The Dynamics of Expected Returns: Evidence from Multi-Scale Time Series Modeling

Daniele Bianchi
University of Warwick

Andrea Tamoni
London School of Economics

9th ECB Workshop on Forecasting Techniques

June 3-4, 2016

Connecting Long- and Short-Horizon Expected Returns

- Find instrument z_t for the expected returns
- Run predictive system

$$
\begin{aligned}
r_{t+1} &= \beta z_t + \varepsilon^r_{t+1} \,, &\qquad \varepsilon^r_{t+1} &\sim N\left(0, \sigma^2_r\right) \\
z_{t+1} &= \phi z_t + \varepsilon^z_{t+1} \,, &\qquad \varepsilon^z_{t+1} &\sim N\left(0, \sigma^2_z\right)
\end{aligned}
$$

Connecting Long- and Short-Horizon Expected Returns

- Find instrument z_t for the expected returns
- Run predictive system

$$r_{t+1} = \beta z_t + \varepsilon^r_{t+1} \ , \qquad \varepsilon^r_{t+1} \sim N\left(0, \sigma^2_r\right)$$
$$z_{t+1} = \phi z_t + \varepsilon^z_{t+1} \ , \qquad \varepsilon^z_{t+1} \sim N\left(0, \sigma^2_z\right)$$

such that

$$E_t r_{t+1} = \beta z_t$$

Connecting Long- and Short-Horizon Expected Returns

- Find instrument z_t for the expected returns
- Run predictive system

$$r_{t+1} = \beta z_t + \varepsilon^r_{t+1} \, , \qquad \varepsilon^r_{t+1} \sim N\left(0, \sigma^2_r\right)$$
$$z_{t+1} = \phi z_t + \varepsilon^z_{t+1} \, , \qquad \varepsilon^z_{t+1} \sim N\left(0, \sigma^2_z\right)$$

such that

$$E_t r_{t+1} + r_{t+2} = \beta \left(1 + \phi\right) z_t$$

Connecting Long- and Short-Horizon Expected Returns

- Find instrument z_t for the expected returns
- Run predictive system

$$r_{t+1} = \beta z_t + \varepsilon^r_{t+1} \ , \qquad \varepsilon^r_{t+1} \sim N\left(0, \sigma^2_r\right)$$
$$z_{t+1} = \phi z_t + \varepsilon^z_{t+1} \ , \qquad \varepsilon^z_{t+1} \sim N\left(0, \sigma^2_z\right)$$

such that

$$E_t \sum_{j=1}^{J} r_{t+j} = \beta \sum_{j=0}^{J-1} \phi^j z_t$$

Connecting Long- and Short-Horizon Expected Returns

- Find instrument z_t for the expected returns

- Run predictive system

$$r_{t+1} = \beta z_t + \varepsilon^r_{t+1} \;, \qquad \varepsilon^r_{t+1} \sim N\left(0, \sigma_r^2\right)$$
$$z_{t+1} = \phi z_t + \varepsilon^z_{t+1} \;, \qquad \varepsilon^z_{t+1} \sim N\left(0, \sigma_z^2\right)$$

  such that

$$E_t \sum_{j=1}^{J} r_{t+j} = \beta \sum_{j=0}^{J-1} \phi^j z_t$$

- This implies that:
  - All relevant information lies at the highest possible frequency
  - Long-run expected returns are inferred mechanically by using short-run estimates

In data, frequency-specific predictability

- The predictive power of forecasting variables studied in the literature varies with the time horizons.

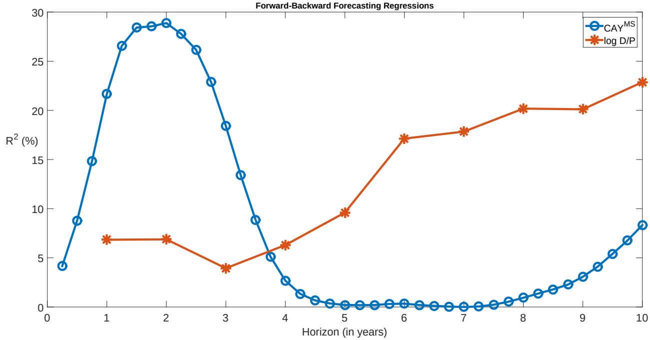

Figure: Horizon-specific predictive regressions. This figure shows the in-sample R^2 of simple linear regressions (with an intercept) of h-period continuously compounded market returns on the CRSP value-weighted index on h-period past consumption-wealth ratio (solid line with circles) and on h-period past log price-dividend ratio (solid line with asterisks). For the consumption-wealth ratio regressions, the sample is quarterly and spans the period 1952Q1-2013Q4. For the dividend-price ratio regressions, the sample is annual and spans the period 1952-2013.

This Paper...

- **Methodology:** We propose a novel econometric framework to coherently encode the scale-specific information from alternative predictors.

This Paper...

- **Methodology:** We propose a novel econometric framework to coherently encode the scale-specific information from alternative predictors.
- **Goal:** Investigate whether the dominant view that long-run expected returns are the result of an aggregation of short-run returns is backed up by empirical evidence.

This Paper...

- **Methodology:** We propose a novel econometric framework to coherently encode the scale-specific information from alternative predictors.
- **Goal:** Investigate whether the dominant view that long-run expected returns are the result of an aggregation of short-run returns is backed up by empirical evidence.
- **Main Result:** We show that expected returns exhibit aggregation properties which differs wildly from those obtained within the class of typical ARMA specifications once conditioning on lower-frequency predictors
  - Long-lasting effects of past shocks on latent expected returns.
  - Implications for forecasting returns and investment decisions.

Related Literature

1. Perfect predictors
   (e.g., Kandel and Stambaugh, 1996; Barberis, 2000; and Campbell and Viceira, 2005)
2. Imperfect predictors
   (e.g., Pastor and Stambaugh, 2009, 2012; Van Binsbergen and Koijen, 2010; Carvalho, Lopes and McCulloch 2015)
3. Mixed-frequency and scale-specific predictability
   (e.g., Ghysels, Santa-Clara and Valkanov, 2004; Corsi, 2009; Lochstoer 2009; Schorfheide and Song 2015; Bandi et al. 2016)
4. Multi-scale time series
   (Mandelbrot, 1974; Calvet and Fisher, 2007; Ferreira et al. 2006 and Ferreira and Lee, 2007)

Multi-Scale Time Series Model

- Objective: The model must
  - combine scale-specific information in a coherent way;
  - subsume the standard AR(1) case;
  - be parsimonious.

Multi-Scale Time Series Model

- Objective: The model must
  - combine scale-specific information in a coherent way;
  - subsume the standard AR(1) case;
  - be parsimonious.
- Framework: couples standard linear ARMA models at different frequencies of observation via a stochastic linkage equation.

Multi-Scale Time Series Model

Key ingredients:

- Latent expected returns; $x_{1:n_x} \sim p\left(x_{1:n_x}\right) = N\left(0, V_x\right).$

Multi-Scale Time Series Model

Key ingredients:

- Latent expected returns; $x_{1:n_x} \sim p\left(x_{1:n_x}\right) = N\left(0, V_x\right)$.
- Scale-specific predictor; $z_{1:n_z} \sim q\left(z_{1:n_z}\right) = N\left(0, Q_z\right)$.

Multi-Scale Time Series Model

Key ingredients:

- Latent expected returns; $x_{1:n_x} \sim p\left(x_{1:n_x}\right) = N\left(0, V_x\right)$.
- Scale-specific predictor; $z_{1:n_z} \sim q\left(z_{1:n_z}\right) = N\left(0, Q_z\right)$.
- The series $x_{1:n_x}$ and $z_{1:n_z}$ are specified at different temporal scales: $m \times n_z = n_x$, with $m > 1$:

$$x_1, x_2, \ldots, x_m, z_1, x_{m+1}, \ldots, x_{2m}, z_2, \ldots$$

Multi-Scale Time Series Model

Key ingredients:

- Latent expected returns; $x_{1:n_x} \sim p\left(x_{1:n_x}\right) = N\left(0, V_x\right)$.
- Scale-specific predictor; $z_{1:n_z} \sim q\left(z_{1:n_z}\right) = N\left(0, Q_z\right)$.
- The series $x_{1:n_x}$ and $z_{1:n_z}$ are specified at different temporal scales: $m \times n_z = n_x$, with $m > 1$:

$$x_1, x_2, \ldots, x_m, z_1, x_{m+1}, \ldots, x_{2m}, z_2, \ldots$$

- Link equation: The z values are averages of non-overlapping groups of m consecutive x;

$$p\left(z_{1:n_z} | x_{1:n_x}\right) = \prod_{s=1}^{n_z} N\left(m^{-1} \sum_{i=1}^{m} x_{(s-1)m+i}, \tau\right) = N\left(A \cdot x_{1:n_x}, \lambda \left(A' V_x A\right)_{11} I\right),$$

where λ measures *the relative increase in uncertainty due to the lack of agreement between time scales.*

Multi-Scale Time Series Model (Cont'd)

- **Problem:** $q(z_{1:n_z}), p\left(z_{1:n_z} | x_{1:n_x}\right)$ and $p(x_{1:n_x})$ are generally inconsistent.

Multi-Scale Time Series Model (Cont'd)

- **Problem:** $q(z_{1:n_z}), p\left(z_{1:n_z}|x_{1:n_x}\right)$ and $p(x_{1:n_x})$ are generally inconsistent.
- **Idea:** Restore consistency by "revising" the marginal of expected returns (see Jeffrey 1957, Jeffrey 1965 and Diaconis and Zabell 1982)

$$q\left(x_{1:n_x}\right) = \int \underbrace{p\left(x_{1:n_x}|z_{1:n_z}\right)}_{\propto\; p(z_{1:n_z}|x_{1:n_x})p(x_{1:n_x})} q\left(z_{1:n_z}\right) dz_{1:n_z},$$

Model Implications: Persistence of Expected Returns

- The consistent, namely revised, distribution for expected returns $x_{1:n_x}$ can now be defined as

$$q\left(x_{1:n_x}\right) = N\left(0, Q_x\right)$$

where $Q_x = V_x - B\left(W - Q_z\right)B'$, with $B = V_x A' W^{-1}$ and $W = AV_xA' + \lambda\left(A'V_xA\right)_{11} I$.

- *The parameter λ controls how much information the low frequency predictor conveys about the expected returns at the high-frequency,*

$$\lim_{\lambda \to 0} AQ_xA' = Q_z \ ,$$
$$\lim_{\lambda \to \infty} Q_x = V_x \ .$$

Model Implications: Persistence of Expected Returns

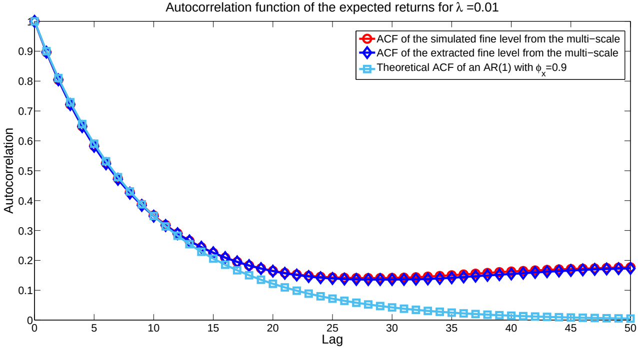

Figure: Multi-scale time series Vs AR(1): autocorrelation functions. This figure shows the theoretical autocorrelation function for a simulated multi-scale time series (red line with circles), the one extracted by conditioning on the lower-frequency series z (blue line with diamonds), and the theoretical autocorrelation of an AR(1) with the same autoregressive parameter as the one used to simulate the multi-scale process (cyan with squares). The parameters $\phi_x = \phi_z = 0.9$ are fixed across panels, $\boldsymbol{\lambda = 0.01}$ and $\sigma_x^2 = \sigma_z^2 = 1$. We simulate $n_x = 720$ with $m = 48$ in all cases.

Model Implications: Persistence of Expected Returns

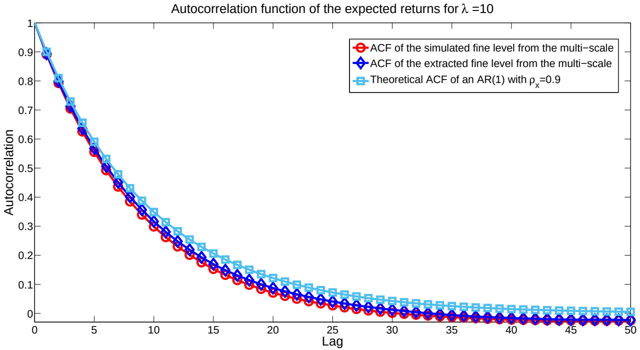

Figure: Multi-scale time series Vs AR(1): autocorrelation functions. This figure shows the theoretical autocorrelation function for a simulated multi-scale time series (red line with circles), the one extracted by conditioning on the lower-frequency series z (blue line with diamonds), and the theoretical autocorrelation of an AR(1) with the same autoregressive parameter as the one used to simulate the multi-scale process (cyan with squares). The parameters $\phi_x = \phi_z = 0.9$ are fixed across panels, $\boldsymbol{\lambda = 10}$ and $\sigma_x^2 = \sigma_z^2 = 1$. We simulate $n_x = 720$ with $m = 48$ in all cases.

Model Implications: Forecasting

- The multi-scale time series model is built in a cascade way from coarse to fine levels of resolution.
- **Step 1:** Predict the observable at the coarse level from the predictive distribution given the current latent state $p\left(z_{n_z+1} \mid x_{1:n_x}, z_{1:n_z}, \Theta\right)$;

Model Implications: Forecasting

- The multi-scale time series model is built in a cascade way from coarse to fine levels of resolution.

- **Step 1:** Predict the observable at the coarse level from the predictive distribution given the current latent state $p\left(z_{n_z+1} \mid x_{1:n_x}, z_{1:n_z}, \Theta\right)$;

- **Step 2:** Now predict the latent expected returns as $p\left(x_{n_x+1} : x_{n_x+m} \mid x_{n_x}, z_{n_z+1}, \Theta\right) \sim N(f_x, F_x)$ where

  $$f_x = r + m^{-1}R1\left(m^{-2}1'R1 + \lambda\left(A'V_xA\right)_{11} I\right)^{-1}\left(z_{n_z+1} - m^{-1}1'r\right)$$

  with $r = x_{n_x}\left(\phi_x, \ldots, \phi_x^m\right)$ and R be the predictive mean and covariance matrix for a not revised AR(1).

- The term $\left(z_{n_z+1} - m^{-1}1'r\right)$ represents the revision due to lower frequency information; if $\lambda \to \infty$ then $f_x = r$.

Simulation: Multiscale Forecasting Vs AR(1)

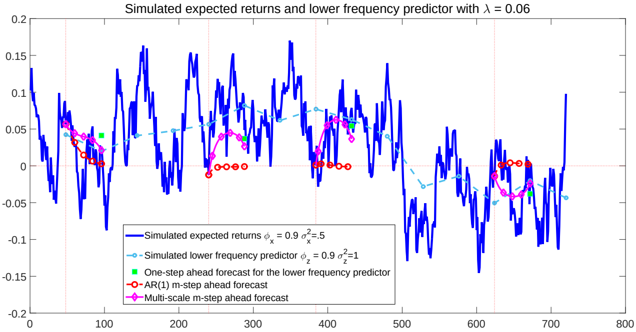

Figure: Comparing forecasts in simulation. Except for the between-levels uncertainty all other parameters are fixed, with $\phi_x = \phi_z = 0.9$, $\sigma_x^2 = .5$ and $\sigma_z^2 = 1$. The figure shows the *m*-step ahead multi-scale forecast (magenta line with diamonds) obtained with $\boldsymbol{\lambda = 0.06}$ compared with the *m*-step ahead (iterated) prediction obtained from an AR(1) with $\phi_x = 0.98$ (red line with circles). The results are obtained by averaging forecasts on 20,000 patterns for predictors and expected returns.

Simulation: Multiscale Forecasting Vs AR(1)

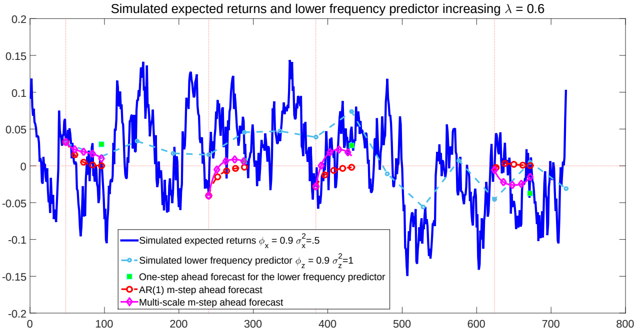

Figure: Comparing forecasts in simulation. Except for the between-levels uncertainty all other parameters are fixed, with $\phi_x = \phi_z = 0.9$, $\sigma_x^2 = .5$ and $\sigma_z^2 = 1$. The figure shows the m-step ahead multi-scale forecast (magenta line with diamonds) obtained with $\boldsymbol{\lambda = 0.6}$ compared with the m-step ahead (iterated) prediction obtained from an AR(1) with $\phi_x = 0.98$ (red line with circles). The results are obtained by averaging forecasts on 20,000 patterns for predictors and expected returns.

Simulation: Multiscale Forecasting Vs AR(1)

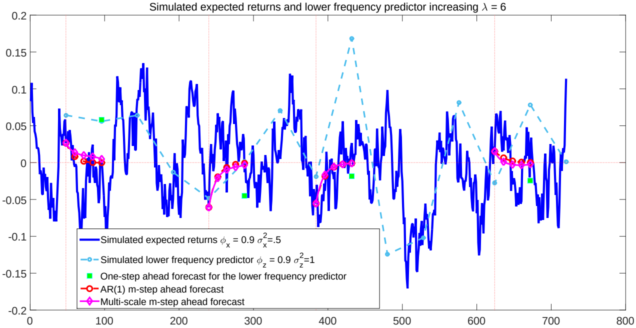

Figure: Comparing forecasts in simulation. Except for the between-levels uncertainty all other parameters are fixed, with $\phi_x = \phi_z = 0.9$, $\sigma_x^2 = .5$ and $\sigma_z^2 = 1$. The figure shows the *m*-step ahead multi-scale forecast (magenta line with diamonds) obtained with $\boldsymbol{\lambda = 6}$ compared with the *m*-step ahead (iterated) prediction obtained from an AR(1) with $\phi_x = 0.98$ (red line with circles). The results are obtained by averaging forecasts on 20,000 patterns for predictors and expected returns.

Estimation Strategy: MCMC Algorithm

- The complete likelihood can be decomposed as

$$p\left(\mathbf{x}_{1:n_x}, \mathbf{z}_{1:n_z} | \phi_x, \lambda, \sigma_x^2, \phi_z, \sigma_z^2\right) = p\left(\mathbf{x}_{1:n_x} | \mathbf{z}_{1:n_z}, \phi_x, \lambda, \sigma_x^2\right) q\left(\mathbf{z}_{1:n_z} | \phi_z, \sigma_z^2\right),$$

Estimation Strategy: MCMC Algorithm

- The complete likelihood can be decomposed as

$$p\left(\mathbf{x}_{1:n_x}, \mathbf{z}_{1:n_z} | \phi_x, \lambda, \sigma_x^2, \phi_z, \sigma_z^2\right) = p\left(\mathbf{x}_{1:n_x} | \mathbf{z}_{1:n_z}, \phi_x, \lambda, \sigma_x^2\right) q\left(\mathbf{z}_{1:n_z} | \phi_z, \sigma_z^2\right),$$

- Metropolis-within-Gibbs algorithm;
  - Conjugate, normal-inverse-gamma marginal priors and standard posterior updates for the predictors $\mathbf{z}_{1:n_z}$.

Estimation Strategy: MCMC Algorithm

- The complete likelihood can be decomposed as

$$p\left(\mathbf{x}_{1:n_x}, \mathbf{z}_{1:n_z} | \phi_x, \lambda, \sigma_x^2, \phi_z, \sigma_z^2\right) = p\left(\mathbf{x}_{1:n_x} | \mathbf{z}_{1:n_z}, \phi_x, \lambda, \sigma_x^2\right) q\left(\mathbf{z}_{1:n_z} | \phi_z, \sigma_z^2\right),$$

- Metropolis-within-Gibbs algorithm;
  - Conjugate, normal-inverse-gamma marginal priors and standard posterior updates for the predictors $\mathbf{z}_{1:n_z}$.
  - Conjugate, normal-inverse-gamma marginal priors and Random-Walk Metropolis-Hasting (MH) steps for posterior updating for the latent expected returns $\mathbf{x}_{1:n_x}$.
- Acceptance rates around 30% of draws for the MH steps, convergence is achieved after 20,000 burn-in draws (Geweke 1992 convergence test).

Empirical Example

- Expected returns are extracted at a monthly frequency conditioning on the joint dynamics of:
  - The log dividend-price ratio sampled at a 4-year horizon, averaged ($n_{dp} = 15, m = 48$);
  - The consumption-wealth ratio sampled at a one-year horizon ($n_{cay} = 60, m = 12$).
- Sample period is 1952:01-2013:12.
- Alternative prior specifications for the parameters.
- The windows m are chosen based on marginal likelihood evidences.

Posterior Estimates of the Multi-Scale Parameters

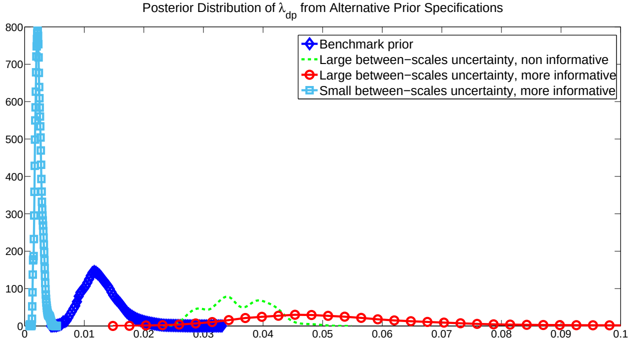

Figure: This figure reports the posterior distributions of the parameters for the latent expected returns extracted by using jointly the annual consumption-wealth and the four-year log dividend-price ratios. The blue line with diamonds shows the posterior distribution obtained from our benchmark prior specification. The light-blue and red-circled lines represent the posterior distribution obtained under tighter and weaker prior elicitation. Estimates are obtained from a sample of 30,000 draws out of 50,000 simulations, storing every other draw. The sample period is 1952:01-2013:12.

Posterior Estimates of the Multi-Scale Parameters

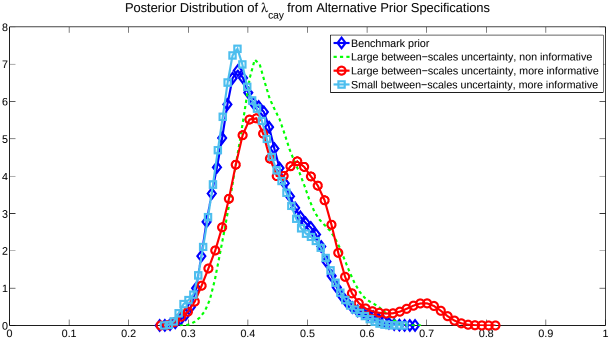

Figure: This figure reports the posterior distributions of the parameters for the latent expected returns extracted by using jointly the annual consumption-wealth and the four-year log dividend-price ratios. The blue line with diamonds shows the posterior distribution obtained from our benchmark prior specification. The light-blue and red-circled lines represent the posterior distribution obtained under tighter and weaker prior elicitation. Estimates are obtained from a sample of 30,000 draws out of 50,000 simulations, storing every other draw. The sample period is 1952:01-2013:12.

Posterior Estimates of the Multi-Scale Parameters

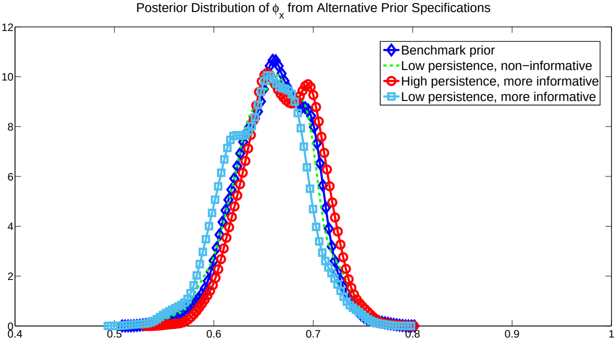

Figure: This figure reports the posterior distributions of the parameters for the latent expected returns extracted by using jointly the annual consumption-wealth and the four-year log dividend-price ratios. The blue line with diamonds shows the posterior distribution obtained from our benchmark prior specification. The light-blue and red-circled lines represent the posterior distribution obtained under tighter and weaker prior elicitation. Estimates are obtained from a sample of 30,000 draws out of 50,000 simulations, storing every other draw. The sample period is 1952:01-2013:12.

Persistence of Expected Returns Vs Sum of two AR(1)

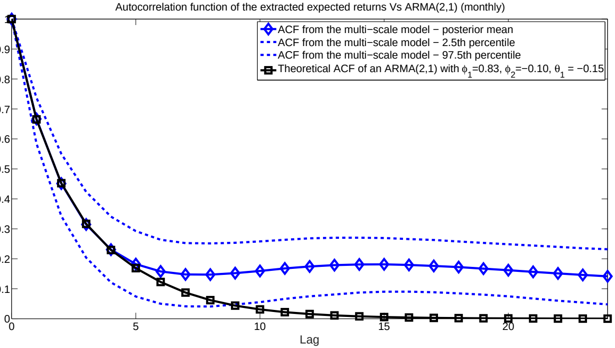

Figure: This figure shows the ACF of expected returns estimated by jointly considering both the annual consumption-wealth ratio ($m = 12$) and the 4-year log dividend-yield ($m = 48$) as scale-specific predictors (solid line with diamonds).

Persistence of Expected Returns Vs Sum of two AR(1)

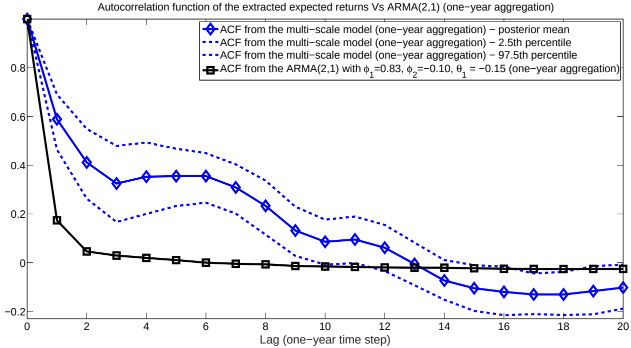

Figure: This figure shows the ACF of expected returns estimated by jointly considering both the annual consumption-wealth ratio ($m = 12$) and the 4-year log dividend-yield ($m = 48$) as scale-specific predictors, aggregated over 1-year (solid line with diamonds).

Persistence of Expected Returns Vs Sum of two AR(1)

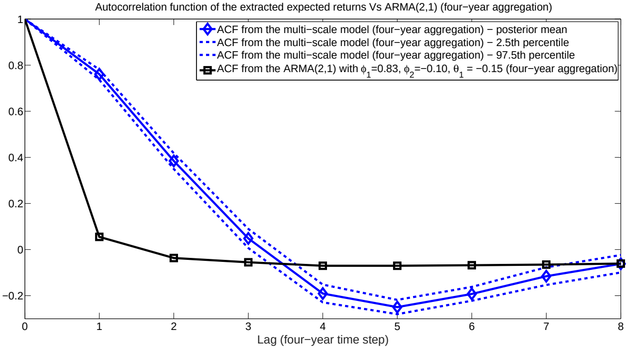

Figure: This figure shows the ACF of expected returns estimated by jointly considering both the annual consumption-wealth ratio ($m = 12$) and the 4-year log dividend-yield ($m = 48$) as scale-specific predictors, aggregated over 4-year (solid line with diamonds).

Persistence of Expected Returns Vs OLS fitted value

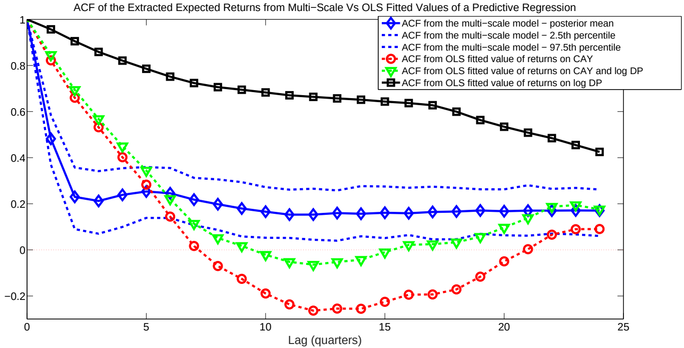

Figure: Comparison with OLS expected returns: autocorrelation functions. This figure shows the autocorrelation functions of the expected returns extracted by using jointly both the annual consumption-wealth ratio ($m = 12$) and the four-year log dividend-price ($m = 48$), in comparison to the ones implied by standard OLS predictive regressions. The Figure displays the posterior average autocorrelation function (solid line with diamonds) along with the 95% confidence intervals (dashed lines). We compare the autocorrelation from the multi-scale against the one implied by the fitted values of a standard predictive regression with the log dividend-price ratio (**black line with squares**), with the consumption-wealth ratio (red line with circles), and with both the log dividend-price and the consumption-wealth ratio (green line with triangles).

Expected Returns Vs OLS fitted value

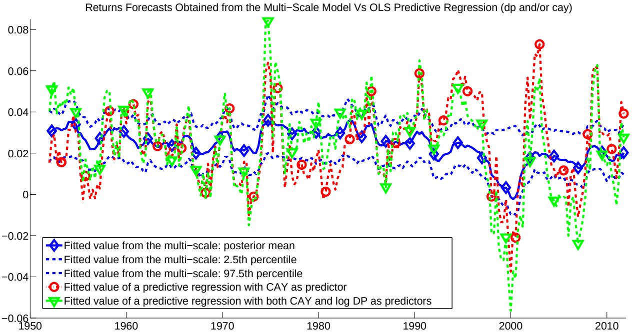

Figure: Comparison with OLS. The expected return process is estimated by jointly considering both the annual consumption-wealth ratio ($m = 12$) and the 4-year log dividend-price ($m = 48$) as scale-specific predictors. We run simple (multiple) regressions of quarterly returns on our extracted series, and on the consumption-wealth ratio (on the dividend-price ratio and the consumption-wealth ratio). The figure compares our expected return series (solid line with diamonds) with the OLS fitted value from the consumption-wealth ratio (red line with circles) as well as with the fitted value from a multiple regression of quarterly returns onto CAY and log DP (green line with triangles).

Forecasting Accuracy of Future Expected Returns

Panel A: Mean Squared Error

	Forecasting Horizon (months)			
Model	$h = 12$	$h = 24$	$h = 36$	$h = 48$
AR(1)	2.346	4.082	4.262	4.877
Sum of two AR(1)	2.321	3.924	4.157	4.766
AR(1) with time-varying mean	3.521	4.178	4.209	4.721
Multi-scale	2.355	3.739	3.796	4.277

Panel B: Log-Predictive Score

	Forecasting Horizon (months)			
Model	$h = 12$	$h = 24$	$h = 36$	$h = 48$
AR(1)	-23.578	-27.627	-27.836	-28.434
Sum of two AR(1)	-23.503	-27.327	-27.631	-28.252
AR(1) with time-varying mean	-31.071	-30.369	-30.391	-30.701
Multi-scale	-23.319	-26.463	-26.603	-27.158

Table: Forecasting accuracy measures. This table reports summary statistics about the forecasting accuracy of future expected returns obtained from our multi-scale time series model. The forecasting performance of the model is compared with: (1) the forecasts obtained from a simple AR(1) fitted on the extracted expected returns; (2) the forecasts obtained from a sum of two independent (at all leads and lags) AR(1) processes; (3) the forecasts obtained from an AR(1) model for which the mean is stationary and is allowed to vary stochastically over time. The latent series of expected returns and corresponding forecasts are obtained from the joint process of the log dividend-price and the consumption-wealth variable *CAY*. Forecasts are produced monthly with an horizon of $h = 48$ months. For the ease of exposition the table reports the results for $h = 12, 24, 36, 48$ months. **Panel A:** Mean Squared Errors obtained from the marginal predictive mean. **Panel B:** Log predictive scores obtained from the marginal predictive likelihood. Draws from the predictive distribution are obtained from a sample of 30,000 draws out of 50,000 simulations, with a thinning size of 20. The sample period is 1952:01-2013:12.

Forecasting Accuracy of Future Expected Returns

Panel A: Mean Squared Error

	Forecasting Horizon (months)			
Model	$h = 12$	$h = 24$	$h = 36$	$h = 48$
AR(1)	2.346	4.082	4.262	4.877
Sum of two AR(1)	2.321	3.924	4.157	4.766
AR(1) with time-varying mean	3.521	4.178	4.209	4.721
Multi-scale	2.355	3.739	3.796	4.277

Panel B: Log-Predictive Score

	Forecasting Horizon (months)			
Model	$h = 12$	$h = 24$	$h = 36$	$h = 48$
AR(1)	-23.578	-27.627	-27.836	-28.434
Sum of two AR(1)	-23.503	-27.327	-27.631	-28.252
AR(1) with time-varying mean	-31.071	-30.369	-30.391	-30.701
Multi-scale	-23.319	-26.463	-26.603	-27.158

Table: Forecasting accuracy measures. This table reports summary statistics about the forecasting accuracy of future expected returns obtained from our multi-scale time series model. The forecasting performance of the model is compared with: (1) the forecasts obtained from a simple AR(1) fitted on the extracted expected returns; (2) the forecasts obtained from a sum of two independent (at all leads and lags) AR(1) processes; (3) the forecasts obtained from an AR(1) model for which the mean is stationary and is allowed to vary stochastically over time. The latent series of expected returns and corresponding forecasts are obtained from the joint process of the log dividend-price and the consumption-wealth variable *CAY*. Forecasts are produced monthly with an horizon of $h = 48$ months. For the ease of exposition the table reports the results for $h = 12, 24, 36, 48$ months. **Panel A:** Mean Squared Errors obtained from the marginal predictive mean. **Panel B:** Log predictive scores obtained from the marginal predictive likelihood. Draws from the predictive distribution are obtained from a sample of 30,000 draws out of 50,000 simulations, with a thinning size of 20. The sample period is 1952:01-2013:12.

Conclusions

- Our multi-scale based expected returns series exhibits aggregation properties which differs wildly from those obtained
  - within the class of typical ARMA specifications, e.g. ARMA(2,1);
  - from standard OLS fitted value.
- Combining information at multiple horizons to construct short-term expected returns has important implications for forecasting
- Framework widely applicable to other problem where multiple frequency are relevant:
  1. short-term rate forecasts: trade-off between long-run targets (inflation, gdp) and higher frequency market fluctuations.
  2. inflation forecasts: long-run unemployment rate, volatility,etc.