Testing Forecast Rationality for Measures of Central Tendency

Timo Dimitriadis
Heidelberg University

Andrew J. Patton
Duke University

Patrick W. Schmidt
University of Zurich

March 2023

What is your point forecast for rainfall tomorrow?

How would you summarize this predictive distribution in a single number?

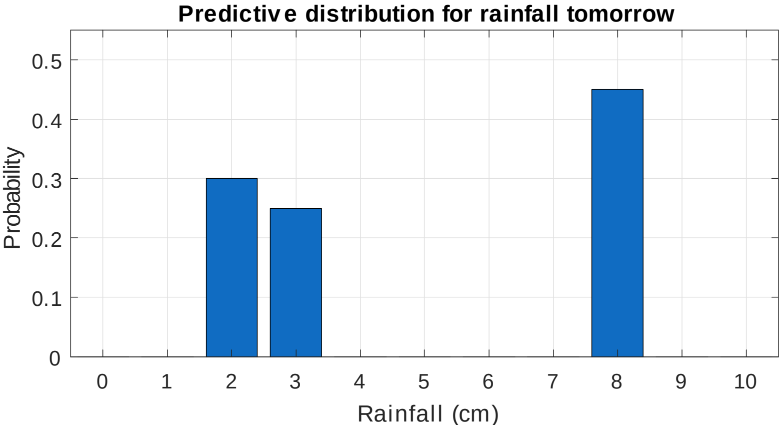

Motivation

- New York Fed's labor market survey:

What do you believe your earnings will be in four months?

- Reasonable answers to this question include the respondent's **expection** of future earnings, or her **median**, or her **mode.**
- When these measures of central tendency coincide, there is no concern about interpreting the point forecast.
- When they may differ, the specific measure used by the respondent may affect its use in other applications, and *testing rationality* becomes difficult.

A general class of central tendency forecasts

- We consider a general class of central tendency measures defined by the mean, median, and mode, and any convex combination of the three:

$$X_t^{\text{central}}(\mathbf{w}) = w_1 X_t^{\text{mean}} + w_2 X_t^{\text{median}} + w_3 X_t^{\text{mode}}$$

- These are the famous "**three Ms**" of statistics.
- The first two Ms are widely studied, but the third:
  - Dalenius (1965, *JRSS*): "The Mode – A Neglected Statistical Parameter"
  - We present new results on the mode to enable us to include it in our analysis.
- Our results can be extended to handle additional measures in the convex combination (e.g., Huber's robust centrality measure, a trimmed mean).

An identification problem

Bank of England forecasts of GDP growth from Feb 2019 (Nov 2018 in parentheses)

Table 5.G Annual average GDP growth rates of modal, median and mean paths[a]

	Mode	Median	Mean
2019	1.2 (1.7)	1.2 (1.7)	1.2 (1.7)
2020	1.5 (1.7)	1.5 (1.7)	1.5 (1.7)
2021	1.9 (1.7)	1.9 (1.7)	1.9 (1.7)

An identification problem, sometimes

Bank of England forecasts of GDP growth from Feb 2018 (Nov 2017 in parentheses)

Table 5.G Annual average GDP growth rates of modal, median and mean paths[a]

	Mode	Median	Mean
2018	1.8 (1.6)	1.8 (1.6)	1.8 (1.6)
2019	1.7 (1.7)	1.8 (1.7)	1.8 (1.7)
2020	1.7 (1.7)	1.7 (1.7)	1.8 (1.7)

An identification problem

- Recall:

$$X_t^{\text{central}}(\mathbf{w}) = w_1 X_t^{\text{mean}} + w_2 X_t^{\text{median}} + w_3 X_t^{\text{mode}}$$

- For symmetric unimodal distributions, the weights on each component are **unidentified**
  - For “nearly” symmetric unimodal distributions, the weights on each component are **weakly identified**
  - For conditionally location-scale distributions, the weights on each component are **partially identified**
- Economic variables may or may not be asymmetrically distributed, and a valid testing approach must accommodate both possibilities.
- We use Stock and Wright’s (2000, *ECMA*) methods for weakly identified GMM to overcome this.

Summary of main findings

1. **New tests of the rationality of mode forecasts**
   - Mode forecast rationality not available in the literature, but essential for considering a test for the class of central tendency forecasts
   - We provide both theoretical and simulation results for these tests

Summary of main findings

1. **New tests of the rationality of mode forecasts**
   - Mode forecast rationality not available in the literature, but essential for considering a test for the class of central tendency forecasts
   - We provide both theoretical and simulation results for these tests
2. **Testing forecast rationality for general measures of central tendency**
   - Allow respondent to use any measure of central tendency
   - Overcome the inherent identification issue that arises

Summary of main findings

1. **New tests of the rationality of mode forecasts**
   - Mode forecast rationality not available in the literature, but essential for considering a test for the class of central tendency forecasts
   - We provide both theoretical and simulation results for these tests
2. **Testing forecast rationality for general measures of central tendency**
   - Allow respondent to use any measure of central tendency
   - Overcome the inherent identification issue that arises
3. **Empirical applications:**
   - New York Fed's survey of consumer expecatations: can reject rationality as mean forecasts, but cannot reject rationality as mode forecasts
   - Fed Board staff's "Greenbook" forecasts: cannot reject rationality as mean forecasts; can reject as median or mode forecasts
   - Random walk forecasts for FX: cannot reject rationality as mean forecasts

Some related work

- **Mode estimation:** Parzen (1962 *AMS*), Eddy (1980 *AoS*), Romano (1988 *AoS*), Kemp et al. (2012, 2019 *JoE*)
    - We draw on this for implementing our test of mode forecast rationality
- **Mode forecasting and elicitability**: Gneiting (2011 *JASA*), Heinrich (2014 *Biometrika*), Fissler and Ziegel (2016 *AoS*)
    - We draw on these in defining asymptotic elicitibility of the mode
- **Survey forecasts and responses:** Manski (2004 *ECMA*), Engleberg, Manski and Williams (2009 *JBES*), Kröger and Peirrot (2019 wp)
    - Motivates our consideration of multiple measures of central tendency
- **Rationality under unknown loss:** Elliott, Komunjer and Timmermann (2005 *REStud*) (EKT), Patton and Timmermann (2007 *JASA*)
    - Like EKT we nest the mean & median as special cases; unlike EKT we consider alternative forecasts only *within the class* of central tendency measures

Outline

Outline

Elicitability of a forecast

- A point forecast (statistical functional) Γ is said to be **elicitable** (Gneiting, 2011 *JASA*) if, for $Y|\mathcal{F} \sim F$, there exists a loss function L where

$$\Gamma(F) = \arg\min_x \; E[L(Y,x)|\mathcal{F}]$$

- Alternatively, it is **identifiable** if there exists an *identification function* V such that

$$0 = E[V(Y,X)|\mathcal{F}] \Leftrightarrow X = \Gamma(F)$$

  - Can think of $V(Y,X)$ as $\partial L(Y,X)/dX$.
  - For the mean, $V(Y,X) = (X - Y) =$ the forecast error
  - For the median, $V(Y,X) = \mathbf{1}\{Y > X\} - \mathbf{1}\{Y < X\}$

- If the forecast is identifiable, then a natural way to test rationality of forecast X is via the first-order condition:

$$0 = E[V(Y,X) \otimes \mathbf{h}] \quad \text{for } \mathbf{h} \in \mathcal{F}$$

Rationality tests for mean and mode forecasts

- **Mean forecast rationality**, defined as

$$H_0^* : X_t = E\left[Y_{t+1}|\mathcal{F}_t\right] \text{ a.s. } \forall\ t$$

is often tested using the famous Mincer-Zarnowitz regression

$$Y_{t+1} = \beta_0 + \beta_1 X_t + e_{t+1}$$

and H_0^* above implies that

$$H_0 : \beta_0 = 0 \cap \beta_1 = 1 \Leftrightarrow H_0 : E\left[\ V_{\text{mean}}\left(Y_{t+1}, X_t\right) \otimes [1, X_t]\ \right] = 0$$

- We implement analogous tests for **mode forecast rationality** by using the asymptotic identification function for the mode introduced below:

$$H_0^* : X_t = Mode\left(Y_{t+1}|\mathcal{F}_t\right) \text{ a.s. } \forall\ t$$

implies

$$H_0 : \lim_{\delta \to 0} E\left[\ V_{\text{mode}}\left(Y_{t+1}, X_t, \delta\right) \otimes [1, X_t]\ \right] = 0$$

Eliciting a mode forecast

- The mode was shown to be **not elicitable** by Heinrich (2014, *Biometrika*), so the "usual" forecast evaluation tests cannot be immediately implemented.

- The *mid-point of the modal interval of length* 2δ **is elicitable**:

$$\begin{aligned} MMP(\delta) &\equiv \arg\max_x \Pr\left[Y \in [x-\delta, x+\delta]\right] \\ &= \arg\max_x E\left[\mathbf{1}\left\{|Y-x| \leq \delta\right\}\right] \end{aligned}$$

- We consider a smoothed version of this:

$$MMP(\delta; K) = \arg\min_x E\left[-\frac{1}{\delta} K\left(\frac{x-Y}{\delta}\right)\right]$$

where K is some kernel function, e.g. the standard Normal PDF.

The mid-point of the modal interval

Red line is the true mode, blue line is the mid-point of the modal interval

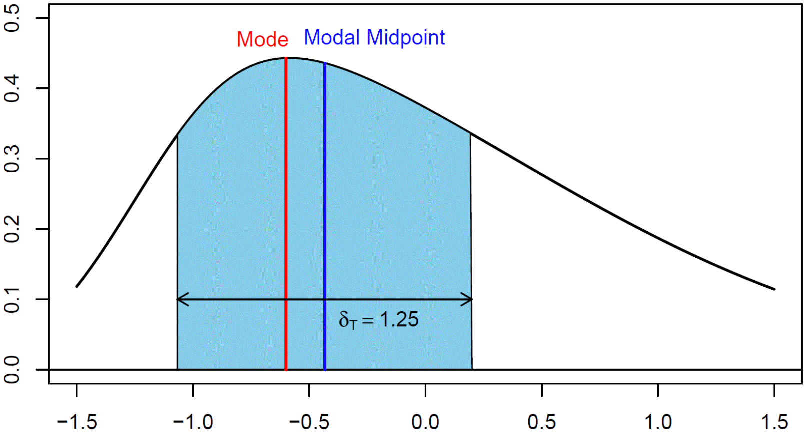

The mid-point of the modal interval

Red line is the true mode, blue line is the mid-point of the modal interval

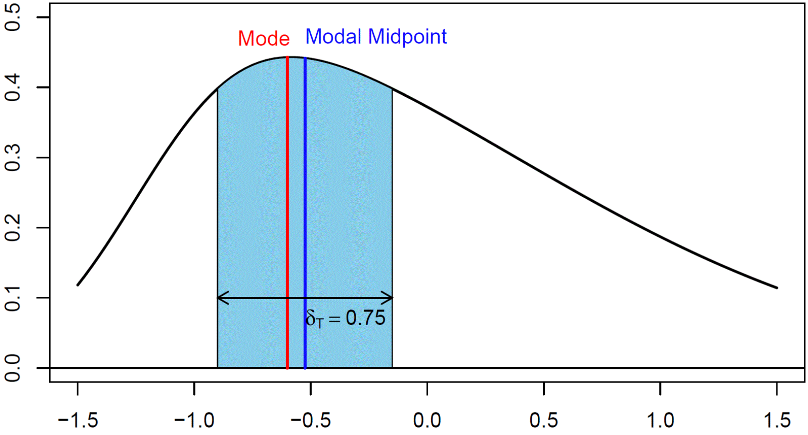

The mid-point of the modal interval

Red line is the true mode, blue line is the mid-point of the modal interval

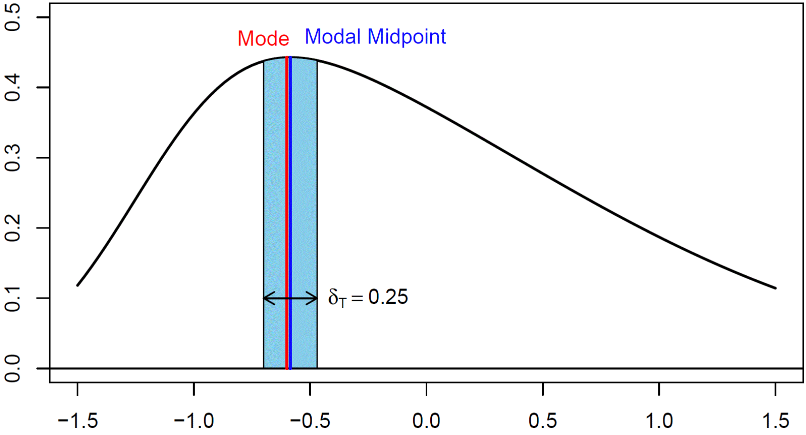

The mode is asymptotically elicitable

Definition: A functional Γ is *asymptotically elicitable* relative to the class of distributions $\mathcal{P}$ if there exists a sequence of elicitable functionals Γ_k such that $\Gamma_k(P) \to \Gamma(P)$ for all $P \in \mathcal{P}$.

Proposition: Let $\mathcal{P}$ denote the class of distributions consisting of absolutely continuous unimodal distributions with bounded density and assume K is is positive, smooth, log-concave, and $\int K(u)\,du = 1$. Then

$$\Gamma_\delta = \arg\min_x \; E\left[-\frac{1}{\delta}K\left(\frac{x-Y}{\delta}\right)\right]$$

is well defined for all $\delta > 0$, and

$$\Gamma_\delta \to Mode(P) \text{ as } \delta \to 0 \text{ for all } P \in \mathcal{P}.$$

Rationality tests for mode forecasts

- We would like to test the null hypothesis:

$$H_0^* : \; X_t = Mode\left(Y_{t+1} | \mathcal{F}_t\right) \; a.s. \; \forall \; t$$

Rationality tests for mode forecasts

- We would like to test the null hypothesis:

$$H_0^* : \; X_t = Mode\left(Y_{t+1}|\mathcal{F}_t\right) \; a.s. \; \forall \; t$$

- Under H_0^* and Assumption 1 (see paper) we have

$$\delta_T^{3/2} T^{-1/2} \sum_{t=1}^{T} V\left(Y_{t+1}, X_t, \delta_T\right) \mathbf{h}_t \xrightarrow{d} N\left(0, \Omega_{\text{Mode}}\right) \text{ as } T \to \infty$$

$$\text{where} \quad V\left(Y_{t+1}, X_t, \delta_T\right) = -\frac{1}{\delta_T^2} K'\left(\; \delta_T^{-1}\left(X_t - Y_{t+1}\right) \; \right)$$

$$\Omega_{\text{Mode}} = E\left[\mathbf{h}_t \mathbf{h}_t' f_t\left(0\right)\right] \int K'\left(u\right)^2 du$$

Rationality tests for mode forecasts

- We would like to test the null hypothesis:

$$H_0^* : \ X_t = Mode\left(Y_{t+1}|\mathcal{F}_t\right) \ a.s. \ \forall \ t$$

- Under H_0^* and Assumption 1 (see paper) we have

$$\delta_T^{3/2} T^{-1/2} \sum_{t=1}^{T} V\left(Y_{t+1}, X_t, \delta_T\right) \mathbf{h}_t \xrightarrow{d} N\left(0, \Omega_{\text{Mode}}\right) \text{ as } T \to \infty$$

$$\text{where} \quad V\left(Y_{t+1}, X_t, \delta_T\right) = -\frac{1}{\delta_T^2} K'\left(\delta_T^{-1}\left(X_t - Y_{t+1}\right) \right)$$

$$\Omega_{\text{Mode}} = E\left[\mathbf{h}_t \mathbf{h}_t' f_t\left(0\right)\right] \int K'\left(u\right)^2 du$$

- Following Kemp et al. (2012 *JoE*, 2019 *JoE*):

$$\hat{\Omega}_{T,\text{Mode}} = \frac{1}{T} \sum_{t=1}^{T} \mathbf{h}_t \mathbf{h}_t' \delta_T^{-1} K'\left(\delta_T^{-1}\left(X_t - Y_{t+1}\right) \right)^2 \xrightarrow{p} \Omega_{\text{Mode}}$$

Outline

Evaluating forecasts of measures of central tendency

- We now seek to test whether we sequence of forecasts is consistent with rationality for *some* measure of central tendency.
- We consider the class of central tendency measures defined by convex combinations of the mean, median, and mode:

$$X_t^{\text{central}}(\mathbf{w}) = w_1 X_t^{\text{mean}} + w_2 X_t^{\text{median}} + w_3 X_t^{\text{mode}}$$

- Intuitively, we will try to find the value(s) for **w** that "rationalize" a given sequence of forecasts.
  - When two or more of these measures are identical, there may be many values of **w** that rationalize the forecasts.

Identification function for a general measure of centrality

- Consider the $(3 \times k)$ matrix of identification functions, weighted and interacted with the instrument vector:

$$\psi_{t,T} = T^{-1/2} \begin{pmatrix} \mathbf{h}_t' \ \widehat{\mathbf{W}}_{T,\text{mean}} & \times & X_t - Y_{t+1} \\ \mathbf{h}_t' \ \widehat{\mathbf{W}}_{T,\text{median}} & \times & 1/2 - \mathbf{1}\left\{X_t \leq Y_{t+1}\right\} \\ \mathbf{h}_t' \ \widehat{\mathbf{W}}_{T,\text{mode}} & \times & \delta_T^{-1/2} K'\left(\frac{X_t - Y_{t+1}}{\delta}\right) \end{pmatrix}$$

where $\widehat{\mathbf{W}}_{T,\bullet}$ are pos def weight matrices that may depend on the data.

- We take a convex combination of these to use in estimation and testing:

$$\phi_{t,T} = \boldsymbol{\theta}' \psi_{t,T}$$

where $\boldsymbol{\theta} \in \Theta \equiv$ the 3D unit simplex.

Estimating the optimal combination weights

- The GMM objective function:

$$S_t(\boldsymbol{\theta}) = \left(\sum_{t=1}^{T} \phi_{t,T}(\boldsymbol{\theta})\right)' \hat{\Sigma}_T^{-1}(\boldsymbol{\theta}) \left(\sum_{t=1}^{T} \phi_{t,T}(\boldsymbol{\theta})\right)$$

  where $\hat{\Sigma}_T(\boldsymbol{\theta})$ is an estimate of the asymptotic covariance matrix of the moment conditions.

- Unlike standard applications, we *cannot* assume that $\boldsymbol{\theta}_0$ is well identified
  - If the distribution is symmetric unimodal, $\boldsymbol{\theta}_0$ is unidentified
  - If the distribution is location-scale, $\boldsymbol{\theta}_0$ is partially identified
  - If the distribution is only weakly asymmetric, $\boldsymbol{\theta}_0$ is weakly identified

- Thus we cannot assume that we can consistently estimate $\boldsymbol{\theta}_0$
  - We use Stock and Wright (2000 *ECMA*) to obtain valid confidence bounds for $\boldsymbol{\theta}_0$ in strongly, weakly, partially, and unidentified cases.

Estimating the optimal combination weights

- Stock and Wright (2000 *ECMA*) provide conditions under which we can obtain asymptotically valid confidence bounds for $\boldsymbol{\theta}_0$ in strongly, weakly, and unidentified cases.

- Consider the set:

$$\hat{\Theta}_T^* = \left\{ \boldsymbol{\theta} \in \Theta \text{ s.t. } S_T(\boldsymbol{\theta}) < \chi_k^{(-1)}(1-\alpha) \right\}$$

  where $\chi_k^{(-1)}(1-\alpha)$ is the $(1-\alpha)$ quantile of the χ_k^2 distribution.

- The set $\hat{\Theta}_T^*$ contains the values of $\boldsymbol{\theta}$ (i.e., the measures of central tendency) that "rationalize" the observed sequence of forecasts and realizations.

- If $\hat{\Theta}_T^*$ is empty, we reject rationality for the *entire class* of measures of central tendency.

Outline

Simulation design

- We simulate data from the following DGP:

$$\begin{aligned} Y_{t+1} &= \mathbf{Z}_t'\zeta + \sigma_{t+1}\varepsilon_{t+1} \\ \varepsilon_{t+1} &\sim \text{Skew } N(0,1,\gamma) \end{aligned}$$

- We vary γ to examine the impact of skewness, and we consider:
  1. iid data
  2. heteroskedastic independent data
  3. homoskedastic AR(1) data
  4. AR(1)-GARCH(1,1) data
- We consider three instrument sets:

$$\mathbf{h}_{1,t} = 1, \quad \mathbf{h}_{2,t} = [1, X_t], \quad \mathbf{h}_{3,t} = [1, X_t, Z_{1,t}]$$

Finite sample size: Mode forecast rationality test

Rejection rates reasonably close to the nominal size of 5.0%.

Testing $$H_0 : \lim_{\delta \to 0} E\left[\, V_{\text{mode}}\left(Y_{t+1}, X_t, \delta\right) \otimes [1, X_t] \,\right] = 0$$

	iid data				**AR-GARCH data**			
	Skewness				*Skewness*			
T	*0*	*0.1*	*0.25*	*0.5*	*0*	*0.1*	*0.25*	*0.5*
100	3.7	3.8	3.9	6.0	3.9	4.3	4.9	5.7
500	5.0	5.5	6.5	7.6	5.4	5.9	6.7	7.6
2000	4.8	5.9	6.2	6.2	5.4	5.5	7.4	6.1
5000	4.9	5.3	6.6	5.9	5.2	6.2	6.7	5.9

Finite sample coverage: Central tendency rationality test

Coverage rates of confidence sets close to 90% nominal level. (AR-GARCH data)

Optimal forecast	**Symmetric data** T				**Skewed data** T			
	100	*500*	*2000*	*5000*	*100*	*500*	*2000*	*5000*
Mean	88.7	87.9	90.6	90.6	88.1	89.7	88.8	90.2
Median	91.2	90.3	90.9	90.1	88.8	87.3	88.5	90.1
Mode	88.9	90.4	89.5	89.8	91.0	91.6	92.6	91.6
Mean-Mode	90.0	90.1	90.5	90.5	91.3	91.7	93.6	92.1
Mean-Median	89.6	90.1	90.0	89.4	90.1	91.4	90.2	90.8
Median-Mode	90.8	90.6	90.8	90.3	91.7	90.9	92.5	90.5
Mean-Med-Mod	89.4	90.0	90.3	90.6	91.7	92.6	93.3	92.3

Finite sample coverage: Central tendency rationality test

Asymmetric data with Mode<Median<Mean. Coverage is >85% "near" the mean vertex.

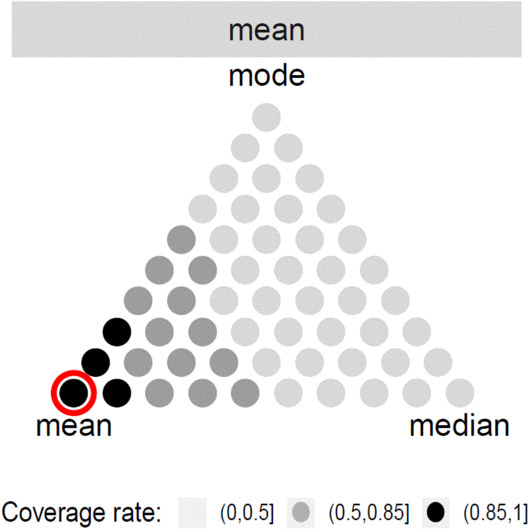

Finite sample coverage: Central tendency rationality test

Symmetric data. Mean rationality => All forecasts are rational. Coverage >85% everywhere.

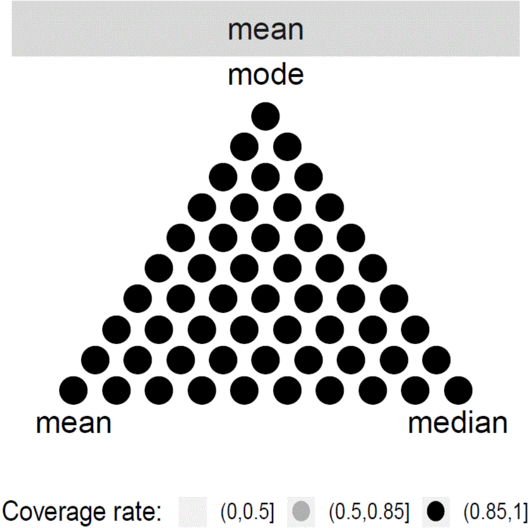

Finite sample coverage: Central tendency rationality test

Asymmetric data with Mode<Median<Mean. Coverage is >85% at the mode vertex.

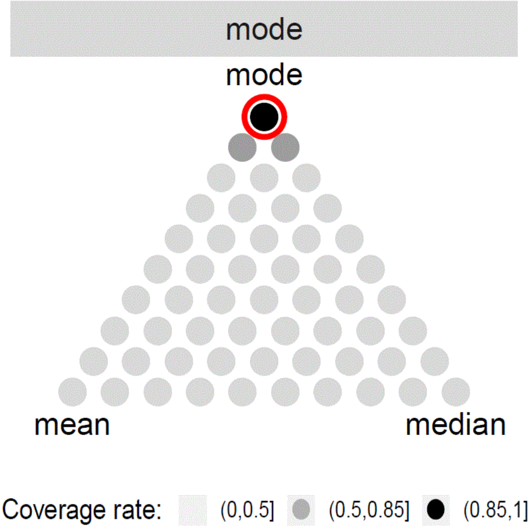

Finite sample coverage: Central tendency rationality test

Asymmetric data with Mode<Median<Mean. Red indicates set of optimal forecasts.

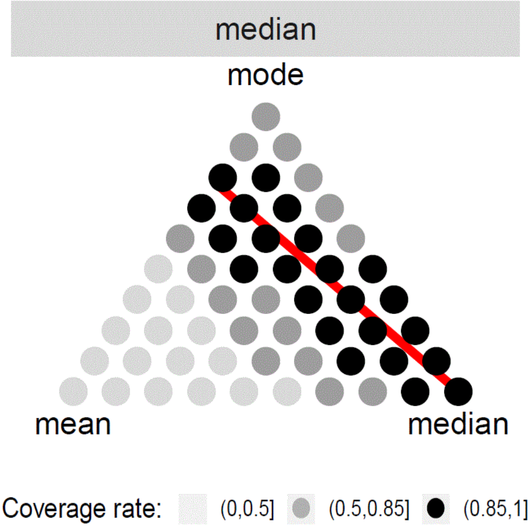

Finite sample coverage: Central tendency rationality test

Asymmetric data with Mode<Median<Mean. Red indicates set of optimal forecasts.

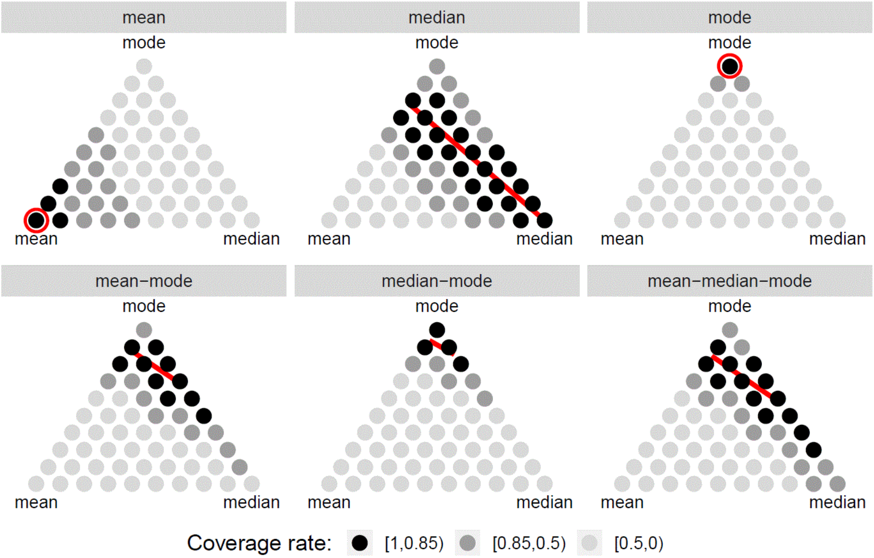

Outline

New York Fed survey forecasts of income

- We use data from the New York Fed's Survey of Consumer Expectations – Labor Market Survey.
  - *"What do you believe your earnings will be in four months?"*
  - Our sample runs from March 2015 to March 2018.
  - We obtain a sample of 3,916 pairs of forecasts and realizations.
- We first test rationality of the entire sample, assuming that all respondents use the same unknown measure of centrality.
- We then stratify our sample using covariates from the survey to allow for cross-respondent heterogeneity in the centrality measure used.

New York Fed survey forecasts of income

Using constant and forecast as instruments (right panel) we can only rationalize as the mode

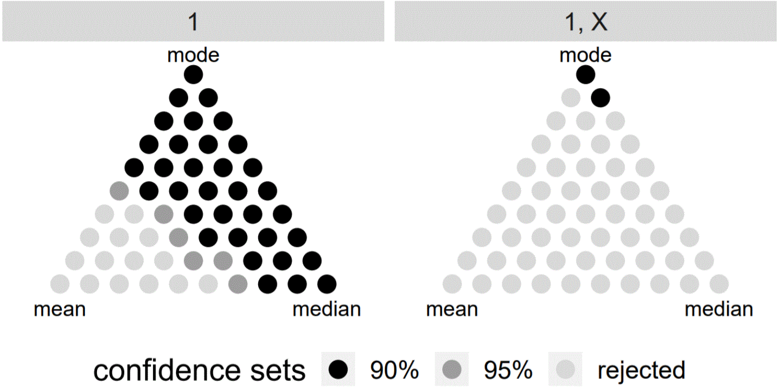

New York Fed survey: Stratify by past income

Forecasts from low-income respondents can only be rationalized as the mode

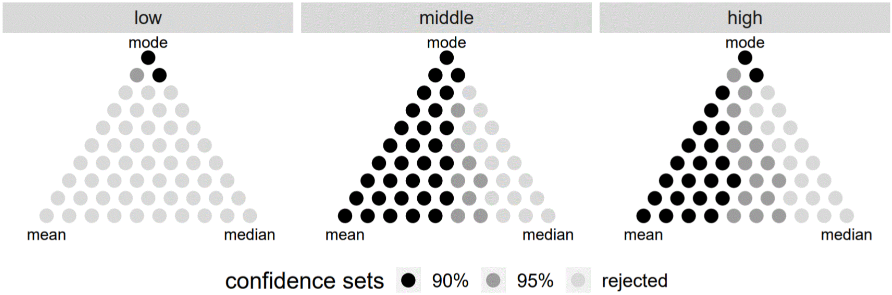

New York Fed survey: Stratify by income and age

Forecasts from younger low-income respondents cannot be rationalized using any measure

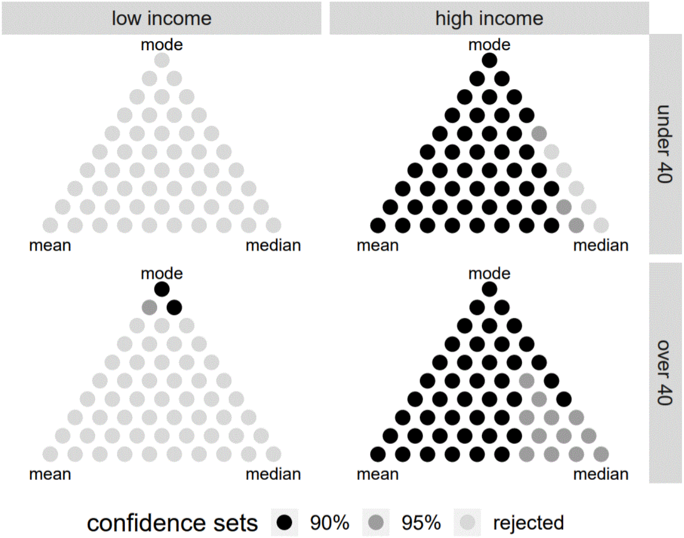

Greenbook forecasts of US GDP growth

- We use one-quarter-ahead "Greenbook" forecasts of US GDP growth produced by the Fed Board
  - Sample is 1967Q2 to 2015Q2, 192 observations.

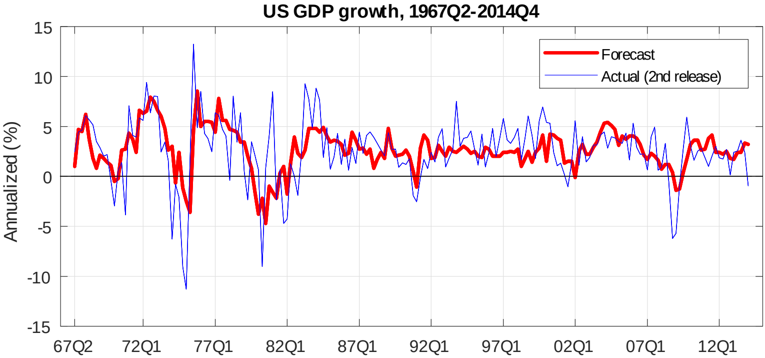

Greenbook forecasts of US GDP growth

Always rationalizable as mean forecasts; also as median forecasts in latest vintage

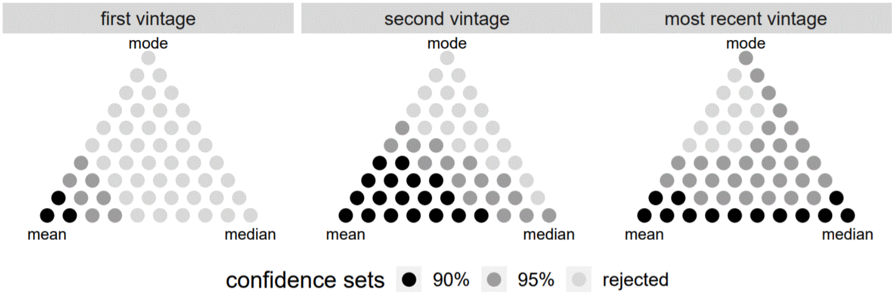

Random walk forecasts of exchange rates

- We revisit the famous result of Meese and Rogoff (1983), that exchange rate changes are unpredictable when using the squared-error loss function.
  - This implies that the lagged exchange rate is an optimal mean forecast.
- We use daily data on USD/EUR, JPY/EUR and AUD/EUR exchange rates
- Sample period is May 2000 to July 2020, which is 5,265 observations.
  - This sample is *way* out-of-sample relative to the Meese and Rogoff paper.

Random walk forecasts of exchange rates

Always rationalizable as mean forecasts; not always for the other centrality measures

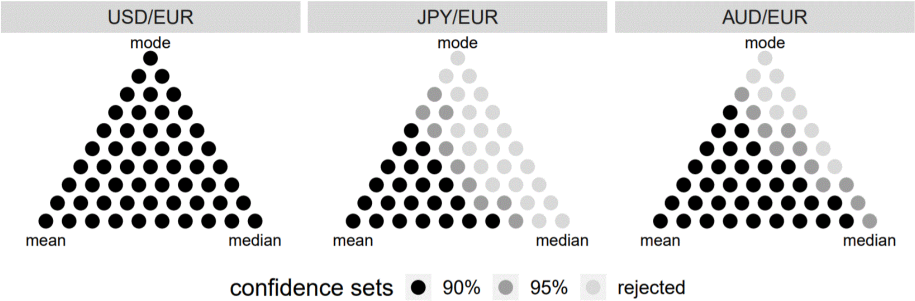

Conclusion

- Economic surveys generally request a *point* forecast, despite calls (e.g. Manski, 2004 *ECMA*) for surveys to solicit distributional forecasts.
- Reasonable people can interpret a request for a prediction of a random variable in a variety of ways
  - Mean, median, mode, perhaps other summary statistics
  - Economic surveys are generally vague about the specific quantity to be reported
- This paper proposes new methods to test the rationality of forecasts of some unknown measure of central tendency
  - Our set of general measures of central tendency is any convex combination of the mean, median, and mode.
  - Our testing approach overcomes an inherent identification problem present for only weakly asymmetric variables.

Summary of results

1. Using NY Fed surveys of individual income forecasts, we find:
   - In aggregate, these forecasts can only be rationalized as **mode** forecasts
   - But forecasts from younger low-income respondents are **not rationalizable** using *any* measure of central tendency
   - Forecasts from high-income respondents are rationalizable using **almost all measures** of central tendency.
2. Using "Greenbook" forecasts of US GDP growth, we cannot reject rationality w.r.t. to the **mean**, but we can w.r.t. to the median and mode.
3. Using random walk forecasts of exchange rates, we find:
   - For USD/EUR, forecasts are not rejected for *any* centrality measure
   - For AUD/EUR and JPY/EUR, forecasts are rational for centrality measures "close" to the mean, but not w.r.t. the median or mode.

Appendix

Assumption 1 (for mode rationality test)

1. The sequence $(Y_{t+1}, X_t, \mathbf{h}_t)$ is stationary and ergodic
2. $E\left[\|\mathbf{h}_t\|^{2+\epsilon}\right] < \infty$
3. $E[\mathbf{h}_t \mathbf{h}_t']$ has full rank for all t
4. The conditional distribution of $\varepsilon_t \equiv X_t - Y_{t+1}$ given $\mathcal{F}_t$ is absolutely continuous with density f_t that is three-times continuously differentiable with bounded derivatives
5. $K : \mathbb{R} \to \mathbb{R}$ satisfies $\int K(u)\, du = 1$, $\int uK(u)\, du = 0$, $\int u^2 K(u)\, du \leq c < \infty$, $\int K'(u)^2\, du \leq c < \infty$, $\sup K(u) \leq c < \infty$, $\sup K'(u) \leq c < \infty$
6. δ_T is a strictly positive, deterministic sequence such that $T\delta_T \to \infty$ and $T\delta_T^7 \to 0$ as $T \to \infty$

Estimating vs. eliciting the mode

- Given a sample of data $\{Y_t\}_{t=1}^T$, Parzen (1962 *AMS*), Eddy (1980 *AoS*), Romano (1988, *AoS*) show how to consistently **estimate** the mode of Y_t :

$$\widehat{Mode}_T = \arg\max_x \frac{1}{T}\sum_{t=1}^{T}\frac{1}{\delta_T}K\left(\frac{x-Y_t}{\delta_T}\right) \xrightarrow{p} Mode \text{ as } T\to\infty,\ \delta_T\to 0$$

- This is related to **eliciting** the (conditional) mode

$$\widetilde{Mode}_t = \arg\min_x \ E\left[-\frac{1}{\delta}K\left(\frac{x-Y_t}{\delta}\right)\middle|\mathcal{F}_t\right] \to Mode_t \text{ as } \delta\to 0$$

- If a functional is elicitable (identifiable) then a natural M- (GMM) estimator exists for it.

- We use the existence of an asymptotic estimator of the mode to obtain asymptotic elicitability of the mode.

Comparison with EKT

- Our approach is related to, but distinct from, that of Elliott, et al. (2005, *REStud*) (EKT) who consider parametric families of loss functions with asymmetry parameter $\theta \in (0,1)$:

$$\begin{aligned} \textit{Lin-Lin} \quad L(Y,X;\theta) &= |Y-X| \; \left(1+(1-2\theta)\mathbf{1}\{Y>X\}\right) \\ \textit{Quad-Quad} \quad L(Y,X;\theta) &= (Y-X)^2 \left(1+(1-2\theta)\mathbf{1}\{Y>X\}\right) \end{aligned}$$

≈ They estimate θ using GMM, and use over-identifying restrictions to test rationality at $\hat{\theta}_T$.

× These families generate the median and mean at $\theta = 1/2$, but for $\theta \neq 1/2$ the optimal forecast is **not** a measure of central tendency.

  - Our alternatives to the median and mean are all measures of centrality.

× Under mild conditions, θ is identified in EKT $\Rightarrow$ no identification issue.

  - Our approach faces a fundamental identification problem, which we address.

Bandwidth selection

- Theory requires $\delta_T \propto T^{-a}$ with $a \in (1/7, 1)$. Optimal rate if $a = 1/7 + \epsilon$.
- Kemp et al. (2019 *JoE*) used the rule:

$$\delta_T = k \times \widehat{MAD}\left[|X_t - Y_{t+1}|\right] \times T^{-0.143}$$

- We found we needed to generalize this to allow the bandwidth to vary with the degree of **skewness**: bandwidth needs to shrink for skewed data
- We use

$$\delta_T = 2.4 \times \exp\left\{-3\left|\hat{\gamma}\right|\right\} \times \widehat{MAD}\left[|X_t - Y_{t+1}|\right] \times T^{-0.143}$$

where

$$\hat{\gamma} = 3\frac{\hat{E}\left[X_t - Y_{t+1}\right] - \widehat{Median}\left[X_t - Y_{t+1}\right]}{\hat{V}\left[X_t - Y_{t+1}\right]^{1/2}}$$

is Pearson's "second coefficient of skewness."

Convex combinations of centrality measures

- We describe our set of centrality measures as a convex combination of the mean, median and mode:

$$X_t^{\text{central}} = \mathbf{w}' \begin{bmatrix} X_t^{\text{mean}} \\ X_t^{\text{median}} \\ \delta X_t^{\text{mode}} \end{bmatrix}$$

- In testing, we take a convex combination of the *identification functions* for the mean, median and mode:

$$T^{1/2}\phi_{t,T} = \boldsymbol{\theta}' \begin{bmatrix} \mathbf{h}_t' \, \widehat{\mathbf{W}}_{T,\text{mean}} & \times & X_t - Y_{t+1} \\ \mathbf{h}_t' \, \widehat{\mathbf{W}}_{T,\text{median}} & \times & 1/2 - \mathbf{1}\left\{X_t \le Y_{t+1}\right\} \\ \mathbf{h}_t' \, \widehat{\mathbf{W}}_{T,\text{mode}} & \times & \delta_T^{-1/2} K'\left(\frac{X_t - Y_{t+1}}{\delta}\right) \end{bmatrix}$$

- The weights vectors $\boldsymbol{\theta}$ and $\mathbf{w}$ lie in the unit simplex, but $\boldsymbol{\theta} \neq \mathbf{w}$ in general.
  - At the vertices, $\boldsymbol{\theta} = \mathbf{w}$, obviously.
  - It is possible to numerically find the mapping from $\boldsymbol{\theta}$ to $\mathbf{w}$, which depends on the weight matrices and instruments. (We do this in our analysis.)