The CARD Forecasting Method

Kevin Sheppard

Advanced Financial Econometrics: Forecasting

May 2020

The Card Forecasting Method

- **C**alibration
  - Forecast by fitting a model to other forecasts
- **A**verage
  - Average forecasts to use in the calibration step
- **R**ho
  - A forecasting method based on autoregressive models
- **D**elta
  - A forecast method using a robust estimate of the trend

References

1. Doornik, J. A., Castle, J. L., & Hendry, D. F. (2020). Card forecasts for M4. *International Journal of Forecasting*, 36(1), 129-134.

Transformation

- If $I_{\ln} = \min\left(X_1, X_2, \ldots, X_T\right) \geq 1$, then the data should be transformed using the natural log.
- Define $Y_t = \ln\left(X_t\right) I_{\ln} + X_t \left(1 - I_{\ln}\right)$to be the transformed series.

Integration $\left(I_\rho\right)$

- Define $I_\rho = I_{[\mathrm{V}[\Delta Y_t] < 1.2\mathrm{V}[Y_t]]}$
  - This value should be 1 for unit-root or highly persistent series. If stationary and short memory, then the overdifferenced series should have a larger variance.
- Define $Z_t = I_\rho \Delta Y_t + \left(1 - I_\rho\right) Y_t$ as the stationarity transformed series

Preliminaries

Seasonality (I_A)

- Define the seasonality indicator

$$I_A = I_{[A > C_{A,10\%}]}$$

  - $C_{A,10\%}$ is the 10% critical value of a $F_{m-1,m(\tau-1)}$
- Transform data using seasonal alignment

$$\begin{bmatrix} Z_{T-m\tau+1} & Z_{T-m\tau+2} & \cdots & Z_{T-m\tau+m-1} & Z_{T-m(\tau+1)} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ Z_{T-2m+1} & Z_{T-2m+2} & \cdots & Z_{T-m-1} & Z_{T-m} \\ Z_{T-m+1} & Z_{T-m+2} & \cdots & Z_{T-1} & Z_T \end{bmatrix}$$

- $\tau = \lfloor T/m \rfloor$ is the number of complete seasonal cycles
- τ rows and m columns
  - $m-1$ observation may be missing

Preliminaries

- The test statistic is defined using an ANOVA

$$A = \frac{m-1}{m(\tau-1)} \frac{\mathrm{V}\left[\bar{Z}_j\right]}{\mathrm{V}\left[Z_{ij}\right] - \mathrm{V}\left[\bar{Z}_j\right]}$$

  - $\bar{Z}_j = {}^1/_\tau \sum_{i=1}^{\tau} Z_{ij}$ are the column averages
  - Column averages should be the same absent seasonality
    - A would be small
  - m is always 12 for monthly data
    - Paper considers multiple seasonalities for high-frequency data (daily, hourly)
- Test is nearly equivalent to regression-based test in SARIMA slides

Seasonal Autoregression (I_R)

- Define the seasonal autoregressive indicator as

$$I_R = I_{\left[m>1 \cap T \geq 3m+1 \cap R > C_{R,10\%}\right]}$$

  - $\cap$ is mathematical *and*
- Test statistic is

$$R = T \frac{\hat{\rho}_m^2}{1 + 2\sum_{i=1}^{m-1} \hat{\rho}_i^2}$$

- $\hat{\rho}_j$ is the j$^{\text{th}}$ autocorrelation of Z_t
- $C_{R,10\%}$ is the 10% critical value from a χ_1^2

Monthly Data If $I_R = 0$ then repeat for $11, \ldots, 2$ using a 1% critical value. If found, set $I_R = 1$ and use the first rejection as the seasonal period m.

Conclusion

CARD makes use of multiple preliminary steps:

- Log positive time series
- Test for a unit root and transform
- Test for seasonal level shifts
- Test for a seasonal autocorrelation

CARD: The Delta Method

Kevin Sheppard

Advanced Financial Econometrics: Forecasting

May 2020

Delta

Nonseasonal Unit Root ($I_\rho = 1$ and $I_A = 0$)

- The simpler of the methods described in the paper
- Let $Z_{(i)}$ be the i value from the sorted Z where $|Z_1| \leq |Z_2| \leq \ldots \leq |Z_n|$
  - $n = T - 1$: one observation lost due to differencing
- Compute

$$d_1 = \frac{1}{n-1} \sum_{i=1}^{n-1} Z_{(i)}$$

$$d_2 = \frac{1}{n-3} \sum_{i=1}^{n-3} Z_{(i)}$$

$$d_r = \frac{1}{6} \sum_{t=T-6}^{T} Z_t$$

- If $T \leq 6$, then set $d_2 = d_1$ and $d_r = d_1$

Delta Forecasts

Nonseasonal Unit Root ($I_\rho = 1$ and $I_A = 0$)

- d_1 is the average change excluding the largest in magnitude change
- d_2 is the average change excluding the **3** largest in magnitude
- d_r is the average of recent changes
- Compute $d_m = \bar{Z}$
- Define $\operatorname{sgnmin}(a, b) = aI_{[ab>0]}$ if $|a| < |b|$ else $bI_{[ab>0]}$
  - The indicator sets the value to 0 if the signs differ **[Typo]**
- If $T > 2m + 1$, define $d_S = m^{-1}\overline{\Delta_m Y_t}$ **[Typo]** and update $d_m = \operatorname{sgnmin}(d_m, d_S)$
- Compute $d_r^* = \operatorname{sgnmin}(d_r, d_m)$
- Forecasts are then

$$\hat{Y}_{T+1} = Y_T + \operatorname{sgnmin}(d_r^*, d_1)$$
$$\hat{Y}_{T+h} = Y_{T+h-1} + \operatorname{sgnmin}(d_r^*, d_2)\,, \; h = 2, \ldots, H$$

Delta

Seasonal Unit Root ($I_\rho = 1$ and $I_A = 1$)

- Use annual means to compute d_1, d_2 and d_r
- The annual mean is a row sum from the seasonal table

$$\bar{Z}_i = m^{-1} \sum_{j=1}^{m} Z_{ij}$$

- Compute d_1 and d_2 using the **ordered values** $\bar{Z}_{(i)}$ $\left(\left|\bar{Z}_1\right| \leq \left|\bar{Z}_{i+1}\right|, i = 1, \ldots, n\right)$
- Compute

$$d_r = \frac{1}{6} \sum_{i=\tau-5}^{\tau} \bar{Z}_i$$

  - If $\tau \leq 6$ then set $d_2 = d_1$ and $d_r = d_1$

Delta Forecasts

Seasonal Unit Root ($I_\rho = 1$ and $I_A = 1$)

- Forecast the series as

$$\hat{Y}_{T+1} = Y_T + \operatorname{sgnmin}\left(d_r^*, d_1\right) + \hat{s}_1$$
$$\hat{Y}_{T+h} = Y_{T+h-1} + \operatorname{sgnmin}\left(d_r^*, d_2\right) + \hat{s}_{h-m\lfloor (h-1)/m \rfloor}, \quad h = 2, \ldots, H$$

- Seasonal estimation makes use of the fact that the final observation is in the final column
- The seasonal effect estimates are estimated using

$$\hat{s}_j^\star = \mathrm{MA}_{3\times 5}\left(Z_{ij}\right)$$

Centered Moving Averages

- $\text{MA}_{3\times 5}$ is read as a centered MA(3) of an MA(5)
- MA(5) has weights ${}^1\!/_5$ over final 5 observations
  - MA(3) then averages last the MA(5) values with weights ${}^1\!/_3$
  - The final three terms from the MA(5) are

$$\begin{bmatrix} \tau+3 & \tau+2 & \tau+1 & \tau & \tau-1 & \tau-2 & \tau-3 \\ \hline {}^1\!/_5 & {}^1\!/_5 & {}^1\!/_5 & {}^1\!/_5 & {}^1\!/_5 & 0 & 0 \\ 0 & {}^1\!/_5 & {}^1\!/_5 & {}^1\!/_5 & {}^1\!/_5 & {}^1\!/_5 & 0 \\ 0 & 0 & {}^1\!/_5 & {}^1\!/_5 & {}^1\!/_5 & {}^1\!/_5 & {}^1\!/_5 \end{bmatrix}$$

  - Combined is trapezoidal weighted average with weights

$$\left[{}^1\!/_3{}^1\!/_5, {}^1\!/_3\left({}^1\!/_5+{}^1\!/_5\right), {}^1\!/_3\left({}^1\!/_5+{}^1\!/_5+{}^1\!/_5\right), \ldots\right] = \left[{}^1\!/_{15}, {}^2\!/_{15}, {}^1\!/_5, {}^1\!/_5, {}^1\!/_5, {}^2\!/_{15}, {}^1\!/_{15}\right]$$

  - Do not have data for rows $\tau+1$, $\tau+2$ and $\tau+3$
  - Weights for $\tau+1$, $\tau+2$ and $\tau+3$ are assigned to τ
    - Random walk forecasting model
  - Effective weights are

τ	$\tau-1$	$\tau-2$	$\tau-3$
${}^9\!/_{15}$	${}^3\!/_{15}$	${}^2\!/_{15}$	${}^1\!/_{15}$

Seasonal Dummy Estimation

Seasonal Unit Root ($I_\rho = 1$ and $I_A = 1$)

- The seasonal forecasting terms are

$$\hat{s}_j = s_j^\star - m^{-1} \sum_{j=1}^{m} s_j^\star$$

- These are the deviation from the average seasonal effect
  - Removes the level from the forecast value

Delta

Non-seasonal Stationary ($I_\rho = 0$ and $I_A = 0$)

- Stationary model uses different methods
- Define mean of r most recent observations as

$$\mu(r) = \frac{1}{\min(T, r)} \sum_{t=\max(T-r,1)}^{T} Z_t$$

- Define $\tilde{m} = \max(2, m)$, then the forecasts are

$$\hat{Y}_{T+1} = \mu(\tilde{m})$$
$$\hat{Y}_{T+h} = {}^1\!/_2 \left(\mu(\tilde{m}) + \mu(6\tilde{m})\right)$$

Note $m = 12$ for monthly data

Delta

Seasonal Stationary ($I_\rho = 0$ and $I_A = 1$)

- Compute $\hat{s}_j$, $j = 1, \ldots, m$ using same method as $I_\rho = 1$ only with a MA$_{7\times 5}$

τ	$\tau - 1$	$\tau - 2$	$\tau - 3$	$\tau - 4$	$\tau - 5$
$^{20}/_{35}$	$^{5}/_{35}$	$^{4}/_{35}$	$^{3}/_{35}$	$^{2}/_{35}$	$^{1}/_{35}$

- Use annual averages in $\mu_A(r)$

$$\bar{Z}_i = m^{-1} \sum_{j=1}^{m} Z_{ij}$$

$$\mu_A = \frac{1}{\min(\tau, r)} \sum_{i=\max(\tau - r, 1)}^{\tau} \bar{Z}_i$$

- The forecasts are

$$\hat{Y}_{T+1} = \mu_A(1) + \hat{s}_1$$

$$\hat{Y}_{T+h} = {}^1/_2\left(\mu_A(1) + \mu_A(6)\right) + \hat{s}_{h - m\lfloor (h-1)/m \rfloor}$$

Conclusion

- The Delta method is a time trend model
- No ARMA-like dynamics are permitted
- Trend is estimated in a highly robust method that shrinks towards zero
- Different procedures for unit root and stationary time series

CARD: The Rho Method

Kevin Sheppard

Advanced Financial Econometrics: Forecasting

May 2020

- The rho method is based around an AR(1)
- The model begins with the transformed series Y_t

Initialization

Set $I_{AR} = I_\rho$ and $I_{UR} = I_{\text{tr}} = 0$. Estimate the model

$$
\begin{aligned}
Y_t =& \mu + \left(\underbrace{\rho Y_{t-1}}_{\text{AR}} + \underbrace{\rho_m Y_{t-m} I_R}_{\text{SAR}} \right) I_{AR} \\
&+ \underbrace{\gamma \lfloor t/m \rfloor I_{\text{tr}}}_{\text{Trend}} + \underbrace{\sum_{j=1}^{m-1} \delta_j I_{S_j} I_A}_{\text{Seasonal}} + \epsilon_t
\end{aligned}
$$

where $I_{S_j} = 1$ if $t - m \lfloor (t-1)/m \rfloor = j$.

Unit Root Retests

If $I_{AR} = 1$ and:

- $\hat{\rho} > 1/2$ and $\hat{\rho} + 2\text{s.e.}\,(\hat{\rho}) > 0.9$ impose $\rho = 1$, set $I_{UR} = 1$ and re-estimate the model using ΔY_t on the LHS.
- $\hat{\rho} < 0$ set $I_{AR} = 0$ and re-estimate the model.

Trend Test

If $I_{UR} = 0$ and $T - k > 10$ where k is the number of parameters in the model, test whether the ***cumulated*** residuals have mean 0 using a t-test with a size of 1%. If the null is rejected, then set $I_{tr} = 1$, and re-estimate the model.

Trend Check

If $I_{tr} = 1$ and $\hat{\rho} < -1/2$, set $I_{tr} = 0$ and re-estimate.

Stationary Forecast

If $I_{UR} = 0$, forecast from the estimated model.

Trend Dampening

If $I_{UR} = 1$, dampen the trend using

$$\tilde{\mu} = \begin{cases} \max\left(0, \hat{\mu} - s\right) & \hat{\mu} > 0 \\ \min\left(0, \hat{\mu} + s\right) & \hat{\mu} < 0 \end{cases}$$

where $s = 1.645\hat{\sigma}/\sqrt{T-1}$ and $\hat{\sigma}$ is the standard deviation from the regression.

Unit Root Forecast

Forecast using $\tilde{\mu}$ in place of $\hat{\mu}$ using the remainder of the estimated parameters.

Conclusion

- The Rho method is fundamentally a Seasonal AR
- Uses a sequential procedure to select a simple but reasonable specification
- Modest robustness when forecasting the trend in unit root models

CARD: Calibrarion and Forecasting

Kevin Sheppard

Advanced Financial Econometrics: Forecasting

May 2020

Calibration

Improving the Delta and Rho forecasts

- The final forecasts come from an "in-sample" fit
- The original data are augmented with H forecasts to produce an augmented dataset

$$\left\{Y_1, \ldots, Y_T, \hat{Y}_{T+1}, \ldots, \hat{Y}_{T+H}\right\}$$

- $\hat{Y}_{T+h}$ are the average of the Delta and Rho forecasts
- The final model has the form

$$Y_t = \mu + \text{AR} + \text{Seasonality} + \text{Breaks} + \epsilon_t$$

- Define the Fourier terms $S_t = \sin\left(2\pi t/m\right)$ and $C_t = \cos\left(2\pi t/m\right)$
- Define the break variable $d_t = I_{[t<T-1/2\min(4m,T+H)]}$

Calibration

Final Model Specification

$$
\begin{aligned}
Y_t =& \mu + \underbrace{\rho Y_{t-1} I_\rho}_{\text{AR}} + \underbrace{(\rho_m Y_{t-m} + \rho_{m+1} Y_{t-m-1})}_{\text{Seasonal AR}} I_R I_\rho I_4 \\
&+ \underbrace{\sum_{j=1}^{m-1} \delta_j I_{S_j}}_{\text{Seasonal}} I_A + \underbrace{\{\delta_s S_t + \delta_c C_t\}}_{\text{Alt. Seasonal (Fourier)}} (1 - I_A) \\
&+ \underbrace{(\gamma_1 d_t + \gamma_2 t d_t I_5)}_{\text{Trend Breaks}} I_6 + \epsilon_t
\end{aligned}
$$

- I_ρ, I_{AR}, and I_A are the same as the start of Rho
- $I_4 = T > 4m$
- $I_5 = I_\rho \times I_{[m \in \{4,12,13\}]}$
- $I_6 = 1$if $m \neq 24$ and $T > 3m$ and $T + H - k > 10$
  - k is the number of regressors excluding $\gamma_\bullet$

Forecast

- The forecasts are the fitted values of $\hat{Y}_{T+1}, \hat{Y}_{T+2}, \ldots \hat{Y}_{T+H}$
- Each horizon will have a different set of forecasts
- The final forecast is:

$$\hat{X}_{T+h|T} = \exp\left(\hat{\hat{Y}}_{T+h|T}\right) I_{\ln} + \hat{\hat{Y}}_{T+h|T} \left(1 - I_{\ln}\right)$$

- If the data was logged, the forecast is the exponential of the expected value of the log
- This is a median forecast under a symmetry assumption

Conclusion

- Averaging is used to combine the Delta and Rho forecasts
- Calibration builds a rich model to fit the the data *and* the average forecasts
- The CARD forecasts are the fitted average forecast values of the calibration model
- Forecasts are finally transformed back to levels using the median, if necessary