Prior and Posterior Predictive Checks

Frank Schorfheide
Department of Economics, University of Pennsylvania

Predictive Checks

- See, for instance, Gelman, Carlin, Stern, and Rubin (1995), Lancaster (2003), Geweke (2005).

- Prior predictive check: does the model have a chance explaining salient features of the data?

- Posterior predictive check: tries to assess the "absolute" fit of the model – similar to classical specification test.

- Recall: posterior odds are designed for relative model comparisons.

Prior Predictive Checks

- Let Y^{rep} be a sample of observations of length T that we could have observed in the past or that we might observe in the future.

- Let's construct a predictive distribution based on our prior knowledge for Y^{rep}:

$$p(Y^{rep}) = \int p(Y^{rep}|\theta) \underbrace{p(\theta)}_{\text{Prior}} d\theta$$

- Let $\mathcal{S}(Y)$ be a sample statistic of interest. From $p(Y^{rep})$ we can derive the predictive distribution of $p(\mathcal{S})$.

- Compute the observed value of $\mathcal{S}$ based on the actual data and assess how far it lies in the tails of its predictive distribution.

Posterior Predictive Checks

- Let Y^{rep} be a sample of observations of length T that we could have observed in the past or that we might observe in the future.

- Let's construct a predictive distribution based on our posterior knowledge for Y^{rep}:

$$p(Y^{rep}) = \int p(Y^{rep}|\theta) \underbrace{p(\theta|Y)}_{\text{Posterior}} d\theta$$

- Let $\mathcal{S}(Y)$ be a sample statistic of interest. From $p(Y^{rep})$ we can derive the predictive distribution of $p(\mathcal{S})$.

- Compute the observed value of $\mathcal{S}$ based on the actual data and assess how far it lies in the tails of its predictive distribution.

Prior (Posterior) Predictive Checks

- Implementation: for $s = 1$ to n_{sim}:

  1. Generate a draw $\theta^{(s)}$ from prior (posterior).

  2. Simulate data $Y^{(s)}$ from model conditional on $\theta^{(s)}$.

  3. Compute $\mathcal{S}(Y^{(s)})$.

Posterior Predictive Checks

- Let Y^{rep} be a sample of observations of length T that we could have observed in the past or that we might observe in the future.

- Derive the sampling distribution of Y^{rep} given the current state of knowledge:

$$p(Y^{rep}|Y) = \int p(Y^{rep}|\theta) \; \underbrace{p(\theta|Y)}_{\text{Posterior}} \; d\theta. \tag{1}$$

- Let $\mathcal{S}(Y)$ be a test quantity and compute predictive distribution for $\mathcal{S}(Y)$.

- Implementation: same as for prior predictive check – replace $\theta^{(s)}$ draws from prior with draws from posterior.

(insert figures)

Figure 8: Posterior Predictive Check

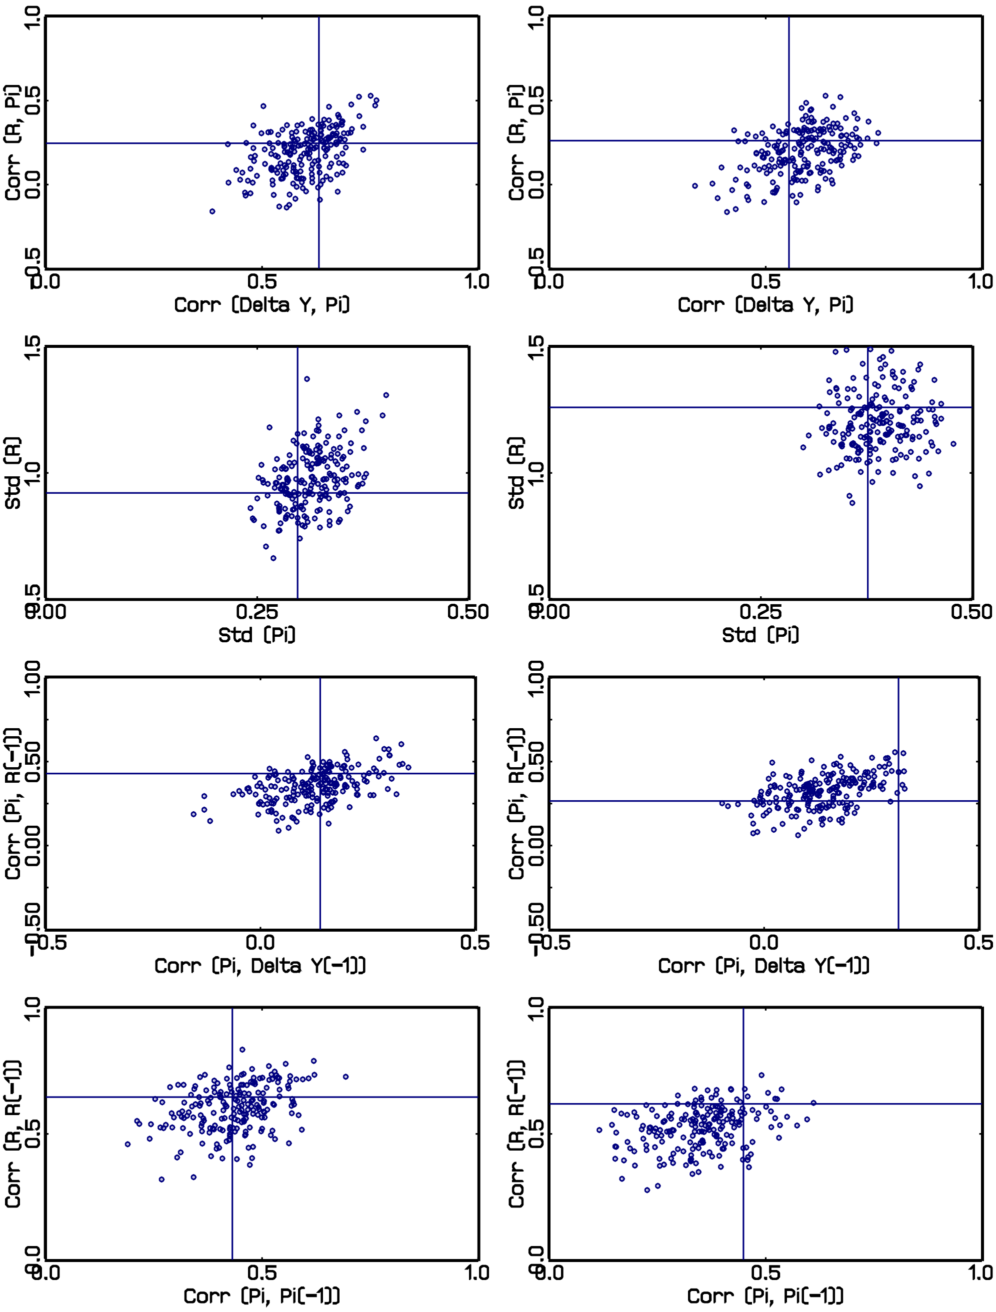

Notes: Output gap rule specification $\mathcal{M}_1$. We plot 200 draws from the posterior predictive distribution of various sample moments. Intersections of solid lines signify the observed sample moments. Left panels: Data Set 1-$\mathcal{M}_1$, no misspecification. Right panels: Data Set 2, misspecification.