



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 II

A THEORETICAL ANALYSIS OF VALUE AND MOMENTUM STRATEGIES*

(Vayanos & Woolley, 2012)

Paper Presentation: Paul Woolley

Senior Fellow, London School of Economics

Discussants: Jakub W. Jurek

Assistant Professor of Economics, Princeton University

Tano Santos**

David L. and Elise M. Dodd Professor of Finance, Columbia Business School

PAUL WOOLLEY: This is work with Dimitri Vayanos, also at the London School of Economics. Because you have me here, there are no equations and it's easy to understand. Dimitri is much more theoretical and mathematical, but he actually explains things far better than I do.

With those disclaimers, I want to talk about a model of asset pricing that recognizes principal-agent problems. We find this explains capital market phenomena that have baffled the efficient market hypothesis and the prevailing paradigm. Moreover, it implies that there has to be a complete reappraisal of asset management strategies. In particular I want to give just one example that shows how long-term savers have become short-termist and lose out.

The Efficient Market Hypothesis

* *Remarks have been edited for clarity.*

** *Remarks delivered by Patrick Bolton on behalf of Tano Santos*

The efficient market hypothesis has been around for 40 or 50 years. It assumes that competition delivers asset prices that reflect fair value, that markets are self-stabilizing and that there is no scope for excess returns for intermediaries. But it is a special and limiting case, rather than general theory. The natural sciences have their own special limiting cases, like vacuums or zero friction, but physics and engineering obviously have to relax those assumptions and go into the real world.

The efficient market hypothesis, however, is doctored lock, stock and barrel by asset management and policy makers. This special limiting case, rather than the general theory of asset pricing, informs for example the use of security indices as benchmarks for performance and for passive investing. It also shows the basis for evaluating and controlling risk. The efficient market hypothesis is also used to optimize asset allocation and as the basis for diversification.

Almost everything it does is actually diametrically opposite to what asset managers should be doing if markets are inefficient. We are in this strange world, where we know the markets are inefficient, which is why we hire managers to exploit the inefficiency. But then all the instructions and delegation is conducted using theory based on the efficient market hypothesis as the instruction manual for how then to proceed.

This gives rise to inappropriate actions and very costly errors. What you really need is a theory to explain mispricing, to inform investors how to proceed in an inefficient market. You do not want them using a manual based on efficiency to conduct themselves in an inefficient market. This leads to inappropriate actions and costly errors by those responsible for investing funds. The efficient market hypothesis has also little to say on the selection of strategies or investment horizon.

In our work, Dimitri Vayanos and I introduce intermediaries. We introduce institutional arrangements by showing that there is delegation, and that delegation creates asymmetric information. The agents have better information; that is why they are hired by the large funds, which we will call the principals. The principals are also agents in a form themselves, but let's call the principals the giant funds: the pension and sovereign wealth funds. They are asset caretakers, rather than asset earners, but they are the principals in this context because they appoint agents and delegate. They delegate because the agents have better information, but they also have different objectives from the asset caretakers. Most importantly, the principals find it difficult to determine the competence and the diligence of the agents.

Principal-agent problems give rise to two phenomena. The first is asset mispricing, and the second is rent capture by agents, but we'll focus on the first today. As Patrick Bolton was saying earlier in his introduction, I was an asset manager for 20 or 30 years. I and my partner thought the markets were becoming more efficient in the 1990s, and then, of course, we had the tech bubble. As value managers, we were underweight in the tech stocks, even though we had some anchor to windward by holding a portion of momentum stocks to try to control risks and to comply with the contracts with our clients. We savagely underperformed the S&P and the FT All-Share

Index over two years, and we were fired by a good portion—or bad portion—of the clients. But we battened down the hatches and tried to ride the storm out. It really did feel like going around Cape Horn. In fact, if it had gone another year, we would have had to start throwing barrels of gunpowder and rum overboard. It was a grim time, but fascinating time, because it was like a scientific experiment with the volume was turned up. We were observing how delegation worked and doing what we thought was right. We followed a value strategy. There was a massive over-valuation of tech stocks. We had some exposure to them to reduce the risks, but we underperformed by 20%. But we stuck to our guns, and later we recovered three times what we'd lost over the next five years.

Most of the money had gone out after we'd done badly, and most of the money came in after we'd done well. I thought that the whole issue of risk momentum and the principal-agent relationship was fascinating. It seemed to me that we were acting rationally and optimally as the fund managers by sticking to value with some exposure to momentum, but the clients were also struggling to understand whether we lost our competence or diligence. We might have had a good record in the past, but this was a new world, a new paradigm, and we were savagely under-performing. They couldn't bear the pain, so a lot of clients departed.

Essentially, this is the basic intuition for the model, in this case in terms of a negative shock. A negative shock hits an asset's fundamentals, and those funds holding the asset realize poor returns. The funds experience outflows, and therefore the manager sells the asset. So in fact, asymmetric information is causing the amplification of shocks. Because it is a negative shock, the price of the asset or group of assets goes down. The fund flows cause the managers to jettison their holdings; they may even do it preemptively because they want to avoid being fired. The outflows are gradual, which is important, because economics and finance have difficulty handling time scales, but in our paper we show that this is a gradual fall or a gradual rise. These things do not happen instantaneously.

Formal Model

This is a formal model conducted in a rational framework; it is not behavioral. Like my experience, everybody is acting rationally and optimally in this model: the investor, the principal, the active fund manager, or the agent. For simplification, we have the investor selecting between an index fund and an active fund manager, but we do have multiple securities. That's hairy enough in mathematical terms, but it all comes out very cleanly. Essentially, we have a single proposition in the model: there is delegation, and it is coupled with uncertainty about ability. The math simply follows that logic through in the environment of the model. That gives rise to a lot of math, but it helps to verify intuition. The math points to things that you would not have expected otherwise.

The important thing we show is that if you introduce delegation, asset prices in general depend on fund flows as well as cash flows. You can think of this in terms relating to stocks or asset classes, currencies, commodities, and so on. The efficient market hypothesis just focuses on the value of the asset being equivalent to the expected future discounted cash flows of the asset. We show that the prices also depend on fund flows, for the reasons I've explained.

I set up a few centers for the study of capital market dysfunctionality, just before the world fell apart. I had always thought that dysfunctionality could be explained with a rational approach, but I thought it would be in bits and pieces. I thought there would be a set of explanations, each for different circumstance. Actually once you introduce delegation, it turns into a unified, tractable, extendable model. It is a new way to look at asset prices. We have set this out in technical papers as well.

Because of my experience in the tech bubble, I felt that if you could explain momentum, you were closer to achieving the objective of a new understanding of asset pricing. After about a year, Dimitri and I alighted on the methodology I just explained. We spent the next two years discovering how many other things it explained. It described the value phenomenon, the growth phenomenon in a rational framework, post-earnings announcement drift, co-movement leads and lags, tracking error, commercial risks, and the list goes on. It is quite extraordinary how illuminating the model is for many phenomena that have not been addressed by the prevailing paradigm.

Implications for Assets Owners and the Long-Term Horizon

Principal-agent problems give rise to asset mispricing and rent capture. The solution actually lies with those who are suffering—the caretakers of capital. Regulation is unlikely to sort out some of the problems that arise because of the bubbles, crashes, and mispricing that derive from the process that I've described and from the way markets function. This has important implications. Currently everybody is basing their risk analysis, benchmarks, diversification, performance, the instructions they give to the managers for the strategies to be adopted on the efficient market hypothesis. But the new model has important lessons about how that behavior should change and what should replace it. For example, it suggests that benchmarks should not be based on the distorted pricing of indexes, both in aggregate and stock by stock and sector by sector, but by underlying cash flows, for example dividends. Similarly, for risk and diversification, the risk analysis should be conducted not on the basis of distorted prices, but on the basis of the returns from the different asset classes.

Let me just focus on one very important aspect: the investment horizon. The efficient market hypothesis is largely silent on horizon. The model that incorporates delegation provides a very helpful answer to an age-old puzzle: is the long run equivalent to the succession of intervening short runs? Expressed another way, should the investor focus on getting the best return each year or the best return over the long

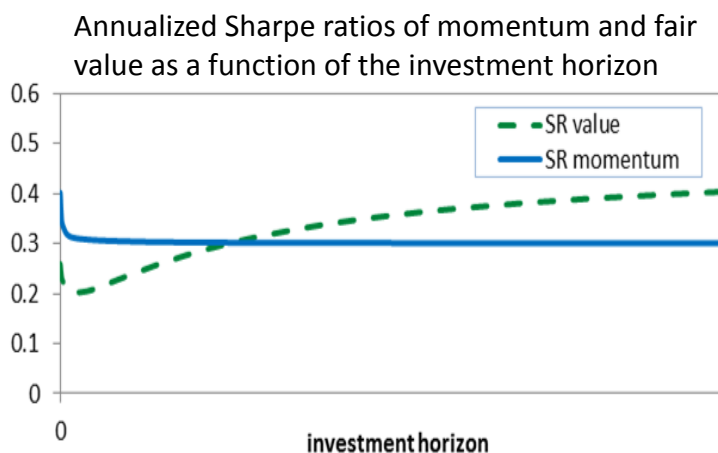
run? I have been fascinated by this puzzle and the issue of efficiency in markets the whole of my career. I'll explain the answer we provide and see what you think.

So do you focus on getting an investment return each year or an investment return over the long run? Horizon addresses not how long you hold a security, but the basis on which you hold it. The answer lies in the choice of strategy and the returns for those strategies. In fact, the whole of asset management is dressed up in many ways, but it boils down to either momentum, which ignores fundamental value and is simply trend-following; or fair value, which disregards fund flows and buys based on cash flows based on the efficient market hypothesis. As we all know, markets often demonstrate trends which persist and are ridden by those seeking excess returns and indeed reducing risk. Most active management is a combination of simply momentum and fair value.

Currently, the only way to explore the returns for these strategies is using historic data. But if you have a model of asset mispricing, you can actually derive the risk-adjusted returns from combinations or stand-alone strategies of momentum and fair value. You can do it not only in one period, but multiple periods with continuous rebalancing and different implementations. To have a model that explains asset mispricing is a powerful tool because it enables you to analyze *a priori* the returns to the different combinations of strategy.

Just for a moment, we are going to insert the investors, the price-taker, into this world. We drop an investor by parachute into a world where prices are turned by cash flows and fund flows. What are the returns? Calculating the returns based on optimal implementation of both strategies in the short-term, the key finding is that momentum dominates fair value. If you have a short horizon, it actually pays you to use momentum. Trend following is the best strategy for short-term investors, assuming they are lucky or good in their timing—that's a big proviso. If you said you were going to shoot me in three months if I haven't outperformed, I would use momentum, because the model demonstrates that trends continue until they end, and then the prices revert to fair value. That is why the momentum investor enjoys the ride for a while and tries—or hopes—to get out in time, and then the fair value starts to click in.

Fair value investors have to exercise patience. For a long-run investor, fair value



dominates momentum for long horizons. The results hinge on differences in risk characteristics. Momentum is a series of uncorrelated bets; the long-run risk of momentum is equal to some of the intervening short-term risk. Fair value

benefits from negative serial correlations of returns. Long-run risk declines over time. Of course, that is well known, but this is now contained in a formalized asset pricing model. The blue horizontal line is the Sharpe ratio of momentum strategy over different horizons. The dotted green line is the Sharpe ratio of value. Because long-run risk is less than the sum of the intervening short run risks, it is a rising trend. You can clearly see that the blue line dominates the green line in the short run, and in the long run it is reversed. If you take nothing else from what I say today, that is worth absorbing. Of course, the shape and the crossover point are sensitive to execution and calibration.

So it is well known that there is a low correlation between momentum and fair value, and therefore it does pay to combine the two. If you are obliged to use one or the other, then clearly it depends on your horizon. The model in effect provides definitive solutions. It shows that the optimal mix of momentum and fair value depends on the term of the liabilities and the horizon. The longer the horizon is, the greater the allocation to fair value.

You might say, well we don't use momentum. But the point is that any principal giving instructions to an agent that incorporates tracking around a volatile benchmark will inevitably use momentum. That is why, in the tech bubble, those who varied the composition of value and momentum had to increase their exposure to momentum simply to buy the stocks that were going up, because they were killing the performance and causing the tracking error guidelines to burst through.

Caveats and Implementation

Fair value is robust to the choice of inputs. Actually, you can use a blunt instrument and be relatively successful in fair value, but momentum is highly sensitive to execution. You need to look back on the holding period just right. Historic data flatters momentum. Momentum in stocks tends to have a period of about 6 months or 18 months; therefore, the annual data always flatters momentum because it is so sensitive to returns. The Sharpe ratio of momentum is a function of the look-back period. Get it wrong, and you get negative returns.

So fair value investors win overall. In fact, momentum investors are always buying after the price has gone up and selling after it has gone down. They are destined collectively to fail, and the fair value investors pick up what they lose.

DISCUSSANTS

On behalf of Tano Santos, PATRICK BOLTON: My friend and colleague, Tano Santos, had an unexpected conflict. I will do my best to convey his comments. Much of Tano's research is on asset pricing and related issues. Let me start with a brief summary of the paper. The summary will include some mathematical symbols—not nearly as much math as in the paper, but a little bit more than the presentation.

The theory of momentum, as Paul emphasized, is a rational theory based on gradual outflows from actively managed funds; that's what makes it a special theory. He highlighted the broad intuition behind the outflows of funds, based on the principal-agent friction.

So what motivates trading by investors here? Why would investors want to trade? Paul has emphasized some of the reasons, such as news about cash flows, rebalancing, shocks to the discount rate, liquidity shocks, and so on. This is summarized in the model as a shock to the opportunity cost of holding shares in the actively managed fund. Then crucially, once an investor faces this shock and wants to rebalance, it takes place gradually. The paper includes a gradual adjustment model, which takes this quadratic form, where the ψ parameter plays a critical role.

On the agent side, the model has investors and managers of actively managed funds. On the asset side, you have three asset classes: a safe asset, an actively managed fund, and an index fund. Investors can allocate their savings into these three asset classes. Then, in short, what happens in the model is that when investors choose to sell or change the composition of their investments in the active fund, they trade with the managers. This is just a device, or a simplified assumption. It is not literally like this in practice.

The challenge with momentum, which this paper takes on, is that if someone sells, and if it is anticipated that returns will be lower for that asset, who is going to buy on the other side? Who wants to pick up an asset where the expectation is that returns will decrease in that asset? That is the challenge for momentum theory to address. The answer here is the bird-in-the-hand effect. This means that the opportunity, which comes from the outflows of funds from the active fund, depresses prices. The opportunity is short-lived, because there is going to be price reversal. The question is when that price reversal will happen. Because you do not know when the reversal will happen, you face risk with respect to it. With the bird-in-the-hand effect, you choose to buy before it is too late. More elegantly put, the manager holds the asset as a hedge against adverse shocks in the investment opportunity set. The paper delivers momentum effects this way and considerable Sharpe ratios.

Let me contrast what the paper does with what has been done before, specifically the behavioral approach to momentum, which is based on the idea of under-reaction to news. Here, there is nothing like that explicitly; I will comment on that a bit more. On the positive side, this paper is an impressive achievement. It is an elegant model that highlights necessary ingredients for momentum to appear, given a reasonable description of the stock markets, which are dominated by institutional trading.

Tano's comments are really more about what next could be explored, so this brings me to the key parameter ψ , which governs the rate at which money flows out. It is calibrated to replicate the flows that we see the data. This is unlike what it does in asset pricing when it comes to, for example, determining the risk premium where the

calibration is done by choosing a particular coefficient of risk aversion for the representative investor. In those exercises, we have a fairly clear idea what, based on other studies, a reasonable number for a coefficient of risk is. Here, it is not obvious what leads us to think that we have a reasonable parameter for ψ . The question is what we can rely on in terms of identifying a good ψ number and whether the theory works well. That is probably where behavioral considerations will have to be brought back in, as Tano's comments suggest.

I want to bring up a couple of other points from Tano's comments. The first point is what Paul alluded to: this model allows you to look at risk-adjusted returns, as opposed to just looking at historical returns. This is what Tano pushes the authors to pursue. Tano looked at some of the data on momentum returns from Ken French's website and ranked ten momentum portfolios. He did this exercise. He found that on the lowest momentum portfolio, you have very large negative a when you do the risk correction. On the highest portfolio you have positive a . This means that there is no risk compensation explanation for momentum returns. Tano suggests that you do this in the model using simulated data. Would the model reproduce what we see in the data? Another way of putting it is that you can decompose momentum returns into two components, covariance of the returns with the market and a dynamic component, and then ask if the model reproduces most of the returns based on the first component or not.

My last comment again uses Ken French's data for a calculation that Tano did. In Tano's exercise, he saw that cap-end β moved a lot over time. Will the model deliver a similar pattern? We also know that regarding momentum returns, there are tail events, such of the crisis of 2007, where you have sudden very massive losses. It is a natural question to ask if we are seeing these returns reflecting some kind of correction for a tail risk that you might be taking on.

JAKUB JUREK: This is a long paper with lots of equations, but I am going to assume that they are all right and focus on some of the practical implications of the theory for what the fund managers in this room are hopefully trying to accomplish.

The paper is about two investment phenomena or themes, value and momentum. Value and momentum essentially appear across a variety of different asset classes, which we as academics like to think of them in the context of equities. You can look at bonds and currencies and see that very similar factors tend to pop up in many of these places. There is a recent paper by Asness and co-authors that effectively explores this across asset classes.

In a broad brush, what are value and momentum? Value and momentum are basically two statements about how prices behave over long horizons. Value essentially says that over long horizons, I expect to have mean reversion, such that volatility at longer horizons is smaller than volatility at shorter horizons. Momentum is the compliment of that; it is about positive autocorrelation. Momentum tells me essentially

that what happened yesterday is likely to happen again, so risk tends to grow more than proportionally with time for some of these assets. The nice thing about value and momentum is that, across asset classes and within asset classes, they are negatively correlated. So you can do a pretty good job of implementing a multi-strategy fund around these two themes.

This paper essentially tries to present a theory of where these effects come from. The way they conceptualize it is by highlighting the channel induced by fund flows. The intuition for the model is quite straight-forward, as Paul alluded to. The idea is that we can allow prices and values to deviate, which is a key component that justifies the existence of investment managers and is essentially a fact of life. Specifically, what happens is that the fund flows amplify these deviations between prices and values. If you have a security that is undervalued and agents are selling it, it becomes more undervalued such that in the short-run you see momentum, but because its price now deviates even further from its fundamental value, you expect to see mean reversion at longer horizons. This is essentially how they generate these two phenomena in their model. The practical implication is that if you look at assets at various horizons, they have different risk and return tradeoffs. In other words, the Sharpe ratio of mean variance and investment frontier is different depending on your investment horizon. This gives a horizon structure to your optimal portfolios. They derive a theory and talk a little about how these Sharpe ratios vary.

It is a very nice paper and really nicely worked out, but if you know that it's nicely worked out and you're an academic, then you know that you have to take a lot of shortcuts, because it is difficult to get things to work out. The paper has to make some assumptions that I do not think are particularly great. One of them is essentially relying on carry utility, which will imply that Jeremy Grantham's investment in stocks was the same when he was a 20 year old, as it is right now. That is not likely to be a true statement, although I don't know enough about Jeremy to say if it is true. When you put this in a general equilibrium context and look over long horizons, this means that risk premium in equilibrium are non-stationary. That is not a very plausible feature for the way that things have looked historically in the world. It is great for theory, but I don't think this is the right way to look at practical implications.

The second matter is whether these effects are quantitatively important. I will try to differentiate between the theory and what we can retain of the theory for practical implementation, versus how important is this in the context of value and momentum. As I mentioned, value and momentum create a term structure of a risk return tradeoff. If you have an asset that is positively autocorrelated in the short horizons—a momentum asset—it becomes riskier over a short interval of time, because it is positively autocorrelated. If you have mean reversion at long horizons, it means the assets are becoming less risky at a longer horizon. If you think about a Sharpe ratio, even if the numerator is not time-varying, the volatility is changing the denominators, so the Sharpe ratio is moving around.

If you are a one-month investor or an investor subject to monthly withdrawals, you might be thinking about a different mean variance frontier than an investor who has a 40-year horizon and is not subject to withdrawals. This essentially creates what I am calling a term structure of the risk and return tradeoff. This has been around in academia in a variety of different flavors since at least the early 1970s in a variety of different examples. It is not in Markowitz's theory, but you can imagine adjusting Markowitz in a static case where you have a certain mean reversion, so therefore means and volatilities measured over different horizons are different. The more common way that people do this in academia uses the dynamic perspective that comes out of Merton's problem, where an investor is hedging time variation in investment opportunities. Whichever perspective you use, you are going to get the same types of conclusions.

Is this quantitatively important? I think the term structure of the risk return tradeoff is there, and it is important. Now, I want to ask, is the term structure of the risk and return tradeoff important in the context of value and momentum? I'm going to look at the value effect in the context of U.S. equities using our standard value factor from Ken French's website.

		MKT	LOW B/M	HIGH B/M	Diff (HML)
Summary statistics for excess returns (zero investment portfolio)	Mean	5.5%	4.7%	10.6%	6.0%
	Volatility	15.7%	18.0%	20.6%	16.2%
	SR	0.35	0.26	0.52	0.37
	Skew	-0.54	-0.20	0.05	0.55
Autocorrelations					
Sums of autocorrelations over various windows	1M-3M	0.07	0.11	0.10	0.13
	4M-12M	-0.02	0.02	-0.03	0.28
	13M-24M	-0.30	-0.28	-0.41	-0.36
	24M-48M	-0.07	-0.10	0.00	-0.08

If you look at the returns for a low book-to-market portfolio—growth firms—over the last 50 years it is about 4.7% in excess of the risk-free rate. High book-to-market firms—value firms—are about 10.6% in excess of the risk-free rate, so value firms tend to outperform growth firms. The outperformance quite large, and it's statistically significant. This is what is frequently referred to as the value factor.

The question is, does the value factor looks like the authors actually say it should look? I did some autocorrelations at short horizons and longer horizons, which are different windows. Indeed at short horizons we observed some positive autocorrelations, so even the value factor has a little bit of momentum in the short run. In the long run, it is pretty strikingly negatively correlated, so there is mean revision. The data confesses to the theory here.

I should emphasize here there is a little bit of a difference in the way that the authors look at stocks. You would sort it into a value portfolio, hold that stock independent of what happened to it, and then study what happens to returns. That is not typically how investors think about value, in that investors re-sort stocks periodically, holding onto the value stocks and shorting the growth stocks in this example. I am thinking about this from the trading perspective of a portfolio that is being rebalanced monthly. This is an important difference between what the theory is trying to address, versus what we perhaps care about in the data.

What about momentum? If I look at momentum, you see that stocks that have

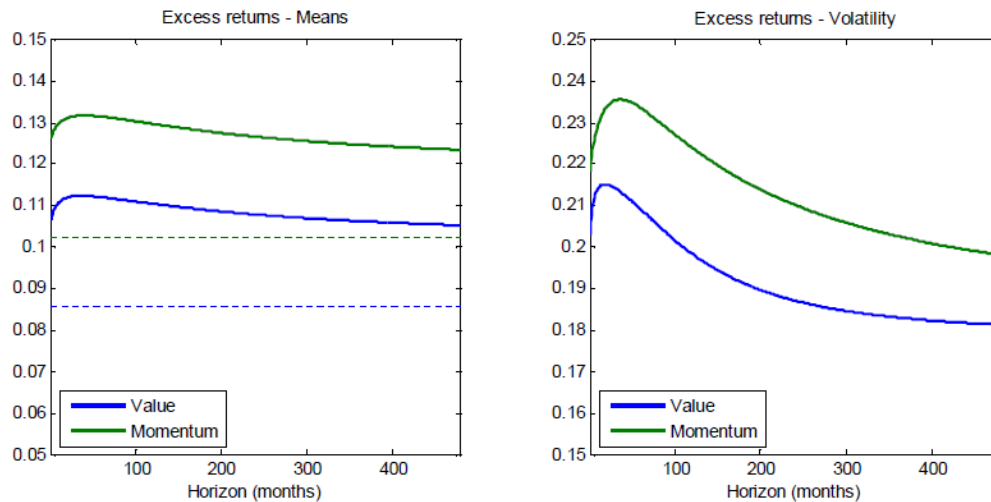
		MKT	DOWN	UP	Diff (UMD)
Summary statistics for excess returns (zero investment portfolio)	Mean	5.5%	-3.2%	13.1%	16.3%
	Volatility	15.7%	28.1%	21.8%	24.3%
	SR	0.35	-0.12	0.60	0.67
	Skew	-0.54	0.69	-0.43	-1.52
Autocorrelations					
Sums of autocorrelations over various windows	1M-3M	0.07	0.14	0.07	0.02
	4M-12M	-0.02	-0.09	-0.10	-0.05
	13M-24M	-0.30	-0.29	-0.28	-0.28
	24M-48M	-0.07	-0.16	0.03	0.18

recently done well continue doing well. They outperform the risk-free rate by 13% and have a massive Sharpe ratio of .6. The losers continue doing poorly. The difference between the winner and the loser portfolio is 16.3% with a Sharpe ratio of .7. That is the momentum effect, and it's strong and quite robust in the U.S. data. Let's look at the autocorrelations: do they confess to the fact that we have positive autocorrelation at short horizons and no mean reversion at long horizons? In fact, we see that there is no positive autocorrelation in the factor at short horizons and at long horizons. Rather than having no autocorrelation, it's actually strongly negatively autocorrelated. In the case of momentum, the autocorrelation structure of momentum is actually completely different from what you see in the theory. I would guess that the difference here stems from the fact that these portfolios are continuously rebalanced into the high value portfolio and the momentum portfolio, versus what they look at the paper, which is freezing a single-value stock at a point in time and looking at what happens to it going forward. It could have drifted from the value to the growth profile, but they are still looking at that particular stock.

I'm going to look at two portfolios, the up portfolio, or the portfolio of winners, and the value portfolio. I'm going to put them into a long horizon portfolio choice framework. Then I will ask what the implications are for the moments and for the means and volatilities, which create the risk return tradeoff, and how these translate into optimal long-horizon portfolios. The right-hand plot has the volatility of the value

The risk-reward term-structure can be characterized using a vector autoregressive model for asset returns (Campbell and Viceira (2005)).

- State-variables: lagged returns $([-3, -1], [-24, -13])$, and Shiller PE ratio (CAPE)
- Value: high book/market portfolio ("high B/M")
- Momentum: winner portfolio ("up")



portfolio and the momentum portfolio. The volatility of the momentum portfolio is greater. So the negative autocorrelation, which is not in the paper but is in the data, means that volatility of the momentum portfolio actually declines with the horizon. The same thing happens for the value portfolio.

In fact they have the same term structure of volatility. On the left-hand side, I have plotted the term structure of the arithmetic mean return. You can see that there's a little bit of variation, but it is not a huge amount. You can see that if you are a one-period investor, you have an expected return of 12.8%, and if you are a 40-year investor, you have a return of 12.4%. This is not quantitatively a huge effect either for value or for momentum on the side of the means. It is a reasonably big effect on the side of the volatility. The way you can think about this, is that both of these portfolios are becoming less risky at long horizons. If you ask what the optimal allocations are as your horizon increases, you want to allocate more both to value and to momentum. But there is no flip where momentum becomes worse at long horizons. It actually has a higher Sharpe ratio up front, so you want more of it. Because in the long run it evolves and they essentially decline at the same rate, you just allocate more to riskier assets as your horizon increases.

To conclude, it is important to think about this risk-return tradeoff and that's an important take away from this paper. We should think about it more broadly, in terms

of all of the risk factors that a sovereign wealth fund or a pension plan is exposed to. Ultimately your asset allocation and the exposure to risk factors in your portfolio primary is the determinate of performance in the long run. There is an entire landscape of risk factors—interest rate risks, equity risks, and downside risks. They come packaged in various flavors and various asset classes, but these are the risks that you are actually bearing. Value and momentum are just a small subset, which is not to say that the term structure-versus-return tradeoff does not kick in everywhere.

So which risks should you worry about? There is a big push into alternatives, and a lot of alternatives specialize in bearing downside risks. It comes in illiquidity, in the form of volatility risks, jump risks, gap risks; however you want to call it, that's essentially what alternatives bear. Let's look at which of these risk factors actually have high Sharpe ratios and which are the ones that we are likely to be most concerned about

Recent returns to tradable risk factors across asset classes (1996:1-2012:6)

- Fixed income: Long 10-year U.S. Treasury, short 30-day T-bill
- Equities: Fama-French (1992) factors
- Alternatives / Downside risk: Excess returns to naked S&P500 put-writing strategy

	B10-RF	MKT-RF	SMB	HML	MOM	Downside Risk
Mean	3.7%	5.5%	3.1%	2.9%	5.5%	10.1%
Volatility	7.3%	17.0%	12.8%	12.2%	20.0%	7.4%
Sharpe Ratio	0.50	0.33	0.25	0.24	0.27	1.36
Skew	0.08	-0.69	0.83	0.04	-1.50	-2.78

as investors. I did an analysis here with a shorter term data, and in the last 15 years, and the excess return on the long bond is about 3.7% with a Sharpe ratio of 0.5. That's probably unlikely to repeat unless people are expecting negative nominal interest rates. The next four factors are the equity risk factors, which have Sharpe ratios of about 0.3. In the most recent data of value and momentum are in fact quite a bit less attractive. This is a little bit like, as Merton calls it, positing a Connecticut Yankee in King Arthur's court. In some sense, it is not fair to look at these factors in the period before they were published. We did not really know about it, although I realize that a lot of investors did. In the most recent history, you see that there is not as much juice in this, versus if you look at the risk return tradeoff in alternatives and for what you get compensated for risk in those asset classes. Here, you see that the Sharpe ratio for the risk that you're bearing in alternatives is an order of magnitude larger than in the other asset classes.

We should think about the risk return tradeoff, but I think the places where it's likely to have most bite is these asset classes that have higher Sharpe ratios.

FURTHER DISCUSSION

AUDIENCE MEMBER: Momentum is not the same across all countries. In Japan, the momentum effect is different, less strong. I wonder whether the fund flows are different in Japan. There might be a nice way to test the theory using some cross-sectional evidence about fund flows in the momentum effect.

WOOLLEY: There has been historically far less momentum, in fact almost no momentum, in Japanese stocks until the last five or eight years. Of course, that is fascinating. It is interesting that in Japan, managers are not hired and fired on the basis of recent performance, which of course is in our model. Our model shows that the principal infers whether the agents are competent or not on the basis of recent performance and then acts accordingly. The last part has not occurred in Japan because the principals cannot bring themselves to confess that they were wrong in appointing the manager, and therefore hold on to underperforming managers. This means that there is no selling of the stocks if they are underperforming, which causes the underperformance in the first place.

Momentum started to creep in during the last decade because you have principals hiring non-Japanese agents and applying Western methods of both delegation and hiring and firing. It is interesting that we have an institutional phenomenon that causes Japan to be quite different from the West, which ties in with there being no momentum effects.

AUDIENCE QUESTION: First of all I agree with your conclusions and recommendations, but I think that the idea that the agency problem somehow comments on market efficiency is a bit of a red herring. Momentum existed long before agents. The South Sea Bubble, on which we have very good data, looks stunningly like the more violent aspects of the tech bubble in duration and magnitude. In fact, the deutsche mark went up more and declined more, and more quickly than the South Sea Bubble. And the South Sea Bubble had no agents; it was just a reflection of the herding instincts of ordinary individuals who can be driven to excess, not just in investing, but in many other instances of crowd madness.

When you move forward to look at the agency problem, you realize that what is happening, almost uniquely in the investment industry, is that we report upwards to decreasing level of sophistication. The most sophisticated well-informed players, the investment managers, report to the hired guns, who used to be very financially illiterate 30 years ago, when I started, but who now are very well informed, though less than the managers themselves. They report to committees at endowment funds and pension funds, which are absolute horrors stocked with complete amateurs; even though they may be hedge fund or private equity whizzes, in terms of a portfolio, they are even more dangerous than random selected people from the street. The threat of being fired passes down the pipeline. It is purely behavioral—the bad behavior of the individuals in their pure natural state.

In 2000, we had the best bunch of clients in the business, and nevertheless the majority of their members on the majority of the committees believed that it was a new golden era, to some extent. They happily fired us and were completely fed up even seeing us. Their very bad instincts overrode the much better instincts of the hired guns, incidentally, and some of them got fired in turn and passed down to us where we got fired. I think this is just a manifestation of the occasionally completely irrational behavior of individuals, and that the agency problem actually can be bypassed.

WOOLLEY: There are many explanations for momentum, of course. The point is, if you can explain it in a rational framework; just assuming everybody is stupid does not get you very far. I do not think GMO was stupid, and I do not think a lot of the people managing money are stupid. But there is a tiering of agency problems, so that the principals themselves are actually agents. The poor old ultimate beneficiaries—as we are in our private capacities—are the ultimate victims of this tiering of responsibility. We are the ultimate victims of a market, in which the bulk of investment is conducted without reference to the cash flows or the assets that are being invested in. Momentum is not just used to add value; it is used to reduce risk. This is all a programmed response given the framework that has been handed down in the efficient market hypothesis.

What we have done is to identify a rational model, which can explain many of the phenomena of markets, without resorting to behavioral explanations. Our work is entirely compatible with the behaviorists' work. We are just trying to devise a model that can give rise to policy implications. Behaviorism by itself has not generated a single policy implication that affects practitioners or policy makers, at least that I can think of. After 25 years, we still have policy makers, for example, imposing mark-to-market constraints that encourage short-termism by funds. We really are in a maelstrom of investing without regard to fundamental value. We need to find a model which explains some of this in a sensible, clear fashion and gives rise to policy implications.

We know there is momentum in markets, or there has been, until 3 years ago. It has actually not been a clever strategy in stock selection. If you have momentum in markets, which takes prices away from fair value, then all the policy implications that we show actually follow. We just happen to have used one particular methodology to show how it gets there in the first place.

AUDIENCE QUESTION: There are two types of intermediaries. One type is like a close-end fund, in which the agent has the discretion to select the portfolio and determine the time horizon for the investment. There, the investors in the fund will not affect their judgment, because they are supposed to do rational investment. Then, there is the mutual fund with the intermediaries that you are talking about here, in which basically the actions of the investors in the fund affect the decision of the agent. In that case, the agent only determines the selection of the fund, and the time horizon is determined by the investors in the fund. If it is an index, they do not even do the first

part. You can view that as completely transparent, basically in that the investors affect the investment at the end. If you substitute what you call a rational model with irrational investing directly in stocks with asymmetric information about the competence of the agent, then at the end it is a similar model to noise trading and to the limit of arbitrage. My question is, at the end, is this just another model of noise trading and the results are due to the limits of arbitrage, because the rational investors do not do their job?

WOOLLEY: As I said, in the short run, it's optimal to use momentum. If you have a short horizon and you have to outperform, you will use the optimal strategy, which is momentum. You will disregard cash flows and use follow fund flows. For example, it is fascinating that hedge funds use momentum to add value, because of their performance fees and short-term performance evaluation. They are very much shorter horizon investors, therefore they use momentum. A fund which is only long, where the principal gives the agent a benchmark of an index and a tracking error, then the manager will also be using momentum. So from two entirely different categories of investing, you will get momentum as the prevailing strategy.