GRESHAM COLLEGE

28 APRIL 2016

THE EFFICIENT MARKETS HYPOTHESIS

PROFESSOR JAGJIT CHADHA

"Many people did foresee the crisis. However, the exact form that it would take and the timing of its onset and ferocity were foreseen by nobody. What matters in such circumstances is not just to predict the nature of the problem but also its timing. And there is also finding the will to act and being sure that authorities have as part of their powers the right instruments to bring to bear on the problem."

Letter from British Academy to H M The Queen, 2009.

"One thing we are not going to have, now or ever, is a set of models that forecasts sudden falls in the value of financial assets, like the declines that followed the failure of Lehman Brothers in September....The main lesson we should take away from the EMH for policymaking purposes is the futility of trying to deal with crises and recessions by finding central bankers and regulators who can identify and puncture bubbles. If these people exist, we will not be able to afford them."

Robert Lucas, 2009.

Introduction

The efficient markets hypothesis (EMH) has taken a 'hell of a beating'[1] in the 9 years since the start of the financial crisis. The very idea that we thought markets were efficient would seem now to beggar belief. Sclerotic and highly volatile markets, as well as all kinds of alleged malpractices, do not seem to correspond to our innate notions of efficiency. It would appear to many casual observers that increasing levels of financial instability has run in tandem with moves to greater degrees of market deregulation and would seem to have been anything other than efficient. But that view of the hypothesis, perhaps, has in mind other notions of efficiency to do with the allocation of scarce capital, the operation or organisation of markets or portfolios that are able to minimise the variance of a return for a given expected payoff.

In fact the efficient markets hypothesis that makes no necessary claims about the social optimality of all financial markets, which may well be constrained or distorted by all kinds of incentives, rigidities, regulation or policy-induced moral hazard. It simply says any individual will use all the information at their disposal to decide on the appropriate price of any given asset and that information has been used, the price will move to reflect that information. The idea of

[1] I borrow this phrase from Bjørge Lillelien who in 1981 was commenting on England's 2:1 loss to Norway in a World Cup qualifier and said: "Lord Nelson, Lord Beaverbrook, Sir Winston Churchill, Sir Anthony Eden, Clement Attlee, Henry Cooper, Lady Diana, vi har slått dem alle sammen, vi har slått dem alle sammen! (we have beaten them all, we have beaten them all!). Maggie Thatcher, can you hear me? Maggie Thatcher ... your boys took a hell of a beating! Your boys took a hell of a beating!"

efficiency in this sense is simply that agents have an incentive to trade any information that they may have so that it becomes publicly available in the new price of the asset. By the process of providing incentives to decant private information into the public domain, the traded price is conditioned on all relevant information, and accurately reflects the payoffs and their probabilities in all states of nature.

The question posed by the Queen as to "why had nobody noticed that the credit crunch was on its way?" is, of course, very important but it is related to the build of risks over a long expansion in economic activity that one might think ought to have been reflected in the changes in some asset prices, particularly those which were vulnerable to these risks, for example, the shares of financial institutions. There are two responses. Financial prices may be informationally efficient and gauge that the probability of a bad event, involving a low or negative return, is non-zero but low and yet, after the fact, if and when that bad event occurs, which it must almost certainly if we run the system for long enough, it may look as though the probability was badly underestimated. Alternatively, the price may not adjust even if the probability of a bad event has risen because it is thought that there is some kind of private or public insurance that limits the extent of the low or negative payoff.

The efficient markets hypothesis really hinges on the incentives to trade all information and to gauge the likely returns in all possible states of nature from holding a given asset. The forecasts we make are this all conditional on states (good, bad or indifferent) and not unconditional with respect to whether a particular state will occur with probability one. The price of that asset will then adjust to reflect the expected payoff from that asset and any premium required for bearing risk. What the EMH does require is some form of rationality in expectations and negligible costs of trading. And so we shall first turn to the formation of rational expectations.

Rational Expectations

When we invest in a financial market, although we might use the history of past payoffs as a guide to expected payoffs from holding an asset,[2] for the calculation of return expectations are actually key because we are interested in the prospective payoff from an asset across all states of nature and thence its implied return, which relates the payoff to the price of purchase. What we are trying to do is to match our subjective priors about these returns to their objective distribution. And under a standard form of expected, or Bernoulli, utility,[3] this expected return is the sum of all the likely payoffs pre-multiplied by the respective probabilities of each state of nature. Once we have the expected or prospective payoff, the price of the asset can jump so that the return compensates you for the risk of that asset. If, for example, we think that two different assets will yield the same payoff, the one with more risk will trade at a lower price and thus give a higher rate of return. We therefore need expectations about payoffs and risk to price an asset, which is the opportunity cost of its purchase, and these derive ultimately from the collection and analysis of information and the development of a number of beliefs.

One lacuna to this story was highlighted by Keynes (1936) in his analogy of asset prices as Beauty Contests. On the one rational paw, what matters to me for the asset price is its actual or fundamental pay-off across all states of nature and that means the actual amount of cash I get in my grubby paw when a good or a bad state occurs. On the other grubby paw, I may care less about the actual cash but what I think that other people think about the likely payoffs. So it is no longer my expectation of the payoffs but my expectation about others' expectations that may matter. Should these higher order expectations dominate the pricing process; the price may deviate from the first order rational level and simply reflect view about views or what we might call fads and fashions.

So that observation leads us to one reasonable question is why are economic forecasts so bad?

[2] Note that I would only use past payoffs as a guide to asset's risk profile against macroeconomic and other events and my own consumption profile.

[3] See Bernstein, 1998.

Everyone 'knows' that forecasters are always wrong, so how can we forecast the correct probability of and payoff from all the possible states of nature. Allow me for a moment to try and take the high ground because there are one set of forecasts that are at least as bad as those provided by economists - the pollsters. The final sets of polls, or forecasts, at the most recent UK general election placed Conservative and Labour neck and neck with both likely to gain 34% of the overall vote.[4] But the results placed the Conservatives some 7% ahead of Labour. The British Polling Council set up an inquiry into this forecasting failure, akin to the various post-mortems in economic forecasting bodies throughout the world after 2007/8, and came to the basic conclusion that the polling sample did not conform particularly well to that of the voters. In terms of the rationality or otherwise of the economic expectations, it was rather like using a long period of non-inflationary stable growth to predict the probability of extreme economic volatility.

There are some other possibilities. Using our economists' notion of using information efficiently, people may simply have changed their minds given the information that was in front of them. If the polls predicted a hung Parliament and an extended period of uncertainty, people may have switched to their second best choice in order to avoid such an outcome and some marginal voters switching may have been sufficient to explain the differences between outcomes and expectations. Indeed this is an example of the Lucas Critique in operation. If I can see the policy and welfare consequences of a given choice or plan and I do not like it, I can reformulate to a new plan that yields a better outcome for me. Let us accordingly be wary of opinion polls in the run-up the referendum in June.

One way we can cut through the problems of sampling and changes in behaviour is to develop a prediction market *per se*, where we might be putting real money on an outcome. Consider uncertainty over a two-horse race where, barring a fall, there are only two outcomes and one horse is called Red and the other Blue. I might phone people to get a view and because all the people who like Red are hard at work, I only manage to get hold of the Blue-supporters, at which point I install Blue as the Nap. But Red is a very good horse and those who know this and are not part of the polling sample, will treat Blue as a Bismarck - likely to be sunk. How can we correct the pollsters? Imagine we buy a £1 payoff for Blue at a price of 50p, then we are placing an evens bet.[5] If I use the polling data as my start point, we might find that I can only buy my £1 payoff for Blue at 80p, or 1/4 on, as it has become the odds-on favourite. At that point, the odds for anyone who has any knowledge of Red, and did not participate in the original poll, are very good at 4/1 against and they will start to fill their boots with bets on Red and the odds will start to turn. Indeed they will continue to do so until the odds reflect the objective or true probability of the outcome. Providing they can bet, the people who know will continue to bet until there odds are the same as the 'true' likelihood of Red winning.

In the US, there have been a set of prediction markets that decant such information from public to private, run in Iowa. And we can see from the charts of the probability of a blue win versus a red win, that despite any political or polling uncertainty, the (near-)rational market had placed Obama as the nap for over a year before the election. Once we introduce money and financial returns into the equations, people may have an incentive to form rational views.

Martingales and Random Walks

The idea we shall start with is that, related to betting or gambling, of a martingale. Here we can imagine a fair game with a random outcome on which we can evaluate the probability but not the exact outcome. The sequence of outcomes is also not history dependent so that the odds do not change over time. The most obvious example is the toss of fair coin where we know that the odds are 50:50 but on any one draw, we certainly cannot know which way the coin will land and even over a small number of tosses, for example 10, we cannot be at all sure that we will have 5 heads and 5 tails. We can thus understand, or even forecast, the process and work out the probability of each outcome

[4] As reported in the Financial Times, earlier this year.

[5] Because this is a fair book, we also buy the £1 payoff for a Red win at 50p. Note that the two 50p in, give us the pot of £1.

but on any one throw the outcome is essentially unpredictable.

Analogously, asset prices, for which we might be able to calculate the probability of returns in different states, may yet still be unpredictable on their next incremental movement. The argument is that if I can calculate the probability of a given return in all states tomorrow, conditioned all information available today, I can work out the correct expectation of tomorrow's asset price. And if having established the correct expectation of the price, today's price must jump to today's expectation of tomorrow's price. Otherwise there would be an unexploited profit opportunity. Today's price has become the best forecast of tomorrow's price, we may know that following tomorrow's coin toss it will almost certainly not be tomorrow's price but because it is a coin toss there is no reason to place a bet one way or the other today!

If all of today's information is contained in today's asset price, then yesterday's information will also be incorporated in today's asset price. But we cannot easily know all the information that was once known, so we can argue that yesterday's information was captured by yesterday's price. So yesterday's price will have no information for today's price or any of our tomorrows. Asset prices are thus simply buffeted by that bit of news that cannot be known today and follow what is called a random walk. Indeed when we look at asset prices that follow a random walk, because their next step is taken from as a random draw from a probability density function, we can start to impose shapes on the series that look like patterns. I can assure you, just like clouds in the sky that look like faces or parts of the anatomy, there are no 'heads and shoulders' waiting to be completed in asset returns.

This lovely bit of logic allowed the efficient markets hypothesis to be tested in a nice juxtaposition of theory and data.[6] The hypothesis has some testable implications when we first suppose that p(t) follows a random walk with its innovations driven by bits of information that we do not know today. Tomorrow's return on the asset will therefore be uncorrelated with today's information set. But because we cannot possibly observe all the information that was available to all traders one any one day, we simply substitute that day's price as the proxy for that day's information. And the observation that future returns are uncorrelated with lagged information develops into a test as to whether returns today or tomorrow are uncorrelated with past returns. And so it turned out to be surprisingly hard to reject that asset prices followed a random walk and so that returns were serially uncorrelated. Where some have found evidence of positive serial correlation in the short run - perhaps as positive or negative news comes in waves and others have found some evidence of lower frequency negative serial correlations but these do not look terribly important to me.[7]

Asset prices seem to be pretty close being informationally efficient. William Sharpe defined this in the following way: A market is efficient with respect to a particular set of information if it is impossible to make abnormal profits (other than by chance) by using this set of information to formulate buying and selling decisions.[8] It is normal then to think of there being three forms of such efficiency: weak form where returns are independent of past asset prices; a semi-strong where asset returns reflect all publicly available information and strong form efficient where returns incorporate all information whether private or public. But if that is the case, why do we invest so much time and so many resources in trying to understand where asset prices will go?

Anomalies

The economics of asymmetric information provided an important attack on the notion of informationally efficient markets. Grossman and Stiglitz (1980) suggested the following paradox: if information is costly to obtain and yet

[6] See Fama's Nobel Prize Lecture, 2013.

[7] See Smithers and Wright, 2002.

[8] Reported in Bailey, 2005.

prices already reflect information, then who would incur the cost of collecting such information? Asset prices cannot then reflect all information and thus asymmetric information matters and asset prices are not efficient. Let us suppose, just to drive the point home, that the return to an asset with access to the full public and private information set is *mu* and if accessing information is costless, any investor can get return *mu* in expectation when they buy the asset. Now suppose that in order to obtain access to the information a cost, *c*, has to be paid, which means the return becomes *mu-c*. Now if the markets already have the information, who would pay cost *c*, and lower their return? But if the cost is not paid how will markets have information? If right, this observation means that changes in the costs of obtaining information or analysing development will have important implications for asset prices because they will act to regulate the flow of information into the market place.

If there are persistently high costs from collecting price sensitive information, then we might be in a better position to understand many of the anomalies that regularly scatter the financial pages. There are number of well-known violations of efficient markets that most text books will list: Calendar effects (January, September, week of the month, Monday blues, hour of the day); Weather (rainy or sunny days); Sports results; Small Firm effects; High earning-price ratio; Closed End Mutual Funds' Pricing; and Initial Public Offerings.[9] But most of these well-known violations of efficient markets that seem to suggest that abnormal profits can be made, which should simply disappear over time or perhaps are very small unless you have Hedge Fund quantities of cash with which to play. And then over the long run any abnormal returns may mostly just pay for the financial intermediary's yacht, leaving you with the market return minus fees at best. The problem with these anomalies is mostly along the lines that if you know you will lose money on a Tuesday what will you do? Hold until Tuesday? No you will sell on Monday and the losses will then happen on Monday, which will mean that you will sell on Friday and so one until the effect just spills away like so much water into a river.

One area where I have tried to understand anomalies is the question of whether bond prices,[10] which should reflect long term trends as well as short term news, tend to over-react to day-to-day news. To remind you, a bond price collects a series of short terms interest expectations over the life of the bond and even though shocks may come along to perturb the expected path of the short run interest rate, we can measure the rate of decay of the shock and hence the expected impact on short term rates and thus long term bond prices and we regularly tend to find significant excess volatility.[11] In some sense this kind of finding harks back to the work of Robert Shiller (1981), who found that equity prices displayed excessive volatility compared to the actual future payoffs. One answer to the Shiller question is that the outcome is one of the many possibilities across many states and any news today may affect all these probabilities, leading current price volatility. There are a number of possible answers to the bond price anomaly: that long term bond prices, which reflect some notion of steady state, may get dislocated by shocks in the absence of a nominal anchor, that risk or term premia may dominate movements in these prices, or that our ability to construct alternate futures is bounded so we overweigh the present or recent past. Indeed I tended to found that more volatile bond prices can be generated if agents themselves adapt very slowly to economic shocks, which implies a perverse result that a stickier real economy may create more volatile asset prices - on which more later.

Surprise, Surprise

One important, emergent literature involves herding in the presence of partial information. The standard example explains both fads and unexploited profit opportunities. Let us suppose that you are a stranger in a new town and are looking for a place to eat and you find two similar-looking restaurants side-by-side, both of which are empty.[12] You

[9] See Bailey, for example.

[10] See Chadha and Holly, 2010.

[11] Blinder (1999) outlines this problem.

[12] For more on these issues, I recommend Chalmley's lovely 2003 book on Rational Herds.

have no information and your smartphone has no charge. Randomly, you go into the one of the left. A few minutes later another random stranger finds herself in the same position and loiters outside the two restaurants. She is faced with the same problem but can see you with a glass of wine in your hand and thinks you know something she does not and also chooses the same restaurant. Very soon there is Party of 1999 proportions going on in the restaurant on the left, even though it is fundamentally the same as the one on the right. We have an informational cascade and a herd. To the extent that the restaurant on the right cannot change its prices and attract customers, this herd may remain but if they can undercut the restaurant next door, it might encourage some learning about the relative merits of these two restaurants and over time a more equal distribution. Assets may be subject to the same types of information cascades, if price changes do not allow news to be bid in.

When we turn to the cross section of views and examine the forecasts of financial market analysts, it can frequently look like there is herding because they mostly seem to publish forecasts close to the median and/or very close to each other. The outcomes, or releases, once again tend to be distributed far from these herded expectations. One explanation is rather obvious and results from our obsession with point forecasts. Consider the role of 2 fair six-sided dice. There are 36 possible combinations of the number on the two dice with eleven possible numerical outcomes each with their own specific probability. Of course, the most likely outcome of seven carries a probability of 1/6. But if we ask 36 people what we think will be the most likely number from a single throw, if they are a Gresham audience, they will all answer 7. The 36 people will mostly be wrong; fives time out of six and journalists can say that the Gresham audience cannot forecast whereas, of course, they can.

Whether because of herding or the focus of expectations on point forecasts, rather than assessment of the full set of subjective prior against objective outcomes, news or announcements, asset prices may jump wildly on news. For example, as discussed at last year's lecture the surprise announcement of operational central bank independence in May 1997 provoked a predictable response. The impact on long run expectations is shown by the fact that bond yields dropped 50 basis points (1/2 of one percent) on the announcement by the incoming government in May 1997. Chadha, McMillan and Nolan (2007) examined the consequence for longer term interest rates, and found that inflation-fighting preferences were highly likely to impact on interest rates and explained well the significant decrease in interest rates. A permanent fall in 1/2% in borrowing costs means that for every £ 1bn borrowed or refinance by the government since 1997, the state has saved some £ 5mn. The value of central bank independence looks clear --it can significantly reduce medium and long nominal rates for both government and private sector liabilities.

Indeed policy-makers, it seems to me vacillate, between implementing the implications of the policy ineffectiveness proposal[13] and working with surprises and assessing their impact on asset prices and signalling intentions in a gradual manner so that asset prices have time to adjust to policy makers' preferences and plans.

Concluding Remarks

There a number of misconceptions about the efficient markets hypothesis and, its building block, rational expectations. Efficient markets in this context are about the exploitation of information for financial gain. The hypothesis says that individuals have an incentive to trade their information, whether the data or news involves surface or deep mining, and in that process remove any excess returns that might accrue from private or public information. The hypothesis does not say that markets cannot lead to socially sub-optimal outcomes involving booms and busts and rational expectations do not imply we can know the future! Given all that it is surprisingly hard to find strong evidence that market returns are not efficient, by which we mean that past returns have essentially no exploitable news for future returns.

[13] See Sargent and Wallace, 1976.

Of course, some studies have found some evidence of past returns or financial data having information for future returns but these effects tend to be small, may be wiped out by trading costs and tend to disappear out of sample. The findings may well be a financial chimera. And yet there are anomalies to explain. Why do financial prices seem to be so volatile in the presence of relatively small bits of news? What kind of uncertainty is being resolved here so that its resolution is so profound? How are all the tiny bits of news aggregated up into the fine-tuning of probabilities and pay-offs? If the aggregation problem is too severe then we might simply over weigh current or recent events. In the cross section, we might just decide to do or forecast what everyone else decides to do, or what I think they will do. Similarly if policy maker opinions evolve over time, we might have to learn and that takes time. Modelling the interaction of many states, views and drip-by-drip news, as well as the interactions of traders with varying levels of market liquidity may help us further understand any asset pricing anomalies.

References

(1) Bailey, R. E. (2005). *The Economics of Financial Markets*, Cambridge University Press.
(2) Bernstein, P. L. (1998). *Against the Gods*, John Wiley & Sons.
(3) Besley, T. and P. Hennessy, (2009), "Letter to H M The Queen", 22nd July 2009, British Academy.
(4) Blinder, A. (1999). *Central Banking in Theory and Practice*, MIT Press.
(5) Chadha, J. S. and M. S. J. Holly, (2010). "Macroeconomic models and the yield curve: An assessment of the fit," *Journal of Economic Dynamics and Control*, vol. 34(8), pp1343-1358.
(6) Chadha, J. S., C. Nolan and P. Macmillan, (2007). "Independence Day For The 'Old Lady': A Natural Experiment On The Implications Of Central Bank Independence," *Manchester School*, vol. 75(3), pp311-327.
(7) Chamley, C. P. (2003). *Rational Herds: Economic Models of Social Learning*, Cambridge University Press.
(8) Fama, E. F. (2013). "Two Pillars of Asset Pricing," Nobel Prize Lecture, December 8, 2013.
(9) Grossman, S. E. and J. E. Stiglitz (1980). "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, vol. 70(3), pp393-408.
(10) Iowa Electronic Markets. http://tippie.uiowa.edu/iem/
(11) Keynes, J. M. (1936). *General Theory of Employment, Interest and Money*, Chapter 12, Macmillan.
(12) Lucas, R. E. (2009). "In defence of the dismal science", *The Economist*, 6th August 2009.
(13) Financial Times, (2016). "General election poll failings blamed on poor sampling", 14th January 2016.
(14) Sargent, T. J. and N. Wallace (1976). "Rational Expectations and the Theory of Economic Policy". *Journal of Monetary Economics*, vol. 2 (2), pp169--183.
(15) Shiller, R. J. (1981). "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review*, vol. 71(3), pp421-436.
(16) Smithers, A. and S. Wright (2002). *Valuing Wall Street: Protecting Wealth in Turbulent Markets*, McGraw-Hill Inc.,US.

Gresham College
Barnard's Inn Hall
Holborn
London
EC1N 2HH
www.gresham.ac.uk