Time-Series Components

- Recall that the optimal point forecast of a series y_{t+h} is its conditional mean

$$\mu_t = \mathrm{E}\left(y_{t+h} \mid \Omega_t\right)$$

- It is useful to decompose this mean into components

$$\mu_t = T_t + S_t + C_t$$

  – T_t = Trend
  – S_t = Seasonal
  – C_t = Cycle

Components

- Trend
  - Very long term (decades)
  - Smooth
- Seasonal
  - Patterns which repeat annually
  - May be constant or variable
- Cycle
  - Business cycle
  - Correlation over 2-7 years
- It is useful to consider the components separately
- We start with the Trend

Trend Forecasting

- A pure trend model has no seasonal or cycle

$$\mu_t = T_t$$

- In a pure trend model, the optimal point forecast for y_{t+h} is $\mu_t = T_t$.
- An actual forecast is an estimate of T_t.

Modeling Trend

- Most trend models are very simple
- Simplest possible trend is a constant

$$T_t = \beta_0$$

- This might seem overly simple, but is appropriate for **stationary** time-series
  - A series not growing or changing over time
  - Many series reported as percentage changes

U.S. Personal Consumption (Quarterly) Monthly Percentage Change

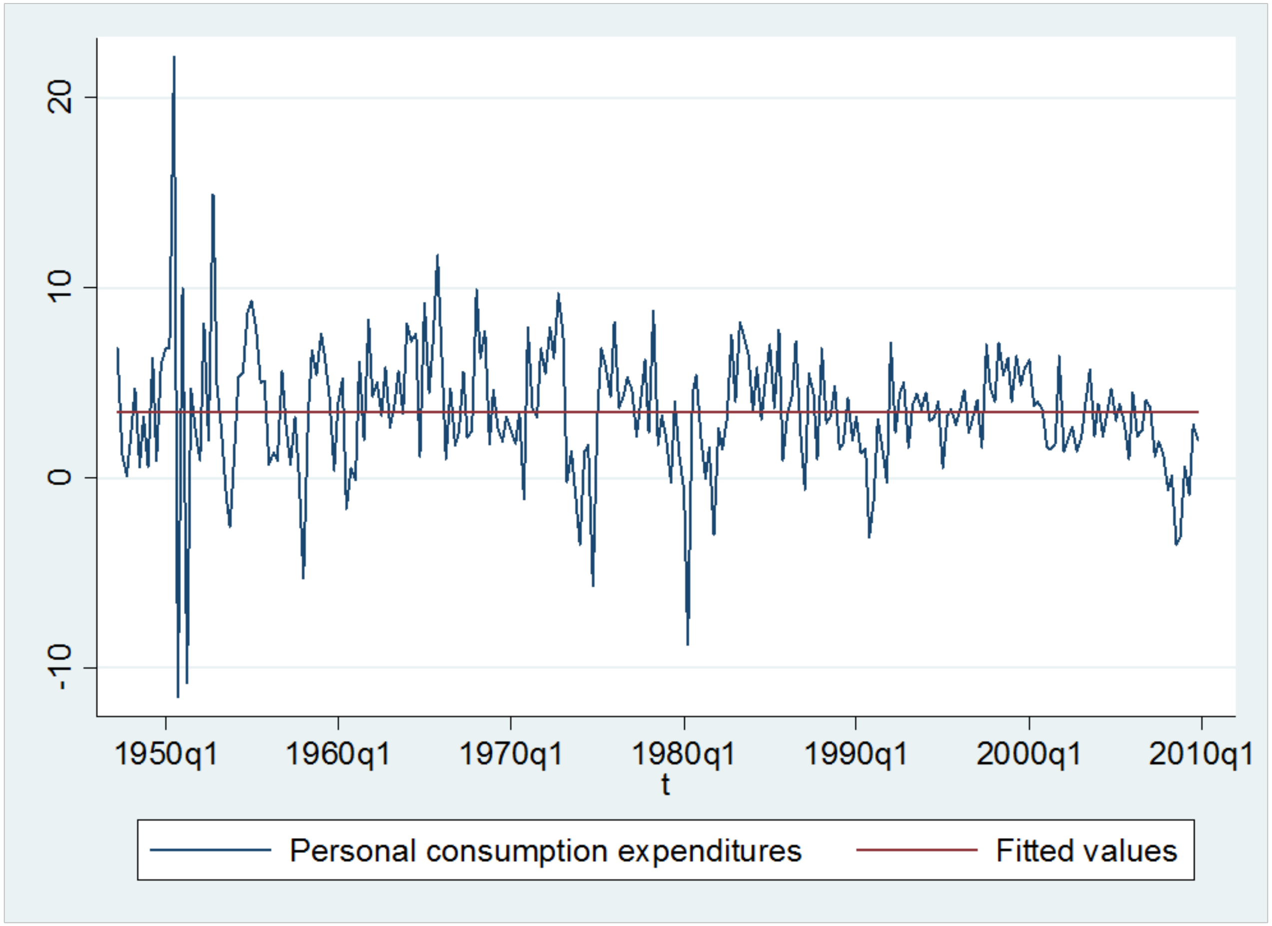

Estimation

- If $E(y_{t+h} \mid \Omega_t) = \mu_t = T_t = \beta_0$ then the optimal forecast is the mean $\beta_0 = E(y_{t+h})$
- The estimate of β_0 is the sample mean

$$b_0 = \frac{1}{T}\sum_{t=1}^{T} y_{t+h}$$

- This is the estimate of the optimal point forecast when $\mu_t = \beta_0$
- b_0 is also the least-squares estimate in an intercept-only model

Estimation

- In STATA, use the **regress** command
- See *STATA Handout* on website
- Sample mean is estimated "constant"

```
. use gdp

. regress consumption

      Source |       SS       df       MS              Number of obs =     251
-------------+------------------------------           F(  0,   250) =    0.00
       Model |           0     0           .           Prob > F      =       .
    Residual |   3017.6648   250   12.0706592           R-squared     =  0.0000
-------------+------------------------------           Adj R-squared =  0.0000
       Total |   3017.6648   250   12.0706592           Root MSE      =  3.4743

------------------------------------------------------------------------------
 consumption |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   3.472112    .219295    15.83   0.000     3.040211    3.904013
------------------------------------------------------------------------------
```

Fitted Values

- Fitted values are the sample mean

$$\hat{y}_t = \hat{\mu}_t = b_0$$

- In STATA use the **predict** command

```
. predict yp
(option xb assumed; fitted values)
```

- This creates a variable “yp” of fitted values

Plot actual against fitted

```
. tsline consumption yp
```

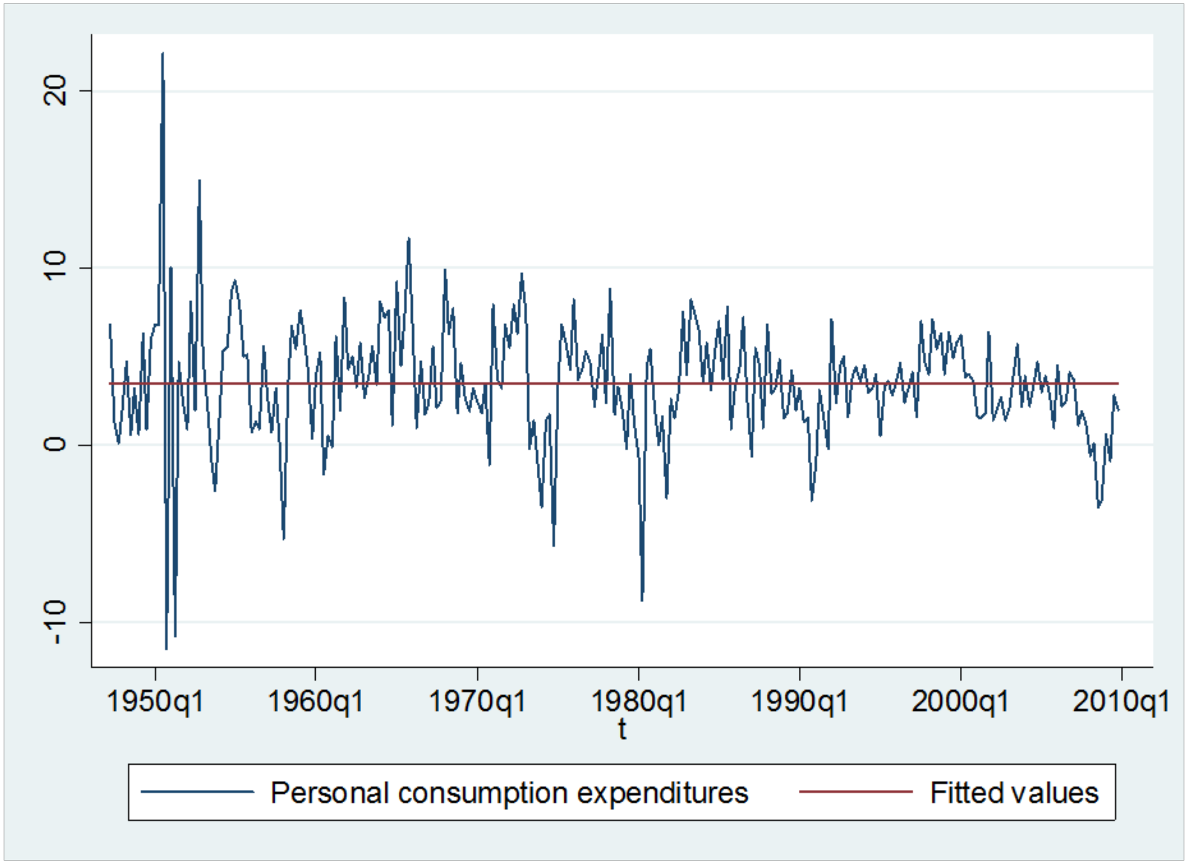

Out-of-Sample

- Point forecasts are the sample mean

$$\hat{y}_{T+h} = b_0$$

- In STATA, use **tsappend** to expand sample, and **predict** to generate point forecasts.

```
. tsappend, add(12)

. predict p if t>tq(2009q4)
(option xb assumed; fitted values)
(251 missing values generated)

. tsline consumption yp p if t>tq(2000q4)
```

Out-of-Sample

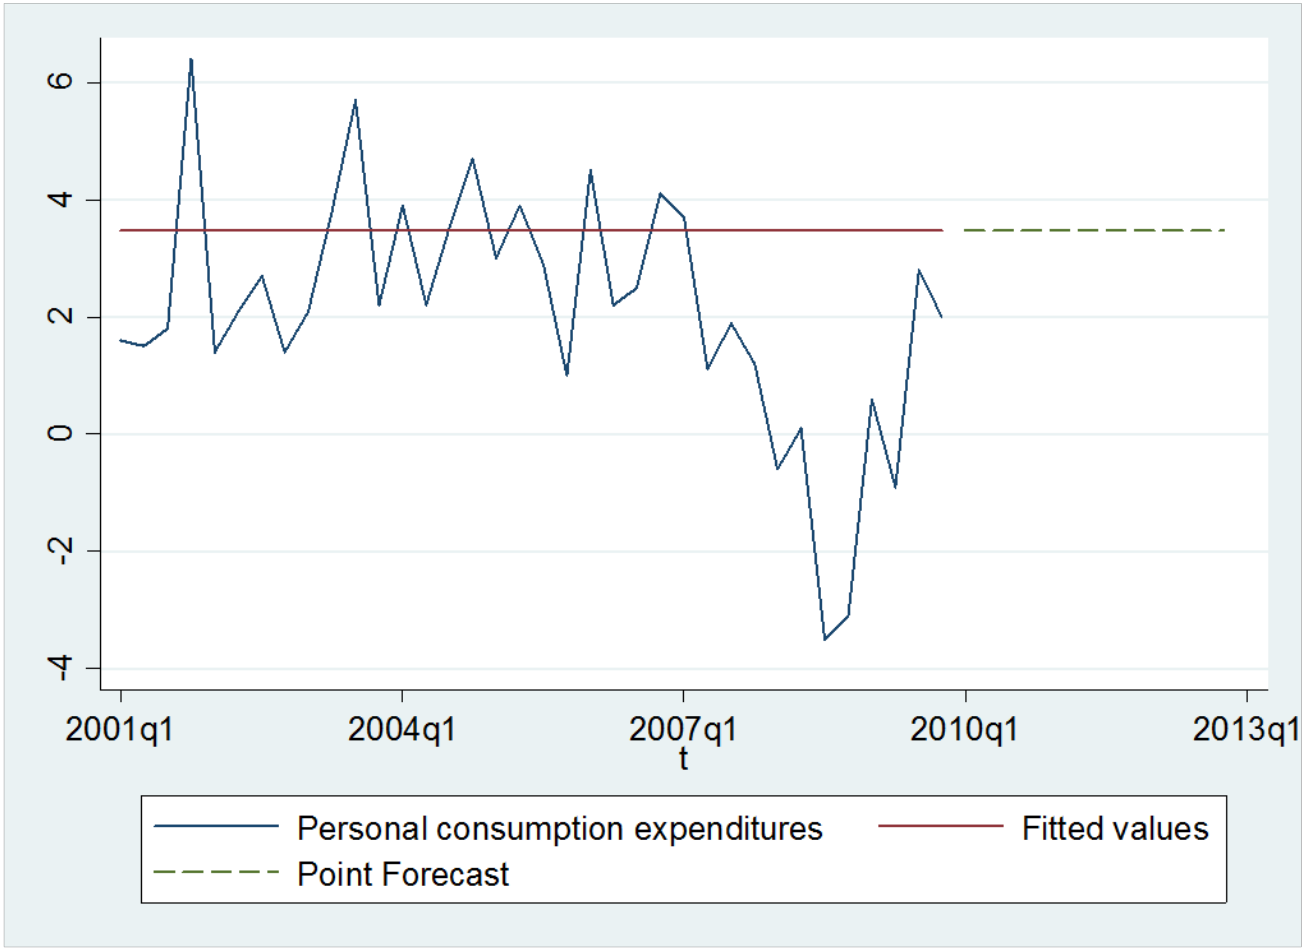

Forecast Errors

- The forecast error e_t is the difference between the realized value and the conditional mean.

$$e_t = y_{t+h} - \mu_t$$

or equivalently

$$y_{t+h} = \mu_t + e_t$$

- We call e_t the forecast error.

Residuals

- The residuals are the in-sample fitted errors.
- The difference between the realized value and the in-sample forecast.

$$\begin{aligned}\hat{e}_t &= y_{t+h} - \hat{\mu}_t \\ &= y_{t+h} - b_0\end{aligned}$$

- In general, it is useful to plot the residuals against time, to see if any time series pattern remains.

Calculate and Plot Residuals

```
. predict e, residuals
(12 missing values generated)

. tsline e
```

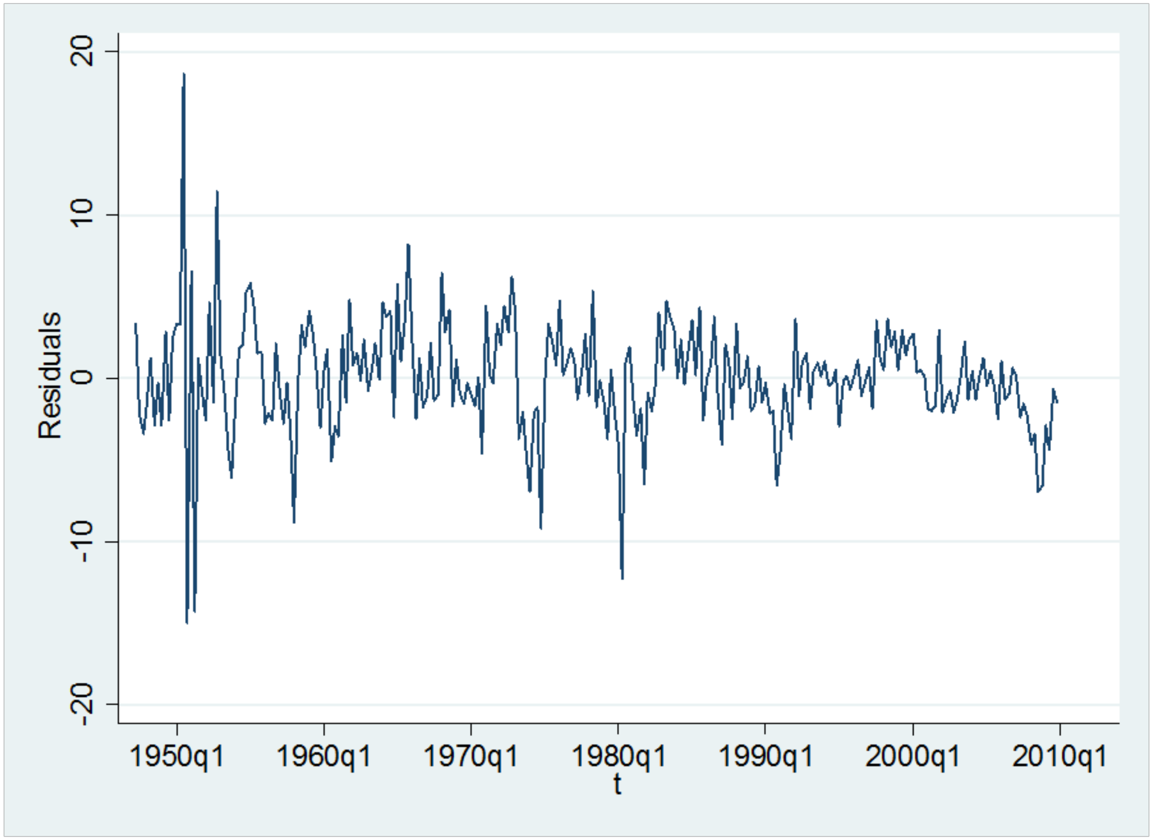

Estimation Uncertainty

- The sample mean

$$b_0 = \frac{1}{T}\sum_{t=1}^{T} y_{t+h}$$

is an estimate of $\beta_0 = E(y_{t+h})$

- The estimation error is

$$\begin{aligned} b_0 - \beta_0 &= \frac{1}{T}\sum_{t=1}^{T} y_{t+h} - \beta_0 \\ &= \frac{1}{T}\sum_{t=1}^{T}\left(y_{t+h} - \beta_0\right) \\ &= \frac{1}{T}\sum_{t=1}^{T} e_t \end{aligned}$$

Estimation Variance

- Under classical conditions,

$$\text{var}(b_0) = \frac{\sigma^2}{T}$$

where $\sigma^2 = \text{var}(e_t)$

- The standard error for b_0 is an estimate of the standard deviation

$$sd(b_0) = \sqrt{\frac{\hat{\sigma}^2}{T}}$$

Forecast Variance

- When the sample mean b_0 is used as the forecast for y_{T+h} then the prediction error is

$$y_{T+h} - b_0 = e_{T+h} + \beta_0 - b_0$$

  which is the sum of the forecast error e_{T+h} and the estimation uncertainty β_0-b_0.
- The forecast variance is

$$\begin{aligned}\operatorname{var}(y_{T+h} - b_0) &= \operatorname{var}(e_{T+h}) + \operatorname{var}(\beta_0 - b_0) \\ &= \sigma^2 + \frac{\sigma^2}{T} \\ &= \left(1 + \frac{1}{T}\right)\sigma^2\end{aligned}$$

Standard Deviation of Forecast

- The standard deviation of the forecast is the estimate

$$s_{T+h} = \sqrt{\left(1 + \frac{1}{T}\right)\hat{\sigma}^2}$$

- This is slightly larger than the regression standard deviation $\hat{\sigma}$

Normal Forecast Intervals

- Let $\hat{y}_{T+h}$ be a forecast for y_{T+h}
- The prediction error is $y_{T+h} - \hat{y}_{T+h}$
- Let s_{T+h} be the st. deviation of the forecast
- If the prediction errors are normally distributed, the (1-α)% forecast interval endpoints are

$$L_{T+h} = \hat{y}_{T+h} + s_{T+h} z_{\alpha/2}$$
$$U_{T+h} = \hat{y}_{T+h} + s_{T+h} z_{1-\alpha/2}$$

where $z_{\alpha/2}$ and $z_{1-\alpha/2}$ are the α/2 and 1-α /2 quantiles of the normal distribution

- *e.g.* $\hat{y}_{T+h} \pm 1.64\, s_{T+h}$ for a 90% interval

```
. predict s, stdf

. generate yp1=yp-1.645*s

. generate yp2=yp+1.645*s

. tsline consumption yp yp1 yp2
```

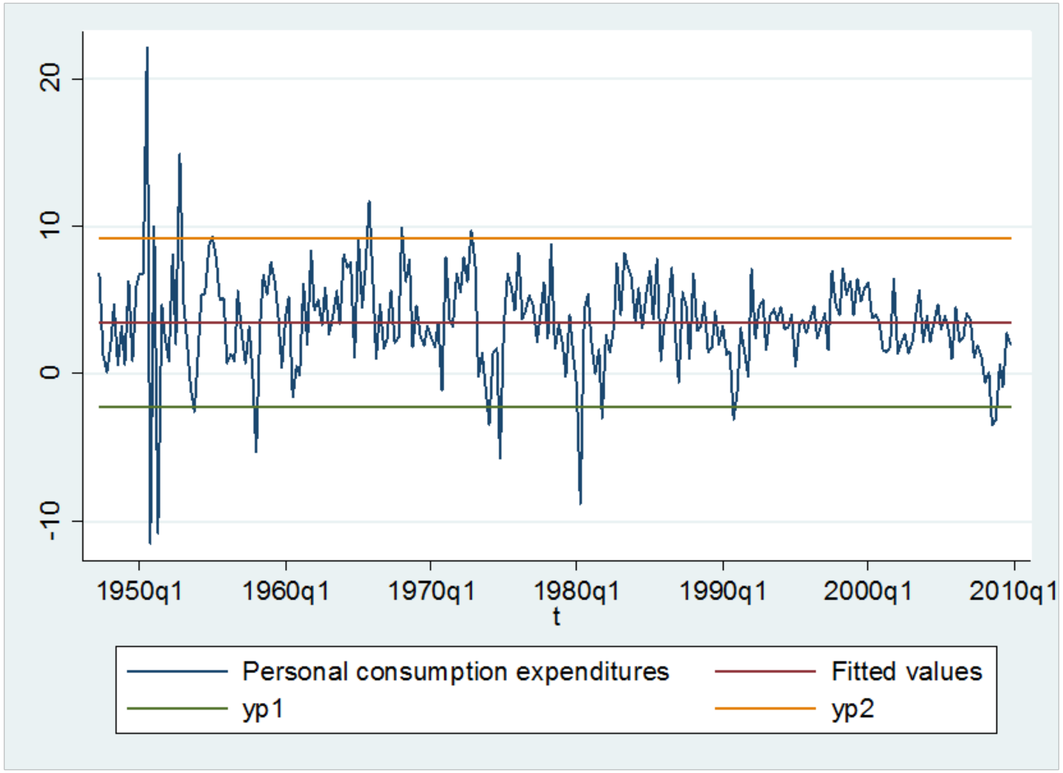

Deficiency of Normal Intervals

- The normal forecast interval is based on the assumption that the prediction errors are normally distributed.
- This requires that the conditional distribution of y_{T+h} be normal, which is rarely valid.
- Instead, we can compute forecast intervals based on the empirical distribution of the forecast residuals.

Empirical Forecast Intervals

- Let $\hat{y}_{t+h}$ be fitted values for y_{t+h} with residuals

$$\hat{e}_t = y_{t+h} - \hat{y}_{t+h}$$

- Let $q_{\alpha/2}$ and $q_{1-\alpha/2}$ be the α/2 and 1-α /2 quantiles of the residuals.
- The (1-α)% forecast interval endpoints are

$$L_{T+h} = \hat{y}_{T+h} + q_{\alpha/2}$$
$$U_{T+h} = \hat{y}_{T+h} + q_{1-\alpha/2}$$

Empirical Forecast Intervals

- The basic method to obtain forecast intervals is the same for any regression model

$$y_{t+h} = \mu_t + e_t$$

- The (1-α)% forecast interval endpoints are

$$L_{T+h} = \mu_T + q_{\alpha/2}$$
$$U_{T+h} = \mu_t + q_{1-\alpha/2}$$

where $q_{\alpha/2}$ and $q_{1\text{-}\alpha/2}$ are the α/2 and 1-α /2 quantiles of the distribution of e_t .

Quantiles

- The *x*'th quantile of a set of numbers is the value q_x such that *x*% are smaller than q_x and (1-*x*)% are larger than q_x.
- You can find q_x by sorting the data.
- In STATA, use the **qreg** command
  - (for quantile regression)

```
. qreg e, quantile(.05)
Iteration  1:  WLS sum of weighted deviations =  500.77838

Iteration  1: sum of abs. weighted deviations =     517.45
Iteration  2: sum of abs. weighted deviations =     213.47

.05 Quantile regression                             Number of obs =       251
  Raw sum of deviations   213.91 (about -5.0721116)
  Min sum of deviations   213.47                    Pseudo R2     =    0.0021

------------------------------------------------------------------------------
           e |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |  -4.672112   1.468371    -3.18   0.002    -7.564066   -1.780157
------------------------------------------------------------------------------

. predict q1
(option xb assumed; fitted values)

. generate yp1=yp+q1

. qreg e, quantile(.95)
Iteration  1:  WLS sum of weighted deviations =  502.46952

Iteration  1: sum of abs. weighted deviations =  507.45001
Iteration  2: sum of abs. weighted deviations =     183.47

.95 Quantile regression                             Number of obs =       251
  Raw sum of deviations   183.47 (about 4.7278881)
  Min sum of deviations   183.47                    Pseudo R2     =   -0.0000

------------------------------------------------------------------------------
           e |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   4.727888    1.77628     2.66   0.008     1.229507    8.226269
------------------------------------------------------------------------------

. predict q2
(option xb assumed; fitted values)

. generate yp2=yp+q2

. tsline consumption yp yp1 yp2
```

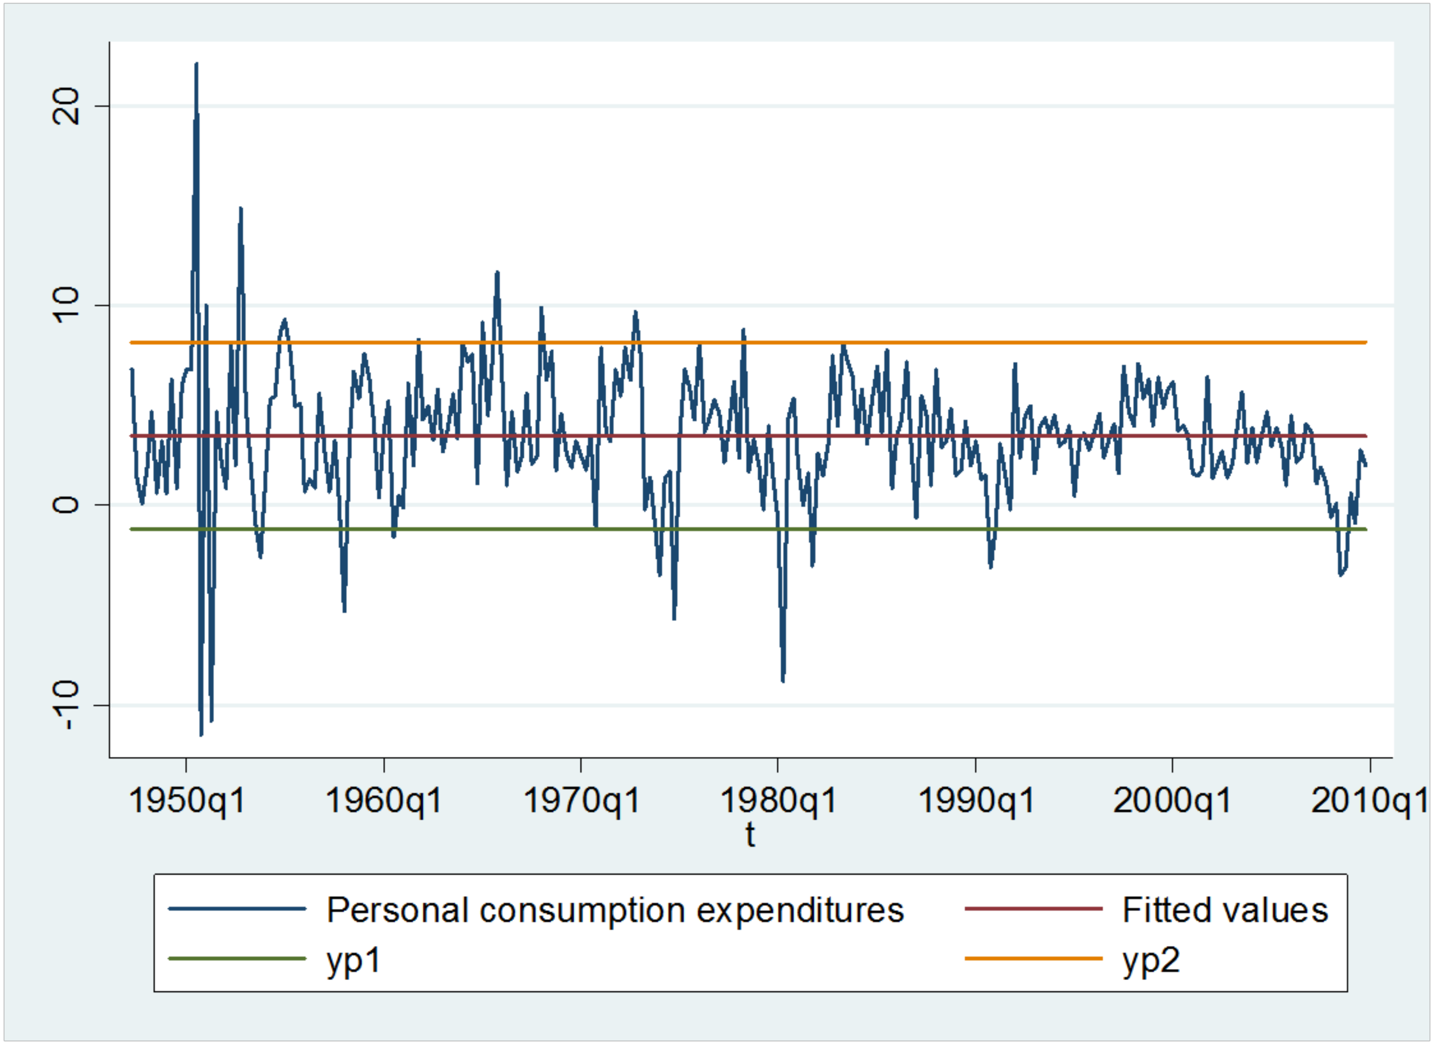
20
10
0
-10
1950q1
1960q1
1970q1
1980q1
1990q1
2000q1
2010q1
t
Personal consumption expenditures
Fitted values
yp1
yp2

Out-of-Sample

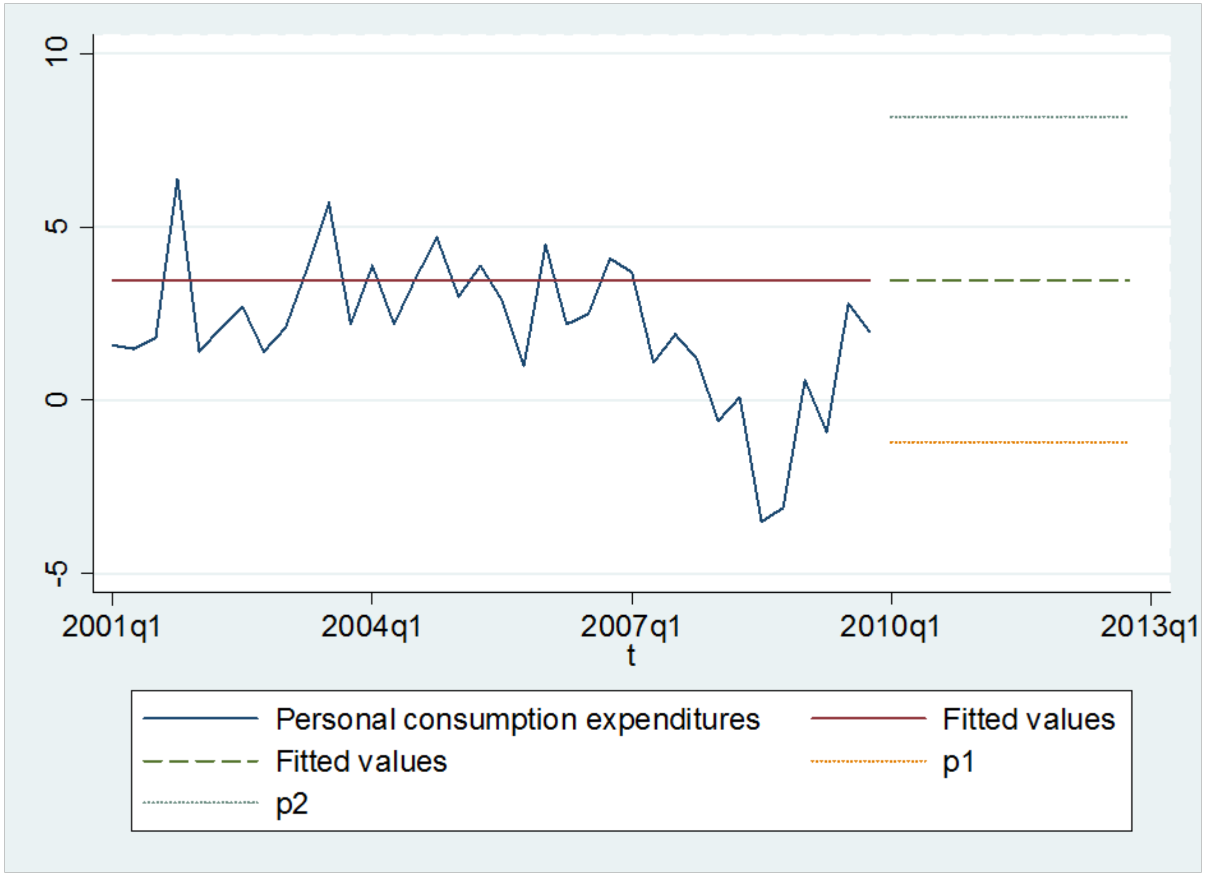

Mean Shifts

- Sometimes the mean of a series changes over time
- It can drift slowly, or change quickly
  - Possibly due to a policy change
- In this case, forecasting based on a constant mean model can be misleading

State and Local Government Spending Percentage Growth Rate (Quarterly)

- Average for 1947-2009: 3.6%
- But this has not been the typical rate in recent years.

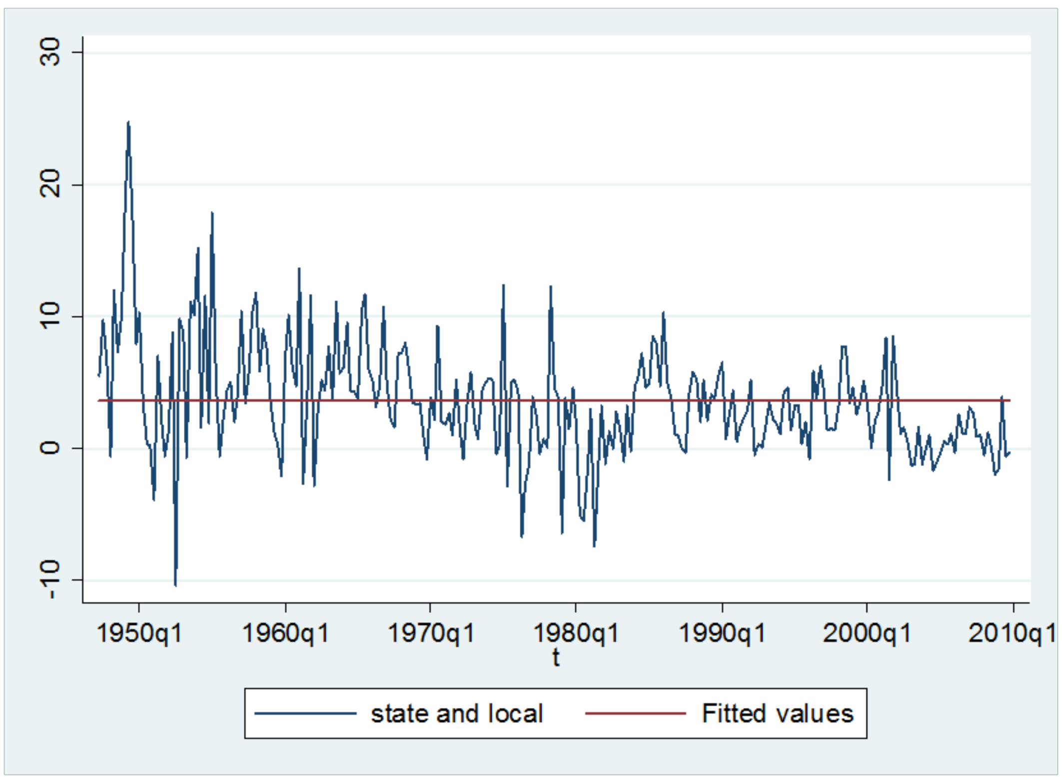

Alternatives

- Subsample estimation
  - Estimate the mean on subsamples
  - Forecasts are based on the most recent
- Dummy Variable formulation

$$\mu_t = \beta_0 + \beta_1 d_t$$
$$d_t = 1(t \geq \tau)$$

- τ is the breakdate
  - The date when the mean shifts
  - The coefficient β_0 is the mean before $t=\tau$
  - The coefficient β_1 is the shift at $t=\tau$
  - The sum $\beta_0+\beta_1$ is the mean after $t=\tau$

Forecast

- Linear Regression y_{t+h} on d_t
- Example
  - State and Local Government Percentage Growth
  - Mean breaks in 1970q1 and 2002q1

```
. regress state d1 d2

      Source |       SS       df       MS              Number of obs =     251
-------------+------------------------------           F(  2,   248) =   26.08
       Model |  853.758681     2  426.879341           Prob > F      =  0.0000
    Residual |  4058.72436   248   16.365824           R-squared     =  0.1738
-------------+------------------------------           Adj R-squared =  0.1671
       Total |  4912.48304   250  19.6499322           Root MSE      =  4.0455

------------------------------------------------------------------------------
       state |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
          d1 |  -3.125867   .5547091    -5.64   0.000    -4.218409   -2.033326
          d2 |  -2.160156   .7995561    -2.70   0.007    -3.734943     -.58537
       _cons |   5.851648   .4240804    13.80   0.000      5.01639    6.686907
------------------------------------------------------------------------------
```

Fitted

- Out-of-sample forecast falls from 3.6% to 0.6%!

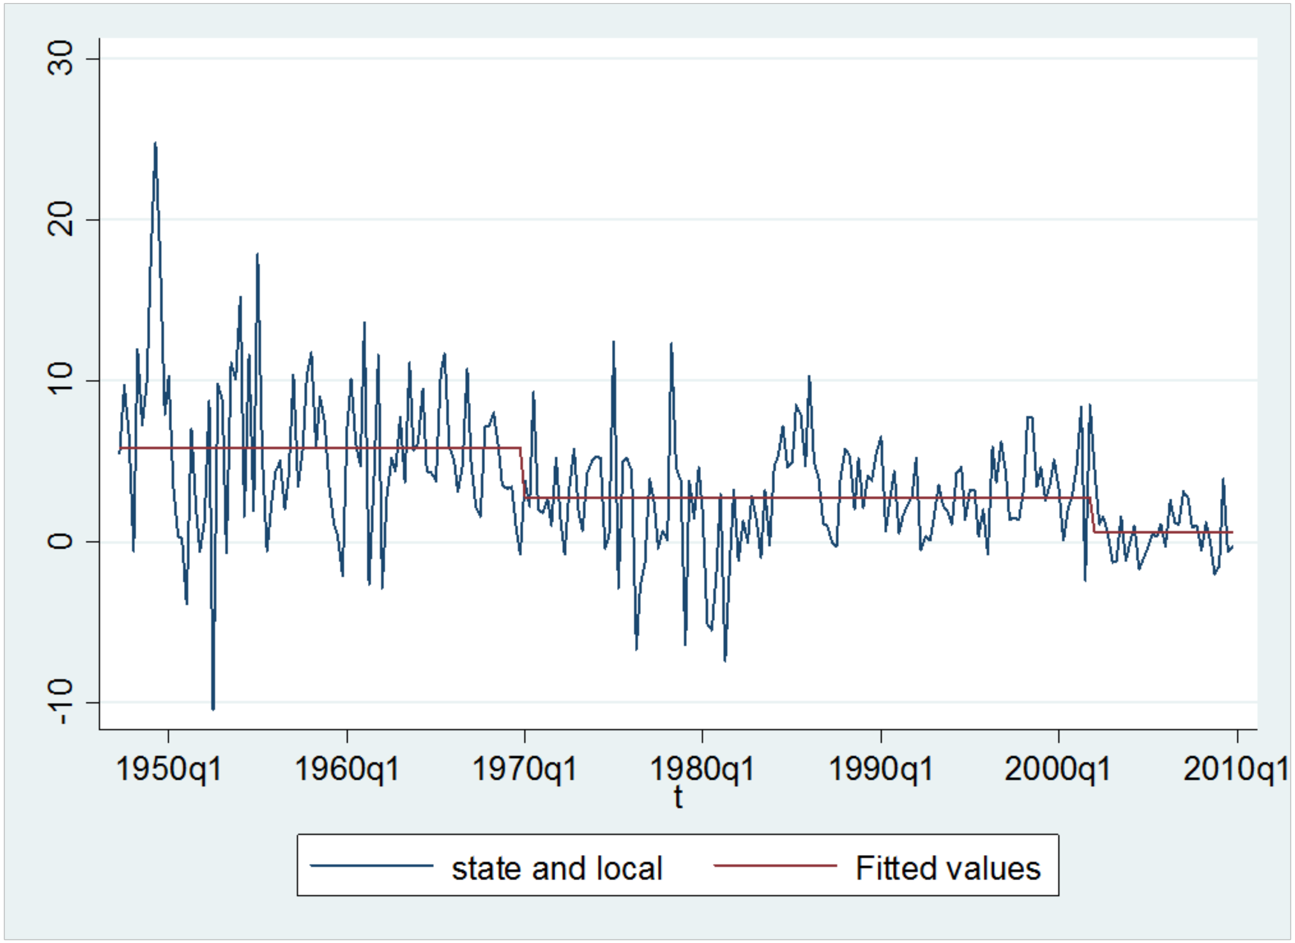

Should you use Mean Shifts?

- Only after great hesitation and consideration.
- Should use shifts and breaks reluctantly and with care.
- Do you have a model or explanation?
- What is the forecasting power of a mean shift?
  - If they have happened in the past, will there be more in the future?
- Yet, if there has been an obvious shift, a simple constant mean model will forecast terribly.

How to Select Breakdates

- Judgmental
  - Dates of known policy shifts
  - Important events
  - Economic crises
- Informal data-based
  - Visual inspection
- Formal data-based
  - Estimate regression for many possible breakdates
  - Select one which minimizes sum of squared error
  - This is the least-squares breakdate estimator