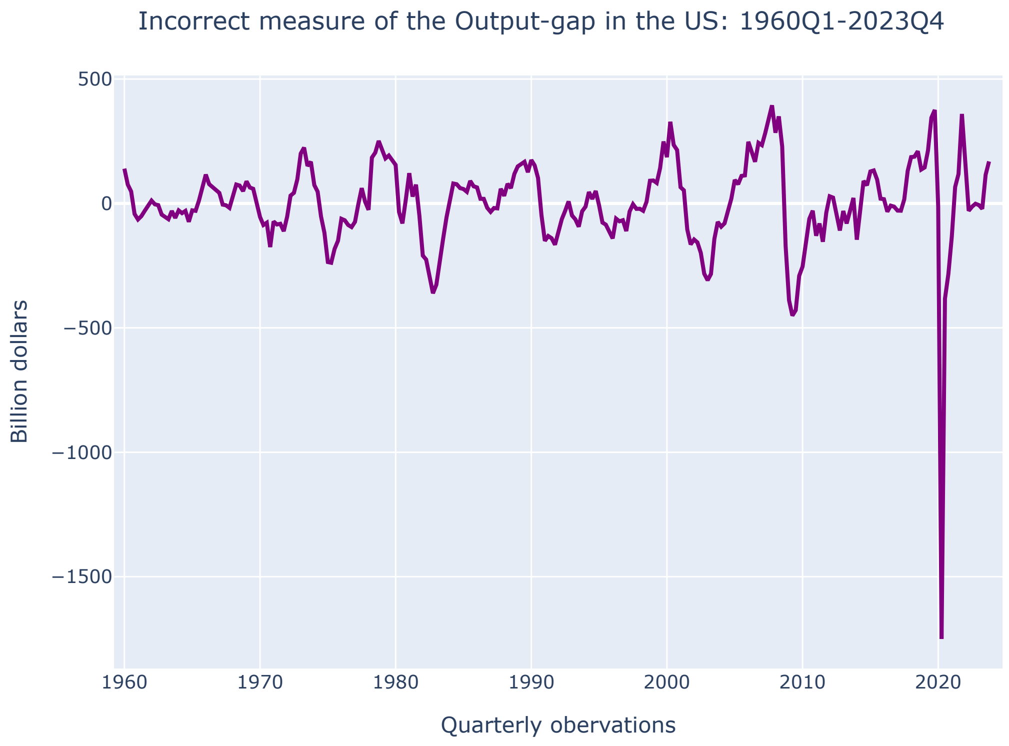

2. Impulse Response Functions

What are IRFs?

- Impulse response functions represent the dynamic impact upon a system of an exogenous “shock” that hits one (or more than one) of its endogenous variables.
- In stationary systems, we expect the shocks to be temporary (not persistent), and over time the system converges to the original state or to some alternative state.
- This depends on the system's structure and the magnitude of the shock.
- Linear systems. The magnitude of the shock may not be highly relevant to
- Nonlinear systems. All macroeconomic models are nonlinear, and in such a case, the magnitude of the shock is of tremendous importance.

An example

- Consider the simplest case, an AR(1):

$$y_{t+1} = a y_t + \varepsilon_{t+1}, \quad \varepsilon_t \sim N(0,1) \tag{2}$$

- Assume that for $t \in (1, n)$:

$$y_1 = 0\ ;\ \varepsilon_2 = 1\ ;\ \varepsilon_t = 0\ ,\ \forall t \neq 2$$

- This implies that at $t = 2$, $y_2 = 1$, but what happens next, if there are no more shocks?
- The IRF of y provides the answer.
- The dynamics of y will depend crucially on the value of a. Six examples:

$$a = \{0, 0.5, 0.9, 0.99, 1, 1.01\}$$

The IRFs of the AR(1) Process

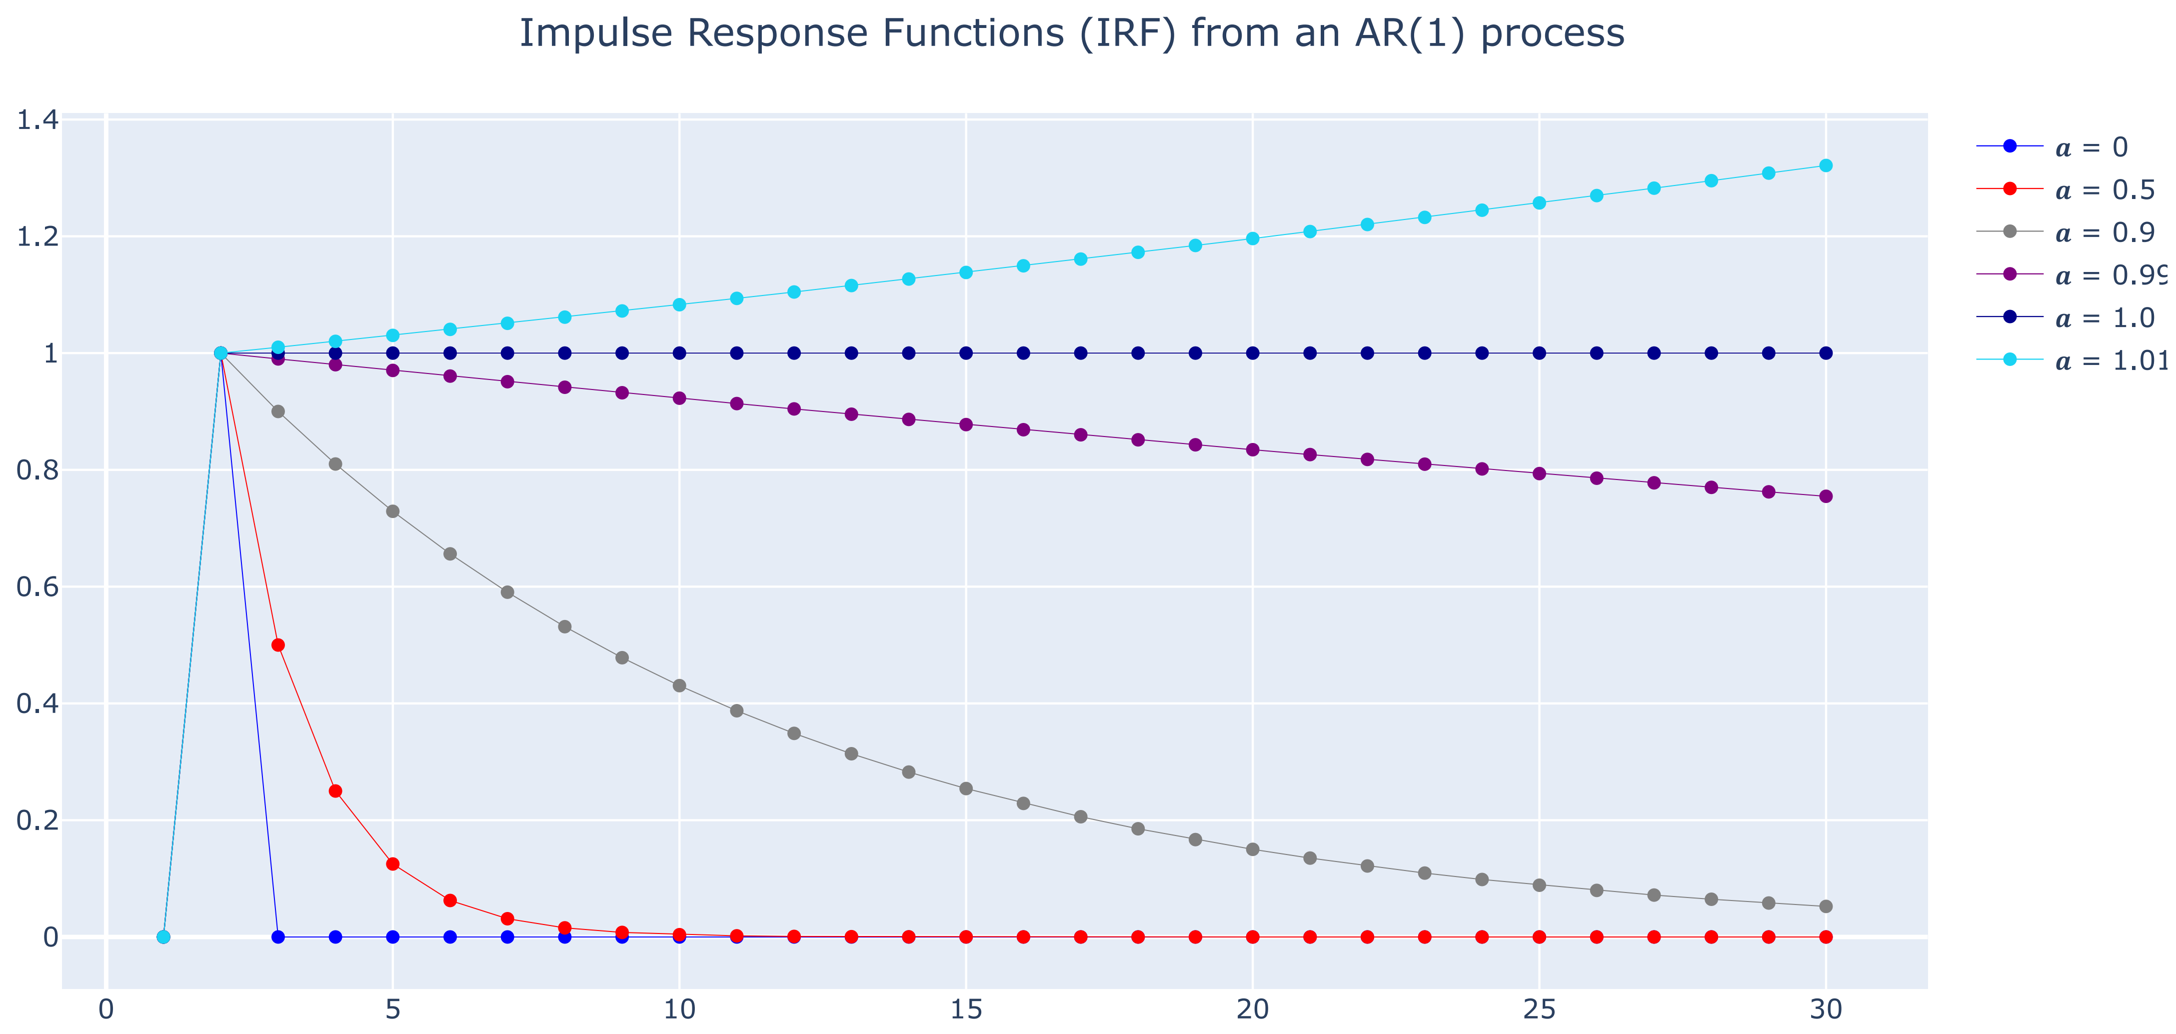

Another example

- Consider a more sophisticated case, an AR(2):

$$y_{t+1} = a y_t + b y_{t-1} + \varepsilon_{t+1}, \quad \varepsilon_t \sim N(0,1)$$

- Assume that for $t \in (1, n)$:

$$y_1 = 0\,;\; y_2 = 0\,;\; \varepsilon_3 = 1\,;\; \varepsilon_t = 0\,,\; \forall t \neq 3$$

- This implies that at

$$t = 3\,,\; \varepsilon_3 = 1 \;\Rightarrow\; y_3 = 1.$$

  What happens next, if there are no more shocks? The IRF of y provides the answer.

- The dynamics of y will depend on the values of a and b. For simplicity consider:

$$b = -0.9\,;\, a = \{1.85, 1.895, 1.9, 1.9005\}$$

The IRFs of the AR(2) Process

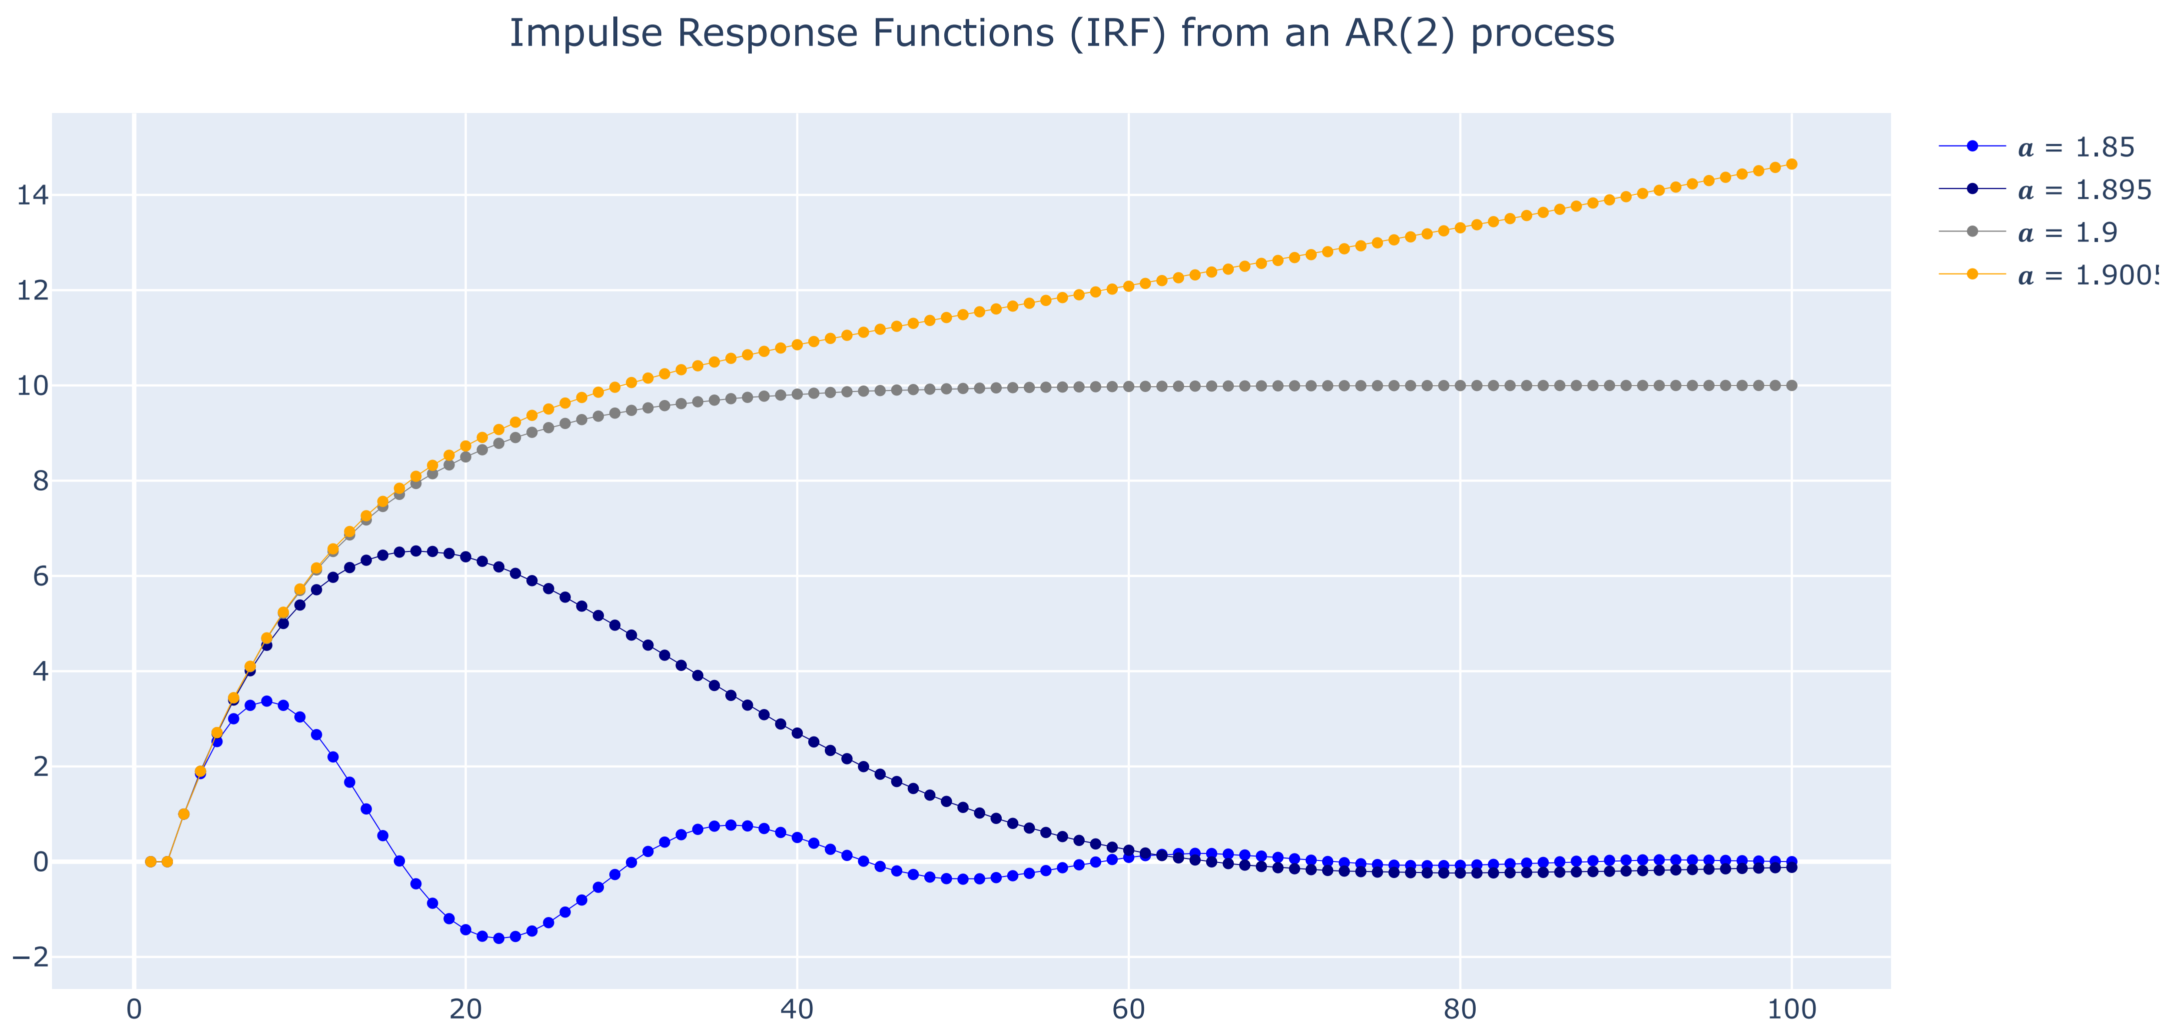

More Sophisticated Examples

A similar reasoning can be applied to our rather more general model:

$$X_{t+1} = A + BX_t + C\varepsilon_{t+1} \tag{3}$$

where B, C are $n \times n$ matrices, while $X_{t+1}, X_t, A, \varepsilon_{t+1}$ are $n \times 1$ vectors.

Consider the following VAR(3) model:

$$X_{t+1} = \begin{bmatrix} z_{t+1} \\ w_{t+1} \\ v_{t+1} \end{bmatrix}$$

In this example we take matrices A, B and C given by:

$$A = \begin{bmatrix} 0.0 \\ 0.0 \\ 0.0 \end{bmatrix}, \quad B = \begin{bmatrix} 0.97 & 0.10 & -0.05 \\ -0.3 & 0.8 & 0.05 \\ 0.01 & -0.04 & 0.96 \end{bmatrix}, \quad C = \begin{bmatrix} 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 \end{bmatrix}.$$

- The initial state of our system (or its initial conditions) are: $z_1 = 0, w_1 = 0$ and $v_1 = 0$, that is:

$$X_1 = [0, 0, 0]$$

- The shock only hits the variable z_t (notice the blue entry in matrix C), and we assume that the shock occurs in period $t = 3$.
- What happens to the dynamics of the three endogenous variables? See next figure.

The IRFs of our VAR(3) Process

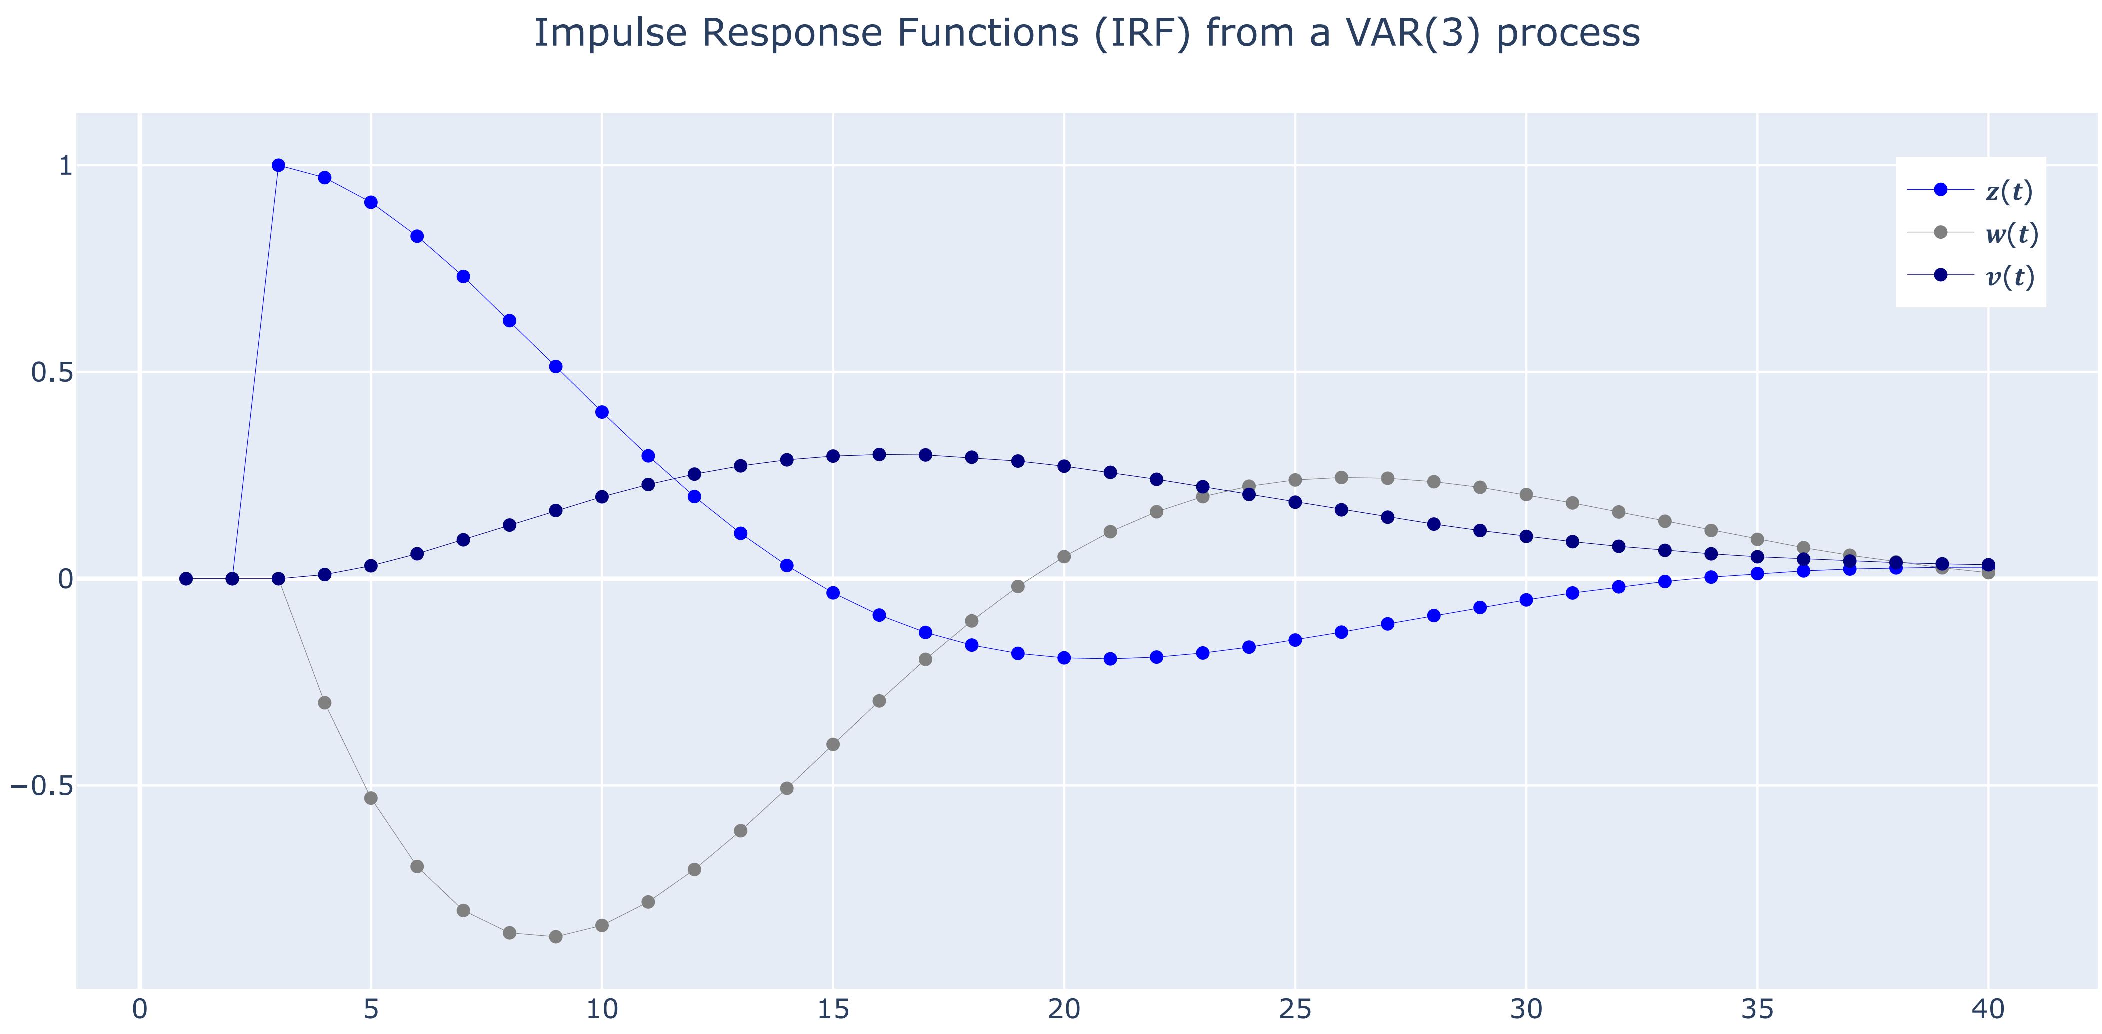

AR(1): A Sequence of Shocks

Consider the same AR(1) as in eq. (2). But now impose a sequence of 200 shocks.

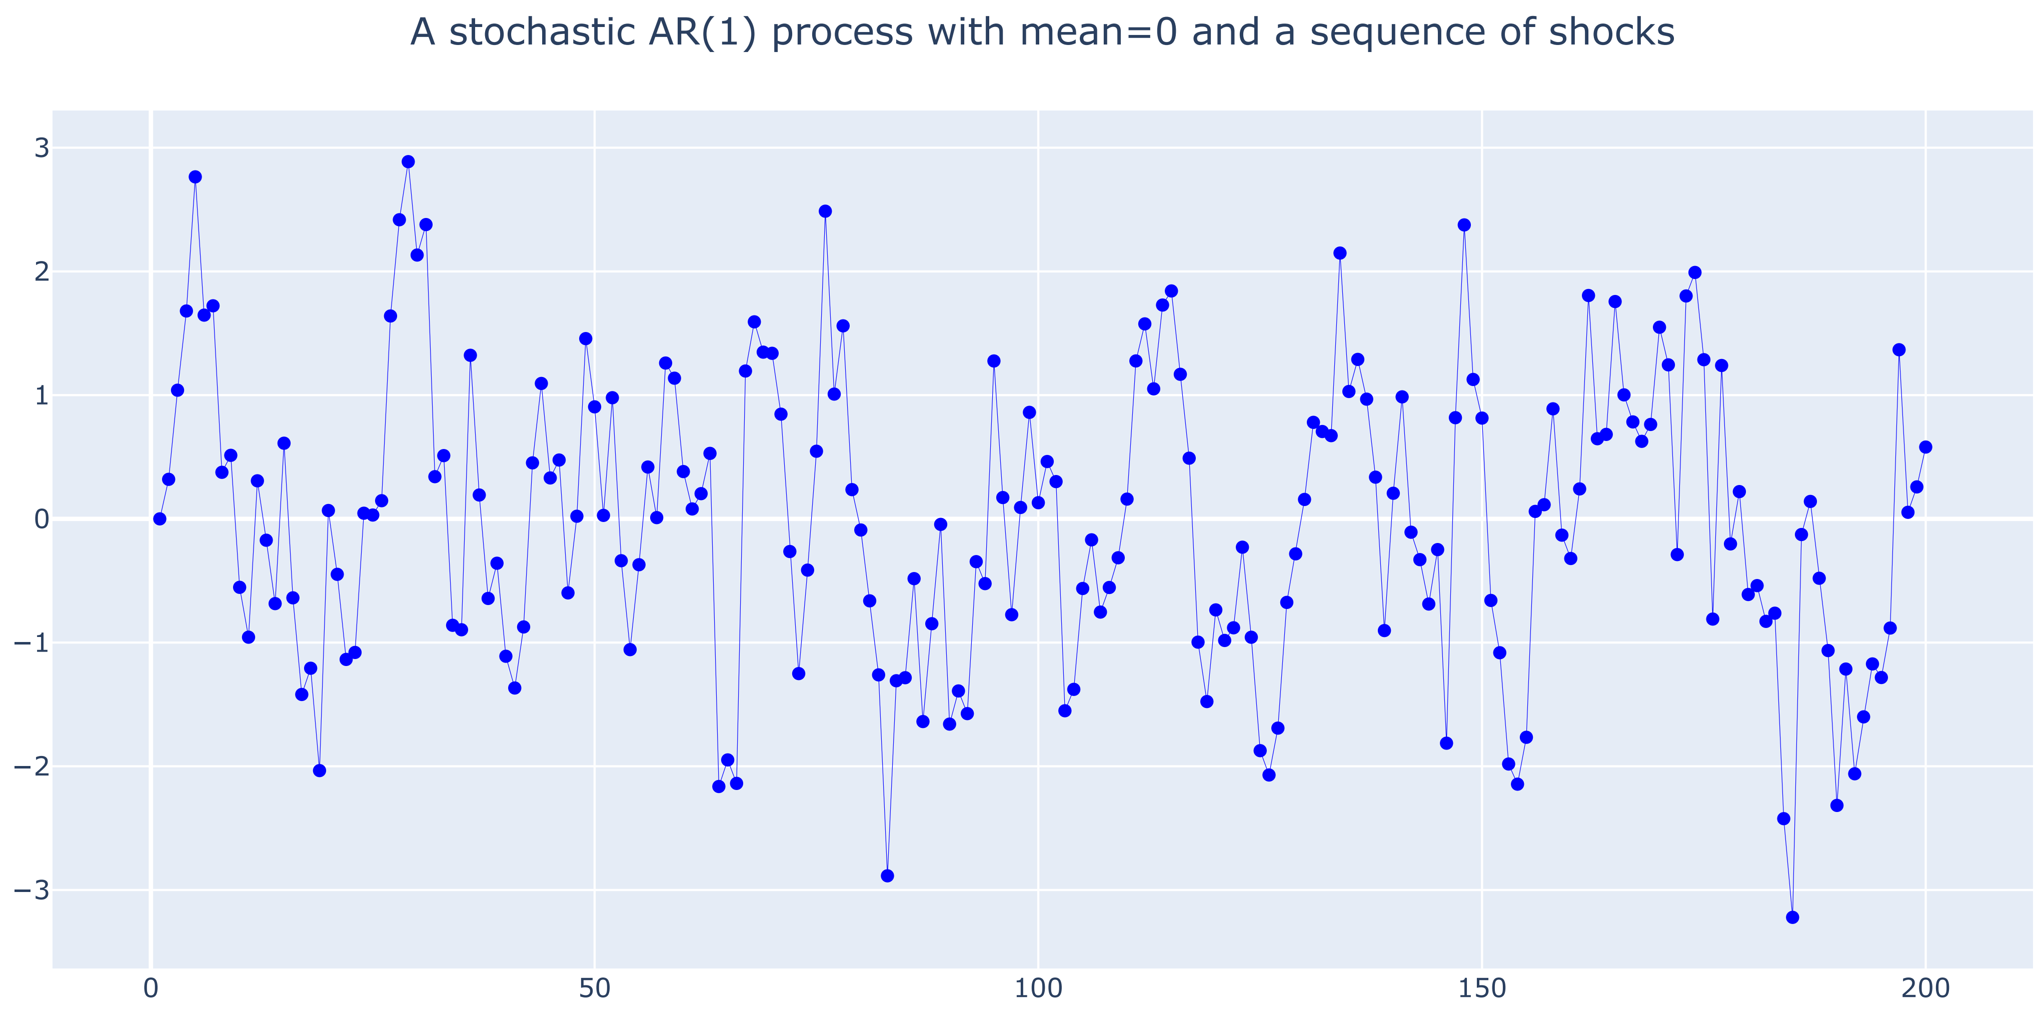

Implications of a Linear Structure

- In the previous examples, the structure of all our models was linear.
- This has a crucial implication:
  The shock's magnitude did not alter the dynamics produced by the shock itself.
  - Only the structure of the model would lead to different outcomes.
- This does not usually occur if the structure of the model is ***non-linear***. In this case, the magnitude of the shock may produce different outcomes even if the system's structure remains the same.
- We do not have time to cover this particular point.
- But be careful: *if the structure of the model is non-linear, large shocks can not be simulated ... in a linearized version of the original system.*

3. Important Problems

Three Major Issues

- There is no perfect filter … but the HP seems the best.
- Measuring Potential GDP (or Natural Unemployment) is difficult:
  - Potential GDP is usually associated with the HP-trend in GDP … but not exclusively.
  - The Natural Rate of Unemployment is largely associated with the HP-trend in unemployment.
- All macroeconomic models are non-linear:
  - Be careful with IRFs that are produced by a linearized version of the model.
  - Big shocks are difficult to be fully represented by linearization.

Limitations of the HP Filter

- New data leads to the rewritting of the history of the economy
- The HP filter is extremely useful but should be used with care

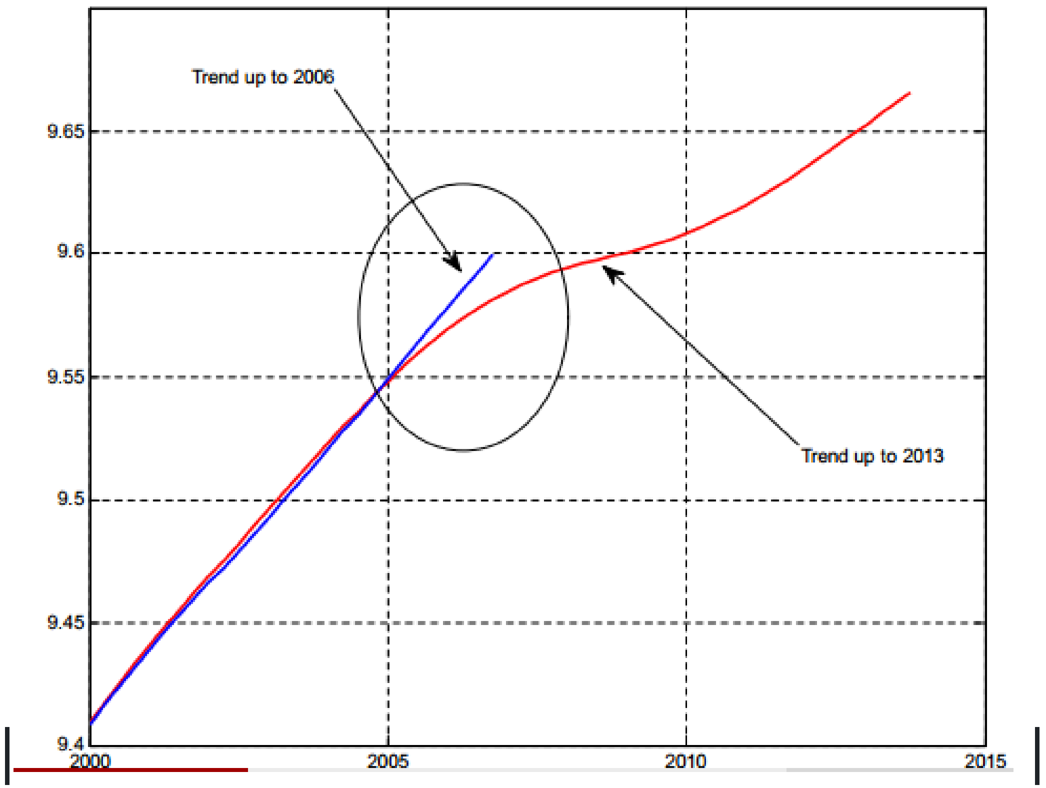

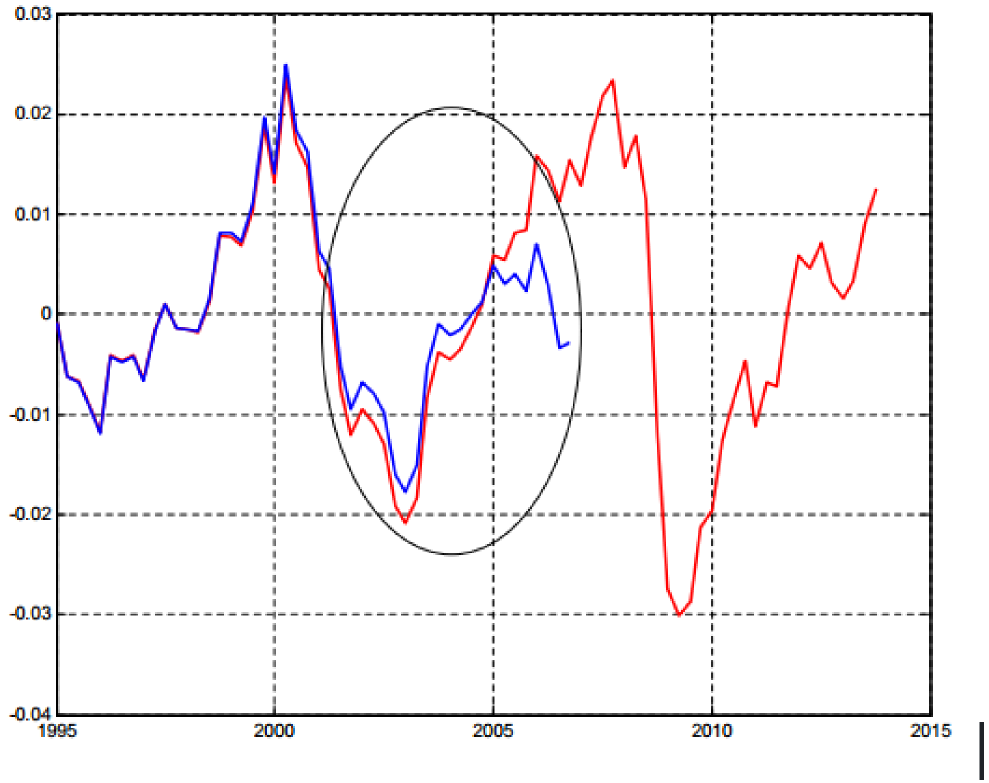

Misuses of the HP Filter

- In 2012, the US economy had an unemployment rate close to 8%, one of the highest rates since WWII.
- The Fed Funds Rate was at 0%, to stimulate the economy.
- The inflation rate was much below the target level (2%) at 0.5% and showing signs of going down.
- James Bullard (the President of the FRB of St. Louis), in a famous speech in June 2012 defended that the US economy had gone back to Potential GDP.
    - Therefore, the Fed should produce a sharp increase in the Fed Funds Rate.
    - He used the HP-filter to substantiate his proposal.

The HP filter according to James Bullard

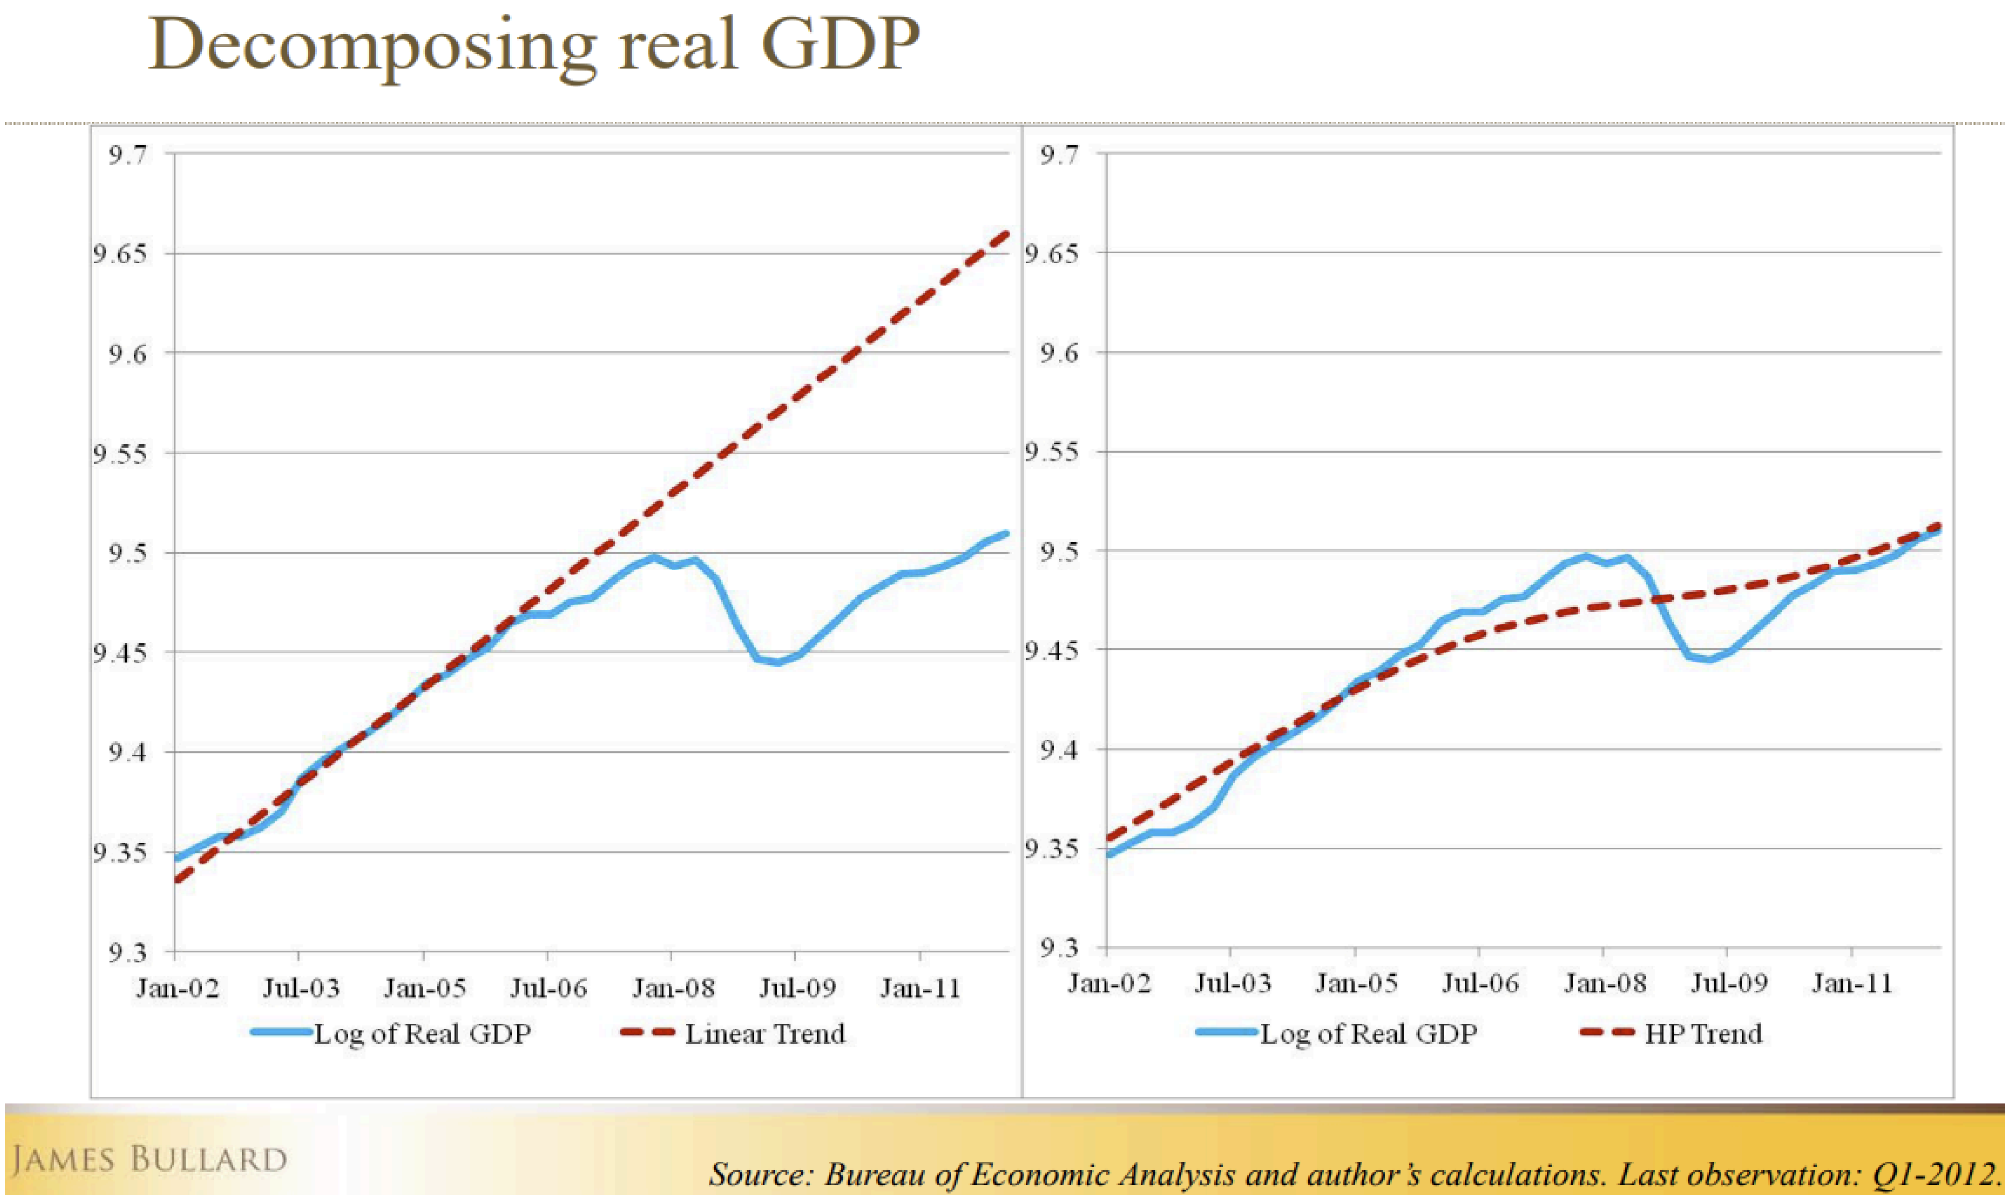

The Output-gap According to the FRB ... St. Louis

The FRB of St. Louis publishes "oficial" US data for Real GDP and Potential GDP.

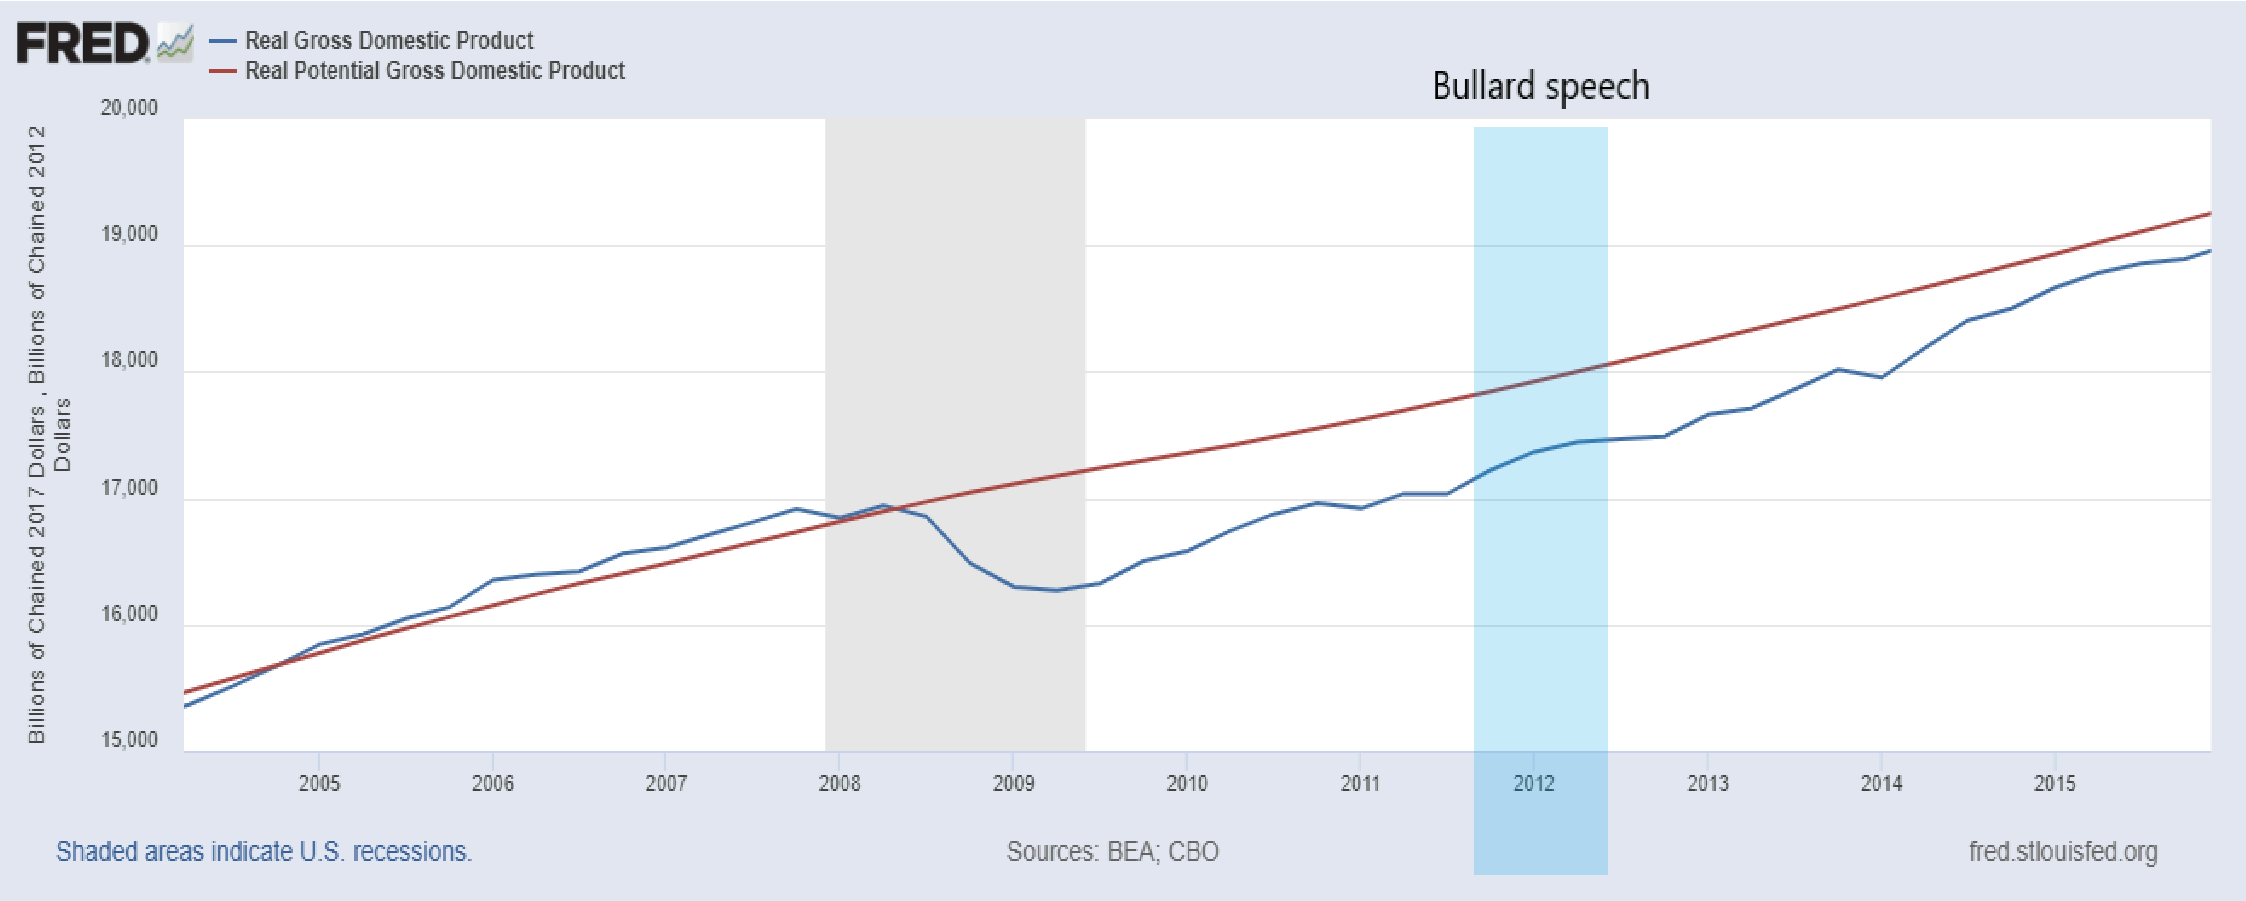

The Natural Rate of Unemployment (NRU)

No, Covid-19 did not raise the NRU; no, an increase in NRU did not anticipate Covid-19!

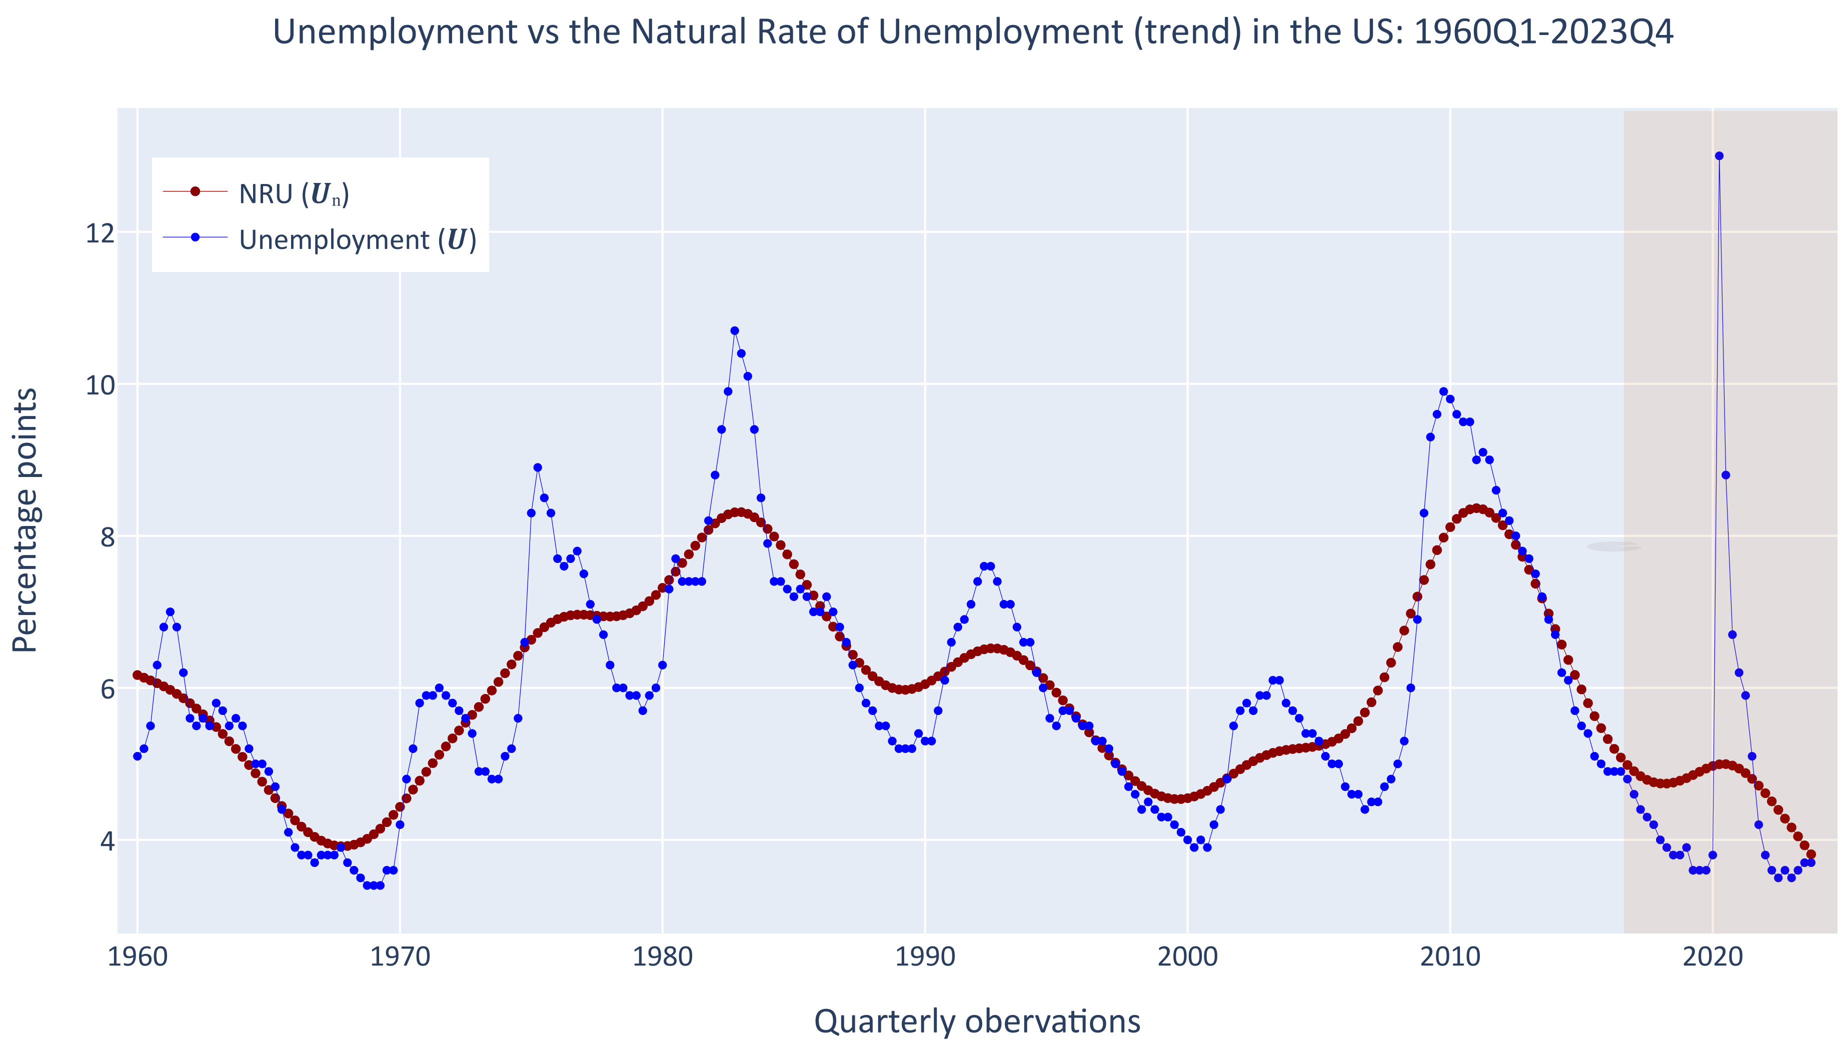

4. Readings

Point 1

- For this point, there is no compulsory reading.
- However, Dirk Krueger (2007). "Quantitative Macroeconomics: An Introduction" (Chapter 2), manuscript, Department of Economics University of Pennsylvania, is well suited for the material covered here.
- This text is a small one (12 pages), easy to read, and beneficial for studying the stylized facts of business cycles, mainly to understand how the Hodrick-Prescott filter is calculated. However, notice that, as mentioned, it is not compulsory reading.

Point 2

- For this point, there is no compulsory reading. However, any modern textbook on time series will cover this subject.
- At an introductory level, see sections 11.8 and 11.9 of the textbook: Diebold, F. X. (1998). *Elements of forecasting*. South-Western College Pub, Cincinnati.
- At a more advanced level, see, e.g., section 2.3.2 of the textbook: Lütkepohl, H. (2007). *New introduction to multiple time series analysis* (2nd ed.), Springer, Berlin.

Point 3

- No textbook covers the topics/controversies mentioned in this section.
- This coursework intends to provide a framework for a better understanding of these controversies at the end of the course.
- All we have to handle is:
  - A little bit of mathematics
  - A little bit of computation
  - A little bit of macroeconomics