Predictive Regression and Robust Hypothesis Testing: Predictability Hidden by Anomalous Observations

Olivier Scaillet

University of Geneva and Swiss Finance Institute

olivier.scaillet@unige.ch

Joint work with Lorenzo Camponovo and Fabio Trojani

Motivation: Literature

- A large literature has investigated whether stock returns can be predicted by economic variables such as, e.g., the dividend yield, labor income, interest rate proxies, variance risk premia,...
    - Rozeff (1984), Fama and French (1988), Campbell and Shiller (1988), Nelson and Kim (1993), Goetzmann and Jorion (1995), Kothari and Shanken (1997), Campbell and Yogo (2006), Ahmihud et al. (2008), Jansson and Moreira (2006), Polk et al. (2006), Santos and Veronesi (2006), Bollerslev et al. (2009), among others.

$\Rightarrow$ Fundamental question for asset pricing models in general.

Motivation: Predictive Regression

- The econometric approach is mostly based on a predictive regression model, such as:

$$y_t = \alpha + \beta x_{t-1} + u_t, \quad (1)$$
$$x_t = \mu + \rho x_{t-1} + v_t, \quad (2)$$

where y_t denotes the stock return and x_{t-1} is a potential predictive economic variable for y_t.

- No predictability hypothesis:

$$\mathcal{H}_0 : \beta_0 = 0 \ ,$$

where β_0 is the true value of the unknown parameter β.

Motivation: Issues

- Endogenous (nearly integrated) predictor x_{t-1} with correlated innovations (u_t, v_t):
  - Several inferential biases arise using standard asymptotic theory.
- Approaches to the problem:
  - Bias-corrected methods: Stambaugh (1999); Amihud, Hurvich and Wang (2008).
  - Near-to-unit-root asymptotics: Lewellen (2004); Torous, Valkanov and Yan (2004); Campbell and Yogo (2006).
  - Resampling methods: Wolf (2000); Choi and Chue (2007); Ang and Bekaert (2007).
- These tests lead to ambiguous results in a number of cases.

Contributions (1): The Lack of Resistance to Anomalous Data

- Monte Carlo simulations show that above testing procedures are dramatically non-resistant to small fractions of anomalous observations in the data.

  ⇒ Anomalous observations can dramatically decrease the test ability to reject the null of no predictability.

- Using the concept of breakdown point, we characterize the resistance properties of (time-series) resampling-based tests of predictability to anomalous observations.

  ⇒ We theoretically confirm the dramatic lack of resistance to outliers of these tests.

Contributions (2): Robust Tests of Predictability

- We develop a novel class of general bootstrap and subsampling tests for time series, which are resistant to anomalous observations in the data.
  - The robust tests are applicable, e.g., to predictive regressions with nearly-integrated multiple predictors and non-linear predictive relations.
  - Monte Carlo analysis confirms their theoretical reliability and the improved resistance over conventional approaches.
  - Beside the sharp robustness improvements, our approach also reduces the resampling computational costs.

Contributions (3): Empirical Analysis

- Robust assessment of the recent evidence on return predictability.
  - ⇒ Single- and multi-predictor models, with well-known predictive variables: dividend yield; variance risk premium; labor income.
- Robust unambiguous evidence of return predictability.
  - ⇒ Robust to choice of predictive variables, sampling frequencies, and prediction horizons.
- Unfrequent clusters of anomalous observations, e.g, in the NASDAQ bubble and during the recent financial crisis.
  - ⇒ Less than 5% of the data, causing the difference in evidence w.r.t. conventional test approaches.
  - ⇒ September 2008 (Lehman Brothers default), October 1987 (Black Monday), September 2001 (Terrorist Attack) are the most influential observations.

Outline

(1) The Lack of Robustness.

- Monte Carlo Analysis.
- Quantile Breakdown Point.

(2) The Robust Approach.

- Robust Resampling Methods.
- Monte Carlo Analysis.

(3) Empirical Analysis.

- Single-Predictor Model.
- Multi-Predictor Model.

(4) Conclusions.

The Lack of Robustness

- Monte Carlo simulations.
- Degree of resistance to anomalous observations of recent tests of predictive relations:
    - Bias-corrected methods; Amihud, Hurvich and Wang (2008).
    - Near-to-unit-root asymptotics; Campbell and Yogo (2006).
    - Bootstrap and subsampling tests; Wolf (2000), Choi and Chue (2007).

Monte Carlo Setting

- $N = 1,000$ random samples $\big((y_1, x_0), \ldots, (y_n, x_{n-1})\big)$, of size $n = 180$, according to predictive regression (1)-(2), with:

$$\beta_0 = 0,\ 0.05,\ 0.1,\ 0.15.$$

- To study the robustness, we also consider replacement outliers random samples $\big((\tilde{y}_1, x_0), \ldots, (\tilde{y}_n, x_{n-1})\big)$, where

$$\tilde{y}_t = (1 - p_t) y_t + p_t \cdot y_{3max},$$

with $Y_{3max} = 3 \cdot \max(y_1, \ldots, y_n)$ and p_t is an iid $0 - 1$ random sequence such that $P[p_t = 1] = 4\%$.

- We test the null hypothesis $\mathcal{H}_0 : \beta_0 = 0$.

Power Curves

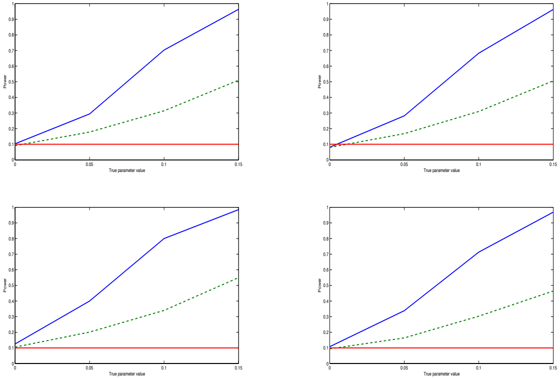

We plot the proportion of rejections of the null hypothesis $\mathcal{H}_0 : \beta_0 = 0$, when the true parameter value $\beta_0 \in [0, 0.15]$. In the top panels we consider Bias-corrected methods (left) and near-to-unit-root asymptotics (right). In the bottom panels we consider the bootstrap (left) and the subsampling (right). We simulate noncontaminated samples (straight line) and contaminated samples (dashed line).

The Lack of Robustness

- To overcome the lack of resistance to anomalous data, one needs to robustify these testing procedures.
  - This task may be hard to achieve for tests based on bias-corrected methods or near-to-unit-root asymptotics.
- By applying resampling methods to more robust statistics do we obtain resistance to anomalous data?
  - For iid data the answer is a definite NO! [e.g., Singh, 1998, Salibian-Barrera and Zamar, 2002, and Camponovo, Scaillet and Trojani, 2012].
  - What about resampling methods for time series data?

Robustness Analysis: Breakdown Point

- $X_{(n)} := (X_1, \ldots, X_n)$: time series sample.
- $T_n := T(X_{(n)})$: statistic with breakdown point $0 < b \leq 0.5$.
- What is the breakdown point?
  - b is the smallest fraction of outliers in sample $X_{(n)}$ such that $T(X_{(n)}) \to \pm\infty$.
  - Measure of (global) robustness; e.g., Donoho and Huber (1983).
- Examples:
  - Sample mean $T_n = \frac{1}{n}\sum_{i=1}^{n} X_i \qquad (b = 1/n)$.
  - Sample median $T_n = med(X_1, \ldots, X_n) \qquad (b = 0.5)$.

Robustness Analysis: Bootstrap/Subsampling Quantile

- $X^*_{(k)} = (X^*_1, \ldots, X^*_k)$ is a subsampling/bootstrap random sample $k = m, n$.

  $\Rightarrow$ The subsampling/bootstrap distribution of $T^*_{n,k} := T(X^*_{(k)})$ consistently estimates the sampling distribution of T_n.

- Given $t \in (0,1)$, the t-quantile of the resampling distribution is defined by:

$$Q^*_t(X_{(n)}) = \inf\{x | P^*(T^*_{n,k} \leq x) \geq t\}$$

  $\Rightarrow$ We characterize the robustness properties of resampling methods through the breakdown point of quantile Q^*_t.

Robustness Analysis: Quantile Breakdown Point

- The breakdown point of the t-quantile Q_t^* is defined by:

$$b_t = \frac{1}{n} \cdot \left[\inf_{\{1 \leq p \leq n/2\}} \left\{ p \middle| \exists Z_p^\zeta \in \mathcal{Z}_p^\zeta : Q_t^*(X_{(n)} + Z_p^\zeta) = +\infty \right\} \right].$$

  - $\mathcal{Z}_p^\zeta$ is the set of n-components outlier samples.
  - $p \in \mathbb{N}$, $\zeta \in \bar{\mathbb{R}}$ are the number and size of outliers in contaminated sample $X_{(n)} + Z_p^\zeta$.

- When breakdown occurs, subsampling/bootstrap inference becomes meaningless!

- We derive quantile breakdown point formulas for the subsampling and the bootstrap (in dependence of parameters n, m, b, t).

Robustness Analysis: Concrete BP Calculations

- Quantile breakdown subsampling and bootstrap.

$n = 120,\ b = 0.5$	$t = 0.95$	$t = 0.99$
Subsampling ($m = 10$)	≤ 0.0500	≤ 0.0500
Subsampling ($m = 15$)	≤ 0.0667	≤ 0.0667
Bootstrap ($m = 10$)	≤ 0.3333	≤ 0.2667
Bootstrap ($m = 15$)	≤ 0.3250	≤ 0.2167

Conclusions on the Lack of Robustness

- Conventional tests of predictability hypotheses are dramatically non-resistant to even small fractions of anomalous observations in the data.

  $\Rightarrow$ Anomalous observations dramatically decrease the test ability to reject the null of no predictability when it is wrong.

- We characterize theoretically the robustness properties of resampling methods through the concept of breakdown point.

  $\Rightarrow$ The theoretical results confirm the dramatic lack of robustness of resampling methods, highlighted in the Monte Carlo experiments.

  $\Rightarrow$ Resampling a robust statistic does not solve the problem.

Robust Approach: General Time Series Setting

- To overcome the robustness problem, we introduce a fast resampling approach for time series.

  ⇒ Fast procedures: e.g., Shao and Tu (1995), Davidson and McKinnon (1999), Hu and Kalbfleisch (2000), Andrews (2002), Goncalves and White (2004), Hong and Scaillet (2006).

  ⇒ Robust procedures for IID data: e.g., Salibian-Barrera and Zamar (2002) Salibian-Barrera, Van Aelst and Willems (2006)-(2007), Camponovo, Scaillet and Trojani (2012).

Robust Approach: The Idea

- $X_{(n)} = (X_1, \ldots, X_n)$ is a random sample on probability space $(\Omega, \mathcal{F}, P)$, indexed by a parameter of interest θ.

- M-estimator: estimator $\hat{\theta}_n$ of θ solves the estimating equation:

$$\psi_n(X_{(n)}, \hat{\theta}_n) := \frac{1}{n} \sum_{i=1}^{n} g(X_i; \hat{\theta}_n) = 0.$$

($T_n := \hat{\theta}_n$ is the statistic of interest.)

- Fast resampling approach: given bootstrap/subsampling sample $X^*_{(k)} = (X^*_1, \ldots, X^*_k)$, compute a suitable linear approximation of $\hat{\theta}^*_k$, instead of the solution of $\psi^*_k(X^*_{(k)}, \theta) = 0$ itself.

Robust Approach: Construction

- Fast Resampling approach.
  - Original sample:
    - θ_0 denotes the true parameter value.
    - $\hat{\theta}_n$ is the solution of $\psi_n(X_{(n)}, \hat{\theta}_n) = 0$.
    - $A_0 := -\big(\frac{\partial \psi_n(X_{(n)},\theta)}{\partial \theta}\big|_{\theta=\theta_0}\big)^{-1}$.
    - Taylor expansion around θ_0:

$$\hat{\theta}_n = \theta_0 + A_0 \psi_n(X_{(n)}, \theta_0) + o_p(1).$$

  - Bootstrap/subsampling sample:
    - $\hat{A}_n$ is a consistent estimator of A_0.
    - Instead of computing $\hat{\theta}^*_k$ as the solution of $\psi^*_k(X^*_{(k)}, \hat{\theta}^*_k) = 0$, consider the linear approximation:

$$\hat{\theta}^*_k \approx \hat{\theta}_n + \hat{A}_n \psi^*_k(X^*_{(k)}, \hat{\theta}_n).$$

Robust Approach: General Properties

- General approach that can be applied to a wide class of resampling methods (subsampling/bootstrap).

- Low computational costs:
  - The conventional approach applied to robust estimators becomes easily unfeasible; see, e.g. Salibian-Barrera and Zamar (2002).
  - The fast approach requires only consistent estimators $\hat{A}_n$ and $\hat{\theta}_n$.
  - Avoid recomputing estimators on each resampled data set.

- General consistency conditions:
  - Bootstrap: Goncalves and White (2004).
  - Subsampling: Hong and Scaillet (2006).

Robust Approach: Robustness Properties

- Robustness of the bootstrap/subsampling quantile:
  - $\hat{\theta}_n + \hat{A}_n \psi_k^*(X_{(k)}^*, \hat{\theta}_n)$ may diverge to infinity only when:
    - $\hat{\theta}_n$ diverges to infinity.
    - $\hat{A}_n$ is a singular matrix.
    - ψ_k^* is not bounded.
  - The quantile breakdown point depends only on the M-estimator $\hat{\theta}_n$, the estimating function ψ_n, and the matrix estimator $\hat{A}_n$.
- ((Robust $\hat{\theta}_n$)+ (Robust $\hat{A}_n$)+($\|\psi_k^*\| < c < \infty$))
  $\Rightarrow$ Robust resampling method!

Robust Approach: Predictive Regression Models

- Analysis of the predictive regression model.

- $z_{(n)} = ((y_1, x_0), \ldots, (y_n, x_{n-1}))$ is the observation sample and $w_t := (1, x_t)$.

- $\hat{\theta}_n^R$ is a robust M-estimator of $\theta := (\alpha, \beta)$, defined for $c > 0$ by:

$$\psi_{n,c}(z_{(n)}, \hat{\theta}_n^R) := \frac{1}{n} \sum_{t=1}^{n} g_c(y_t, w_{t-1}, \hat{\theta}_n^R) = 0,$$

where (bounded) estimating function g_c is given by:

$$g_c(y_t, w_{t-1}, \theta) = (y_t - \theta' w_{t-1}) w_{t-1} \cdot \min\left(1, \frac{c}{||(y_t - \theta' w_{t-1}) w_{t-1}||}\right).$$

Predictive Regression Model: Robustness Properties and Findings

- Let b_t be the t-quantile breakdown point of our robust bootstrap or subsampling approach. Then,

$$b_t = 0.5, \qquad t \in (0,1),$$

  i.e., we obtain a maximal resistance to anomalous observation for each t-quantile!

- We verify the theoretical robustness properties of our approach through Monte Carlo simulations.

Power Curves

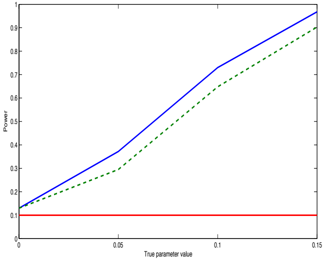

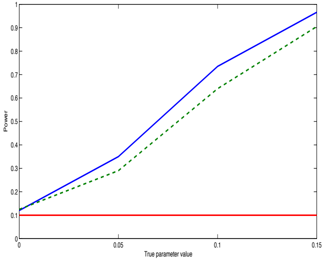

We plot the proportion of rejections of the null hypothesis $\mathcal{H}_0 : \beta_0 = 0$, when the true parameter value $\beta_0 \in [0, 0.15]$. We consider the robust bootstrap (left column) and the robust subsampling (right column). We simulate noncontaminated samples (straight line) and contaminated samples (dashed line).

Conclusions on the Robust Approach

- For the class of M-estimators, we develop a fast and robust resampling approach that implies quantile resistance to anomalous data points.

- The approach is general and can be applied to robustify a wide class of resampling approaches (subsampling/bootstrap).

- Our resampling approach inherits directly the robustness properties of the bounded estimating functions used in robust estimation.

- Using the robust approach, we obtain robust resampling tests of predictability hypotheses in predictive regression settings.

Empirical Analysis

- We consider both single-predictor and two-predictor models. In particular, we study the predictive ability of variables, such as:
  - The Dividend Yield,
  - Variance Risk Premium proxies,
  - Labor Income proxies,

for predicting future stock market returns.

Single-Predictor Model

- Data: S&P 500 index data (1871-2008) from Shiller (2000).
  - The one-period real total return is defined as

$$R_t = (P_t + d_t)/P_{t-1},$$

    - P_t is the end of month real stock price.
    - d_t is the real dividend paid during month t.
  - The annualized dividend series D_t is defined as

$$D_t = d_t + (1 + r_t)d_{t-1} + \cdots + (1 + r_t)\ldots(1 + r_{t-10})d_{t-11},$$

  where r_t is the one-month Treasury-bill rate.

- Simple predictive regression model:

$$\ln(R_t) = \alpha + \beta\left(\frac{D_{t-1}}{P_{t-1}}\right) + \epsilon_t.$$

Confidence Intervals: Nonrobust Tests

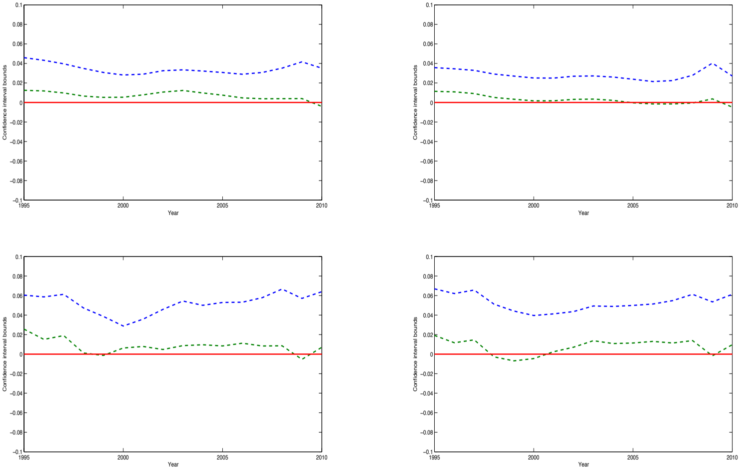

90% confidence intervals for parameter β_0. In the top panels we present the evidence for the Bias-corrected method (left) and for the near-to-unit-root asymptotics (right). In the bottom panels we consider the bootstrap (left) and the subsampling (right).

Confidence Intervals: Robust Tests

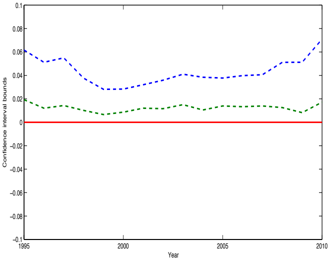

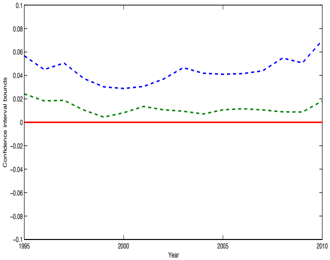

90% confidence intervals for parameter β. We consider the robust bootstrap (left) and the robust subsampling (right).

Huber Weights: The Anomalous Observations

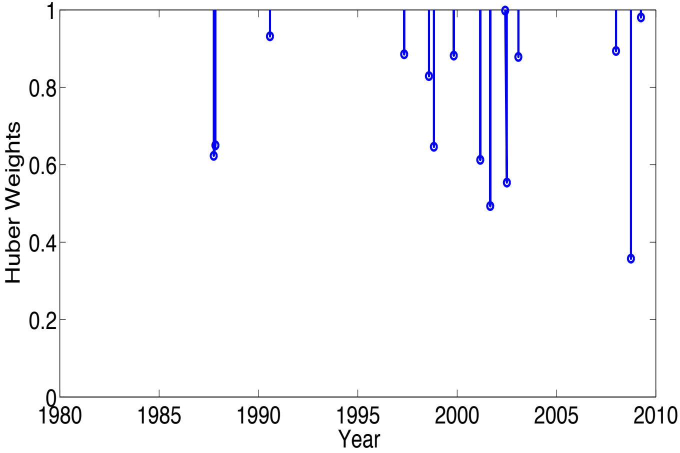

Huber weights for the predictive regression model in the period 1980-2010.

Properties of Anomalous Observations

- In the whole sample period 1980-2010, the proportion of anomalous observations is less than 4.8%.

- Most influential data points:
  - October 2008: Lehman Brothers default on September 15 2008.
  - October 2001: Terrorist attack on September 11 2001.
  - November 1987: Black Monday on October 19 1987.

Two-Predictor Model: Bollerslev et al. (2009)

- We study the predictive ability of dividend yields and variance risk premium proxies for the two-predictor model in Bollerslev, Tauchen and Zhou (2009):

$$\frac{1}{k}\ln(R_{t+k,t}) = \alpha + \beta_1 \ln\left(\frac{D_t}{P_t}\right) + \beta_2 VRP_t + \epsilon_{t+k,t},$$

where $k = 4$ and

- $\ln(R_{t+k,t}) := \ln(R_{t+1}) + \cdots + \ln(R_{t+k})$.
- $VRP_t := IV_t - RV_t$.

Confidence Intervals: Dividend Yields

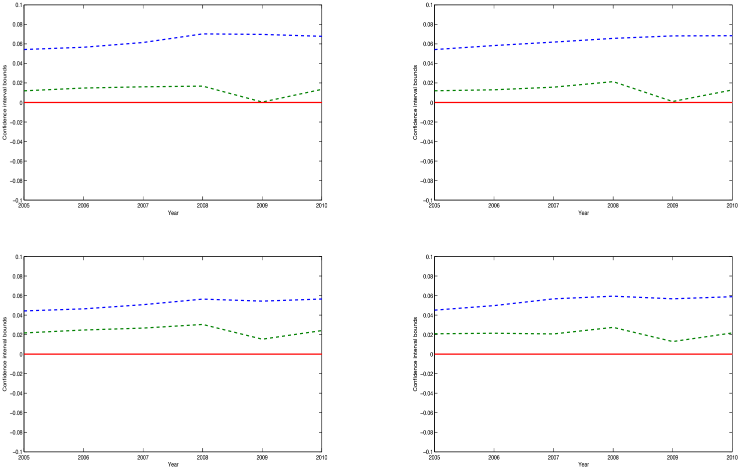

90% confidence intervals for parameter β_1. In the top panels, we consider the nonrobust bootstrap (left) and the nonrobust subsampling (right). In the bottom panels, we consider the robust bootstrap (left) and the robust subsampling (right).

Confidence Intervals: Variance Risk Premia

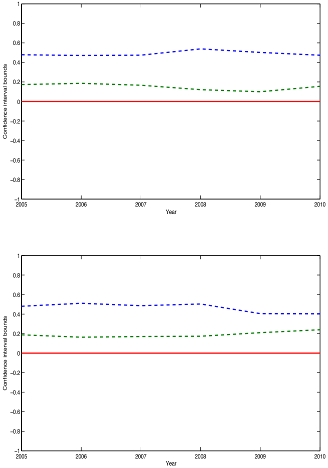

90% confidence intervals for parameter β_2. In the top panels, we consider the nonrobust bootstrap (left) and the nonrobust subsampling (right). In the bottom panels, we consider the robust bootstrap (left) and the robust subsampling (right).

Huber Weights: The Anomalous Observations

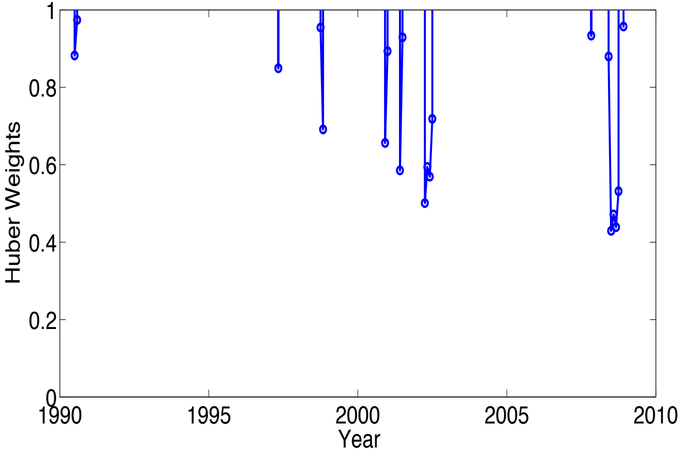

Huber weights for the predictive regression model in the period 1990-2010.

Properties of Anomalous Observations

- In the whole sample period 1990-2010, the proportion of anomalous observations is less than 5.6%.

- Most influential observations:
  - October 2008: Lehman Brothers default on September 15 2008.
  - August 2002: Dot-Com bubble collapse.
  - October 2001: Terrorist attack on September 11 2001.

Two-Predictor Model: Santos and Veronesi (2006)

- We study the predictive ability of dividend yield and labor income proxies for the two-predictor model in Santos and Veronesi (2006):

$$\ln(R_t) = \alpha + \beta_1 \ln\left(\frac{D_{t-1}}{P_{t-1}}\right) + \beta_2 s_{t-1} + \epsilon_t,$$

  where $s_{t-1} = w_{t-1}/C_{t-1}$ is the share of labor income to consumption.

- We consider quarterly returns on the value weighted CRSP index, which includes NYSE, AMEX, and NASDAQ stocks.

Confidence Intervals: Dividend Yields

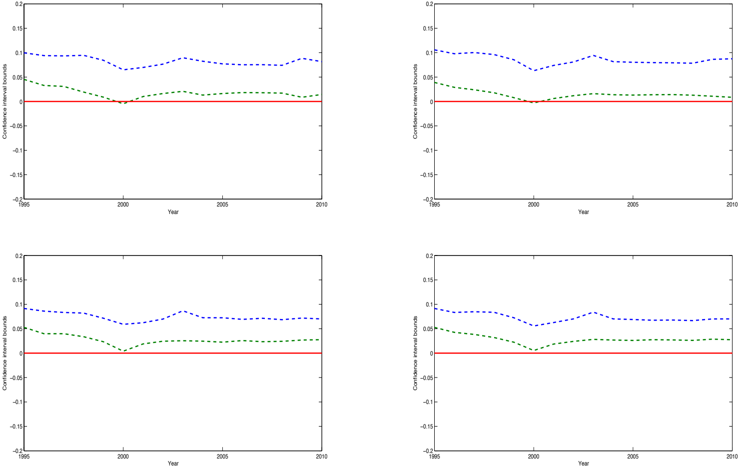

90% confidence intervals for parameter β_1. In the top panels, we consider the nonrobust bootstrap (left) and the nonrobust subsampling (right). In the bottom panels, we consider the robust bootstrap (left) and the robust subsampling (right).

Confidence Intervals: Labor Income

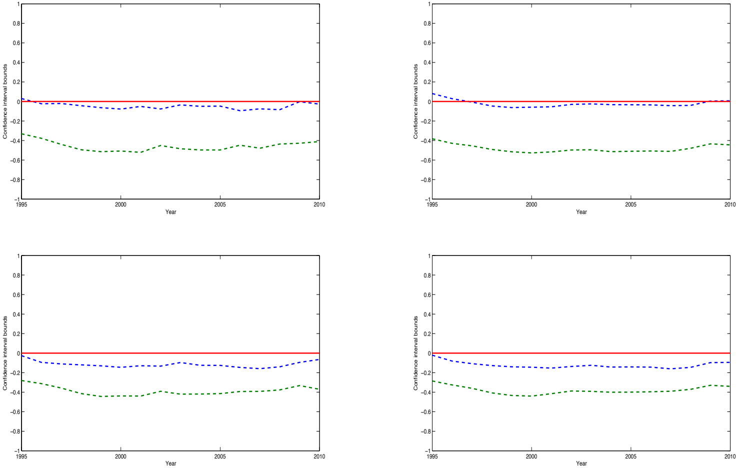

90% confidence intervals for parameter β_2. In the top panels, we consider the nonrobust bootstrap (left) and the nonrobust subsampling (right). In the bottom panels, we consider the robust bootstrap (left) and the robust subsampling (right).

Huber Weights: The Anomalous Observations

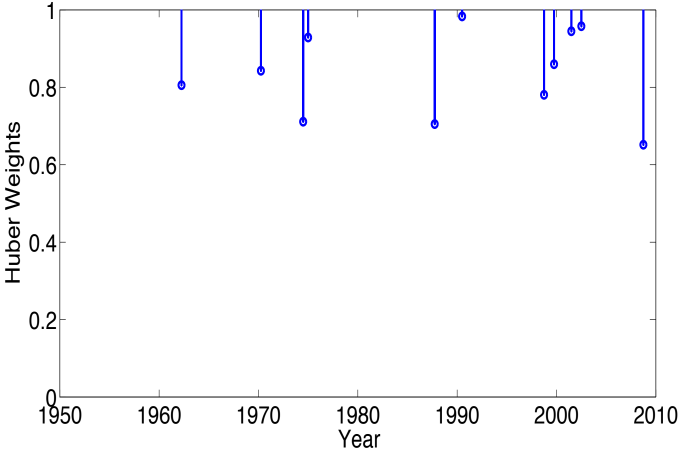

Huber weights for the predictive regression model in the period 1950-2010.

Properties of Anomalous Observations

- In the whole dataset sample period 1950-2010, the proportion of anomalous observations is less than 4.2%.

- Most influential observations:
  - 4Q 2008 - 1Q 2009: Lehman Brothers default on September 15 2008.
  - 4Q 1987 - 1Q 1988: Black Monday on October 19 1987.
  - 4Q 1973 - 1Q 1974: Oil Crisis.

Conclusions on the Empirical Study

- The dividend yield is a robust predictive variable of market returns.
  - $\Rightarrow$ Unambiguous robust significance at 5% significance level, for each subperiod, sampling frequency and forecasting horizon.
  - $\Rightarrow$ No or weak predictive ability in several subperiods using conventional tests, due to small fractions of anomalous data.
- The variance risk premium is a robust predictive variable of future market returns at quarterly forecasting horizons.
- Labor income is a robust predictive variable of market returns at quarterly frequencies.
  - $\Rightarrow$ Unambiguous robust significance at 5% significance level in all subperiods.
  - $\Rightarrow$ No or weak predictive ability in several subperiods using conventional tests, due to small fractions of anomalous data.

Final Conclusions

(1) Conventional tests of predictability hypotheses are dramatically non-resistant to small fractions of anomalous observations.
   - ⇒ Anomalous data points reduce the test ability to reject the null when it is violated.

(2) Novel class of bootstrap and subsampling procedures implying inference with desirable robustness/resistance properties.

(3) Robust resampling tests consistently detect predictability structures that are uncovered by conventional methods.

(4) In the data, predictability by the dividend yield, the variance risk premium or labor income is unambiguous and robust.
   - ⇒ The different findings of conventional tests are caused by unfrequent anomalous observations.
   - ⇒ Anomalous observations, linked to particular events not modeled by the predictive regression, hide predictability features.