



LONG-TERM INVESTING: AN OPTIMAL STRATEGY IN SHORT-TERM ORIENTED MARKETS

December 3, 2012

*Organized by the Committee on Global Thought at Columbia University
& the Sovereign Wealth Fund Research Initiative*

http://cgt.columbia.edu/conferences/long-term_investing/

Paper Presentation III

NATURAL EXPECTATIONS, MACROECONOMIC DYNAMICS, AND ASSET PRICING*

(Fuster, Herbert & Laibson, 2011)

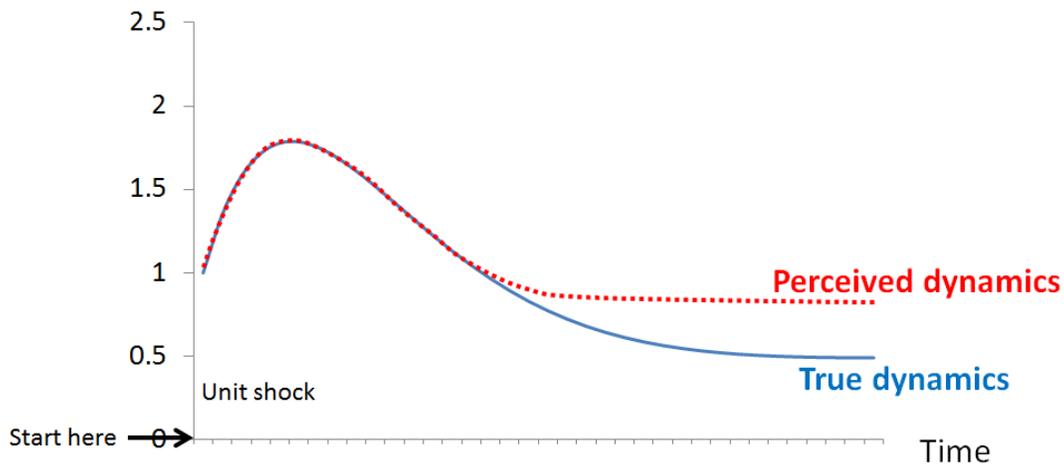
Paper Presentation: David Laibson
*Robert I. Goldman Professor of Economics,
Harvard University*

Discussants: William Goetzmann
*Edwin J. Beinecke Professor of Finance and
Management Studies, Yale School of
Management*
Jan Svejnar
*James T. Shotwell Professor of Global Political
Economy, School of International and Public
Affairs, Columbia University*

DAVID LAIBSON: I am going to tell you about some joint work with Andreas Fuster, who is now with the Federal Reserve Bank of New York, and Benjamin Hebert, a doctoral student at Harvard. This paper has two fundamental assumptions. The first is that macro fundamentals—think of that as earnings—are hump shaped in their dynamics. By that we simply mean that when an impulse arises in the aggregate economy, that impulse tends to have some momentum in the short run and then have some mean reversion in the long run. For now, we are not going to worry about the amount of mean reversion. The first critical assumption we want to make is that there is an up and down pattern in the data. We will back that up with empirics later.

The second assumption is that agents in the economy underestimate the degree of mean reversion. To illustrate that idea, imagine that the true dynamics are here in blue, and earnings after an impulse tend to go up and then come down on average.

* *Remarks have been edited for clarity.*



Agents making forecasts in this environment tend to under-appreciate the full degree of mean reversion. Agents plotted in red anticipate an up and down pattern, with not quite as much down in the long run than they should expect on average. Why might they make that mistake? Why is the long-run mean reversion hard to see or anticipate? First, it comes at long lags, so these agents would have to anticipate something that will play out over many years—10 or 15 years. Secondly, it is the sum of many small effects, so it does not come in one big bump; it comes slowly over time on average, over the long run. Thirdly, if you estimate using the kinds of models that we have been estimating on this planet for the last 100 years, you will miss the mean reversion. If you introduce five years of lag with quarterly data, you would have 20 lagged quarters and you would not pick up mean reversion.

What happens in an economy like this? As an overview of what we are going to see, I will start with fundamentals. Everything is symmetric there, so it could be a good or bad news story that drives the cycle; I'll tell the good news story. Secondly, agents will over-estimate the long-run persistence of that good news; they will think that the good news is going to stay around forever and will fail to appreciate the full degree of the mean reversion of that good news. Asset prices will respond to this belief about the persistence of the good news, and agents will price assets as if the good news was going to be highly persistent. As a consequence, not only will asset prices rise, but consumption and investment will begin to rise. Over the long run, however, because on average agents under-appreciate the degree of mean reversion in macro fundamentals like earnings, the average assets prices, consumption, and investment will fall as that mean reversion plays out slowly over the long run. We will see an up and down cycle in the economy and in asset prices.

Let me quickly summarize the consequences of this framework. First, agents who estimate simple models to try understand this economy are going to appreciate and foresee the momentum in the short run. They will foresee that earnings tend to keep drifting in the same direction after a good shock. But those agents will fail to see the long-run mean reversion, so there will end up being an extrapolation bias in this

economy. Asset returns will be excessively volatile, rising and falling as the underlying fundamentals tend to systematically rise and fall. We will see returns that are negatively predicted by lag returns, by Campbell and Shiller-style lag price-to-earnings ratio, and by high consumption growth in the past. When the economy is booming today, that is a warning sign of below-average excess returns in the future.

Real economic activity will also have amplified cycles. Consumption and investment will have cycles of going up and down, so there will be long-run volatility in the real economy, not just in financial markets. The equity premium is going to be too large, although long-run equity returns are going to co-vary only weakly with long-run consumption growth. Agents in this world will perceive a very risky equity market because they perceive that there is a lot of persistent volatility in the real economy. If they fully appreciated the degree to which what goes up tends to come down, they would better appreciate that equities are less risky than they seem to be to these agents. Consequently agents will be unwilling to hold equities at appropriate prices; they will bid down the price of equities and create an anomalously high equity premium.

If you are rational in this economy, you have a great opportunity because you can take advantage of the mistakes of these agents. You can lever up and hold more equities, because you realize that equities are not quite as risky as they appear to be to everyone else. Moreover, you will have a counter-cyclical investment policy. As the rest of the economy is convinced that the economy is completely sour, you're going to lever up and deepen your position in equities. Then, of course, you will weaken your exposure to equities during the boom periods, when other agents mistakenly think that the economy is persistently healthy; you will foresee that the mean reversion will probabilistically set in.

There is a large literature on which this paper is building. We are not the first to think about extrapolation, but we are going to offer a particular parametric approach and calibration exercise.

Model

I'm going to omit a lot of the Greek details and just give a few equations to set the stage. We have a standard framework in which agents maximize a discounted sum of utilities.

CARA habit preferences (Alessie and Lusardi)

$$\sum_{t=0}^{\infty} \delta^t u(c_t, c_{t-1}) = \sum_{t=0}^{\infty} \delta^t \left\{ -\frac{1}{\alpha} \exp(-\alpha[c_t - \gamma c_{t-1}]) \right\}$$

δ is the discount factor and u is the utility function. Utility is a function both of my current consumption, c_t , and lag consumption, c_{t-1} . Economists call that a habit; it is simply there to slow down consumption dynamics, and has no effect on the asset pricing in this constant, absolute risk-averse world. How do we formalize that? On the right side of the model is a constant absolute risk aversion utility function. a is the parameter that captures risk aversion, and γ is the parameter that slows down consumption adjustment, the so-called habit.

There are only two assets in this very simple, stripped down economy. There is an equity tree.

Dynamic budget constraint for wealth, w_t

$$w_t = (w_{t-1} - c_{t-1} - \theta_{t-1}P_{t-1})R + \theta_{t-1}(d_t + P_t)$$

θ_t represents shares purchased at date t

d_t represents per share dividend at date t

P_t represents share price at date t

Equities here are purchased in quantity θ_{t-1} at price P_{t-1} . This is the amount of equities that I've purchased in period $t-1$. c_{t-1} is the amount of consumption that I have undertaken in period $t-1$, and w_{t-1} is the amount of wealth that I had in period $t-1$. The difference is simply my liquid wealth in period $t-1$, that grows at a gross rate of return R . My total wealth, w_t , is my liquid wealth grossed up by the rate of return plus the amount of equity shares that I have, θ_{t-1} , multiplied by the dividends I get from my equities, d_t , and the price of those equities, P_t , at date t .

To make things tractable, we assume an elastic supply of foreign capital available at gross rate of return of R . This is the Chinese economy happily lending to us at gross rate of return R . We also assume that foreign agents are happy to lend but not happy to hold domestic capital, so there is a home bias in equity investment.

Now how do we model the earnings process, the dynamics of fundamentals? Let's think about dividend growth as $\Delta \ln D$ or Δd , as the sum of historical dividend growth rates.

Growth in dividends ($\Delta \ln D = \Delta d$) is captured by an autoregressive model with p lags:

$$\Delta d_t = \theta_1 \Delta d_{t-1} + \theta_2 \Delta d_{t-2} + \dots + \theta_p \Delta d_{t-p} + \varepsilon_t$$

This is going to produce the cycles of up and down patterns. We call this kind of process an AR(p) model in dividend growth, because it is an auto-regressive model with little p lags, represented by $\theta_p \Delta d_{t-p}$. These lags derive the dividend process and for now assume that is exogenous.

The crux of the psychology in this behavioral economics model is the following assumptions. Everything else in the model is completely standard in a classical finance sense; this is the only exotic piece. We are going to assume that dividend growth is truly driven by an auto-regressive process with 40 lags. The process has a lot of memory, 40 lags, to determine that up and down cycle.

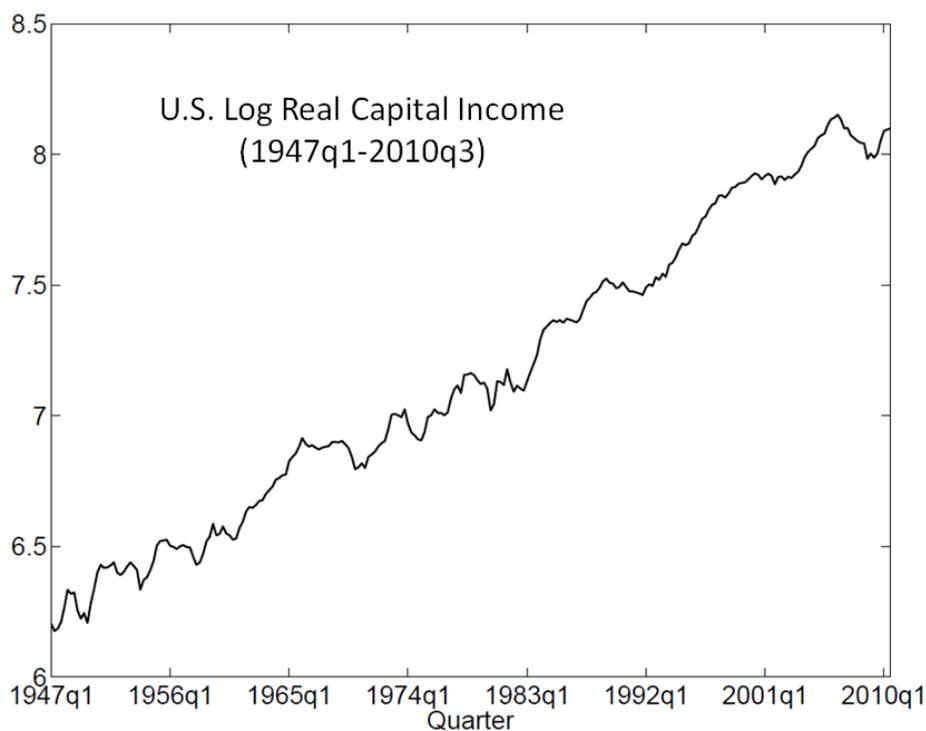
But agents in the economy do not think of it that way. They think of the world as being driven by some smaller number of lags. Let's say that they think of the smaller number of lags being p , because we will study all the different cases where agents oversimplify the world. We take a very complicated world and treat it as a simpler object than it actually is. Here that simplification process is going from 40 lags, which is what we should be considering, to going to p lags, where p is less than or equal to 40.

$$\Delta d_t = AR(40) \quad \text{Data generating process}$$

$$\Delta d_t = AR(p), p \leq 40 \quad \text{Natural Expectations}$$

This is the simplified model that we take to the data and will use to form our forecast.

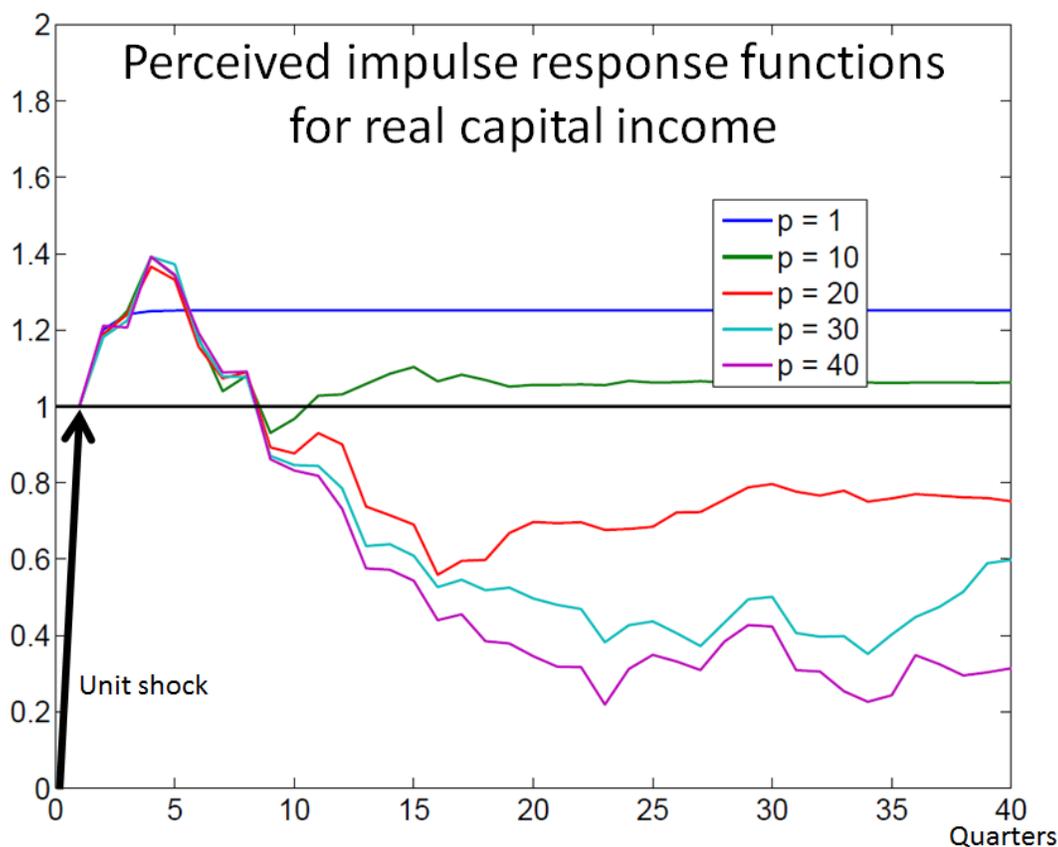
We will act as if the imperfectly rational agents in this economy will assume that this autoregressive model with p lags is the true model, but in fact they are making a mistake. They are neglecting the fact that there are more lags that they should be taking into account, which are being ignored in their oversimplified analysis.



U.S. NIPA (BEA): net operating surplus of private enterprises.

The data that we're going to work with is capital income data adjusted for inflation for the U.S. economy, which is available from 1947, quarter 1 to the present day. At the time the paper was written, we had data through 2010, quarter 3. As you can

see this data shows a fair bit of mean reversion; if I fit a straight line to this process, we would not deviate far from it at any given point in time. This process does not look like it has a very strong unit root in the language of time series econometrics, but it has a lot of mean reversion over the long run.



Let's estimate an auto-regressive growth processes on this data and begin with the most naïve, simplest model we possibly can. Let's begin with a model with only one lag. $p = 1$. An agent comes to this data naively saying that the only thing that matters for growth is the growth rate of dividends one period ago; all history beyond that does not matter. That very naïve agent, taking this data available back to 1947 Q1 and estimating using her model on the data, will determine that after a shock there is a little more momentum and then the economy levels out. That is the consequence of estimating that auto-regressive model, and again, this an average effect of a shock. She would estimate that average effect using this 50- or 60-year period of data.

The more lags she adds to her estimating model, the more she begins to see a very different picture emerge. Let's take a look at 10 lags, the green line; when she adds 10 lags to her estimating model, she notices that what goes up keeps going up and then comes down a little bit. If she adds 20 lags, what goes up goes up even more with trending and momentum, and then comes down a lot, to the red line. If she adds 30 lags, it goes up and down, and with 40 lags, up and down even more. As she adds more and more lags—an analysis that most econometricians would not do because they

would say the additional lags are not significant, as they do not show up with a high enough coefficient to matter—she would find that the mean reversion was getting stronger and stronger. With 40 lags she sees very powerful mean reversion in the long run dynamics of this economy, and with one lag she sees absolutely no mean reversion.

To anticipate where this will go, the red line, with 20 lags, will fit the data for the U.S. economy very well. Agents in this economy think that there is this much mean reversion in the data but who are missing the fact that there is even more mean reversion in the true economy. They think that this is the amount of mean reversion, and they are wrong because in fact there is deeper mean reversion. That is going to create all sorts of cycles and asset pricing anomalies, which we will trace out next.

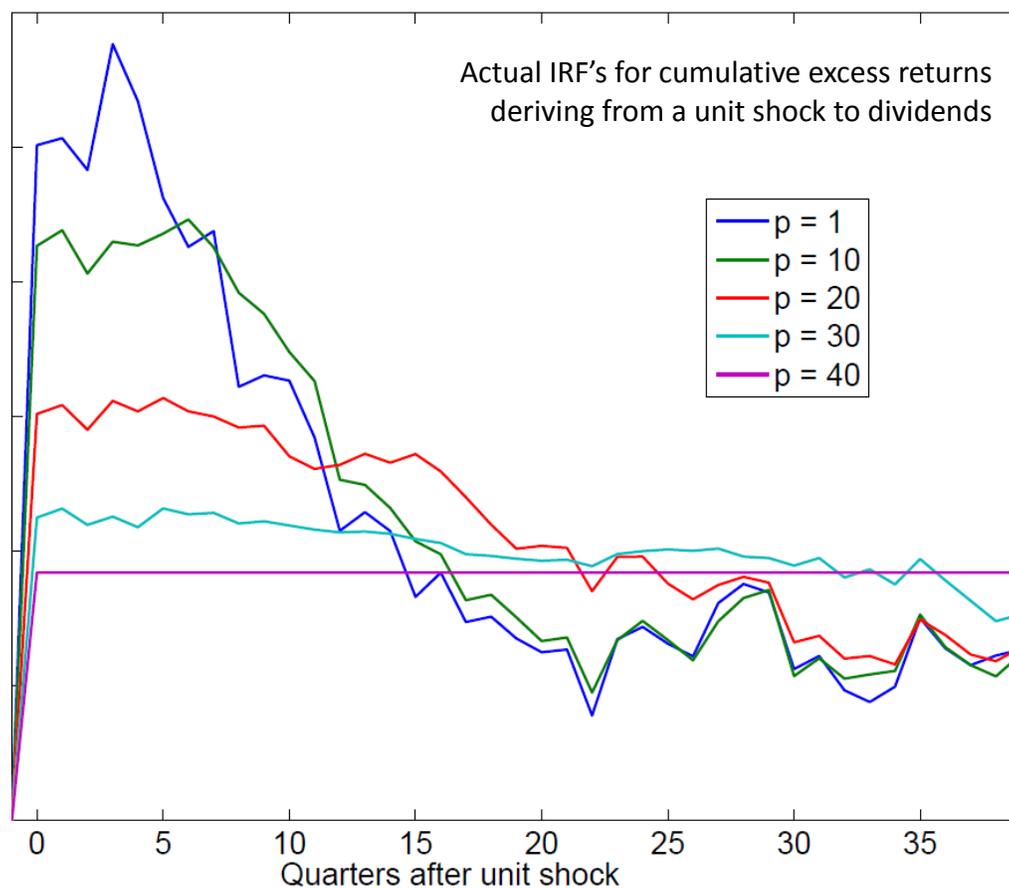
How do we calibrate the model? The true data generating process (DGP) we assume to be a 40 lagged auto-regressive model. We estimate that using the actual data from the U.S. economy that is available quarterly, but we assume that agents in this world—forecasters at banks, in the U.S. government, in households—do not realize that it is 40 lags. They think that it is p lags, and they will estimate oversimplified model with too few lags. We have standard calibrations for the other parameters.

Model Calibration	
True DGP	$\Delta \ln d \square \text{AR}(40)$ estimated from NIPA
Perceived DGP	$\Delta \ln d \square \text{AR}(p)$ estimated from NIPA
$R = 1.0025$	gross risk-free rate (quarterly)
$\delta R = 1$	discount factor
$\gamma = 0.9$	habit weight
$\alpha = \frac{4}{c \left(1 - \frac{\gamma}{R}\right)}$	local relative risk aversion of 4

Asset Pricing

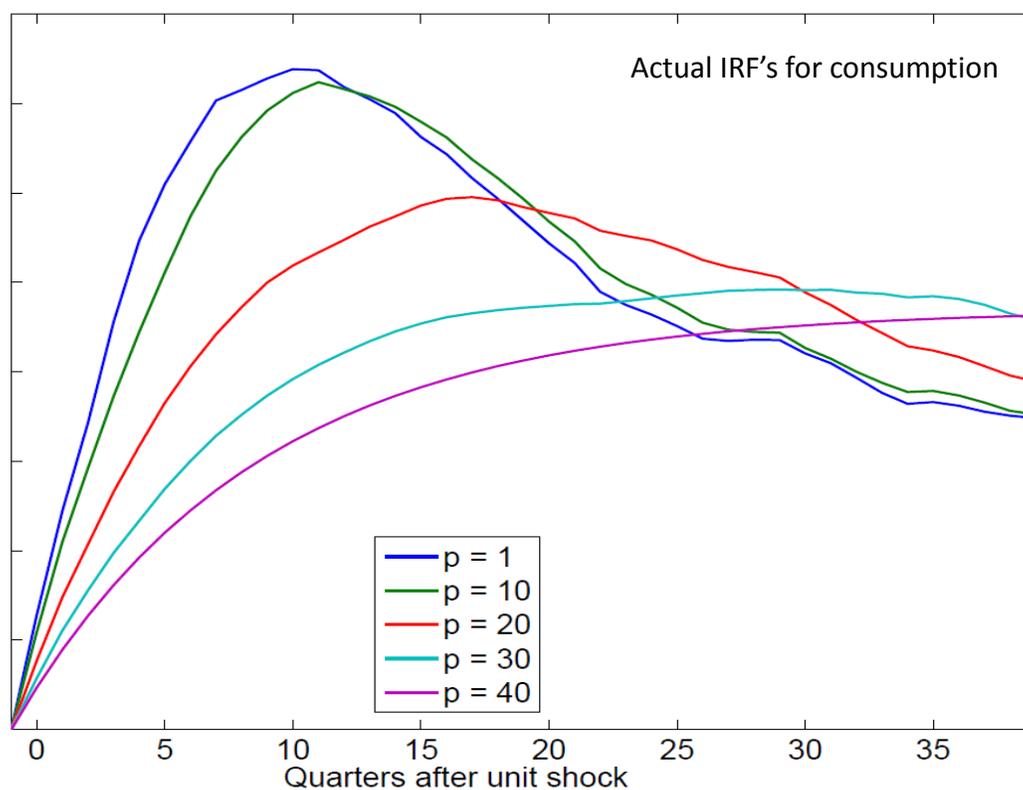
What happens for asset pricing in this world? Let's study a shock, which is derived from a unit shock to dividends, in this environment. Let's first think about the rational case.

If I included 40 lags in my model, using the right model to estimate the data, when a shock comes, prices go up, dividends rise and then there's no additional average excess return after that. On average, excess returns in the purple case are not forecastable after the shock.



Now let's think about the opposite world, where the agent does not use the correct 40-lag model but instead has an oversimplified view of the world, only estimating a one-lag model. In that case, asset prices jump up a lot after the shock because agents think it is going to persist. There will only be momentum effects; what was good yesterday will get even better tomorrow, and there will be no long run mean reversion. Asset prices race higher after the good news, and then because the agents do not anticipate the mean reversion, that asset price will slowly on average, over 10 years, mean revert a fair degree. About 3/4 of the initial return will be given back stochastically, randomly over the long run.

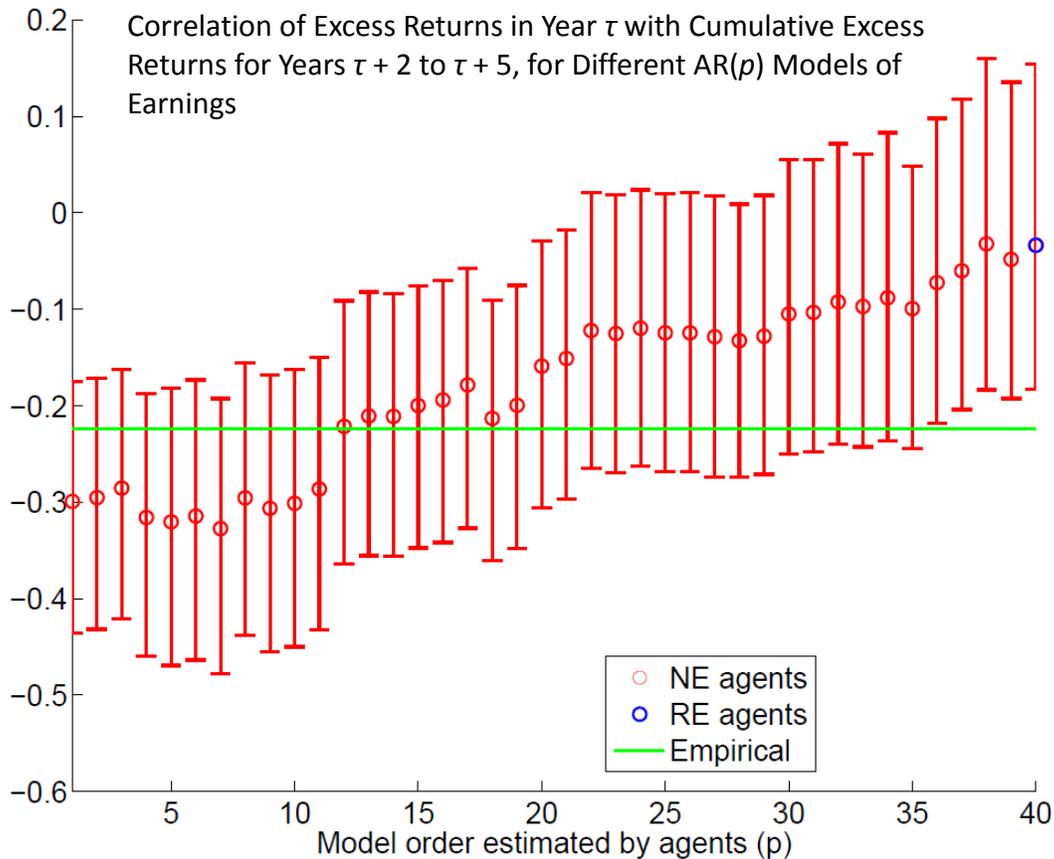
As the agents include more and more lags—1 lag, 10 lags, 20 lags, 30 lags, 40 lags—they get closer to the rational asset pricing. Let's take a look at the 20 lag case. Here an impulse or shock to dividends in the form of good news raises asset prices a lot, but because the model does not have enough lags, there is some mean reversion. About 1/3 of the original impulse is eventually given back in terms of asset price mean reversion. That is the story that will play out in this economy. Agents who over-react to good news and then slowly over the long run, as new news comes in, the prices begin to fall back. There is a little bit of mean reversion not just in dividends, but also in the asset prices that agents have set and in the consumption choices that they make.



Here is consumption. I want to show you the dynamics for the rational case, in which agents perfectly anticipate all of the dynamics in the underlying economy. In that case, represented by the purple line of 40 lags, when good news comes in, consumption slowly asymptotes up. It moves slowly. Because of the habit in this economy, there is an incentive to smooth adjustment in the household's consumption profile, so we slowly asymptote on average to a new higher level of consumption.

But when agents use the wrong models that do not reflect the full degree of mean reversion, they overreact. They consume too much in the short run. Then, as the bad news starts coming in and they realize that there was more negative earnings news than they anticipated, consumption falls back down. This is the cycle that emerges in the real economy. Good news comes in, people overreact and then as they discover they've overreacted, they tend to mean revert back in the consumption and investment dynamics for economies like this.

Let's take this to the data and ask if the model explains the asset pricing anomalies in the time series data. I will use annual per capita consumption data and excess return data from 1929 to the present. We will also use the price-to-earnings ratio that Shiller and Campbell like to use, which is the S&P 500 divided by inflation and adjusted 40 quarters of earnings (P/E ratio) over 40 quarters, inflation adjusted. We are going to ask if we can reproduce the relationships that emerge in the data.

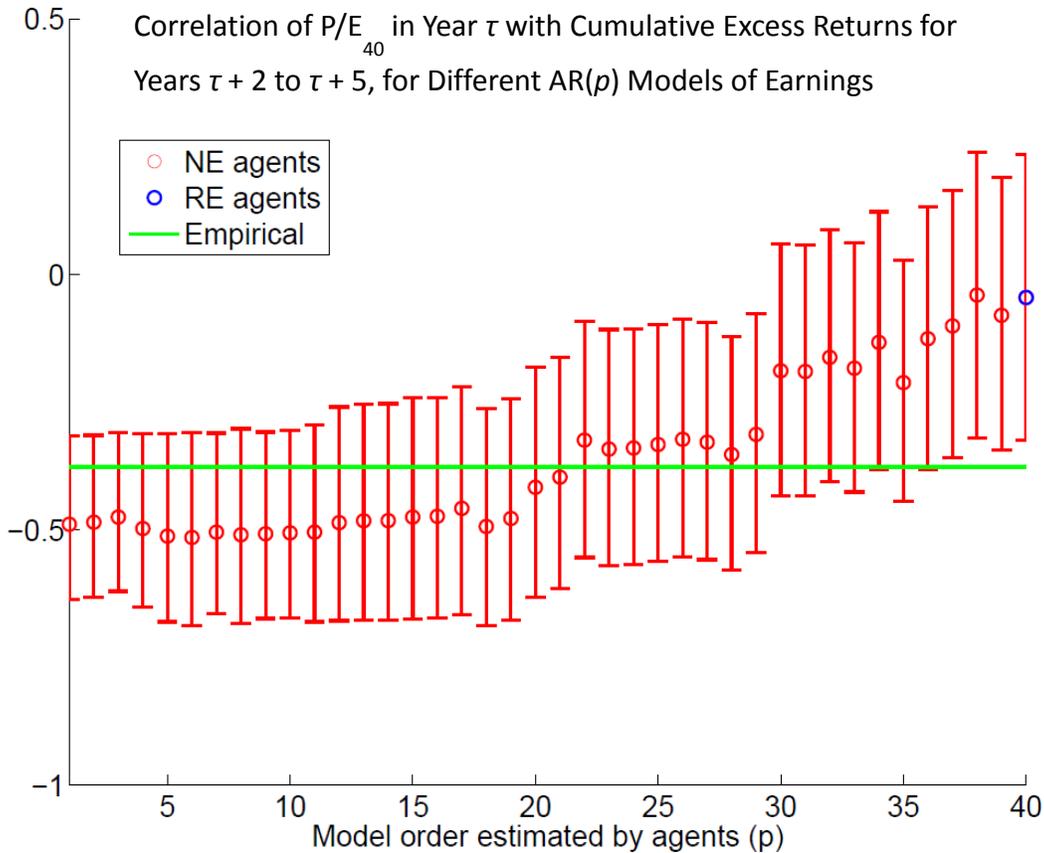


The sophistication of the agents in the economy is on the horizontal axis. If they are estimating models with one lag, they are pretty stupid. If they are estimating models with 40 lags, they are perfectly rational. I will warn you that it will turn out again and again that an estimate models with about 20 lags, where agents will miss about 1/3 of the mean reversion, is going to explain the data.

The vertical axis shows the correlation of excess returns in year t with cumulative excess returns in years $t + 2$ to $t + 5$. We simulate that correlation in different kinds of universes where agents are either very stupid, very sophisticated or somewhere in the middle. Those simulations are the red lines, showing 95% confidence bands. The mean is the circle in the middle. The green line is the actual empirical data.

You should be immediately be surprised and intrigued because in the actual empirical data there is a sharp negative correlation between excess returns at year t and excess in years $t + 2$ to $t + 5$. The model does a good job of picking that out. At around 20 lags, we explain that negative correlation, because good news today is associated with bad news in the future. People are surprised to find the mean reversion coming in over the next five years. As a consequence, there is a negative correlation in asset returns that the model predicts. If we make the agents perfectly rational, we end up with a zero correlation—the blue circle which should be at zero—because if agents are perfectly rational there is no equilibrium correlation between excess returns today and excess returns in the future. Asset prices are always equal to fundamentals and hence asset

prices are a random walk; there is no predictability. We are first able to explain this negative correlation in excess returns by assuming that agents are imperfectly appreciating the degree of mean reversion.



This is the correlation of price-to-earnings ratios in year t with cumulative excess returns for year $t + 2$ to $t + 5$, again for different autoregressive models that are held by the agents in the economy. Stupid agents are on the left side of the x-axis, and sophisticated agents are on the right. Again, we find that agents in the middle are roughly able to predict the patterns that we see in the data, which is a very strong negative correlation between a price-to-earnings ratio today and excess returns in the future. That negative correlation is about -0.4.

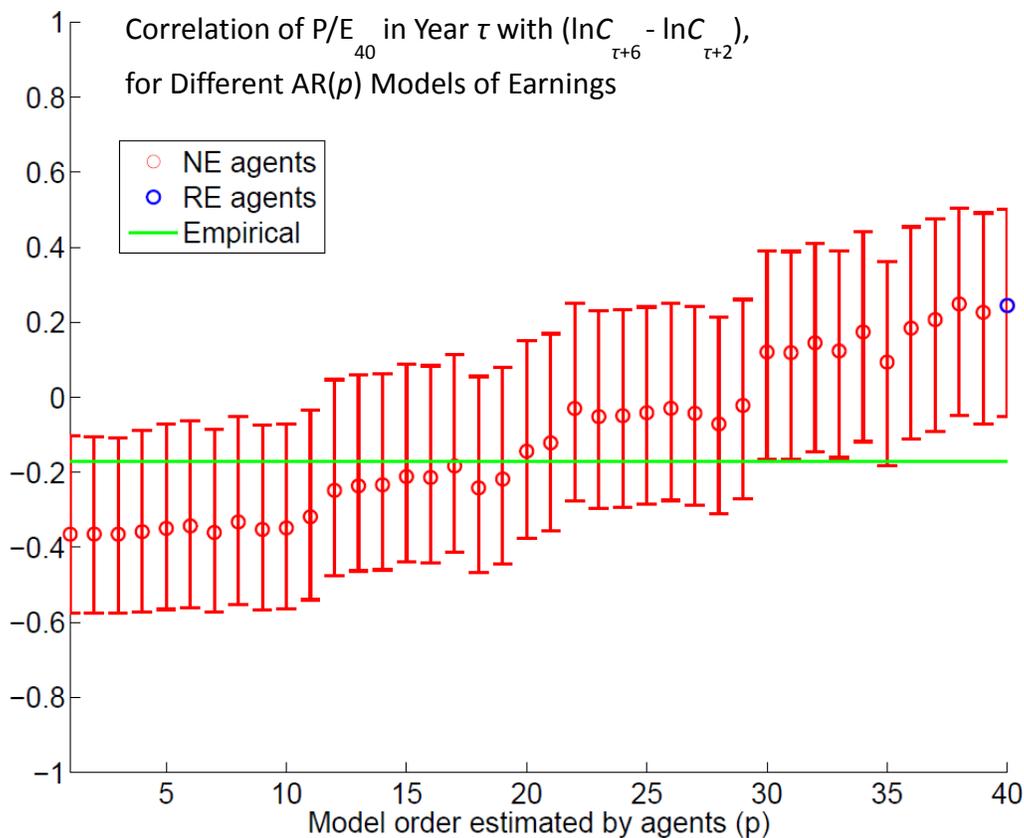
How do we explain that with our 20 lag models? Agents are doing a somewhat good job of anticipating the degree of mean reversion but they are missing a lot of the mean reversion. When a price-to-earnings ratio is very high today, it is going to anticipate unexpectedly negative earnings news in the future, and agents are not anticipating all of the mean reversion. That is why a high price-to-earnings ratio today produces negative excess returns later. From a calibrational perspective, we are able to match the -0.4 correlation. Again, the rational model will not do it, because the rational model wants the correlation to be zero; that's very natural. The rational model

fundamentals are all priced exactly to the price of equities, so consequently there is no correlation between the price of equities today and excess returns in the future.

Consumption Growth

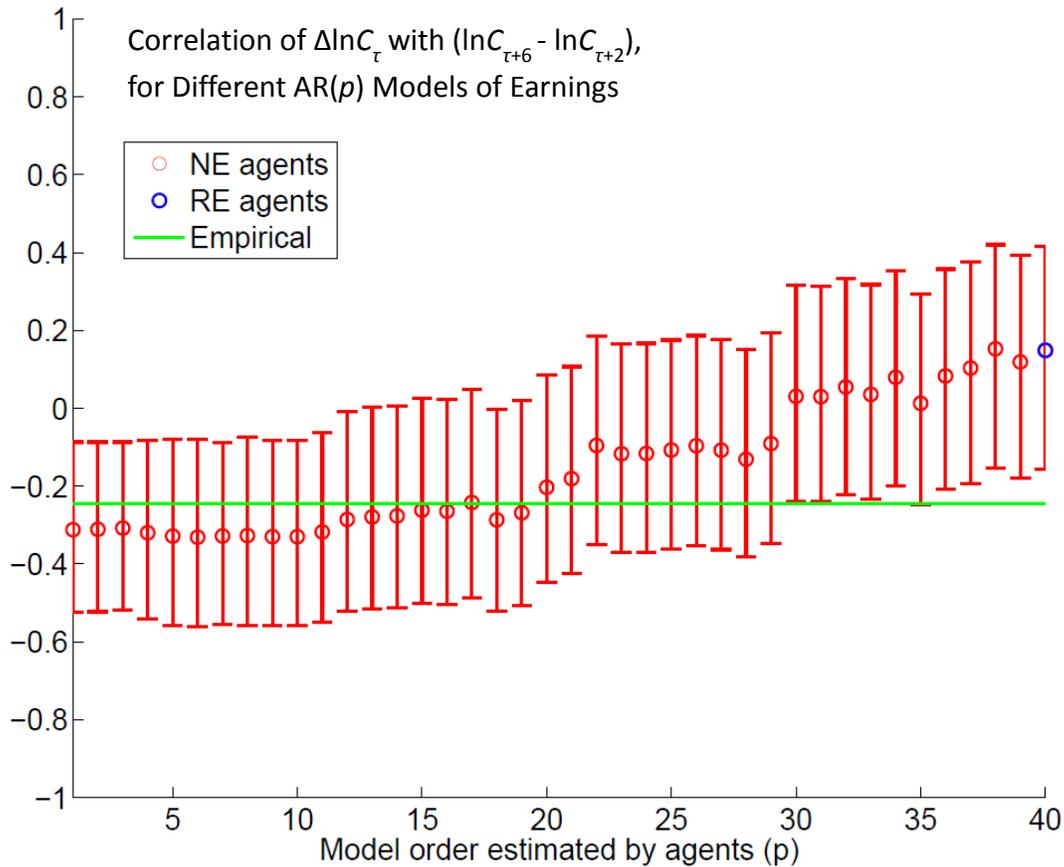
What about consumption growth? When the economy is booming, as evidenced by rapidly rising consumption, it anticipates negative excess returns in the future. That correlation is about -0.3. We would expect that correlation to be zero in the rational case, when agents are perfectly sophisticated. Again, when agents are half way to sophistication, they turn out to be relatively close to what we want to see in the data, which is a strong negative correlation between consumption growth today and excess returns in the future.

We can spin all these things around and ask how historical events like high P/E ratios, high equity returns, and high consumption growth predict future consumption growth. Rather than predicting future excess returns, we can also predict future consumption growth.



We can look at the correlation of the P/E ratio in year t with consumption growth between year $t + 2$ and $t + 6$, once again assuming agents at intermediate levels of sophistication do a good job at matching that negative correlation between a high price-to-earnings ratio today and below-average consumption growth in the future.

The rational model gets it wrong. For the rational model, you want to see a positive correlation between price-to-earnings ratios today and consumption growth in the future, working through the fact that consumption only slowly responds to fundamentals because of the habit. We do not see a positive correlation in the data; in fact, we see a negative correlation and again, the model does a good job of picking that out.



What about the correlation between consumption growth today and consumption growth in the future? The model predicts cycles of increases in consumption and decreases in consumption. That cyclical pattern in consumption produces this negative correlation between consumption growth today and consumption growth in the future. That negative correlation is once again well picked up by models around $p = 20$. It is not predicted if you assume that the agents in the economy are perfectly rational. With perfect rationality, we end up with a positive correlation in consumption growth because of the slow adjustment in consumption. That positive correlation is very much counterfactual; the actual correlation in the data is around -0.25.

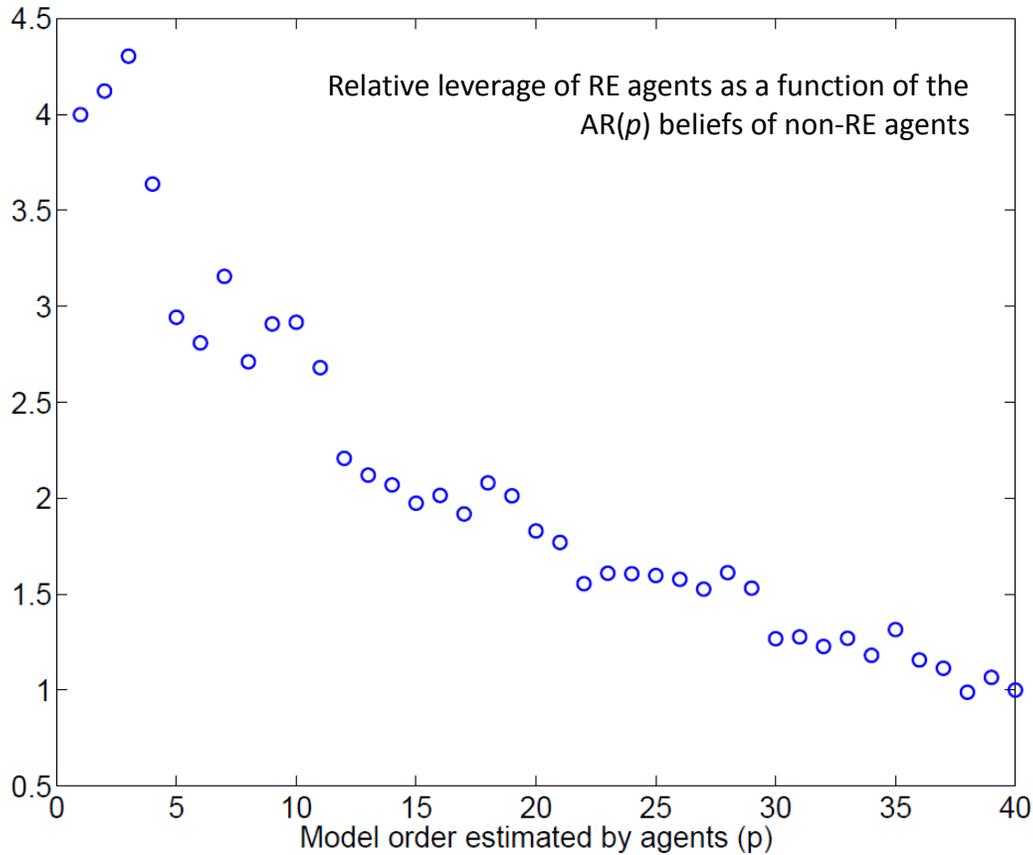
Equity Premium

In this world, agents radically misperceive the riskiness of equities. Agents think that when there is good news, equity prices are going to go up and stay up, and when there is bad news, equity prices are going to go down and stay down. This means that agents think there is a lot of volatility in equity markets. If they were more rational, they would realize that good news and bad news tend to mean revert, so equity prices should move very little since they are the discounted value of the stream of future dividends. Agents think that equities are going to be much riskier than they actually are, because agents are overestimating the degree of persistence in the earnings shock and the degree to which consumption will eventually respond to those shocks. The actual risk measure in this economy is about 1/9 of the level of the perceived risk level, because agents are overestimating both the long run volatility of consumption by a factor of three, and the long-run volatility of equity returns by a factor of three.

Implications for Rational Agents

If you understand these relationships, then you want to hold equities, because you realize that everyone else in the world is inappropriately fleeing this perceived risky asset, when it is in fact much less risky in the long run than the agents perceive it to be in the equilibrium.

In our calibration where non-rational agents make forecasts using 20 lags, the rational agents (RE agents) should hold about 50% more equities than normal agents. They should be levered up relative to the rest of the population and they should have a counter-cyclical strategy. A counter-cyclical strategy here means that whenever the economy is booming, the rational agents lighten up equity exposure, because they realize that other agents are over-anticipating good earnings events in the future. Likewise, when there is a bad set of events in the earnings process, rational agents should increase their exposure to equities, because rational agents now recognize that other agents are mistakenly inferring that the bad news is persistent when probabilistically it will be transitory.



This plot shows what the rational agent should do in this economy as a function of what the non-rational agents are doing. The horizontal axis again represents the p value of the non-rational agents. The more sophisticated the non-rational agents become, the less equity the rational agents should hold, because there is no mistake to exploit. But as the non-rational agents become less and less sophisticated, the rational agents begin to exploit their mistakes and lever up in their equity exposure. For example with a p value of 20, the rational agent should be holding about 1.5 to 2 times as much equity per capita as the non-rational agents. The rational agents exploit the mistaken beliefs of the non-rational agents, and are able to take advantage of their mistakes by holding an asset that everyone else is mistakenly scared of.

Conclusions

In this paper, we assume that fundamentals follow hump-shaped dynamics, with short-run momentum and long-run mean reversion. We assume that agents use simple models that turn out to miss much of the mean reversion in the long-run. This means that they fail to anticipate the degree to which good news today is followed by the giving up of that good news in the future. If they fail to appreciate that pattern, they are going to misprice assets like equities. So we find that equity prices are overly responsive to news events. Equity prices exhibit excess volatility and long-run mean reversion. We

see predictable cycles in consumption and investment. In the paper there are lots of predictions about covariance dynamics that the model gets right, and that are quite paradoxical and show up in the data. Equity will be perceived as many times riskier than it actually is by non-rational agents. They will be terrified of equities, when in fact they should recognize that equities are a lot less scary than they think. Rational investors should hold more equity than typical investors, and rational investors should invest counter-cyclically; in other words, when times are good lighten up, when times are bad, double down.

DISCUSSANTS

WILLIAM GOETZMANN: I got very excited about this paper, so some of what I have to say is just plunging ahead to see how much of this matches up against data I have access to. This is a great idea, and I thought it was beautifully executed. I think the nice thing about it is the elegance, and simplicity of this crafting of this behavioral limitation. We have lots of behavioral models, and they can get pretty complicated. This model really nicely says that the world has long cycles that people do not see, and that behavior leads to all sorts of predictions. Many of the predictions in this paper match what we have already seen in the U.S. data over time. I will present a bit of challenging evidence about beliefs, and then some supporting evidence.

Since this is Columbia University, I want to pay homage to Graham and Dodd and remind everybody that this is a very old question and a very old model that goes back at least to these value investors of the past. As Graham says, "It must be remembered that the automatic or normal economic forces militate against the indefinite continuance of a given trend." This is one of the secrets of Graham and Dodd: don't behave like a 20 *p* idiot. It could be 40 *p*. Then you would be getting it all wrong, and you would be handing your money over to Graham and Dodd, because after all they made hay out of this as investors.

One of the interesting questions posed in this paper is how beliefs are formed. In other words, who are these agents that are looking at 20-quarter lags and 40-quarter lags? This is rather abstract. I don't do my econometric analysis for my investment portfolio with 40 lags in it, because typically you do not have that much data and it is very hard to estimate. Most people and most investment managers, to some extent, outsource their predictions. Even if they have an in-house economist, they look at what other economists are saying about the economy. Actually in the U.S. we have some information about what economists predict is going to happen with all sorts of things. I am primarily going to look at GDP, because it is the data I had available and I have been interested in it.

The basic model is that economists make the forecasts, they tell people their forecasts, and people believe them and act on them. The Livingston survey is a panel of

economist that goes back to 1947, so we can look at the median forecast of the future changes in GDP through time. They essentially forecast a six-month horizon of changes in GDP, so we can then ask if these forecast were short-sighted. In other words, if there is some structure to their forecast compared with the reality, we can see whether they're making mistakes.

This is based on a paper by Campbell and Diebold (2009). One of the things you would expect is that these forecasts are noisier at the beginning of the series when they did not have a lot of macroeconomic data to estimate 20 lags or 40 lags, and that seems to be the story in Campbell and Diebold. Maybe those economists were using some kind of econometric model, but back then it was really hard to estimate these models to begin with. I want to focus on the forecast error. The error takes the economists' forecast and compares it to the realization. You would really like them to be centered on zero, but you see deviations from that. Campbell and Diebold have documented a positive autocorrelation in the structure between the estimated forecast GDP and the actual GDP. That tells you that there might be some basic flaw that is not taking into account the full lags. That is nice empirical evidence, but I think there's room here to play with that data.

I want to talk a little bit about some counter-cyclical allocation evidence. Campbell and Diebold find in their analysis that when GDP forecasts are low, for the next six months, the future realized returns are high. This is more empirical evidence that you can take to your model and see whether it is confirmatory or contradictory. It is consistent with a premium for holding stocks in a low consumption state, so it would be nice to know how that fits with your model. On the other hand, the sign on the CAY variable is positive, so they have their own interpretation of time-bearing risks versus time-bearing risk aversion.

Now I'm going to leap to some data. I want to talk about learning the true model. In 1947, how much data did we have? Not very much. Working with Sebastien Pouget at Toulouse and David Labrie at Bordeaux, we have collected data about the Bazacle Company that extends over for hundreds of years, from 1372 up until 1949. This was a grain milling company that was located in Toulouse. When I read this paper, I said, this is great, because now I can go to a situation where there is as much macro data as you could possibly want and it is out of sample, because it is completely different than the U.S. economy, so we can check if people are making the forecast we expect. I did some yield regressions and autoregressive estimates. We constructed dividend yields using the price process and dividend process back to the 1300s. I wanted to see whether there might some empirical evidence in support of these very long horizon models. I estimated to the full 10 years, but I found that the autoregressive coefficients were significant out to 7 on the change of the dividends. What we are finding is that there is long memory in the changing dividend process, so if you were just going to focus on the four or five year, you would miss these three lags and you would be the same kind of challenged agent that we see in this paper.

Let's take a look at regression of returns on dividend yields, and then future dividends on current dividends. In other words, do dividend yields predict future dividends and future returns? The results were really a big shock. We have a strongly significant relationship between the three-year and five-year horizon returns with high explanatory power. This actually looks a lot like what you get in the U.S. data. There is a little bit less explanatory power for the dividend predictions themselves. There are all sorts of caveats on this, of course. I just did this in the last week or so to try to take this model to the data in a situation where we might expect that if this was prevalent it would develop over the life of this one particular company.

JAN SVEJNAR: One way to look at this paper is as an alternative theory of how agents behave, or how we take on the rational expectations model. It is pretty fundamental. There is support using the U.S. annual and quarterly data. I find the results quite credible, even if not yet fully conclusive. These are reasonable conclusions, and obviously there are issues for discussion and future research.

The fundamental question is how an economy behaves, particularly if the fundamentals are this hump-shaped type that exhibit plausible momentum in the short-run and partial mean reversion in the long-run. Here, I think the authors are very good about saying that this is plausible, and not saying that this is true. Indeed, they are showing quite well how one should treat it at least as plausible. The agents here do not know how the real world fundamentalists behave, so they base their beliefs on this approximation of it, which are these parsimonious models that are fitted to the available data. So we have a war of natural expectations against rational expectations, so to speak. These really are two competing worlds, so you can see that this is a pretty fundamental piece of research in terms of how we view the economy behaving, at least the U.S. and maybe the French economy.

A summary of the five major points: we begin with the pickup of the short-term momentum in the fundamentals but then failure, at least completely, to capture the long-run mean reversion. Then there is the volatility of asset prices and partial mean reversion, so we have overreaction. The real economic activity that would come out of it, then, has more amplified cycles than there were before. You have these large equity premiums, because the agents are perceiving equities as being way too risky. The results in the paper are presented in this comparative perspective, so you always have the natural expectations agents compared with the rational expectation agents. The rational expectations agents are the sophisticated agents, and they will in the end hold far more equity than the natural expectations agents who hold these parsimonious expectation-type models.

The methodological approach is the assumption that the natural expectations agents who populate the world behave as if using this parsimonious model fits the available data. These agents are then embedded together with the hump shaped

fundamentals that are taken as plausible into a consumption model with asset pricing. The methodology uses a particular type of preference and constant absolute risk aversion. You also have the habit model to slow the adjustment and the consumption to have the smoother curvature that we saw in David's presentation. The authors then derive and calibrate the model and perform the simulations in the empirical evaluation.

There is a real appeal of these simple models. It was not in the verbal presentation, but is nicely described in the paper, that you can have different ways to come to this. You can have a statistical approach, where there is a tradeoff between the goodness or fit of the model and parsimony. The major emphasis is probably placed on the psychological considerations, motives, and behavioral underpinnings of why the agents might behave like this. But there is other literature that is discussed as well. I just selected one piece that is interesting to describe the methodological distinction here relative to the literature. For example, if you would think the bias here is driven by the fact that not enough lags are being included in the forecast variables, let's say 20 lags instead of 40, the other literature might look at it from the standpoint of omitting relevant variables from the forecasting equation.

I think there is one big challenge here, so although I think the paper is terrific I will focus on one thing. We looked at the difference in the point estimates, but David also presented the associated confidence interval bands. This is basically obtained from a boot-strap-type exercise that assesses the statistical significance and the potential bias of the authors' results. You get these incredible predictions that these five sets of predicted effects are confirmed qualitatively in the data, but there is this 'but'. The authors are very up-front about this, since they acknowledge that the results in a sense are statistically weak. In other words, there are large standard errors. The precision of the estimate is not as strong as one would like to have, presumably due to the fact that there is only a limited span of data. The authors are very good about going as far back as they can with the annual data to 1929. They go quite far back with the quarterly data. Given what they do within the model, you basically get a strong indication of these results, but if you take the results strictly on statistical grounds you cannot reject the classical models that would be consistent with rational expectations. In a way, they are saying that they should be able to do it, because the paper is terrific.

There is a question of what could be done about this. Are there alternative tests that one could use? Could one try another, maybe simpler model that would in some sense give a stronger way to contrast the rational expectation with something? Obviously, additional data would be terrific. In the end, this is what will answer the question as whether the world behaves according to rational expectations or is it natural expectations.

This is an extremely fine paper, with an obviously very important new concept that could be path-breaking. It is a very nicely developed, elegant theory, and the presentation is very good. It is just a pity that it is not a conclusive test. Although it is definitely a valuable base for discourse and further research.

FURTHER DISCUSSION

AUDIENCE QUESTION: This morning we had a paper contrasting momentum and value strategies. One of the big messages of the paper was that if you are an investor with a very short horizon, you are better off pursuing a momentum strategy than a value strategy. Here we see that the rational investor is pursuing a value strategy, but there is no word anywhere about momentum. Yet, when we see the patterns in prices going up and only slowly mean reverting, it suggests that if you were rational, you would ride the bump first and only later exit. It seems like if you are rational, you should begin with a momentum strategy and then follow up with a value strategy.

LAIBSON: Let me take that question apart in a few different ways. First, an empirical observation: momentum patterns are very strong in the cross-section. When one thinks about excess returns, a cross-sectional momentum strategy—where you go long the stocks that have done well recently, and short the stocks that have done badly recently and hold that portfolio for six months—does very well. This paper is about time series evidence, as opposed to the cross-sectional patterns. In the time series, there is only the weakest evidence for momentum effects. Let's say you adopt a strategy where, when the whole U.S. stock market is doing well, you increase your leverage and hold more stocks; and when the whole U.S. stock market is doing badly, you decrease your leverage and hold fewer stocks. Even in the back sample, that strategy basically does not perform. To the extent we believe in momentum, in the existing empirical finance literature, we believe in momentum really only in the cross-section. There is very, very little momentum in the time series.

This paper actually explains why you might expect there to be very little momentum in the time series, which is that it is very easy for even naïve forecasters using very simple models to see the momentum patterns in aggregate variables. We basically explained that there is not much momentum in the time series returns, because the thing you would be exploiting—the momentum in aggregate earnings or aggregate GDP—is in essence transparent to even second-rate forecasters. Hence, for an aggregate strategy—meaning an asset class strategy, for example equities versus bonds—I do not recommend momentum strategies. I certainly agree that there is overwhelming evidence in favor of momentum, not in the aggregate, but in terms of cross-sectional strategies. If I have already decided that I am going to hold equities, then, sure, maybe I want to build portfolios that heavily weight momentum in the next six months and then mean reversion in the next six years.

AUDIENCE QUESTION: Supposed the true model is $p = 30$, and I over-fit and estimate with $p = 40$. Why do I do that? Because David Laibson told me I should! What happens then? Is this symmetric, and does it lead people to overestimate mean reversion? Not necessarily like I said, but there must be a statistical theory that says what occurs if p is too great.

LAIBSON: Letting go of the particular numbers, if I'm wrong in asserting that there is a lot of mean reversion out there because these models with many lags are the wrong models. And in fact, the right models are the models with fewer lags, then you are going to draw incorrect inferences. We agree that there is not enough data in the U.S. to say definitively what the right structure is, and a lot of the conclusions that one draws depend upon the parametric structure that you put on the estimation process. If I incorporate models that allow me to have sinusoidal asymptotic dynamics, then I can easily pin down powerful mean reversion by picking the right model class. If I pick a different model class, I may end up wanting to pick a model with only a few lags.

A major research question is trying to figure out what the true dynamics are. I think we need to go beyond the kinds of auto-regressive moving-average models I have been talking about today and that impose very little structure. Instead we should contemplate models with more structure where we have priors that are hump-shaped dynamics. When I write down a hump shape dynamic prior, I can estimate a model with 4 parameters rather than 40 to learn what those run asymptotic dynamics look like. I think we will make a lot more progress in the next phase of research by moving out of the prior list framework, like the auto-regressive moving-average models, and moving into models where we put so much structure on it that with only 4 or 5 parameters, we can estimate the long-run mean reversion because we have the sinusoids built into our impulse response functions.

You point to a big open question. Not only should we be looking at different model classes, but we should also be looking at data from many economies. I see powerful evidence of mean reversion when I look at developed economies around the world. You can look at the incredible comebacks after World War II. Cities in Japan that were absolutely destroyed had the same relative GDP within the Japanese economy that they had before the war. They basically came back to exactly that relative position within the Japanese economy. Cities that were absolutely decimated by bombing in World War II and cities that were completely ignored by bombing ended up with the same fraction of GDP 30 or 40 years after the war. There are so many different examples of bounce-backs after crisis, natural disaster, and war. So many countries return to their familiar position in the hierarchy of nations, in terms of output per capita. This makes me think that there is a lot that we can do looking across different countries' experience to identify finer estimates of these mean reverting patterns. There is lots of work to be done.

AUDIENCE QUESTION: For the sake of the argument, let me defend the vision of what you call a second-rate or silly forecaster. Your model of 40 lags fits the last century very well. Now looking ahead for the next century to come, what are the best forecasts for the next 50 to 100 years? Your result that 40 is the best lag for the 20th century actually encourages me to think that next century will be 20. Why? Because the speed at which information gets integrated into the price has at least doubled and is going to decrease further. I would actually defend that $p = 20$ is the best estimate for the next century.

LAIBSON: That's a cool argument and an interesting way to go. I'm working on another paper with Benjamin Hebert and Michael Woodford, where we actually derive the optimal number of lags. 40 was just illustrative, and there was nothing optimal about it. We noticed in the historical data that the asymptotes seem to occur with that many lags, but there was nothing magical about 40.

I am not an econometrician, but all of the advice that formal statisticians give about picking the number of lags in these kinds of models, is based on short-run forecasting goals, for example the Bayesian information criterion or the Schwarz information criterion. In all of these frameworks, the algorithms for answering how many lags we should have, are optimized with the assumption that your goal is to do short-run forecasting. All of those answers tend to be that you should use a few lags.

The point of this conference, as I understand it, is to think about long-run investing. If you are thinking about long-run investing you probably want to think about long-run forecasting. If your goal is to do long-run forecasting, you have to throw out all the things that you were taught in econometrics classes, because they are addressed to the short-run forecasting question. We are deriving long-run forecasting norms, or optimal algorithms for picking the number of parameters in a model, and then it turned out that we want to use a huge number of lags. For example, depending upon calibrations, you might want a number of lags equal to something like a third of the number of data points in your data set. It is one third of n , where n is the number of data points that you have available to you. Of course, you do not use enough lags to cover your whole database, because then you have an over-fitting problem that has gone out of control. But as a first approximation, do not think that the number of lags should be some number like 10, 20 or 30; instead think of it as some fraction of your total sample where the fraction might be one-quarter or one-third.

Then you start to realize that we are way off in terms of the norms that our profession has adopted. If you go into the data research group of any major bank in New York, they estimate using a maximum of 15 or 20 lags. In econometrics textbooks, you see a maximum of 8 lags max. In the economics literature and in the journals, you see 8 or 12 lags, max. I think we are off by an order of magnitude in terms of the number of lags we should be using when doing long-run forecasting.