Comments on "Volatility, the Macroeconomy and Asset Prices, by Ravi Bansal, Dana Kiku, Ivan Shaliastovich, Amir Yaron, (BKSY) and "An Intertemporal CAPM with Stochastic Volatility" John Y. Campbell, Stefano Giglio, Christopher Polk, and Robert Turley (CGPT)

John H. Cochrane

NBER Asset Pricing Meeting, April 13 2012

These are two, good hard papers. I have 20 minutes to discuss them both. An apology to the authors, and an ad to the audience: there's lots of good stuff in here that I won't have time to mention.

These papers focus on two very popular ingredients of macro-finance models: Volatility, and recursive utility / long-run risks. That makes me in some sense a singularly unqualified discussant. I've been a sceptic of both ingredients, and haven't worked in either area. If they had rare disasters too, we'd have a trifecta of discussant ignorance!

So, the most useful thing I can do, I think, is to express some of that scepticism, and then discuss how the papers do or don't move my (fairly diffuse) priors. Not being an expert at this kind of sausage, however, I won't be able to comment that knowledgeably on the details of the sausage machines in both papers.

Some sceptic's questions

My first sceptic's question:

- Do long run risks in recursive utility really matter?

We can express this question by looking at the discount factor, or by looking at the expected return - beta representation,

$$
\begin{aligned}
(E_{t+1} - E_t)\ \ln m_{t+1} &= -\gamma \Delta c_{t+1} + (1-\gamma) \sum_{j=1}^{\infty} \beta^j \Delta c_{t+j} \\
E(R^e_{t+1}) &= cov(R^e_{t+1}, \Delta c_{t+1})\gamma + cov(R^e_{t+1}, \sum_{j=1}^{\infty} \beta^j \Delta c_{t+j})\,(\gamma - 1)
\end{aligned}
$$

Since the first term doesn't do much, the second term is almost everything in these models.

Are people really afraid of news of long-run future consumption growth *holding today's consumption growth constant?* Is this *bad news* the crucial state variable for "bad times?" Is it really true that the average investor says, "I'm afraid of holding stocks, because they do badly when we get bad news about consumption growth 20 years from now, even if we're in a boom now?

Keep in mind the crucial orthogonalization in this view of the world. I'm a big fan of long-run risks – the idea that the economy is driven by bad news about long-run future investment opportunities, taxes, regulations, etc. But in the usual view, bad news about the future causes investors to reduce consumption immediately; so it is reflected in lower consumption today or other measurable indicators of higher marginal utility today. Not in the recursive utility / long run risk world. Here, you can have bad news about the future, yet you do not reduce consumption today. The news itself lowers marginal utility. You eschew stocks that do badly along with this bad news, even if today remains boom times by any other measure.

Maybe, but I hope not. Like a time-varying risk of a rare disaster, the view that the major risk in the stock market is ephemeral news about the far-off future orthogonal to today's economic fortunes may be true, but it seems to lead us to an awfully untestable theory. We certainly can't make a lot of fun of "sentiment" explanations. Or "the Gods are angry" explanations either.

We can summarize a huge empirical adventure in finance by the idea that there is some sort of "second factor," having to do with business cycles, and that this second factor drives marginal utility. The conditional CAPM, habits, outside income (human capital), and so on and so on offer theories to account for these observations. Is bad news about the far off future really the "second factor," having nothing to do with current economic stress? Count me a sceptic.

Again, I'm a big fan of "long run" thinking. I think finance should arrive at a price vs. payoff focus, not a focus on one-month returns. I think a long-run perspective frees us from many of the timing problems in macro data. It's sensible that concerns about the long run affect today's prices and risk premiums. But worries about the long run should be reflected in current marginal utility and economic conditions. Where I have trouble is the idea that pure bad news, not reflected in current conditions, is the crucial state variable.

Moreover

- Is there really much variation in $\sum_{j=1}^{\infty} \beta^j \Delta c_{t+1+j}$ not reflected in current circumstances?

For the long-run risks / recursive utility model to matter, consumption must deviate substantially from a random walk. There must be important variation in news about long-run future consumption growth. "We can assume it's there and you can't prove it isn't" is fairly unsatisfying.

These papers are about volatility, leading to another set of questions:

- Does time-varying consumption volatility $\sigma_t(\Delta c_{t+1})$ (Kandel/Stambaugh; Recursive tradition) or time-varying risk aversion (say by habits, leverage,etc.) generate time-varying expected returns $\sigma_t \ln m_{t+1}$?

Recursive utility does not naturally generate time-varying risk premiums and return forecastability. These have been grafted on the models by supposing that the conditional *variance* of consumption varies as well, as Kandel and Stambaugh explored for power utility.

Well, does it? Is there *enough* variation in the conditional variance of consumption growth? Remember, DP predictability implies that expected returns vary as much as the mean expected return. $\sigma_t(\Delta c_{t+1})$ must vary over time by a factor of two!

(The alternative is that "risk aversion" translates a roughly homoskedastic macroeconomy to a highly heteroskedastic discount factor. Habits do this;

$$(E_{t+1} - E_t) \ln m_{t+1} = -\gamma \left[1 + \lambda \left(s_t\right)\right] \Delta c_{t+1}.$$

Leverage and many other mechanisms do so as well.

That's the fundamental question: Does our economy have time varying *risk*, or time-varying *risk aversion?)*

To get much action, these papers need even more; they need *persistent* variation in conditional variance. So, another sceptic's question:

- Is there much persistence to conditional volatility? Is there a lot of variation in $\sum_{j=1}^{\infty} \beta^j \sigma_t \left(\Delta c_{t+1+j}\right)$?

Answers – consumption volatility

BKSY address my questions about consumption.

So, how do BKSY address the big question: Is there much persistent variation in conditional volatility?

1. They construct the realized variance of monthly industrial production $RV_t = \frac{1}{12} \sum_{j=0}^{11} \Delta ip^2_{t-j/12}$
2. They forecast RV_{t+1} from a VAR using annual data from 1930. The important part of the VAR is the volatility forecast equation (bold = t>2)

	Δc_t	Δy_t	r_t	pd_t	RV_t	R^2
RV_{t+1}	**-0.007**	**-0.005**	0.001	-0.001	**0.291**	0.33
	(0.001)	**(0.001)**	(0.001)	(0.001)	**(0.004)**	

3. They assume that the conditional variance of consumption growth is the same as the VAR forecast of realized industrial production growth volatility, $\sigma_t^2(\Delta c_{t+1}) = E_t RV_{t+1}$! Here I really part company with the paper. The section is titled a "Volatility-Based Permanent Income Model." Well, the whole point of the permanent income model is that people smooth income shocks when choosing consumption, not that consumption always equals income. $\Delta c_t = (E_t - E_{t-1})\, r\beta \sum \beta^j y_{t+j}$. In what general equilibrium (what technology, with recursive preferences) do people choose consumption to have the same volatility as labor income?
4. Even so, the VAR shows little *persistence,* little *long-run volatility* $\sum_{j=1}^{\infty} \beta^j \sigma_t^2 \left(\Delta c_{t+1+j}\right)$. The major forecast of future volatility is current volatility, with a 0.291 autoregression, so volatility spikes die out quickly.

5. I have to stifle many other complaints about this VAR. Consumption and income are cointegrated! Mixing pre and postwar data is an obvious way to generate the appearance of time-varying volatility. The authors kindly sent me the data, and I was all hot to replicate the VAR, investigate the effect of prewar data, and so on. By midnight last night, I had gotten as far as plotting the data, which is where I will stop:

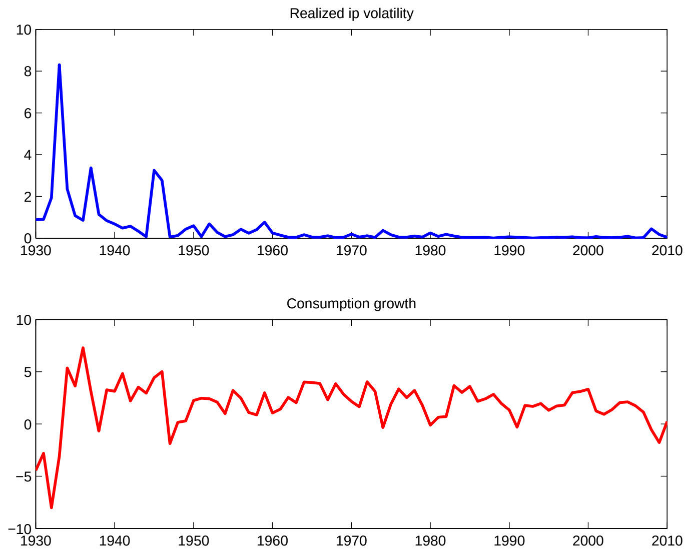

The top line is their realized volatility of industrial production growth. The bottom line is the consumption growth series. The assumption is that the top line measures the volatility of the bottom line. Once we look at postwar data (collected by completely different procedures, as Christina Romer pointed out), It's hard to see any changes in volatility. OK, maybe it's there, but this is supposed to be the central defining fact that generates factor of two time varying risk premiums. Let's just say my priors haven't moved that much.

To be fair, our purposes are different. They want a rough and ready calibration to feed in to a simulation model, an exercise I am not covering for lack of time. I'm looking for solid evidence to move a prior away from iid consumption growth. Lots of shortcuts are acceptable for the former purpose that are not for the latter. But in the same vein, an exercise designed for the former purpose does little to shift my priors on the latter.

CGPT, return volatility and the value effect.

CGPT examine the volatility of returns. It's not the same thing – I'd be willing to bet that most return volatility is "discount rate" volatility and thus not easily connected to consumption or other cashflow volatility.

CGPT's thesis is that persistent changes in return volatility are the missing "second factor," that the value premium in particular is earned for the exposure of value stock returns to changes in volatility.

As before, let me phrase the discussion in terms of my sceptic's doubts:

- Does volatility matter to long run investors? If so, why?

- Is volatility really the missing second factor for which hml is proxy? As opposed to, say, expected returns or other ICAPM state variables, or non-market businesses, non traded income, or human capital, as Fama and French (1996) suggested?

We can phrase these questions in the standard Merton portfolio theory / ICAPM framework: The optimal portfolio weight is driven by market-timing (conditional mean-variance) and hedging demands:

$$w_t = \frac{1}{\gamma_t}\Sigma_t^{-1}E_t(R_{t+1}^e) + \beta_{R,z}\frac{\eta_t}{\gamma_t}$$

where risk aversion and state variable aversion are value-function derivatives,

$$\gamma_t = -\frac{WV_{WW}(W,z)}{V_W(W,z)}; \eta_t = \frac{V_{Wz}(W,z)}{V_W(W,z)}$$

and expected returns correspond to multiple factors,

$$E_t(R_{t+1}^e) = cov_t(R_{t+1}^e, R_{t+1}^{em})\gamma_t^m - cov_t(R_{t+1}^e, z_{t+1})\eta_t^m$$

where m denotes market-wide averages and z a state variable, either for investment opportunities or for outside income. So

- Is η large for $z_t = \sigma_t^2$? Is $cov_t(hml, \sigma_t^2)$ large?

There's a deeper question bugging me: how in the world do we account for market portfolio weights given huge (and not perfectly correlated) variation in conditional means and conditional variances? Is it time-varying risk aversion, or neatly offsetting time-varying hedging demands? Or can conditional means and variances possibly scale perfectly?

However, CGPR is explicitly not about this question: They are about unconditional means and covariances and thus the average hedging component only.

Now, on the answer to the question, do long-term investors care about volatility? Here my priors lead to some doubt.

Here is one example. An earlier Campbell paper (also Jessica Wachter) showed that long-term bonds are the right riskless asset for a very risk averse long-run investors. I made this graph a while ago ("Discount rates") to illustrate: No matter what happens to the bond price or to interest rate volatility, the bond pays off at the desired horizon. From the Mertonian point of view, bonds are very desirable because bad returns covary strongly with rises in the state variable, yield.

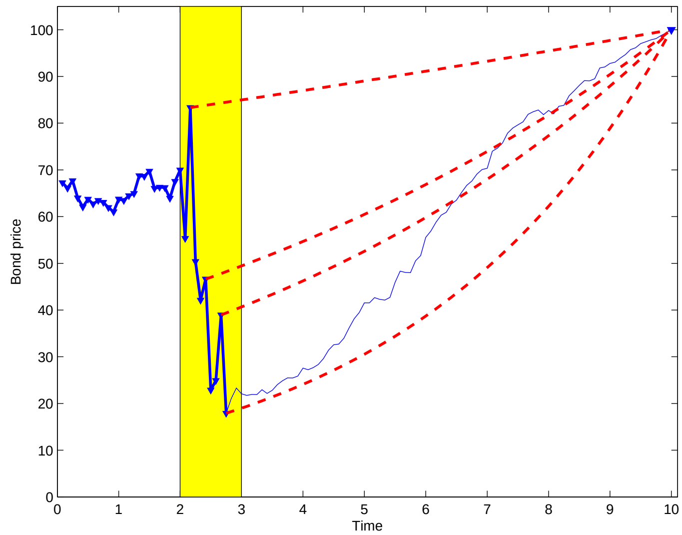

But in this example, *The long run bond investor does not care about volatility.* The right state variable is yield, and yield only. $V(W, yield)$, not $V(W, \sigma^2(yield))$. There is an important lesson: not everything that describes investment opportunities is automatically a state variable. How much does the stock investor really care about volatility per se? Of course it's a different problem, but this is a good cautionary example.

This is a great example of the importance of the long run. I conclude that thinking about long-run bond investing as one period mean/variance problem with state-variable hedging is nuts.

Really, the example is an instance of another sort-of theorem: "Merton state variables disappear from long-run investment problems." The right answer here is to buy and hold and don't do any dynamic anything!

A second useful example is the $Max\ EU(W_T)$ investor. This investor correctly ignores "short term" changes in volatility because it does not much affect $\sigma^2(R_{0\to T}) = \sigma^2(\sum r_{t+j})$

Thus "does the investor care" about volatility is clearly related to "is there a lot of persistent variation, variation in long-run volatility? " η is a value function derivative, and value functions depend on the environment as much as they do on preferences.

Ok, enough guessing. The right way to answer the question "do investors care about volatility" is to write down and solve a Merton problem, hopefully without losing track of economic intuition along the way.

CGPR do that, or at least the first part. They use the recursive utility machinery. To me, this machinery is so involved – but so familiar to its authors – that it has started to resemble a Rube Goldberg contraption.

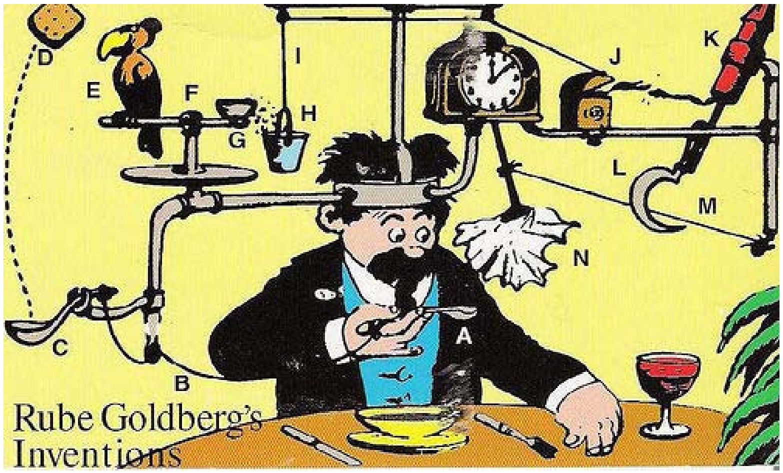

Let's just reflect on a few of the assumptions

The market return = the wealth portfolio (real estate, outside bushiness, human capital)

Consumption is total consumption – there are no other goods in utility index. Not even leisure.

The VAR captures all investor information. This is not a theory that handles "conditioning down" implications

A wealth of approximations

"Substituting out consumption," as the CAPM does and as most uses of Merton do. Is this an advantage or a disadvantage? Substituting out consumption from $c_t = kW_t$ allowed us to ignore the CAPM's prediction of 18% consumption volatility for a generation, until Mehra and Prescott rediscovered it.

At any rate, I think we should only whine softly about assumptions. The meat of the paper is that with all these assumptions, we do indeed derive that the representative investor cares about volatility – CGPR derive η(my V_{Wz}):

$$\begin{aligned} E_t r_{i,t+1} - r_{ft}.. &= \gamma\ cov_t(r_{i,t+1}, N_{CF,t+1}) + cov_t(r_{i,t+1,} N_{DR,t+1}) - \frac{1}{2}\omega(\gamma)\ cov_t[r_{i,t+1}, N_{V,t+1}] \quad (18) \\ N_{V,t+1} &= (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^{j-1} \sigma^2_{t+j}(r_{t+j}) \end{aligned}$$

1. This investor does cares about volatility. And there is a nice single-parameter restriction for how much – the coefficient ω in front of the covariance of return with volatility is a function of the risk aversion coefficient.

2. This investor cares about *long-run* volatility $\sum_{j=0}^{\infty} \rho^{j-1} \sigma^2_{t+j}(r_{t+j})$ my intuition was correct, and we have to believe in persistent volatility for this effect to matter.

Long-run volatility.

So, is there variation in *long-run* volatility that long-run investors / recursive utility care about? Here's what the paper does to document this issue.

First, they create realized return volatility RVAR. Then, they forecast RVAR with its own past and a vector of the usual subjects, including the value spread VS_t. (2 standard error successes in boldface)

	r_t	$RVAR_t$	PE_t	TY_t	DEF_t	VS_t	R^2
$RVAR_{t+1}$	-0.017	**0.30**	0.013	-0.002	**0.024**	0.001	10%
	0.021	**0.061**	0.007	0.002	**0.006**	0.008	

Next, they define EVAR as this fitted value foreast of realized volatility, keeping all the regressors,

$$EVAR_t = E_t(RVAR_{t+1})$$

This graph, from the paper, plots realized volatility and EVAR.

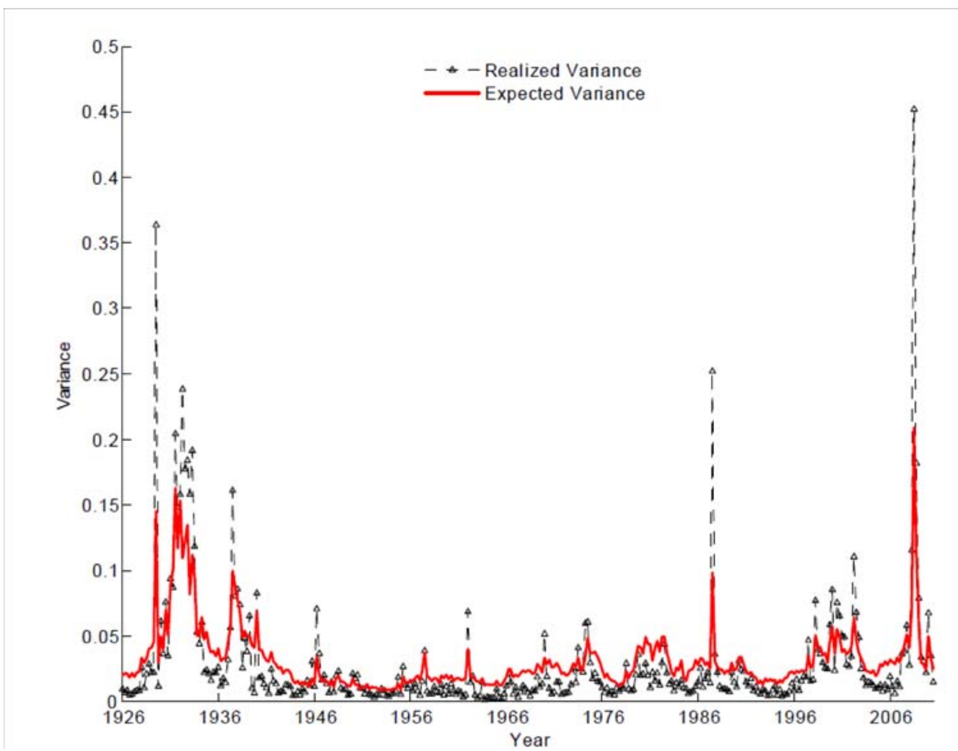

I read this graph as a dramatic *failure* in the quest to document persistent variation in forecastable volatility. The red line is supposed to forecast the black line *one period in advance*. The fact that they line up so well means that most of the forecast $EVAR_t = E_t RVAR_{t+1}$ is just 0.30 $\times RVAR_t$.

The graph shows basically that realized variance has a very short half life, and that its forecasts are strongly and spuriously affected by spikes in realized returns that do not persist.

Next, the paper constructs a VAR for EVAR, using the same regressors. I don't really understand what the point of forecasting realized variance twice is in this way, but that's what they do.

Here is the VAR. Again, bold face means a t of 2.0 or greater, normal face means not.

	r_t	$EVAR_t$	PE_t	TY_t	DEF_t	VS_t
r_{t+1}	0.12	0.66	-0.054	0.007	-0.029	-0.017
(se)	(0.082)	(0.93)	(0.039)	(0.009)	(0.028)	(0.047)
$EVAR_{t+1}$	-0.004	**0.34**	0.012	-0.001	**0.018**	0.005
(se)	(0.005)	**(0.085)**	(0.007)	(0.001)	**(0.004)**	(0.008)
PE_{t+1}	**0.19**	0.57	**0.96**	0.007	-0.024	-0.004
(se)	**(0.079)**	(0.88)	**(0.037)**	(0.008)	(0.027)	(0.044)
TY_{t+1}	-0.16	2.91	-0.002	**0.85**	0.099	0.044
(se)	(0.37)	(4.01)	(0.160)	**(0.039)**	(0.13)	(0.20)
DEF_{t+1}	**-0.45**	2.23	-0.033	-0.003	**0.87**	0.035
(se)	**(0.20)**	(1.82)	(0.080)	(0.020)	**(0.064)**	(0.10)
VS_{t+1}	0.066	0.97	-0.010	-0.005	-0.001	0.93
(se)	(0.073)	(0.74)	(0.033)	(0.008)	(0.025)	(0.041)

This is really where the bodies are buried, so to speak.

Evar decays very quickly, with an 0.34 autoregression. The other forecsaters are much more persistent, with the dividend yield being the champ at 0.96.

Now, when we use this VAR to forecast *long run* volatility $N_{Vt} = (E_t - E_{t-1}) \sum \rho^{j-1} EVAR_{t+j}$, we start raising this matrix to bigger and bigger powers. Obviously, the persistent variables become more and more important. Thus the forecast of $\sigma_t^2(r_{t+10})$ is a linear function of the very persistent forecasting variables, PE, TY, and DEF, and will load very little on EVAR itself.

This strikes me as very dangerous. PE and TY do not significantly forecast EVAR on their own. Yet when we take these point estimates and raise them to larger and larger powers, PE and TY will be economically important state variables for "forecasting long-run variance."

Aside from the statistical issue, this fact raises an important economic issue. Covariances with shocks to "expected long-run variance" will be, mechanically, covariances with shocks to PE, TY, and DEF. We have lots of Mertonian and non-Mertonian stories for shocks to those state variables to matter! At a minimum, they are all here because they forecast mean returns! Maybe this has nothing to do with volatility and long-run risk at all!

In sum, I remain suspicious that there is substantial long-run volatility, or that this procedure measures it.

As a positive suggestion, the mean-return forecasting business complements the dangerous business of raising VARs to large powers with direct estimates of long-run forecastability. This robustness check would be a good idea in this case. Do PE, TY and DEF significantly forecast long-run volatility? If not, maybe some of them should be dropped from the VAR?

(In the discussion, Torben Anderson brought up the important point that the literature on long-run volatility forecasting does find significant long-run variation in volatility, and doesn't use linear models based on past realized returns in order to avoid the spikes in the above graph. Good point.)

OK, so much for "do long run investors care about volatility" and "is there significant variation in long run volatility / significant persistence in volatility?" Now, lets look at the question,

- Is volatility the extra state variable that explains the value effect?

I thought it worthwhile to look at huge recent data point! Here is the VIX (realized returns track the vix almost one for one here, so this is not about a volatility risk premium) along with the cumulative market return and cumulative hml return. (I added rf to each and then cumulated)

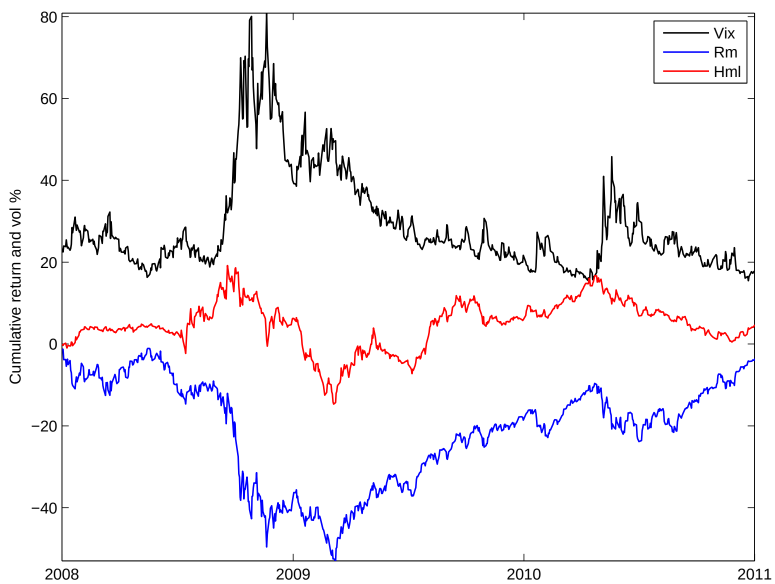

The negative correlation between vix and market premium is really astonishing.

Issue 1: I have been puzzled by a big issue in this episode: why do investors hold the market as σ rises from $0.18^2 = 0.0324$ to $0.80^2 = 0.64$? Standard portfolio theory

$$w_t = \frac{1}{\gamma_t}\frac{E_t(R^e_{t+1})}{\sigma^2_t(R^e_{t+1})} + \beta_{R,z}\frac{\eta_t}{\gamma_t}$$

says to get out! Is there time-varying hedging demand to make holding the market sensible? Can mean returns really rise from 3.2% to 64%? But that's not a question for today, as we are looking only at unconditional moments.

Issue 2: Are value stocks really correlated with volatility? The market itself seems correlated with volatility, but not value! The idea that value is strongly associated with volatility just does not scream at you from this graph!

Again, though, this paper does it seriously, not just plotting data points and telling stories. It estimates and evaluates the model.

The main result is a standard OLS cross-sectional regression pricing the Fama-French 25 portfolios. The result is well encapsulated by the following graph.

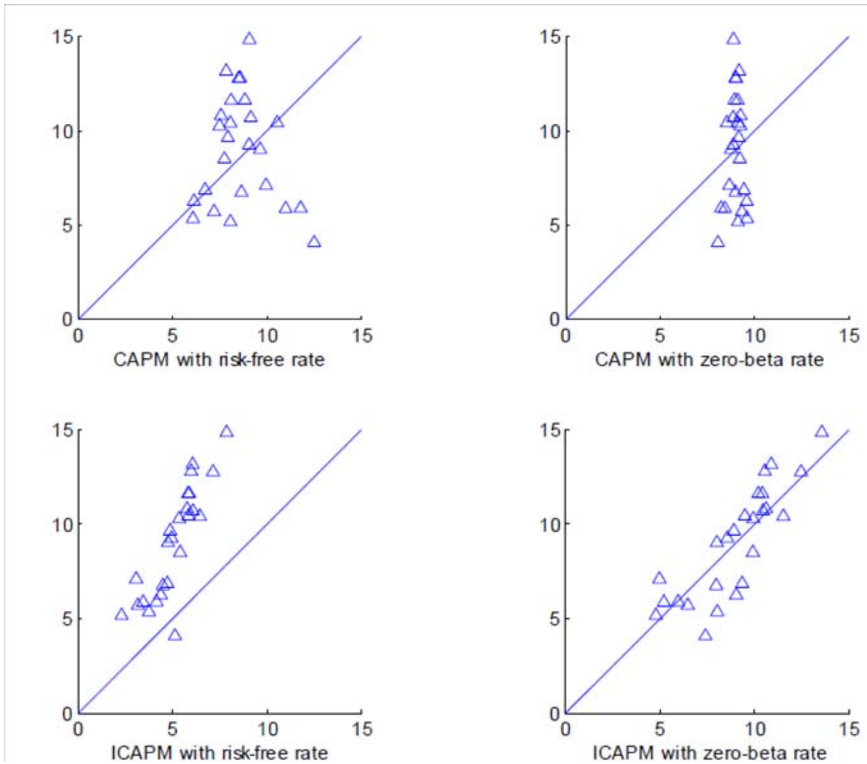

$$E_t r_{i,t+1} - r_{ft}.. = \gamma \; cov_t(r_{i,t+1}, N_{CF,t+1}) + cov_t(r_{i,t+1,}N_{DR,t+1}) - \frac{1}{2}\omega(\gamma) \; cov_t[r_{i,t+1}, N_{V,t+1}] \quad (18)$$

Here I do need to whine a little bit. Our collective standards for asset pricing tests are better than this.

Producing plots like this is a fun industry. As an example, I just was rereading Malloy, Moskowitz, and Vissing Jorgeson who made almost the same plot with stockholder consumption. Lettau and Ludvigson (conditional CAPM), Roussanov, and many others have fit models that look this good. Fine, but producing a 60% R^2 doesn't move my priors much that σ^2 is *the* extra state variable above all the others.

Panel A: Mean returns versus consumption covariances

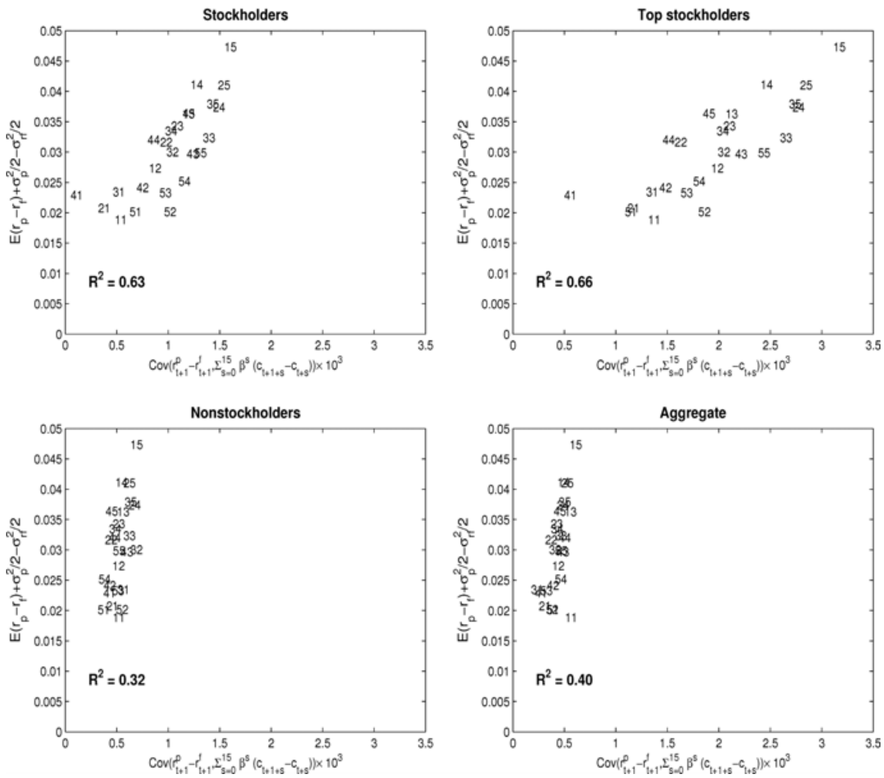

(In the conference I asked for a show of hands who had produced a Fama French 25 plot that looked visually as good. Many hands went up. ,

More deeply, we all now understand that the Fama French 25 are not a "cross section" of returns. The point of Fama and French is that there are basically three degrees of freedom; the 25 portfolios follow a nearly perfect three factor model of *variance*

$$R^{ei} = b_i rmrf + h_i hml + s_i smb;\ R^2 = 0.95$$

and so, by APT logic, follow a nearly perfect three factor model of expected returns

$$E(R^{ei}) = b_i E(rmrf) + h_i E(hml) + s_i E(smb).$$

We now all understand that once you price hml, you price the 25 (unless this is a paper about the small growth anomaly, which it is not.) We also can recite the pitfalls of cross-sectional R^2 in our sleep. (Though to its credit, this paper has fewer degrees of freedom than most.)

Fama and French's variance observation is also important in this context. What is the time-series R^2 of this model? Do volatility innovations soak up anywhere near as much variance as does hml? I suspect not, in part because the authors tell us the model is not rejected. Rejections are based on $\alpha'\Sigma^{-1}\alpha$. I bet the "non-rejection" corresponds to a large Σ, as Fama and French's dramatic rejection corresponded to a large Σ.

In sum, the real question should be "how volatility innovations N_V explain hml," and "why do all the value stocks move together ex-post."

The free zero beta rate is a standard and suspicious trick. It really means abandoning bonds as test assets, and abandoning GLS (efficient) estimates which say to force the cross sectional line through the risk free rate as any other observed factors. The excess zero beta rate is big here, 4.55%, not some small liquidity premium.

You can make the CAPM look good on the value premium by assuming a negative market premium and a huge intercept. But this is a sign of a misspecified model. Here's an example: The CAPM on the FF 25 for a postwar sample. If you include the risk free rate (red or dark blue lines) it's a disaster. If you do an OLS cross sectional regression with a free intercept, you fit the green line.

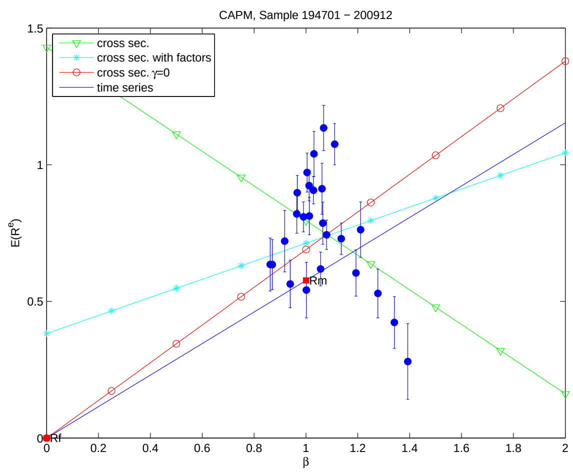

and this obviously produces a much nicer actual vs. predicted graph

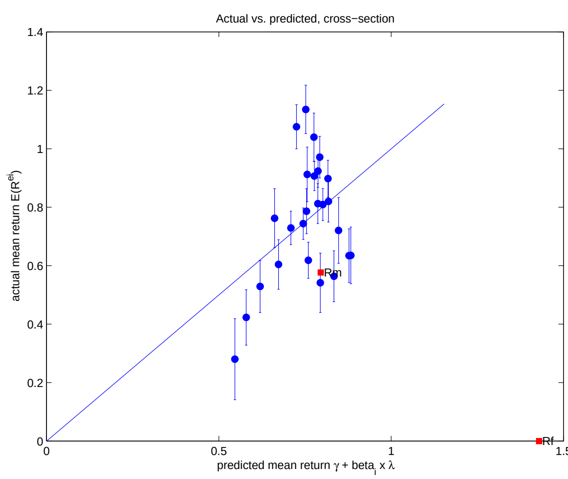

Thus it seems especially worrisome in this context, that in order to get this great fit, the paper estimates a negative premium on volatility risk!

Comparing the two papers, BKSY have a *positive* premium for volatility risk, which seems a bit more intuitive. I understand the claim that a negative premium is consistent with the math of this paper, though have yet to digest it. Why do people *like* high volatility states?

Anyway, I'd feel more comfortable if all parties in the debate could agree on the right sign of volatility risk!

Betas are as much a part of asset pricing tests as are mean returns. Here are the betas in the later sample

$\widehat{\beta}_V$	Growth		2		3		4		Value		Diff	
Small	0.61	[0.32]	0.37	[0.26]	0.24	[0.24]	0.19	[0.22]	0.02	[0.32]	-0.59	[0.12]
2	0.69	[0.29]	0.42	[0.26]	0.24	[0.23]	0.19	[0.25]	0.08	[0.27]	-0.60	[0.12]
3	0.67	[0.28]	0.36	[0.24]	0.27	[0.22]	0.13	[0.25]	0.15	[0.19]	-0.52	[0.14]
4	0.63	[0.25]	0.36	[0.23]	0.19	[0.26]	0.17	[0.27]	0.11	[0.27]	-0.53	[0.12]
Large	0.47	[0.22]	0.36	[0.17]	0.18	[0.19]	0.12	[0.25]	0.14	[0.21]	-0.33	[0.09]
Diff	-0.14	[0.14]	-0.01	[0.13]	-0.06	[0.09]	-0.07	[0.09]	0.12	[0.14]		

Look side to side, along value, not up and down. The size effect in average returns is almost exactly explained by the CAPM and poses no puzzle. Here we see a nice pattern – the betas vary considerably from growth to value. (Though the standard errors are an order of magnitude larger than Fama and French's, suggesting again that N_V is poorly correlated with HML and explains little return variance)

But here are the betas in the earlier sample:

$\widehat{\beta}_V$	Growth		2		3		4		Value		Diff	
Small	-0.72	[0.29]	-0.79	[0.24]	-0.85	[0.26]	-0.82	[0.25]	-0.88	[0.25]	-0.16	[0.13]
2	-0.50	[0.17]	-0.52	[0.22]	-0.60	[0.20]	-0.62	[0.22]	-0.82	[0.25]	-0.32	[0.12]
3	-0.48	[0.20]	-0.38	[0.15]	-0.53	[0.19]	-0.55	[0.19]	-0.85	[0.27]	-0.37	[0.13]
4	-0.21	[0.13]	-0.35	[0.17]	-0.43	[0.18]	-0.58	[0.24]	-0.87	[0.28]	-0.65	[0.19]
Large	-0.22	[0.14]	-0.23	[0.14]	-0.44	[0.21]	-0.67	[0.27]	-0.68	[0.18]	-0.46	[0.16]
Diff	0.50	[0.21]	0.56	[0.14]	0.41	[0.17]	0.16	[0.14]	0.19	[0.12]		

It's a little worrisome that the spread in betas has pretty much vanished. But it's much more worrisome that the *sign of the betas has changed!*

If this is a stable phenomenon, how can the sign of betas change so dramatically? If this time series structure changes from day to night, why is the VAR stable?

More deeply, where do betas come from anyway? What is there about the economics of value firms that makes them rise or fall when return volatility rises or falls?

Bottom line:

I'm still a big fan of long-run thinking. But I'm not convinced yet that recursive utility is crucial; that shocks to long-run news unrelated to current events are important state variables, that volatility is a crucial missing state variable (not, say, expected returns or nontraded income), or that volatility is very persistent.

But these are deep papers, and there is a lot to learn. Just because my priors haven't shifted a lot doesn't mean that they wouldn't if I understood the papers better.