Forecasting stock returns with a structural estimate of risk aversion

Ian Dew-Becker

Harvard

11/8/2011

Introduction

- Thanks for coming
- Section of my job market paper on forecasting returns
- Most of this is covered in the appendix

The basic result

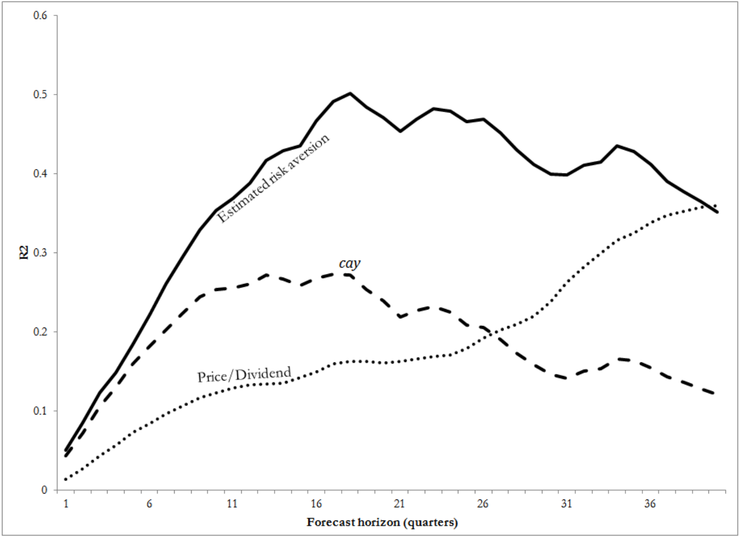

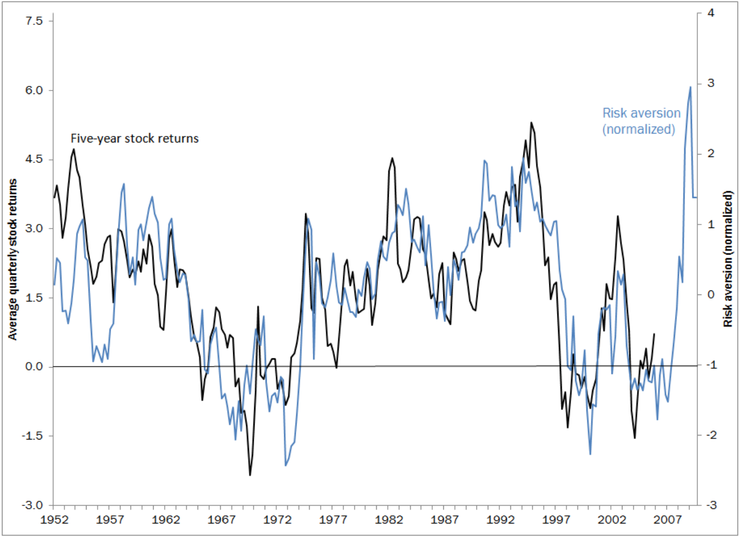
Five-year stock returns
Risk aversion
(normalized)
Average quarterly stock returns
Risk aversion (normalized)
7.5
6.0
4.5
3.0
1.5
0.0
-1.5
-3.0
4
3
2
1
0
-1
-2
-3
1952
1957
1962
1967
1972
1977
1982
1987
1992
1997
2002
2007

Introduction

- Estimated risk aversion comes from a model
- It forecasts stock returns embarrassingly well
- Discuss today whether this is right

Outline

1. The model
2. Basic forecasting regressions
3. Out-of-sample results
4. Robustness to parameter choices

Household preferences

- Epstein–Zin preferences with time-varying risk aversion

$$
\begin{aligned}
V_t &= \left\{ (1-\beta)\, C_t^{1-\rho} + \beta \left[E_t V_{t+1}^{1-\alpha_t} \right]^{\frac{1-\rho}{1-\alpha_t}} \right\}^{\frac{1}{1-\rho}} \\
\alpha_t &= (1-\phi)\,\bar{\alpha} + \phi \alpha_{t-1} - \lambda \left(\log V_t^A - E_{t-1} \log V_t^A \right)
\end{aligned}
$$

  Motivated from a model with habit formation
- If we can measure the history of V_t^A, can measure α_t.

Measuring risk aversion

- Define the *Wealth Portfolio* as the asset that pays C_t as its dividend
- For Epstein-Zin preferences,

$$W_t = V_t^{1-\rho} C_t^{\rho} / (1-\beta)$$

intuitively, $W = V / MU_C$

$$V_t \propto W_t \left(W_t / C_t \right)^{\frac{\rho}{1-\rho}}$$

$\Longrightarrow$ With V_t, we can calculate risk aversion

$$\alpha_{t+1} = (1-\phi)\,\bar{\alpha} + \phi \alpha_t - \lambda \log \left(V_{t+1}^A / E_t V_{t+1}^A \right)$$

- Check for whether the assumed process for α_t fits returns

Measuring risk aversion

$$\alpha_t = (1-\phi)\bar{\alpha} + \phi\alpha_{t-1} - \lambda\left(\frac{1}{1-\rho}\Delta wc_{t-j} + \Delta c_{t-j} - \overline{\Delta w}\right)$$

$$\alpha_t = \bar{\alpha} - \lambda\sum_{j=0}^{\infty}\phi^j\left(\frac{1}{1-\rho}\Delta wc_{t-j} + \Delta c_{t-j} - \overline{\Delta w}\right)$$

Calibrate ρ and ϕ; $\rho = 2/3$; $\phi = 0.96$ (half-life $=$ 17 quarters)
Ignore $\bar{\alpha}$ and λ

Measuring risk aversion

$$\alpha_t = (1-\phi)\,\bar{\alpha} + \phi\alpha_{t-1} - \lambda\left(\Delta w_t + \frac{\rho}{1-\rho}\Delta wc_t - \overline{\Delta w}\right)$$

- How do we measure $\overline{\Delta w}$?
- If the model is true, full-sample mean should work
    - Could also regress Δw_t on α_t or wc_t
- Or, for $E_{t-1}\Delta w_t$ use sample mean of Δw_j for $j < t$
    - Should this matter versus full-sample mean?
    - Just a level shift in α_t

Relationship with the wealth/consumption ratio

$$\alpha_t = \bar{\alpha} - \lambda \sum_{j=0}^{\infty} \phi^j \left(\frac{1}{1-\rho} \Delta wc_{t-j} + \Delta c_{t-j} - \overline{\Delta w} \right)$$

$$\phi \sum_{j=0}^{\infty} \phi^j \Delta x_{t-j} = x_t - (1-\phi) \sum_{k=0}^{\infty} \phi^k x_{t-k}$$

- α_t overweights recent innovations to *wc*
- Similar to detrending wc_t (effect is small)

Data

- Two methods:
  - Measure consumption and wealth
  - Using innovations to the level of technology as innovations to V_t

Lettau and Ludvigson (2001)

- Financial wealth from flow of funds; labor income from the BEA
- Suppose

$$\begin{aligned} w_t &= \chi f_t + (1-\chi) h_t \\ h_t &= i_t + \mu_t \end{aligned}$$

- w_t wealth; f_t financial wealth; h_t human wealth; i_t labor income; μ_t stationary error
  - Everything in logs
- Observe f_t and i_t,

$$w_t = \chi f_t + (1-\chi) i_t + (1-\chi) \mu_t$$

Lettau and Ludvigson (2001)

- Under balanced growth (or looser theories) consumption and wealth are cointegrated
  - Stationary consumption growth and discount rates
- Suppose $c_{ND,t} = \eta c_t + \zeta_t$

$$\begin{aligned} c_t - w_t &= \varepsilon_t \\ \eta^{-1} c_{ND,t} - \chi f_t - (1-\chi)\, i_t &= \underbrace{\varepsilon_t + (1-\chi)\,\mu_t + \eta^{-1}\zeta_t}_{\text{stationary}} \end{aligned}$$

- Estimate the cointegrating vector for c_t, f_t, and i_t
- χ is then identified (as long as η is unconstrained)
- Lettau and Ludvigson (2001) call residual *cay*

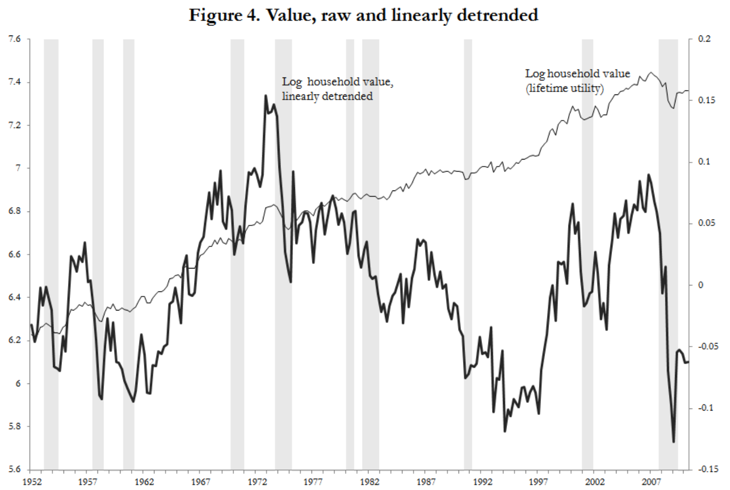

Note: Household value (lifetime utility) is measured using data on wealth and consumption from Sydney Ludvigson's website. The thin line is the absolute level of value (left-hand axis); the thick line is value linearly detrended. Both variables are measured in logs. Grey bars are NBER-dated recessions

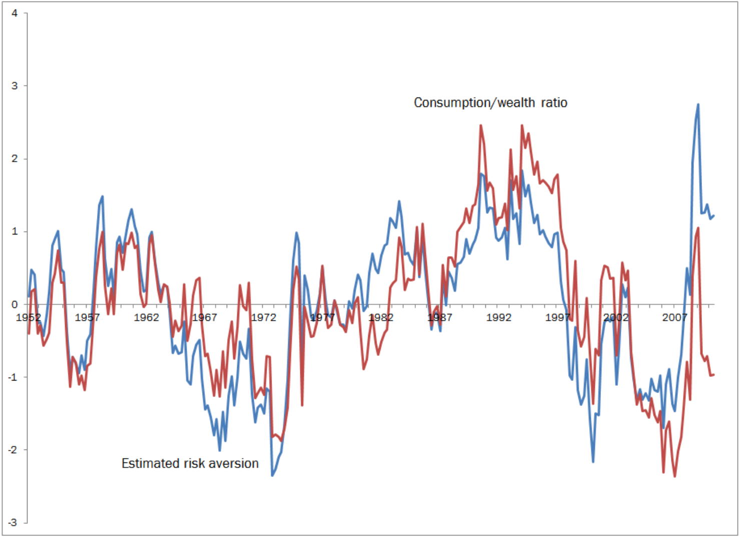
4
3
2
1
0
-1
-2
-3
1952
1957
1962
1967
1972
1977
1982
1987
1992
1997
2002
2007
Consumption/wealth ratio
Estimated risk aversion

Out-of-sample testing

- Compare my forecast and *cay* to a constant-mean forecast
- Under the null (no predictability), RMSE is lower for constant mean than anything else
- Need to adjust for this difference
- Clark and McCracken (2001, 2005); Clark and West (2007)

Out-of-sample testing

- Estimate VECM on data 0 to t
- Construct $\hat{\alpha}$
- On data up to date t, estimate $b_{1,t}$ and $\varepsilon_{t,j}$ as

$$
\begin{aligned}
r_{k,k+j} &= b_0 + b_{1,t}\hat{\alpha}_k \\
\varepsilon_{t,j}^{\hat{\alpha}} &\equiv r_{t,t+j} - b_0 - b_{1,t}\hat{\alpha}_t
\end{aligned}
$$

- Use $\varepsilon_{t,j}^{\hat{\alpha}}$ to construct OOS test statistics, $1 \leq j \leq 20$

Out-of-sample testing

- Denote $\varepsilon_{t,j}^{iid}$ as error from a constant-mean model for returns
- $\varepsilon_{t,j}^{iid}$ is encompassed by $\varepsilon_{t,j}^{\hat{\alpha}}$
- Test statistic is

$$\begin{aligned} f_{t,j} &= \left(\varepsilon_{t,j}^{iid}\right)^2 - \left(\left(\varepsilon_{t,j}^{\hat{\alpha}}\right)^2 - \left(\varepsilon_{t,j}^{iid} - \varepsilon_{t,j}^{\hat{\alpha}}\right)^2\right) \\ &= 2\varepsilon_{t,j}^{iid}\left(\varepsilon_{t,j}^{iid} - \varepsilon_{t,j}^{\hat{\alpha}}\right) \\ OOS_j &= \frac{(T-R)^{-1/2}\sum_{t=R}^{T} f_{t,j}}{S\left(f_{t,j}\right)} \end{aligned}$$

where $S\left(f_{t,j}\right)$ is a long-run variance estimate

Out-of-sample test statistics

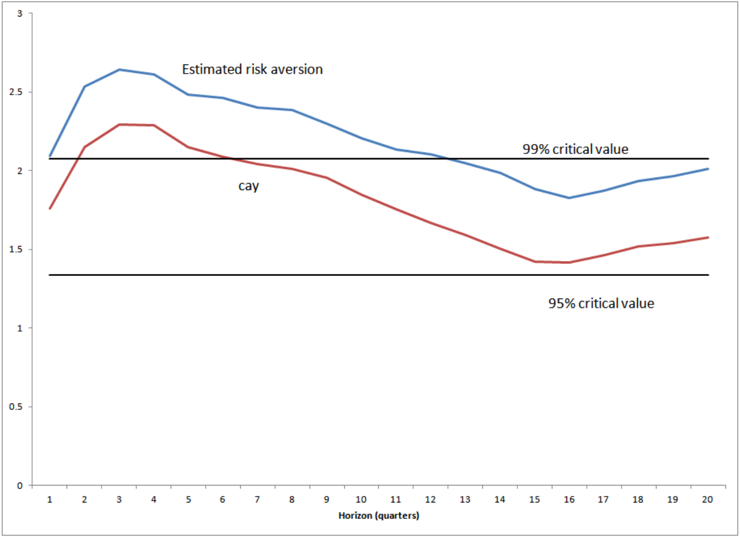

Stambaugh bias

$$r_t = b_0 + b_1 \alpha_{t-1} + \varepsilon_t$$

- Finite sample results for OLS assume $E_t \varepsilon_t \alpha_{t-j} = 0 \ \forall \ j$
- We can only assume $E_t \varepsilon_t \alpha_{t-j} = 0 \ \forall \ j > 0$
- Innovations to r_t are correlated with innovations to α_t
  - Good news lowers risk aversion
- Asymptotically, OLS is consistent; biased in small samples
  - See also Mankiw and Shapiro (1986); Cavanagh, Elliott, and Stock (1995)

Bootstrapping

- Suppose consumption and wealth follow random walks; returns i.i.d.

$$\begin{aligned} \Delta f_t &= \varepsilon_{f,t} \\ \Delta i_t &= \varepsilon_{i,t} \\ \Delta c_t &= \varepsilon_{c,t} \end{aligned}$$

1. Draw bootstrap samples of $\{\varepsilon_{f,t}, \varepsilon_{i,t}, \varepsilon_{c,t}, r_t\}$
2. Estimate vecm
3. Construct α_t^{boot}
4. Calculate OOS test statistics

- Hard test: lots of estimation error

Bootstrapped critical values

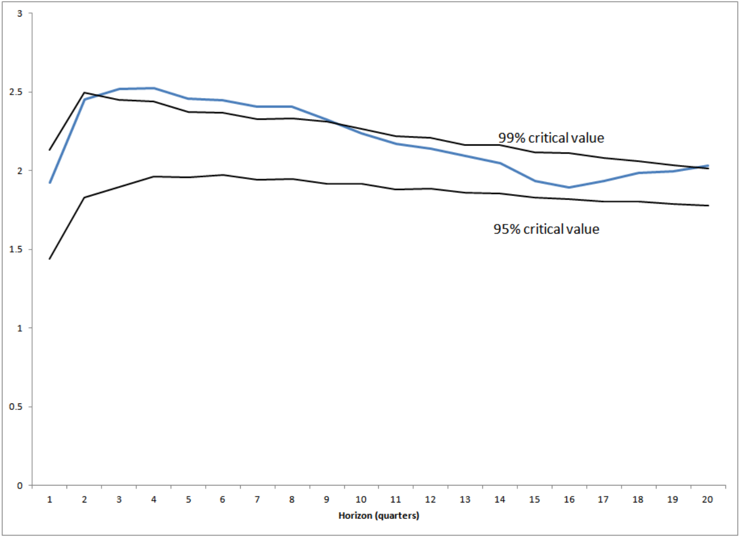

Estimated risk aversion versus cay

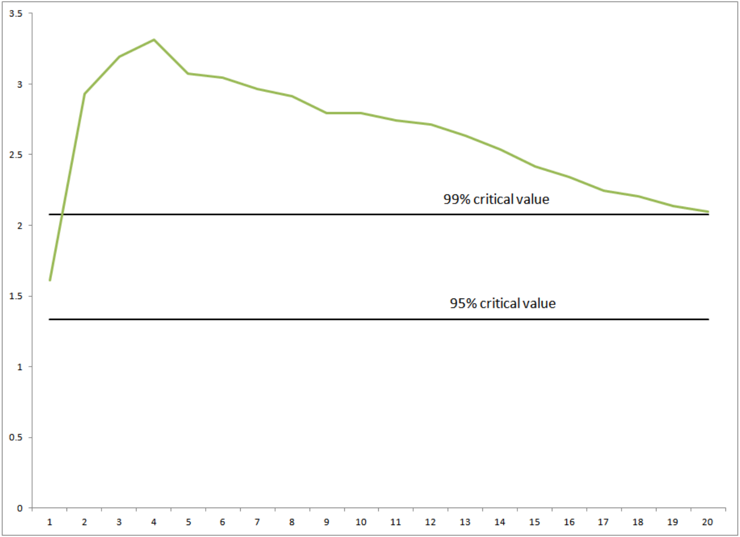

Robustness

- Everything up to here fixes ϕ and ρ exogenously
- ρ from micro evidence and aggregate facts
- ϕ might have look-ahead bias
- Look at R^2 for varying values of ρ and ϕ

Varying persistence of risk aversion

phi^4	1q	5q	10q	20q
0.7	**0.04**	**0.17**	0.31	0.40
0.75	**0.04**	0.19	0.33	0.43
0.8	0.05	0.21	0.36	0.47
0.85	*0.06*	*0.23*	*0.39*	*0.49*
0.9	0.06	0.25	0.41	0.50
0.95	0.07	0.27	0.43	0.48
0.99	0.07	0.27	0.42	0.43
cay	0.05	0.18	0.27	0.28

- Robust to any beliefs about ϕ

Varying EIS

EIS	1q	5q	10q	20q
10	**0.05**	0.21	0.35	0.50
3	0.05	0.22	0.38	0.52
1.5	*0.06*	*0.023*	*0.39*	*0.49*
1.1	0.05	0.20	0.33	0.38
2/3	**0.01**	**0.03**	**0.05**	**0.01**
1/5	**0.01**	**0.04**	**0.06**	**0.19**
1/10	**0.01**	**0.06**	**0.09**	**0.24**
cay	0.05	0.18	0.27	0.28

- Need EIS>1
- Common result in this literature

Risk aversion measured from technology growth