



Università Commerciale
Luigi Bocconi

Lecture 2: Forecasting stock returns (part 1)

Prof. Massimo Guidolin

Advanced Financial Econometrics III

Winter/Spring 2020

Overview

- The objective of the predictability exercise on stock index returns
- Predictability and the efficient market hypothesis
- How much predictability can we expect?
- In-sample vs. Out-of-Sample Predictability
- The role of economic restrictions
- Forecast combinations
- Diffusion indexes
- Regime shifts

The key point

- The objective of the literature consists of OOS forecasting of the equity premium, i.e., realized excess returns
 - There is now some limited evidence that stock returns are to some extent predictable
-
- A number of studies, such as Goyal and Welch (2008, RFS), argue that, despite extensive in-sample evidence of excess return predictability, popular predictors fail to outperform the simple historical average benchmark in OOS tests
 - Recent studies, however, indicate that better forecasting strategies deliver statistically and economically significant OOS gains
 - Economically motivated model restrictions
 - Forecast combination
 - Use of diffusion indexes
 - Regime shifts
 - This significant OOS evidence of predictability has implications for asset pricing models and asset management strategies

The key point

- The objective of the literature consists of OOS forecasting of the equity premium, i.e., realized excess returns
 - There is now some limited evidence that stock returns are to some extent predictable
-
- A number of studies, such as Goyal and Welch (2008, RFS), argue that, despite extensive in-sample evidence of excess return predictability, popular predictors fail to outperform the simple historical average benchmark in OOS tests
 - Recent studies, however, indicate that **better forecasting strategies deliver statistically and economically significant OOS gains**
 - Economically motivated model restrictions
 - Forecast combination
 - Use of diffusion indexes
 - Regime shifts
 - There is significant OOS evidence of predictability has implications for asset pricing models and asset management strategies

The key point

- The objective of the literature consists of OOS forecasting of the equity premium, i.e., realized excess returns
 - There is now some limited evidence that stock returns are to some extent predictable
-
- A number of studies, such as Goyal and Welch (2008, RFS), argue that, despite extensive in-sample evidence of excess return predictability, popular predictors fail to outperform the simple historical average benchmark in OOS tests
 - Recent studies, however, indicate that **better forecasting strategies deliver statistically and economically significant OOS gains**
 - Economically motivated model restrictions
 - Forecast combination
 - Use of diffusion indexes
 - Regime shifts
 - There is significant OOS evidence of predictability has implications for asset pricing models and asset management strategies

The key point

- We should not expect the amount of excess return predictability to be substantial
-
- Why should we expect predictability to be modest?
 - Stock returns inherently contain a large unpredictable component, so that **the best forecasting models will explain only a small part of high-frequency (e.g., monthly) stock returns**
 - Competition among traders implies that once successful forecasting models are discovered, they will be readily adopted by others
 - The widespread adoption of successful forecasting models can then cause stock prices to move in a manner that eliminates the models' forecasting ability
 - Rational asset pricing theory posits that stock return predictability can result from exposure to time-varying aggregate risk
 - Only to the extent that previously successful forecasting models consistently capture time variation in aggregate risk premiums, they will likely remain successful

The key point

- We should not expect the amount of excess return predictability to be substantial
-
- Why should we expect predictability to be modest?
 - Stock returns inherently contain a large unpredictable component, so that **the best forecasting models will explain only a small part of high-frequency (e.g., monthly) stock returns**
 - Competition among traders implies that once successful forecasting models are discovered, they will be readily adopted by others
 - The widespread adoption of successful forecasting models can then cause stock prices to move in a manner that eliminates the models' forecasting ability
 - **Rational asset pricing theory posits that stock return predictability can result from exposure to time-varying aggregate risk**
 - Only to the extent that previously successful forecasting models consistently capture time variation in aggregate risk premiums, they will likely remain successful

Predictability and the Efficient Markets Hypothesis

- The EMH does not imply that stock returns should not be predictable, but **any predictability ought to be justified by exposure to systematic risk factors**
- Common misconception that stock return predictability is contrary to market efficiency
 - The canonical random walk model implies that future stock returns are unpredictable on the basis of currently available information
- Yet, while the random walk model is consistent with market efficiency, so is a predictable return process, since predictability is consistent with exposure to aggregate risk
- It is only when the risk-adjusted return—after further adjusting for transaction costs and other trading frictions (e.g., liquidity and borrowing constraints, research costs)—becomes economically positive can we say that the market is inefficient
- However theory does impose certain bounds on the maximum degree of return predictability consistent with market efficiency

Predictability and the Efficient Markets Hypothesis

- The EMH does not imply that stock returns should not be predictable, but **any predictability ought to be justified by exposure to systematic risk factors**
- Common misconception that stock return predictability is contrary to market efficiency
 - The canonical random walk model implies that future stock returns are unpredictable on the basis of currently available information
- Yet, while the random walk model is consistent with market efficiency, so is a predictable return process, since predictability is consistent with exposure to aggregate risk
- It is only when the risk-adjusted return—after further adjusting for transaction costs and other trading frictions (e.g., liquidity and borrowing constraints, research costs)—becomes economically positive can we say that the market is inefficient
- However **theory does impose certain bounds on the maximum degree of return predictability consistent with market efficiency**

Predictability and the Efficient Markets Hypothesis

- To the extent that return predictability exceeds these bounds, it can be interpreted as evidence of market inefficiency
 - This may derive from information processing limitations and/or the types of psychological influences emphasized in behavioral finance
- Stock return predictability is typically examined via the following predictive regression model:

$$r_{t+1} = \alpha + \beta x_t + e_{t+1}$$

where r_{t+1} is the time-(t + 1) return on a broad stock market index in excess of the risk-free interest rate and x_t is a variable used to predict the equity premium (such as the dividend-price ratio)

- Because a valid stochastic discount factor (SDF, or state-price density or pricing kernel) m_{t+1} , satisfies $E(R_{j,t+1}m_{t+1}|I_t) = 1$ ($j = 1, \dots, N$), where $R_{j,t+1}$ is the gross return on asset j
- Assuming that the risk-free rate R_f is constant, Ross (2005) shows that the regression R^2 has an elegant **upper bound**:

$$R^2 \leq R_f^2 \text{Var}(m_t)$$

Predictability and the Efficient Markets Hypothesis

- To the extent that return predictability exceeds these bounds, it can be interpreted as evidence of market inefficiency
 - This may derive from information processing limitations and/or the types of psychological influences emphasized in behavioral finance
- Stock return predictability is typically examined via the following predictive regression model:

$$r_{t+1} = \alpha + \beta x_t + e_{t+1}$$

where r_{t+1} is the time-($t + 1$) return on a broad stock market index in excess of the risk-free interest rate and x_t is a variable used to predict the equity premium (such as the dividend-price ratio)

- Because a valid stochastic discount factor (SDF, or state-price density or pricing kernel) m_{t+1} , satisfies $E(R_{j,t+1}m_{t+1}|I_t) = 1$ ($j = 1, \dots, N$), where $R_{j,t+1}$ is the gross return on asset j
- Assuming that the risk-free rate R_f is constant, Ross (2005) shows that the regression R^2 has an elegant **upper bound**:

$$R^2 \leq R_f^2 \text{Var}(m_t)$$

Predictability and the Efficient Markets Hypothesis

- To the extent that return predictability exceeds these bounds, it can be interpreted as evidence of market inefficiency
 - This may derive from information processing limitations and/or the types of psychological influences emphasized in behavioral finance
- Stock return predictability is typically examined via the following predictive regression model:

$$r_{t+1} = \alpha + \beta x_t + e_{t+1}$$

where r_{t+1} is the time-($t + 1$) return on a broad stock market index in excess of the risk-free interest rate and x_t is a variable used to predict the equity premium (such as the dividend-price ratio)

- A valid stochastic discount factor (SDF, or state-price density or pricing kernel) m_{t+1} , satisfies $E(R_{j,t+1}m_{t+1}|I_t) = 1$ ($j = 1, \dots, N$), where $R_{j,t+1}$ is the gross return on asset j
- Assuming that the risk-free rate R_f is constant and using SDF, Ross (2005) shows that the regression R^2 has an elegant **upper bound**:

$$R^2 \leq R_f^2 \text{Var}(m_t)$$

How Much Predictability Can We Expect?

- Rational asset pricing models suggest that **we should expect very limited predictability in monthly data**, of 1% at most
 - Under an annualized risk-free rate of 3.5%, an annualized standard deviation of 20% for the U.S. aggregate stock market, and an upper bound on market risk aversion equaling five times the observed VIX, the R^2 bound is approximately 8% for monthly returns
 - The SDF corresponds to the representative investor's intertemporal marginal rate of substitution in consumption-based models
 - This bound is too loose to be binding in applications, e.g., Zhou (2010, EL) reports monthly R^2 s of less than 1% for individual predictive regressions based on ten popular variables
- Under special assumptions on the SDF, Kan and Zhou (2007, JoBus) report much tighter bound,
$$R^2 \leq \rho_{z,m_0}^2 R_f^2 \text{Var}[m_t(z_t)]$$
where ρ_{z,m_0} is the correlation between the predictors and the SDF
- With ρ_{z,m_0} ranging from 0.10–0.15 the R^2 bounds are a fraction of 1% so that many empirical papers violate these bounds

How Much Predictability Can We Expect?

- Rational asset pricing models suggest that **we should expect very limited predictability in monthly data**, of 1% at most
 - Under an annualized risk-free rate of 3.5%, an annualized standard deviation of 20% for the U.S. aggregate stock market, and an upper bound on market risk aversion equaling five times the observed VIX, the R^2 bound is approximately 8% for monthly returns
 - The SDF corresponds to the representative investor's intertemporal marginal rate of substitution in consumption-based models
 - This bound is too loose to be binding in applications, e.g., Zhou (2010, EL) reports monthly R^2 s of less than 1% for individual predictive regressions based on ten popular variables
- Under special assumptions on the SDF, Kan and Zhou (2007, JoBus) report much tighter bound,
$$R^2 \leq \rho_{z,m_0}^2 R_f^2 \text{Var}[m_t(z_t)]$$
where ρ_{z,m_0} is the correlation between the predictors and the SDF
- With ρ_{z,m_0} ranging from 0.10–0.15 the R^2 bounds are a fraction of 1% so that many empirical papers violate these bounds

In-Sample vs. OOS Predictability

Even an apparently small degree of return predictability can translate into substantial utility gains for a risk-averse investor

- Predictive models that claim to explain a large part of stock return fluctuations thus imply massive market inefficiencies and the availability of substantial risk-adjusted abnormal returns
- Although we should expect a limited degree of stock return forecastability, it is important to realize that a little goes a long way
- Even an apparently small degree of return predictability can translate into substantial utility gains for a risk-averse investor who does not affect market prices (e.g., Kandel and Stambaugh, 1996, JF; Xu, 2004, JEF; Campbell and Thompson, 2008, RFS)
- Most popular variables found to predict U.S. stock returns?
 - The dividend-price ratio
 - The earnings-price ratio
 - Nominal interest rates
 - Interest rate spreads

In-Sample vs. OOS Predictability

Even an apparently small degree of return predictability can translate into substantial utility gains for a risk-averse investor

- Predictive models that claim to explain a large part of stock return fluctuations thus imply massive market inefficiencies and the availability of substantial risk-adjusted abnormal returns
- Although we should expect a limited degree of stock return forecastability, it is important to realize that a little goes a long way
- **Even an apparently small degree of return predictability can translate into substantial utility gains for a risk-averse investor** who does not affect market prices (e.g., Kandel and Stambaugh, 1996, JF; Xu, 2004, JEF; Campbell and Thompson, 2008, RFS), see part 2
- Most popular variables found to predict U.S. stock returns?
 - The dividend-price ratio
 - The earnings-price ratio
 - Nominal interest rates
 - Interest rate spreads

In-sample vs. OOS Predictability

Even an apparently small degree of return predictability can translate into substantial utility gains for a risk-averse investor

- Predictive models that claim to explain a large part of stock return fluctuations thus imply massive market inefficiencies and the availability of substantial risk-adjusted abnormal returns
- Although we should expect a limited degree of stock return forecastability, it is important to realize that a little goes a long way
- Even an apparently small degree of return predictability can translate into substantial utility gains for a risk-averse investor who does not affect market prices (e.g., Kandel and Stambaugh, 1996, JF; Xu, 2004, JEF; Campbell and Thompson, 2008, RFS)
- Most popular variables found to predict U.S. stock returns?
 - The dividend-price ratio
 - The earnings-price ratio
 - Nominal interest rates
 - Interest rate spreads

In-sample vs. OOS Predictability

- Inflation
 - Dividend payout ratio
 - Corporate issuing activity
 - Consumption-wealth ratio
 - Stock market volatility
 - Labor income
 - Aggregate output
 - Output gap
 - Oil prices
 - Lagged industry portfolio returns
 - Accruals
- The evidence for U.S. aggregate stock return predictability is predominantly in-sample
 - Goyal and Welch (2008) show that OOS forecasts based on the bivariate regression model fail to consistently outperform the simple historical average benchmark forecast in terms of MSFE

In-sample vs. OOS Predictability

- Inflation
 - Dividend payout ratio
 - Corporate issuing activity
 - Consumption-wealth ratio
 - Stock market volatility
 - Labor income
 - Aggregate output
 - Output gap
 - Oil prices
 - Lagged industry portfolio returns
 - Accruals
- The evidence for U.S. aggregate stock return predictability is predominantly in-sample
 - Goyal and Welch (2008, RFS) show that OOS forecasts based on the bivariate regression model fail to consistently outperform the simple historical average benchmark forecast in terms of MSFE

The Role of Economic Restrictions

Imposing sign and “sum-of-parts” restrictions on predictive regressions has been shown to improve their performance

- Goyal and Welch (2008) also find that a multiple regression forecasting model that includes all potential predictors—the “**kitchen sink**” forecast—performs much worse than the historical average
- It is well known that, due to in-sample over-fitting, highly parameterized models typically perform very poorly OOS
- The first approach for improving forecasting performance imposes **economically motivated restrictions** on predictive regressions
- Campbell and Thompson (2008, RFS) recommend **sign restrictions**
 - Such restrictions on fitted r_{t+1} and β reduce parameter estimation uncertainty and help to stabilize predictive regression forecasts
 - They find that restricted regression forecasts based on a number of economic variables outperform the historical average forecast
- Ferreira and Santa-Clara’s (2011, JFE) **sum-of-the-parts method** from the standard decomposition in capital gain + dividend yield:
$$R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t} = CG_{t+1} + DY_{t+1}$$

The Role of Economic Restrictions

Imposing sign and “sum-of-parts” restrictions on predictive regressions has been shown to improve their performance

- Goyal and Welch (2008) also find that a multiple regression forecasting model that includes all potential predictors—the “kitchen sink” forecast—performs much worse than the historical average
- It is well known that, due to in-sample over-fitting, highly parameterized models typically perform very poorly OOS
- The first approach to improving forecasting performance imposes **economically motivated restrictions** on predictive regressions
- Campbell and Thompson (2008, RFS) recommend **sign restrictions**
 - Such restrictions on fitted r_{t+1} and β reduce parameter estimation uncertainty and help to stabilize predictive regression forecasts
 - They find that restricted regression forecasts based on a number of economic variables outperform the historical average forecast
- Ferreira and Santa-Clara’s (2011, JFE) **sum-of-the-parts method** from the standard decomposition in capital gain + dividend yield:

$$R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t} = CG_{t+1} + DY_{t+1}$$

The Role of Economic Restrictions

Imposing sign and “sum-of-parts” restrictions on predictive regressions has been shown to improve their performance

- Goyal and Welch (2008) also find that a multiple regression forecasting model that includes all potential predictors—the “kitchen sink” forecast—performs much worse than the historical average
- It is well known that, due to in-sample over-fitting, highly parameterized models typically perform very poorly OOS
- The first approach for improving forecasting performance imposes **economically motivated restrictions** on predictive regressions
- Campbell and Thompson (2008, RFS) recommend **sign restrictions**
 - Such restrictions on fitted r_{t+1} and β reduce parameter estimation uncertainty and help to stabilize predictive regression forecasts
 - They find that restricted regression forecasts based on a number of economic variables outperform the historical average forecast
- Ferreira and Santa-Clara’s (2011, JFE) **sum-of-the-parts method** from the standard decomposition in capital gain + dividend yield:

$$R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t} = CG_{t+1} + DY_{t+1}$$

The Role of Economic Restrictions

- Simple algebra and a few accounting identities show that

$$\log(R_{t+1}) = gm_{t+1} + ge_{t+1} + dp_{t+1}$$

where gm_{t+1} (ge_{t+1}) is the log growth rate of the price-earnings multiple (earnings), and dp_{t+1} is the log of 1 + the dividend-price ratio

- Since price-earnings multiples and dividend-price ratios are highly persistent and nearly random walks, reasonable forecasts of gm_{t+1} and dp_{t+1} based on information through t are zero and dp_t
- Earnings growth is nearly entirely unpredictable, apart from a low frequency component, so that they employ a 20-year moving average
- Their sum-of-the-parts equity premium forecast is then given by

$$\hat{r}_{t+1}^{SOP} = \overline{ge}_t^{20} + dp_t - r_{f,t+1}$$

- The sum-of-the-parts forecast is a predictive regression forecast that restricts the slope coefficient to unity for $x_{i,t} = dp_t$ and sets the intercept to $\overline{ge}_t^{20} - r_{f,t+1}$
 - Monte Carlo simulations indicate that the sum-of-the-parts forecast improves upon conventional predictive regression forecasts by substantially reducing estimation error

The Role of Economic Restrictions

- Simple algebra and a few accounting identities show that

$$\log(R_{t+1}) = gm_{t+1} + ge_{t+1} + dp_{t+1}$$

where gm_{t+1} (ge_{t+1}) is the log growth rate of the price-earnings multiple (earnings), and dp_{t+1} is the log of 1 + the dividend-price ratio

- Since price-earnings multiples and dividend-price ratios are highly persistent and nearly random walks, reasonable forecasts of gm_{t+1} and dp_{t+1} based on information through t are zero and dp_t
- Earnings growth is nearly entirely unpredictable, apart from a low frequency component, so that they employ a 20-year moving average
- Their sum-of-the-parts equity premium forecast is then given by

$$\hat{r}_{t+1}^{SOP} = \overline{ge}_t^{20} + dp_t - r_{f,t+1}$$

- The sum-of-the-parts forecast is a predictive regression forecast that restricts the slope coefficient to unity for $x_{i,t} = dp_t$ and sets the intercept to $\overline{ge}_t^{20} - r_{f,t+1}$
 - Monte Carlo simulations indicate that the sum-of-the-parts forecast improves upon conventional predictive regression forecasts by substantially reducing estimation error

Forecast Combinations

Several combination schemes of individual predictive regression forecasts significantly beat the historical average forecast

- Combining forecasts across models often produces a forecast that performs better than the best individual model
 - Forecast combinations can be viewed as a diversification strategy that improves forecasting performance in the same manner that asset diversification improves portfolio performance
- The predictive power of individual models can vary over time, so that a given model provides informative signals during certain periods but predominantly false signals during others
- If the individual forecasts are weakly correlated, a combination of the individual forecasts should be less volatile, thereby reducing risk and improving forecasting performance in environments with substantial model uncertainty and parameter instability
- Rapach et al. (2010, RFS) find that combinations of individual predictive regression forecasts significantly beat the historical average forecast

Forecast Combinations

Several combination schemes of individual predictive regression forecasts significantly beat the historical average forecast

- Combining forecasts across models often produces a forecast that performs better than the best individual model
 - Forecast combinations can be viewed as a diversification strategy that improves forecasting performance in the same manner that asset diversification improves portfolio performance
- The predictive power of individual models can vary over time, so that a given model provides informative signals during certain periods but predominantly false signals during others
- If the individual forecasts are weakly correlated, a combination of the individual forecasts should be less volatile, thereby reducing risk and improving forecasting performance in environments with substantial model uncertainty and parameter instability
- Rapach et al. (2010, RFS) find that combinations of individual predictive regression forecasts significantly beat the historical average forecast

Forecast Combinations

Combination forecast can be interpreted as a “shrinkage” forecast that circumvents in-sample over-fitting problems

- Cremers (2002, RFS) and Dangl & Halling (2012, JFE) use Bayes model averaging to account for model uncertainty in predictive regressions
- It can be beneficial to “tilt” the combining weights toward certain individual: Rapach et al. (2010) show that simple and discounted MSFE combination forecasts of the quarterly U.S. equity premium consistently outperform the historical average
- It is curious that the simple combination forecast performs much better than the kitchen sink forecast, since both approaches entail the estimation of many slope coefficients
- Rapach et al. (2010) show that simple combination forecast can be interpreted as a “shrinkage” forecast that circumvents in-sample over-fitting problems
- Simple combination forecast replaces the slopes in the kitchen sink forecasts, $r_{t+1} - \bar{r} = \sum_{i=1}^K \beta_i^{KS} (x_{i,t} - \bar{x}_i) + \varepsilon_{t+1}$ with $(1/K)\hat{\beta}_i$ in

Forecast Combinations

Combination forecast can be interpreted as a “shrinkage” forecast that circumvents in-sample over-fitting problems

- Cremers (2002, RFS) and Dangl & Halling (2012, JFE) use Bayes model averaging to account for model uncertainty in predictive regressions
- It can be beneficial to “tilt” the combining weights toward certain individual: Rapach et al. (2010) show that simple and discounted MSFE combination forecasts of the quarterly U.S. equity premium consistently outperform the historical average
- It is curious that the simple combination forecast performs much better than the kitchen sink forecast, since both approaches entail the estimation of many slope coefficients
- Rapach et al. (2010) show that simple combination forecast can be interpreted as a “shrinkage” forecast that circumvents in-sample over-fitting problems
- Simple combination forecast replaces the slopes in the kitchen sink forecasts, $r_{t+1} - \bar{r} = \sum_{i=1}^K \beta_i^{KS} (x_{i,t} - \bar{x}_i) + \varepsilon_{t+1}$ with $(1/K)\hat{\beta}_i$ in

Diffusion Indices

$$\hat{r}_{t+1} = \bar{r} + (1/K) \sum_{i=1}^K \hat{\beta}_i (x_{i,t} - \bar{x}_i)$$

- This stabilizes the forecast via two channels: (1) reducing estimation variability by substituting the bivariate regression slope estimates for the multiple regression estimates; (2) shrinking the forecast toward the historical average forecast by pre-multiplying each slope coefficient by $1/K$
- Ludvigson and Ng (2007, JFE) explore diffusion indexes to improve equity premium forecasting
- The diffusion index approach assumes a latent factor model for the potential predictors: $r_{t+1} = \alpha_{DI} + \beta'_{DI} f_t + \varepsilon_{t+1}$

f_t is a q -vector of latent factors and β_{DI} is a q -vector of loadings

- A strict factor model assumes that the disturbance terms are contemporaneously and serially uncorrelated
- An “approximate” factor model permits a limited degree of contemporaneous and/or serial correlation in the residuals

Diffusion Indices

$$\hat{r}_{t+1} = \bar{r} + (1/K) \sum_{i=1}^K \hat{\beta}_i (x_{i,t} - \bar{x}_i)$$

- This stabilizes the forecast via two channels: (1) reducing estimation variability by substituting the bivariate regression slope estimates for the multiple regression estimates; (2) shrinking the forecast toward the historical average forecast by premultiplying each slope coefficient by $1/K$
- Ludvigson and Ng (2007, JFE) explore diffusion indexes to improve equity premium forecasting
- The diffusion index approach assumes a latent factor model for the potential predictors: $r_{t+1} = \alpha_{DI} + \beta'_{DI} f_t + \varepsilon_{t+1}$

f_t is a q -vector of latent factors and β_{DI} is a q -vector of loadings

- A strict factor model assumes that the disturbance terms are contemporaneously and serially uncorrelated
- An “approximate” factor model permits a limited degree of contemporaneous and/or serial correlation in the residuals

Diffusion Indices

Diffusion indices are built on the basis of latent factor models in which a small number of factors summarize the true variables

- **Co-movements in the predictors are primarily governed by fluctuations in the relatively small number of factors**
 - For either the strict or approximate factor model, the latent factors can be consistently estimated by principal components
- Estimates of the latent factors then serve as regressors in the predictive regression model: $r_{t+1} = \alpha_{DI} + \beta'_{DI} f_t + \varepsilon_{t+1}$
 - All K predictors, $x_{i,t}$ ($i = 1, \dots, K$), contain relevant information for r_{t+1}
 - However individual predictors can also provide noisy signals
 - Rather than using the $x_{i,t}$ variables directly, we first identify the important common fluctuations in the potential predictors thereby filtering out the noise in the individual predictors
- The factor structure thus generates a more reliable signal from a large number of predictors to employ in a predictive regression

$$\hat{r}_{t+1}^{DI} = \hat{\alpha}_{DI,t} + \hat{\beta}'_{DI,t} \hat{f}_{t,t}$$

Diffusion Indices

Diffusion indices are built on the basis of latent factor models in which a small number of factors summarize the true variables

- **Co-movements in the predictors are primarily governed by fluctuations in the relatively small number of factors**
 - For either the strict or approximate factor model, the latent factors can be consistently estimated by principal components
- Estimates of the latent factors then serve as regressors in the predictive regression model: $r_{t+1} = \alpha_{DI} + \beta'_{DI} f_t + \varepsilon_{t+1}$
 - All K predictors, $x_{i,t}$ ($i = 1, \dots, K$), contain relevant information for r_{t+1}
 - However individual predictors can also provide noisy signals
 - Rather than using the $x_{i,t}$ variables directly, we first identify the important common fluctuations in the potential predictors thereby filtering out the noise in the individual predictors
- The factor structure thus generates a more reliable signal from a large number of predictors to employ in a predictive regression

$$\hat{r}_{t+1}^{DI} = \hat{\alpha}_{DI,t} + \hat{\beta}'_{DI,t} \hat{f}_{t,t}$$

Regime Switching Predictability

A regime switching model captures the **instability in predictive regressions**

- This approach recognizes that data-generating processes for stock returns are subject to parameter instability
- One strategy for modeling breaks is based on Markov switching predictive models:
$$r_{t+1} = \alpha_{S_{t+1}} + \beta'_{S_{t+1}} x_t + \sigma_{S_{t+1}} u_{t+1}$$

where S_{t+1} is a first-order Markov-switching process representing the state of the economy

- S_{t+1} can take integer values between 1 and m , corresponding to the state of the economy, where the transition between states is governed by an matrix with typical element,

$$p_{ij} = \Pr(S_t = i | S_{t-1} = j) \quad (i, j = 1, \dots, m)$$

- Since the state of the economy is unobservable, the model cannot be estimated using conventional regression techniques
- The EM algorithm can be used to estimate the parameters via MLE and make inferences regarding the state of the economy

Regime Switching Predictability

A regime switching model captures the **instability in predictive regressions**

- This approach recognizes that data-generating processes for stock returns are subject to parameter instability
- One strategy for modeling breaks is based on Markov switching predictive models:

$$r_{t+1} = \alpha_{S_{t+1}} + \beta'_{S_{t+1}} x_t + \sigma_{S_{t+1}} u_{t+1}$$

where S_{t+1} is a first-order Markov-switching process representing the state of the economy

- S_{t+1} can take integer values between 1 and m , corresponding to the state of the economy, where the transition between states is governed by an matrix with typical element,

$$p_{ij} = \Pr(S_t = i | S_{t-1} = j) \quad (i, j = 1, \dots, m)$$

- Since the state of the economy is unobservable, the model cannot be estimated using conventional regression techniques
- The EM algorithm can be used to estimate the parameters via MLE and make inferences regarding the state of the economy

Regime Switching Predictability

Regime switching predictability models implicitly diversify across state-specific predictability patterns

- A forecast of r_{t+1} for $m = 2$ is given by

$$\hat{r}_{t+1}^{MS} = \Pr(S_{t+1} = 1|I_t)(\hat{\alpha}_{1,t} + \hat{\beta}'_{1,t}x_t) + \Pr(S_{t+1} = 2|I_t)(\hat{\alpha}_{2,t} + \hat{\beta}'_{2,t}x_t)$$

- Implicitly, we diversify across forecasts from two possible regimes
- In periods where it is difficult to determine next period's state, approximately equal weights are placed on the two regimes
- If there is strong evidence based on data through t on one regime, much more weight is placed on that regime forecast
 - Guidolin and Timmermann (2007, JEDC) estimate a multivariate 4-state MS model for U.S. aggregate stock and bond returns via MLE
 - Characterizing the four states as “crash,” “slow growth,” “bull,” and “recovery,” they present statistical evidence favoring four regimes
 - Real-time asset allocation decisions yield some utility gains relative to asset allocation decisions based on constant expected excess returns

Regime Switching Predictability

The historical average is sufficient during “normal” times, while economic variables provide useful signals during recessions

- Henkel et al. (2011, JFE) estimate a two-regime MSVAR includes the DP ratio, short-term interest rates, term spread, and default spread
- They estimate their model via Bayesian methods and find that in-sample predictability is highly concentrated during recessions
- Stock return forecasts outperform the historical average benchmark in terms of MSFE and OOS return predictability is concentrated during cyclical downturns
 - Instead of parameters switching among a small number of states via MS, **time-varying parameter (TVP) models** allow for parameters to continuously evolve, so that each period is viewed as a new regime
- Dangl and Halling (2009, JFE) specify the following TVP model:

$$\begin{aligned} r_{t+1} &= \alpha_t + \beta_t' x_t + \varepsilon_{t+1}, & \varepsilon_t &\sim N(0, \sigma^2), \\ \begin{pmatrix} \alpha_t \\ \beta_t \end{pmatrix} &= \begin{pmatrix} \alpha_{t-1} \\ \beta_{t-1} \end{pmatrix} + \begin{pmatrix} w_{1,t} \\ w_{2,t} \end{pmatrix} & w_t &\sim N(0, W_t). \end{aligned}$$

Regime Switching Predictability

The historical average is sufficient during “normal” times, while economic variables provide useful signals during recessions

- Henkel et al. (2011, JFE) estimate a two-regime MSVAR includes the DP ratio, short-term interest rates, term spread, and default spread
- They estimate their model via Bayesian methods and find that in-sample predictability is highly concentrated during recessions
- Stock return forecasts outperform the historical average benchmark in terms of MSFE and OOS return predictability is concentrated during cyclical downturns
 - Instead of parameters switching among a small number of states via MS, **time-varying parameter (TVP) models** allow for parameters to continuously evolve, so that each period is viewed as a new regime
- Dangl and Halling (2012, JFE) specify the following TVP model:

$$\begin{aligned} r_{t+1} &= \alpha_t + \beta_t' x_t + \varepsilon_{t+1}, & \varepsilon_t &\sim N(0, \sigma^2), \\ \begin{pmatrix} \alpha_t \\ \beta_t \end{pmatrix} &= \begin{pmatrix} \alpha_{t-1} \\ \beta_{t-1} \end{pmatrix} + \begin{pmatrix} w_{1,t} \\ w_{2,t} \end{pmatrix} & w_t &\sim N(0, W_t), \end{aligned}$$

Utility-Based Metrics

- The intercept and slope coefficients in the predictive regression evolve as (driftless) random walks
- The TVP model can be estimated using the Kalman filter and MLE
- A return forecast based on the TVP model is $\hat{r}_{t+1}^{TVP} = \hat{\alpha}_{t,t} + \hat{\beta}'_{t,t}x_t$.
- Dangl and Halling (2012) employ Bayesian estimation methods and find that forecasts significantly outperform the historical average
- The **OOS gains are concentrated during recessions**
- In line with Leitch and Tanner (1991, AER), researchers frequently analyze return forecasts with profit- or utility-based metrics
- In these exercises, stock return forecasts serve as inputs for ad hoc trading rules or asset allocation decisions derived from expected utility maximization problems
- A leading utility-based metric for analyzing U.S. equity premium forecasts is the average utility gain for a mean-variance investor:

$$a_{i,t} = \left(\frac{1}{\gamma} \right) \left(\frac{\hat{r}_{i,t+1}}{\hat{\sigma}_{i,t+1}^2} \right)$$

Utility-Based Metrics

- The intercept and slope coefficients in the predictive regression evolve as (driftless) random walks
- The TVP model can be estimated using the Kalman filter and MLE
- A return forecast based on the TVP model is $\hat{r}_{t+1}^{TVP} = \hat{\alpha}_{t,t} + \hat{\beta}'_{t,t}x_t$.
- Dangl and Halling (2009) employ Bayesian estimation methods and find that forecasts significantly outperform the historical average
- The OOS gains are concentrated during recessions
- In line with Leitch and Tanner (1991, AER), researchers frequently analyze return forecasts with profit- or utility-based metrics
- In these exercises, stock return forecasts serve as inputs for ad hoc trading rules or asset allocation decisions derived from expected utility maximization problems
- A leading utility-based metric for analyzing U.S. equity premium forecasts is the average utility gain for a mean-variance investor:

$$a_{i,t} = \left(\frac{1}{\gamma}\right) \left(\frac{\hat{r}_{i,t+1}}{\hat{\sigma}_{t+1}^2}\right)$$

Utility-Based Metrics

The difference in realized utility gains from mean-variance strategies can be interpreted as a portfolio management fee

- Over the evaluation period, the investor realizes the average utility

$$\hat{v}_i = \hat{\mu}_i - 0.5\gamma\hat{\sigma}_i^2$$

where we use sample mean and variance of realized ptf. returns

- If the investor instead relies on historical average (using the same variance forecast), she allocates the portfolio share

$$a_{0,t} = \left(\frac{1}{\gamma}\right) \left(\frac{\bar{r}_{t+1}}{\hat{\sigma}_{t+1}^2}\right)$$

and, over the forecast evaluation period, realizes the average utility

$$\hat{v}_0 = \hat{\mu}_0 - 0.5\gamma\hat{\sigma}_0^2$$

- The difference between \hat{v}_i and \hat{v}_0 represents the utility gain (certainty equivalent return, CER) from using the predictive regression in place of the historical average forecast
- CER can be interpreted as the portfolio management fee that an investor would be willing to pay to have access to the information in the predictive regression forecast

Utility-Based Metrics

The difference in realized utility gains from mean-variance strategies can be interpreted as a portfolio management fee

- Over the evaluation period, the investor realizes the average utility

$$\hat{v}_i = \hat{\mu}_i - 0.5\gamma\hat{\sigma}_i^2$$

where we use sample mean and variance of realized ptf. returns

- If the investor instead relies on historical average (using the same variance forecast), she allocates the portfolio share

$$a_{0,t} = \left(\frac{1}{\gamma}\right) \left(\frac{\bar{r}_{t+1}}{\hat{\sigma}_{t+1}^2}\right)$$

and, over the forecast evaluation period, realizes the average utility

$$\hat{v}_0 = \hat{\mu}_0 - 0.5\gamma\hat{\sigma}_0^2$$

- The difference between $v\text{-hat}_i$ and $v\text{-hat}_0$ represents **the utility gain (certainty equivalent return, CER) from using the predictive regression in place of the historical average forecast**
- CER can be interpreted as the portfolio management fee that an investor would be willing to pay to have access to the information in the predictive regression forecast

An Empirical Application

- A number of papers detect sizable utility gains for MV investors who rely on equity premium forecasts based on economic variables
- Consider an application based on forecasting the monthly U.S. equity premium using updated data from Goyal and Welch (2008) spanning 1926:12–2010:12
- Fourteen popular economic variables serve as candidate predictors:
 - Log dividend-price ratio: log of a twelve-month moving sum of dividends paid on the S&P 500 index minus the log of stock prices
 - Log dividend yield]: log of a twelve-month moving sum of dividends minus the log of lagged stock prices
 - Log earnings-price ratio
 - Log dividend-payout ratio
 - Stock variance
 - Book-to-market ratio for the DJIA
 - Net equity expansion: ratio of a twelve-month moving sum of net equity issues by NYSE-listed stocks to the total end-of-year market capitalization of NYSE stocks

An Empirical Application

- A number of papers detect sizable utility gains for MV investors who rely on equity premium forecasts based on economic variables
- Consider an application based on forecasting the monthly U.S. equity premium using updated data from Goyal and Welch (2008) spanning 1926:12–2010:12
- Fourteen popular economic variables serve as candidate predictors:
 - Log dividend-price ratio: log of a twelve-month moving sum of dividends paid on the S&P 500 index minus the log of stock prices
 - Log dividend yield: log of a twelve-month moving sum of dividends minus the log of lagged stock prices
 - Log earnings-price ratio
 - Log dividend-payout ratio
 - Stock variance
 - Book-to-market ratio for the DJIA
 - Net equity expansion: ratio of a twelve-month moving sum of net equity issues by NYSE-listed stocks to the total end-of-year market capitalization of NYSE stocks

An Empirical Application

- Treasury bill rate
 - Long-term yield on government bonds
 - Long-term return on government bonds
 - Term spread: long-term yield minus the Treasury bill rate.
 - Default yield spread: difference between BAA- and AAA-rated corporate bond yields
 - Default return spread: long-term corporate bond return minus the long-term government bond return
 - Inflation (INFL): calculated from the CPI (all urban consumers)
- Use 1926:12–1956:12 as the initial in-sample estimation period, so that we compute out-of-sample forecasts for 1957:01–2010:12 (648 observations)
 - Forecasts employ a recursive (or expanding) estimation window
 - In the figure, the solid line in each panel depicts the difference in cumulative square errors for the historical average forecast vis-à-vis the predictive regression forecast

An Empirical Application

- Treasury bill rate
 - Long-term yield on government bonds
 - Long-term return on government bonds
 - Term spread: long-term yield minus the Treasury bill rate.
 - Default yield spread: difference between BAA- and AAA-rated corporate bond yields
 - Default return spread: long-term corporate bond return minus the long-term government bond return
 - Inflation (INFL): calculated from the CPI (all urban consumers)
- Use 1926:12–1956:12 as the initial in-sample estimation period, so that we compute out-of-sample forecasts for 1957:01–2010:12 (648 observations)
 - Forecasts employ a recursive (or expanding) estimation window
 - In the figure, the solid line in each panel depicts the difference in cumulative square errors for the historical average forecast vis-à-vis the predictive regression forecast

An Empirical Application

Overall

Economic variable	R_{OS}^2 (%)	Δ (annual %)
-------------------	----------------	---------------------

Panel A: Unrestricted predictive regression forecasts

$\log(DP)$	-0.05 [0.10]	0.87
$\log(DY)$	-0.37 [0.07]	1.18
$\log(EP)$	-1.88 [0.28]	0.57
$\log(DE)$	-2.04 [0.97]	-0.44
<i>SVAR</i>	0.32 [0.17]	-0.11
<i>BM</i>	-1.74 [0.31]	-0.72
<i>NTIS</i>	-0.91 [0.41]	-0.21
<i>TBL</i>	-0.01 [0.09]	1.53
<i>LTY</i>	-1.17 [0.12]	1.29
<i>LTR</i>	-0.08 [0.20]	0.57
<i>TMS</i>	0.06 [0.16]	1.15
<i>DFY</i>	-0.04 [0.59]	0.29
<i>DFR</i>	-0.01 [0.38]	0.39
<i>INFL</i>	-0.09 [0.50]	0.34

An Empirical Application

Economic variable	Overall		Expansion		Recession	
	R_{OS}^2 (%)	Δ (annual %)	R_{OS}^2 (%)	Δ (annual %)	R_{OS}^2 (%)	Δ (annual %)
<i>Panel A: Unrestricted predictive regression forecasts</i>						
$\log(DP)$	-0.05 [0.10]	0.87	-1.24 [0.42]	-1.47	2.41 [0.00]	11.87
$\log(DY)$	-0.37 [0.07]	1.18	-2.28 [0.40]	-1.98	3.56 [0.00]	16.17
$\log(EP)$	-1.88 [0.28]	0.57	-2.21 [0.31]	-0.41	-1.20 [0.38]	4.99
$\log(DE)$	-2.04 [0.97]	-0.44	-1.26 [0.80]	0.06	-3.67 [0.97]	-2.73
<i>SVAR</i>	0.32 [0.17]	-0.11	-0.02 [0.50]	-0.37	1.01 [0.16]	1.08
<i>BM</i>	-1.74 [0.31]	-0.72	-2.56 [0.44]	-2.01	-0.04 [0.28]	5.12
<i>NTIS</i>	-0.91 [0.41]	-0.21	0.50 [0.03]	0.72	-3.82 [0.94]	-4.71
<i>TBL</i>	-0.01 [0.09]	1.53	-0.84 [0.30]	0.24	1.71 [0.10]	7.58
<i>LTY</i>	-1.17 [0.12]	1.29	-2.37 [0.38]	-0.21	1.32 [0.11]	8.38
<i>LTR</i>	-0.08 [0.20]	0.57	-0.85 [0.63]	-0.35	1.52 [0.05]	4.74
<i>TMS</i>	0.06 [0.16]	1.15	-0.40 [0.34]	0.01	1.00 [0.09]	6.49
<i>DFY</i>	-0.04 [0.59]	0.29	-0.06 [0.64]	0.00	-0.01 [0.48]	1.48
<i>DFR</i>	-0.01 [0.38]	0.39	0.12 [0.25]	0.15	-0.28 [0.48]	1.53
<i>INFL</i>	-0.09 [0.50]	0.34	0.10 [0.22]	0.19	-0.48 [0.66]	1.16

An Empirical Application

Overall

Economic variable	R_{OS}^2 (%)	Δ (annual %)
-------------------	----------------	---------------------

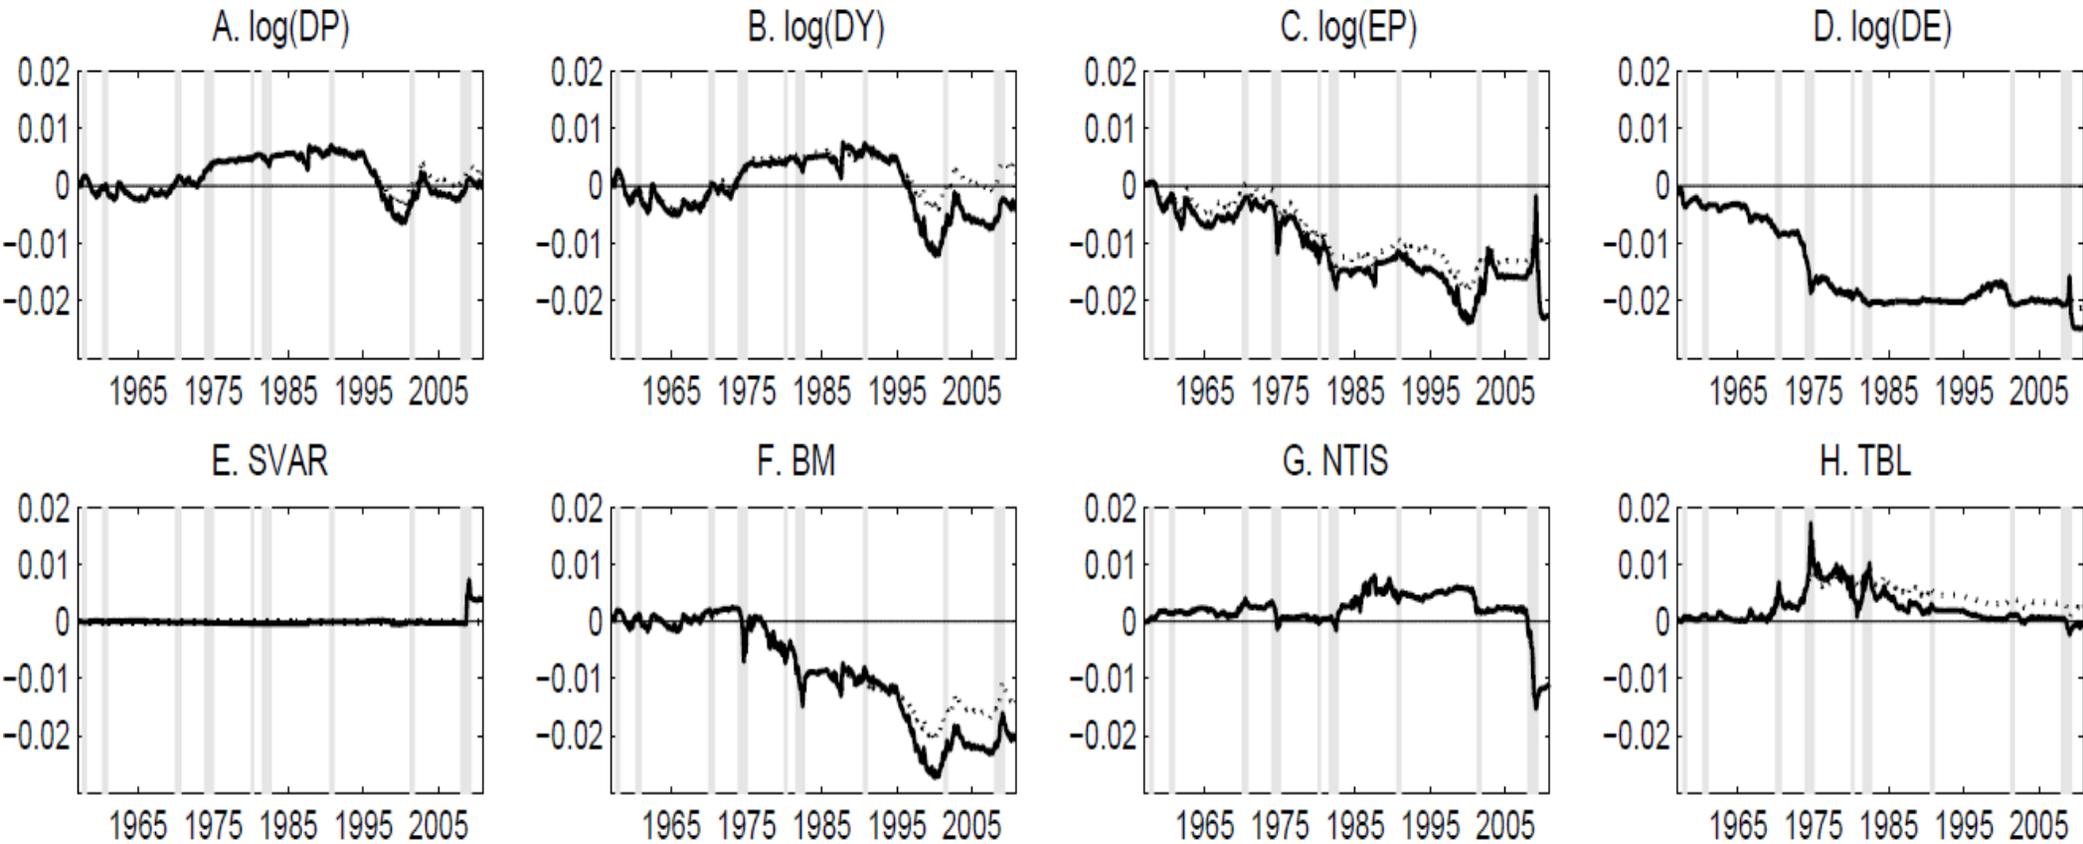
Panel B: Predictive regression forecasts with Campbell and Thompson (2008) restrictions

$\log(DP)$	0.15 [0.07]	0.87
$\log(DY)$	0.17 [0.04]	1.18
$\log(EP)$	-0.82 [0.24]	0.57
$\log(DE)$	-1.74 [0.98]	0.01
<i>SVAR</i>	0.00 [—]	-0.26
<i>BM</i>	-1.17 [0.29]	-0.72
<i>NTIS</i>	-0.91 [0.41]	-0.21
<i>TBL</i>	0.21 [0.10]	1.53
<i>LTY</i>	-0.01 [0.09]	1.29
<i>LTR</i>	0.22 [0.12]	0.63
<i>TMS</i>	0.12 [0.15]	1.15
<i>DFY</i>	-0.01 [0.50]	0.29
<i>DFR</i>	-0.16 [0.49]	0.56
<i>INFL</i>	-0.06 [0.46]	0.34

An Empirical Application

Economic variable	Overall		Expansion		Recession	
	R_{OS}^2 (%)	Δ (annual %)	R_{OS}^2 (%)	Δ (annual %)	R_{OS}^2 (%)	Δ (annual %)
<i>Panel B: Predictive regression forecasts with Campbell and Thompson (2008) restrictions</i>						
$\log(DP)$	0.15 [0.07]	0.87	-0.92 [0.38]	-1.47	2.36 [0.00]	11.87
$\log(DY)$	0.17 [0.04]	1.18	-1.33 [0.40]	-1.98	3.26 [0.00]	16.17
$\log(EP)$	-0.82 [0.24]	0.57	-1.19 [0.30]	-0.41	-0.06 [0.32]	4.99
$\log(DE)$	-1.74 [0.98]	0.01	-1.19 [0.78]	0.01	-2.88 [0.98]	0.00
<i>SVAR</i>	0.00 [-]	-0.26	0.00 [-]	-0.36	0.00 [-]	0.18
<i>BM</i>	-1.17 [0.29]	-0.72	-1.68 [0.40]	-2.01	-0.13 [0.30]	5.12
<i>NTIS</i>	-0.91 [0.41]	-0.21	0.50 [0.03]	0.72	-3.82 [0.94]	-4.71
<i>TBL</i>	0.21 [0.10]	1.53	-0.25 [0.27]	0.24	1.16 [0.10]	7.58
<i>LTY</i>	-0.01 [0.09]	1.29	-0.67 [0.29]	-0.21	1.36 [0.07]	8.38
<i>LTR</i>	0.22 [0.12]	0.63	-0.47 [0.52]	-0.36	1.64 [0.03]	5.15
<i>TMS</i>	0.12 [0.15]	1.15	-0.42 [0.37]	0.01	1.23 [0.06]	6.49
<i>DFY</i>	-0.01 [0.50]	0.29	-0.03 [0.55]	0.00	0.01 [0.45]	1.48
<i>DFR</i>	-0.16 [0.49]	0.56	0.09 [0.27]	0.15	-0.68 [0.66]	2.51
<i>INFL</i>	-0.06 [0.46]	0.34	0.10 [0.22]	0.19	-0.38 [0.63]	1.16

An Empirical Application



An Empirical Application

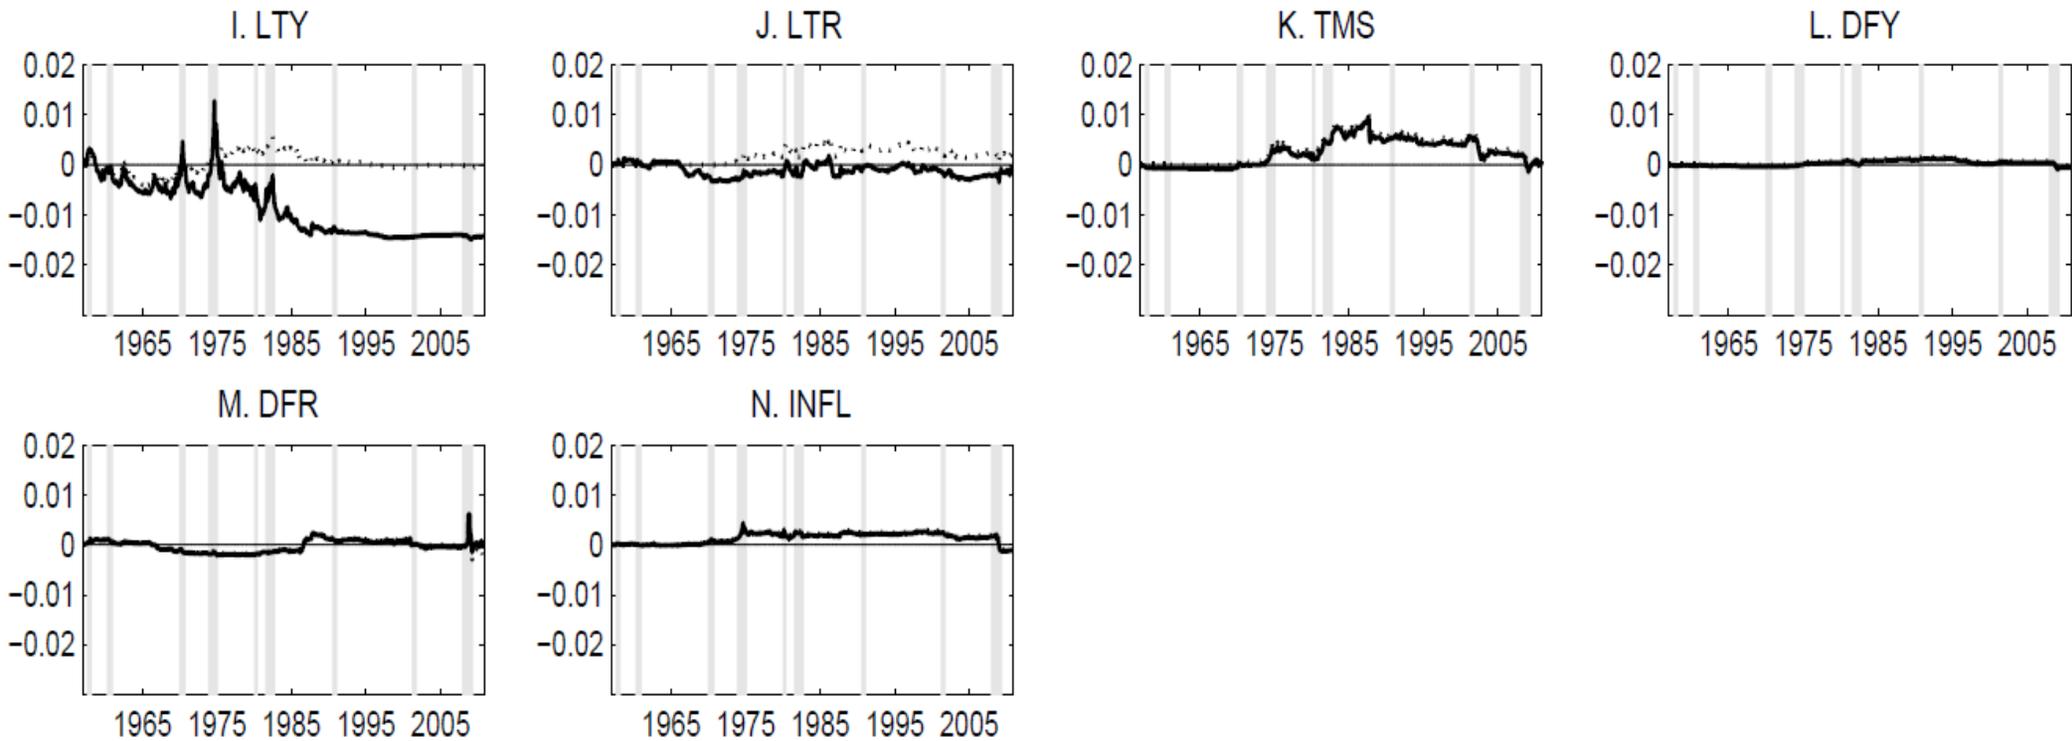


Figure 1 Differences in cumulative square forecast errors, U.S. equity premium out-of-sample forecasts based on individual economic variables, 1957:01–2010:12. Solid (dotted) lines in each panel depict differences in cumulative square forecast errors for the historical average forecast relative to the unrestricted predictive regression forecast (predictive regression forecast with Campbell and Thompson (2008) restrictions imposed) based on the economic variable given in the panel heading. Vertical bars depict NBER-dated recessions.

An Empirical Application

Method	Overall		Expansion		Recession	
	R_{OS}^2 (%)	Δ (annual %)	R_{OS}^2 (%)	Δ (annual %)	R_{OS}^2 (%)	Δ (annual %)
<i>Panel A: Monthly forecasts</i>						
Kitchen sink	-8.43 [0.42]	0.24	-9.41 [0.68]	-1.60	-6.38 [0.28]	8.94
SIC	-5.61 [0.99]	-1.77	-5.80 [1.00]	-3.24	-5.21 [0.79]	5.00
POOL-AVG	0.44 [0.03]	1.25	0.12 [0.21]	0.41	1.10 [0.01]	5.14
POOL-DMSFE	0.51 [0.02]	1.52	0.08 [0.25]	0.40	1.39 [0.01]	6.76
Diffusion index	0.68 [0.01]	1.65	-1.00 [0.27]	-1.46	4.15 [0.00]	16.42
Sum-of-the-parts	0.93 [0.01]	2.47	0.29 [0.13]	0.27	2.24 [0.01]	12.87
<i>Panel B: Quarterly forecasts</i>						
Kitchen sink	-29.82 [0.76]	0.01	-25.89 [0.43]	-2.07	-34.97 [0.85]	8.22
SIC	-19.16 [0.95]	-3.64	-15.27 [0.59]	-4.40	-24.28 [0.97]	-1.31
POOL-AVG	0.73 [0.09]	1.74	0.40 [0.20]	0.98	1.15 [0.13]	4.42
POOL-DMSFE	1.09 [0.05]	2.27	0.23 [0.26]	0.68	2.22 [0.05]	8.41
Diffusion index	2.10 [0.01]	3.40	-3.58 [0.23]	-1.26	9.56 [0.00]	21.81
Sum-of-the-parts	2.02 [0.03]	3.61	0.24 [0.22]	0.63	4.35 [0.03]	15.34

Notes: R_{OS}^2 measures the percent reduction in mean square forecast error for the forecasting method given in the first column relative to the historical average benchmark forecast. Brackets report p -values for the Clark and West (2007) *MSFE-adjusted* statistic corresponding to $H_0: R_{OS}^2 \leq 0$ against $H_A: R_{OS}^2 > 0$. Average utility gain (Δ) is the portfolio management fee (in annualized percent return) that an investor with mean-variance preferences and risk aversion coefficient of five would be willing to pay to have access to the forecasting method given in the first column relative to the historical average benchmark forecast. 0.00 indicates less than 0.005. R_{OS}^2 statistics and average utility gains are computed for the entire 1957:01–2010:12 forecast evaluation period and separately for NBER-dated expansions and recessions.