Paper topic suggestions for PhD Students

These are suggestions for class papers or second year papers. I have not researched their novelty to make sure nobody has done them before (and I welcome pointers to papers that have). Some may be appropriate as thesis topics as well. This list is just beginning, so check back as I add more topics!

Cross section of Stock returns

1. Lamont finds that E/D forecasts earnings growth and returns in the time series, even controlling for P/D or book/market. Does the same hold in the cross section? I.e., does cross sectional variation in E/D forecast cross sectional variation in earnings growth and average returns? Does it still do so controlling for size and value (B/M?) (Note: there are now funds trading on this idea. They buy high D/E stocks, on the idea that this is a good signal of future "growth" in earnings and price. But of course the price should build this in ex ante!) Address this by 1) running regressions $R_{i,t+1} = a + b(B/M_{it}) + c(E/D)_{it} + ... + \varepsilon_{it+1}$ (time dummies can let you focus on cross section) and 2) forming portfolios. If so, does the 3 factor model account for E/D sorts, or do we need a E/D factor too?

2. We conventionally sort stocks into 10 to 25 portfolios based on some characteristic, e.g. size, beta or book to market, and then perform asset pricing tests on these portfolios. This is nuts – every econometrics textbook says don't group the data. Instead, we should be using the information in the characteristic as an instrument, and using the full cross section of individual stock returns. Figure out how do to it. Redo some standard tests – CAPM, three factor model, consumption based model – as illustrations.

3. Vuolteenaho's linearization assumes all assets have a *common* unconditional B/M ratio. Find a linearization that allows variation in the constant, so you can say "asset i has a higher unconditional mean B/M than asset j because xyz."

4. No one (least of all us!) has tried to estimate and test the Campbell Cochrane habit utility function on a cross section of stock returns, like the Fama French 25 size and B/M portfolios. Does it work? How does it compare to other macro models like Lettau Ludvigson (cay) or Schneider Piazzesi Tuzel (housing)

5. Redo Jung and Shiller NBER WP 9348 in a more refined manner. Use portfolios to avoid selection bias. Use B/M and E/P as well as D/P. Use all dividends, not excluding liquidation etc. (p.8). Check whether B/M forecasts returns as well as dividend growth.

6. How much of Jung and Shiller's results comes from dividend smoothing? For example, the Schlumberger story is that in the great depression prices *and earnings* plunged, but dividends did not. D/P was high, but dividends eventually gave in and plunged.

But one would know this from P/E. It could be that D/E here forecasts the crash in future D, but that P/E only forecasts returns! Do like Jung and Shiller with P/E and D/E.

7. The fact that so much of B/M (P/D) forecasts dividend or earnings growth means that we must be able to clean up B/M as a forecaster of returns and forecast returns *even better* than Fama French. Do so! (Start with Piotroski, "Value Investing" http://gsbwww.uchicago.edu/fac/joseph.piotroski/research/.)

8. Carhart finds that losing funds keep losing, even after taking in to account fees, expenses, and reasonable trading costs. Maybe this can explain it: People start taking money out of losing funds, and so the losing funds have to liquidate securities. If they're holding illiquid stuff, they have to sell at "fire sale" prices. This causes further bad returns, that would not be captured by Carhart. (Google the story of Heartland funds as a suggestive example.) To see if this works, you have to find the flows out of Carhart's losing funds.

Predictability of stock returns

1. Figure out the accounting of share repurchases, and produce an adjusted dividend number, or other way to adjust the price dividend regressions for repurchases. (The conventional formula does work with repurchases – if the firm repurchases half the stock, the remaining stock gets twice as many dividends per share. But does CRSP data account for this correctly? And for the effects of new issues?)

2. Read p.424 on consumption, Can you see the extreme predictability that p.424 suggests you should see from consumption and P/C stationary? Does D/C, E/C forecast ΔD ΔE as they should? (This involves getting the data definitions right. You don't want value of a dollar invested in 1926, as there is no reason that should be cointegrated with GDP.)

3. Read Lettau Ludvigson on cay, Ribeiro on d/labor income, and Lamont on d/e. Surely there is *one* common trend in all these series – consumption, labor income, earnings, dividends, prices. You have to get the accounting right to see it. Get the single common trend representation to work.

4. Statisticians love to reject return predictability and dividend predictability. But we can't reject both! Estimate return and dividend predictability from d/p *jointly*, imposing the condition that they must sum up.

5. How much of the evidence of return predictability from d/p comes from predicting negative excess returns? Try some specifications like $R_{t+1} - R^f_{t+1} = e^{a+b(d_t - p_t)} + \varepsilon_{t+1}$ and see how the R^2 is affected.

6. Extend the Monte Carlo the distribution of d/p forecasts of returns and dividend growth. Check Campbell's and Lewellen's recent papers on the distribution of these statistics.

7. More generally, in my QJE paper, I never got around to checking if d/p predicts consumption and income, or if c/y predicts stock returns. In a growing economy, consumption, income, investment, stock value, dividends, etc. all should share a single common trend. Since consumption is much more stable ($\sigma = 1\%$) than dividends, C/D should forecast a lot of dividend growth, and the permanent component of stock prices should be much less volatile than d/p regressions suggest. Properly measured, P/C should be a much stronger forecast variable than P/D.

Taylor rules

1. Regressions: Clarida, Gali, and Gertler and many others estimate Taylor rules of the rough form $i_t = \bar{\imath} + \phi\pi_t + \beta y_t$.They find that $\phi < 1$ for the US in the 1970s, but $\phi > 1$ since 1980. This is considered evidence that the Fed learned how to run monetary policy to give price stability.

   (a) Yet this must depend crucially on the y variable. If you run $i_t = \bar{\imath} + \phi\pi_t$ you get $\phi = 1$ by Fisher equation logic. How sensitive is the result on ϕ to sensible choices for y? (See Athanasoulis, AER on one choice)

   (b) What if you use specifications like Piazzesi's that give much better fit to i? In particular, what if we include long term interest rates in the description of Fed policy. Does it make sense to do this?

   (c) Do we see $\phi > 1$ outside the US?

   (d) Woodford's *interest and prices* tells us that optimal monetary policy should have a time varying intercept $\bar{\imath}_t$. Can you estimate a Taylor rule that conforms to Woodford's optimal monetary policy?

   (e) With $i_t = \bar{\imath}_t + \phi_\pi\pi_t$, if the Fed is *optimally* setting $\bar{\imath}_t$ as suggested by Woodford, then the right hand variable (π) is massively correlated with the errors ($\bar{\imath}_t$), so estimates of Taylor coefficient are massively biased. Figure out how to deal with this.

   (f) Worse, in the standard equilibrium idea for Taylor rules, we don't see the the Fed reacting to inflation because the inflation never happens. Look at artificial time series of a Taylor rule economy. Do OLS regressions recover ϕ_π in those economies? NO! Is estimating a Taylor rule in Taylor rule economies like trying to see if surpluses respond to debt in Fiscal economies – you can't estimate off-equilibrium reaction functions from time series of an economy in equilibrium?

2. Theory: The Taylor rule in an open economy. Woodford (*Interest and prices)* does a wonderful exposition of Taylor rule dynamics and optimal monetary policy in a closed economy. In an open economy, domestic real interest rates must equal the world real interest rate so "real interest rate" policy is impossible. Central banks in open economies pay more attention to exchange rates. Formulate an exchange rate based Taylor rule that provides a determinate price level in an open economy (at least as

well as the conventional Taylor rule does so for a closed economy.) Corresponding to Woodford's characterization of optimal monetary policy in a closed economy – roughly, set the Taylor intercept to the current "natural" real rate – what is the characterization of optimal monetary policy in an open economy?

3. Theory: Half of the Taylor rule theory says you should have explosive dynamics to give price stability, (see Woodford) and half say you should have stable dynamics (see Taylor). Reconcile these views.

NOte: Date: Tue, 25 Nov 2003 13:24:44 -0500 (EST)

From: Jose Tessada <tessada@MIT.EDU>

Subject: paper topics file

To: john.cochrane@gsb.uchicago.edu

Message-id: <Pine.GSO.4.55L.0311251320390.27211@buzzword-bingo.mit.edu>

MIME-version: 1.0

Content-type: TEXT/PLAIN; charset=US-ASCII

Original-recipient: rfc822;john.cochrane@gsb.uchicago.edu

X-UID: 22351

Dear Professor Cochrane:

Looking for some info on the internet I found your file with suggestions

for paper topics.

You ask about estimations of taylor rules for other countries. Apart from

some estimates Clarida et al published in the early 2000s, I have seen a

couple of papers estimating Taylor functions for Latin american countries.

I was an RA in one of them. It was published in a conference proceedings

volume by the Central Bank of Chile in 2002. The author is Vittorio Corbo.

I guess the article must be available online, either as working paper or

the final published version.

Best regards

Jose Tessada

MIT

Other asset pricing

1. If real rates are constant and nominal bonds vary from inflation expectations, then short term bonds have the lowest real risk even for long horizon investors. If inflation is constant and real rates vary, then long term bonds have the lowest risk for long horizon investors. Which is it? Did this change from the 70s – high inflation variability, low real rate variability (?) to the 90s, with the opposite pattern?

2. Are bond risk premia *real* risk premia, or premia for holding inflation risk?

3. As a specific instance, Fama and Bliss find that the expectations hypothesis works pretty well at long horizons, but poorly at a one year horizon. Is this because long horizon yield curve movements reflect inflation, while short horizons reflect less persistent changes in real rates, and you get a premium for holding real rate risk but not for holding inflation risk?

4. We can't all rebalance. Hence, after stock i goes up, stock i's expected return must *rise* so that we now want to hold that stock in its, now larger, proportion of the market portfolio. *Stock returns can't be i.i.d.* (Note: I assume constant second moments here, and this presumes a short run in which shares are constant. With linear technologies, rates of return can be constant and you invest more in the ith production process) Is this enough to account for momentum?

   This got done by "Two Trees." Still, Two trees gets the autocorrelation too low by a factor of 10. (See "new facts in finance" for the link between autocorrelation and momentum, as well as the discussion in Two trees.) But two trees uses a log utility function, and clearly risk aversion needs to be cranked up more, not least of all to get the equity premium right. Doing the same exercise as two trees but with higher risk aversion, and potentially with epstein zin or habits so you can have high risk aversion without a time-varying risk free rate would be the right way to get at the question

5. Epstein Zin and habits in the utility function allow you to separate *interetemporal substitution* from *risk aversion.* They allow people to be quite averse to substituting consumption across states of nature, while allowing them to be quite willing to substitute consumption across time. The usual separable utility ties the two together, $U = \sum_{states\ i} \pi_i \sum_{dates\ t} \beta^t u(c_{i,t})$, gives the same willingness to substitute across states i as time t. In standard micro, everything you can do for utility you can do for production functions. But our current production functions treat outcomes over *time* and across *states* in very simple forms. The typical production function $y_{it} = \lambda_i f(k_{t-1})$ (λ=shock, i = state index) allows pretty smooth substitution over *time* but absolutely no substitution across *states* – it's Leontief here. More generally, devices such as adjustment costs that have been used to introduce less smoothing across states also give less smoothing over time and hence more interest rate variability, just like raising the risk aversion coefficient. (See Urban Jermann's paper in the JME.)

   Paper topic: do like Epstein Zin or habits for production functions. Separate the ability to smooth over time and the ability to smooth over states. (The Hansen-Sargent robust control machinery may prove useful here. I also have an old working paper titled "rethinking production under uncertainty that may help – though since I gave up on the approach in that paper, maybe not.)

1. Merge the fiscal theory of the price level and optimal distorting taxation including an inflation tax. (Sims' Dollarization in Mexico paper is a good place to start)

2. The Q theory of investment thinks shocks are to preferences. The consumption based model thinks shocks are to production. Where are the shocks? How does the Q theory look in general equilibrium if there are shocks to future productivity?

3. Many people have recently suggested that "asset price inflation" should be included in measures such as the CPI. This is conventionally not done. The CPI includes the rental value of housing, for example, not its price, and does not include interest rates or stock prices at all. Leaving aside the question whether the Fed should respond or target "asset price inflation" there is one sound economic motivation. OUr index theory presumes a static consumer. It measures the change in cost of consuming a given consumption bundle at date t. People don't live one period however. We could instead think of inflation as the change in cost of purchasing the *lifetime* consumptoin stream. For this purpose, interest rates and stock prices do matter. If goods prices do not change but interest rates go down, the time t cost of purchasing a given lifetime consumption stream has risen. (Of course the time t value of a labor income stream has also risen.) High stock prices work the same way.

   Explore this idea to construct a price index that reflects the cost of purchasing a lifetime stream of consumption. You will have to think of incomplete (or incompletely measured) markets. For perfect foresight, interest rates are sufficient intertemporal prices. To handle uncertainty, you will have to think about stock prices which can provide some "consumption insurance" and make assumptions about market prices of unmeasured risks. Calculate a "lifetime income" index to go with your "lifetime inflation" index . Do these measures correlate better with exchange rates? How do they compare to the CPI?