

# The Undercover Economist

By Tim Harford

## Chapter 6

### Rotten Investments and Rotten Eggs

#### **What's this?**

This is chapter six of the second edition of my book *The Undercover Economist*. When I published the new edition, by far the most substantial change was the addition of this chapter, which is all about the financial crisis.

#### **Why are you publishing a big extract from your book?**

Lots of people wrote to say they'd purchased the first edition and wondered whether they needed to buy the second edition too. This chapter gives you most of what's changed, and that seems only reasonable.

#### **Can I send this to friends?**

Feel free to distribute this document, but please do not change it or use it commercially. Issued under a Creative Commons BY-NC-ND License:

<http://creativecommons.org/licenses/>

#### **If I want to buy *The Undercover Economist*, what's the story?**

Thanks, feel free: <http://timharford.com/books/undercovereconomist/>

# 6

## Rotten Investments and Rotten Eggs

In the first edition of *The Undercover Economist*, this chapter detailed the dot-com bubble and bust of the late 1990s and early 2000s: in a wave of excitement about the internet's potential to transform business models, shares in red-hot high-tech start-ups were traded up to absurd levels before inevitably crashing back to earth again. After that first edition was published, however, the banking crisis that began in 2007 made the dot-com crash look like a minor inconvenience.

We'll come back to the dot-com crash in due course, but it would be impossible to produce this second edition without going undercover again to see if we can explain just what went so hideously wrong with the banks. And this is quite a challenge, because as well as being deeply serious, the banking crisis was also mind-bogglingly complicated. So in an attempt to understand it, let's start with something altogether simpler: a carton of eggs.

In Cumbria, in the north of England, early in 2010, a lady named Fiona Exon bought a carton of six eggs and discovered that all six of them had double yolks. Newspapers reported this as beating odds of 1,000,000,000,000,000 to one. (Some of them also reported this as 'a trillion to one'. That's old school. These days we'd call a number with eighteen zeroes a 'quintillion'.) The mathematics seemed straightforward. According to the egg-watchers to whom the newspapers spoke (the British Egg Information

Service – who knew?), the chance of any given egg having two yolks is one in a thousand. The chance of any two given eggs both having double yolks would therefore appear to be, from multiplying the two probabilities together, one in a million. Three in a row would be a one in a billion chance; four would be a trillion, five a quadrillion, and six double-yolked eggs in a row would be a one in a quintillion chance. If that calculation is right, then if each and every person in the world bought six eggs each morning, we'd expect to see a carton of double-yolk eggs being sold somewhere on earth roughly every four centuries.

The trouble is, after Mrs Exon's apparently astonishing discovery was reported in the media, lots more people popped up to announce that the same thing had happened to them. I was invited onto national radio to explain the mathematics, and a number of listeners phoned in to say they'd also once bought a carton of eggs that turned out to have six double-yolks. It seemed they were rare, but not *that* rare.

What had gone wrong with the calculation? Perhaps double-yolk eggs are more common than the British Egg Information Service thinks. But let's be charitable to the British Egg Information Service and assume that it provides accurate information about British eggs. (It would be a minor tragedy if it did not.) The problem is that the newspapers made an insidious assumption: that double-yolk eggs do not come in clusters. Allowing for the fact that they might, we get a very different equation. In an extreme case, where double-yolk eggs *always* come clustered with other double-yolk eggs, you will either see no double-yolk eggs in a carton (with probability 999 in 1000) or you will see six double-yolk eggs. Even though the chance of a single double-yolk egg hasn't changed, the chance of seeing six together has increased a million billion fold, simply because the eggs tend to come in clusters.

All this would be simply a curiosity unless large amounts of money were riding on it. So imagine that Fiona Exon's friendly local egg retailer had come up with a clever marketing wheeze and offered a million-pound jackpot to any customer who bought a carton full of double-yolk eggs. Sitting with his calculator, and working from odds of a thousand to one per egg, he would have concluded there was effectively no chance that he'd ever have to pay out. If the offer enticed in a few more egg buyers a week, it was easy money.

Unless his calculations were wrong, of course, and the eggs came in clusters. Then it would be only a matter of time before the retailer bankrupted himself.

How might double-yolk eggs come in clusters? As it turns out, there's a human explanation. A caller to the radio show I was on explained that they used to work in an egg packing plant. They and their colleagues could easily identify double-yolk eggs by holding them up to the light. They used to put several double-yolk eggs in a carton together to take home for breakfast. And if each of them had all the eggs they could stomach, the double-yolk carton would be placed back on the production line to be shipped to the likes of Fiona Exon – to the horror of our imaginary calculator-wielding retailer.

All this is roughly – very roughly – what happened in the financial crisis. Bear with me, and we'll discover how.

## **Rotten eggs and rotten investments**

Essentially, the banking crisis came about because banks and other financial institutions were taking very big bets on the chance of events that, like Fiona Exon's discovery, seemed too unlikely to contemplate. They were wrong,

because they had made the wrong assumptions about the process producing these events.

The financial egg cartons were called mortgage-backed securities, and the rotten eggs that filled them were the now-infamous subprime loans. Banks bought these risky – or ‘subprime’ – mortgages from ‘originators’, companies who initially made the loans. Then the banks repackaged them into financial products, which provided the rights to a stream of mortgage repayments and the risks that mortgage would not be repaid. The repackaging was often repeated many times. (The first package was called a residential mortgage-backed security, or RMBS. The repackaged versions were called collateralised debt obligations, or CDOs, then CDOs-squared and CDOs-cubed.)

These RMBSs, CDOs and the rest were tremendously complex products, but their aim was simple enough: to create financial assets that promised a safe and predictable income, rather like the one our fictional egg retailer thought he’d come up with. This predictable income could then be traded between banks, pension funds and other financial institutions. If the underlying risks had been misunderstood, however, the value of these financial assets would collapse, threatening the companies that owned them with bankruptcy.

So how did these strange mortgage-backed products work? Let’s switch the analogy from double-yolk eggs to rotten ones. Imagine that a bunch of mortgages is simply a basket full of eggs, some of which are rotten – say 5 per cent. The bank repackages the eggs into cartons of six, keeping one egg and selling the rest. The price of those other eggs reflects the risk of who gets the rotten eggs. It turns out that, assuming the rotten eggs don’t cluster, the chance that at least one egg will turn out to be rotten is 27 per cent. If the first egg is

rotten, the bank is first in line and takes the hit. But it's a risk they're happy to take, because they've sold the rest of the carton at a profit.

The buyer of the second egg – the 'junior investor' – is next in line if there's another rotten egg. Do the sums, and the chance of a second bad egg comes out at just over 3 per cent. The junior investor gets a decent discount on the price of the egg for taking this fairly small but still meaningful risk. Senior investors buy the remaining eggs, and as the chance of any further bad eggs is very low – only about 0.2 per cent – they pay top dollar. They're expecting a safe and predictable income; in the very unlikely event that there's a third rotten egg, they'll get a nasty surprise.

The repackaging makes some sense, because some investors prefer to pay for a guarantee of safety while others would prefer to get a cheap egg and take a risk. But why repackage two or three or four times, creating CDOs-squared, CDOs-cubed and so on? There is a reason. The first egg is quite likely to be rotten – not very interesting for repackaging purposes. The final four eggs are almost certainly fresh and don't need any repackaging. But the second egg is more interesting. It is fairly safe but not that safe. So a bank could take six 'second eggs' and repackage them into another carton. The bank would take the first rotten egg, new senior investors would queue up for the four guaranteed-fresh eggs, and there would be another 'second egg' which might or might not be rotten. This process can be – and was – repeated many times. In each case, it all seemed at the time like a perfectly rational way to ensure that investors with different preferences could pay a fair price for the different levels of risk they were taking on.

Now, here comes the catch with this repackaging. If the banks had got their sums wrong – no prizes at this stage for guessing whether they had – then each level of repackaging dramatically amplified the effect of that mistake.

Remember we assumed that the risk of a rotten egg picked at random from the basket of eggs was 5 per cent. This turned out to imply that, in the first carton of eggs, the chance of a second rotten egg is just 3 per cent. The CDO, then, is made of eggs with a probability of being bad of 3 per cent. The chance that any one of these eggs is bad turns out to be about 18 per cent, and the chance that a second egg in the CDO is off is about 1.5 per cent. The CDO-squared is made of eggs with a probability of being rotten of 1.5 per cent, and the chance that the second egg in the CDO-squared is bad is less than one third of one per cent. That's a very small risk.

What happens if the chance of a rotten egg isn't 5 per cent at all, but really 10 per cent? Crunch the numbers and it turns out that the chance the second egg in the RMBS is rotten hasn't fallen from 5 to 3 per cent, it's *risen* from 10 to 11.5 per cent. The fact that it's risen is more important than how much it's risen – because now each successive repackaging will further increase the risk. The second egg in the CDO now has a 15 per cent chance of being bad, rather than a 1.5 per cent chance. It's ten times as risky. The second egg in the CDO-cubed becomes almost 2500 times riskier. The risk seemed laughable; suddenly, it's no joke.

Much the same analysis applies if the eggs come in clusters (although the mathematics are harder). Why did the banks' mathematicians assume that mortgage defaults wouldn't cluster together as much as they did? Well, one borrower might go through a divorce, another lose her job, and a third fall ill. But one person's divorce is nothing to do with a stranger's health, so mortgage defaults caused by these problems should come in ones and twos, not clusters. As long as that assumption held true, the re-re-re-packaged loans were extremely safe, for much the same reason that Fiona Exon apparently faced a

one in a quintillion chance of picking up a carton with six double-yolk eggs in it.

But if the mortgage defaults come in clusters, disaster looms. Every egg could end up stinking; the safest-seeming mortgage products could end up suffering awful losses. The banks' mathematicians realised, of course, that some degree of clustering was possible. But clustering is a difficult thing to measure,\* and they apparently based their estimates on a limited run of data from a time when house prices were rising across the United States. And they therefore seem to have missed the scenario in which mortgage defaults would cluster together ferociously: a national slump in the price of housing.

With hindsight, it is blindingly obvious why such a slump would lead to mortgage defaults. Subprime mortgages depended – in a way no other mortgage in history ever had – on the value of houses continuing to rise. Subprime mortgages were typically offered with generously low introductory interest rates for a fixed period, after which both borrower and lender assumed that they would be cancelled and replaced with a new mortgage. The idea was that, as the house would by then be worth more, lenders would be happy to

---

\* Here's an example from the financial analyst and commentator, Paul Wilmott. Imagine that you're trying to measure the clustering (more formally, the correlation) between the price of shares in two rival companies, Nike and Adidas. At first they're not very closely correlated. Then imagine that one of Adidas's sponsorship deals goes wrong and the brand is tarnished by association with some wayward celebrity. Adidas shares would fall and Nike shares would rise as they picked up market share – a negative correlation. Then imagine there's a craze for branded sportswear thanks to a great World Cup. The price of shares in both companies rises. So the correlation could be small or large, positive or negative, depending on what's happening in the world. If you looked only at the two share prices moving, you would have very little idea what to make of it all. And if you looked only at data for a limited period of time, you might leap to entirely the wrong conclusion.





insurers such as AIG, a once-boring insurance company that found itself at the heart of the crisis after insuring far more bad loans than it could possibly afford. The credit default swaps made banks even more careless about taking on too much risk. If catastrophic losses to an individual bank are vanishingly unlikely, the swaps look like a sensible precaution. If catastrophic losses to multiple banks are sooner or later inevitable, they are an unwittingly suicidal way to ensure that apparently separate companies are all bound together and will all be dragged down together – which is exactly what happened.

Despite all these layers of explanation, if I had to point to one single innovation that sparked the crisis, it would be those endlessly repackaged cartons of rotten subprime eggs. It was an accident waiting to happen – indeed, bound to happen – as soon as the global economic equivalent of Fiona Exon came along.

### **Cheap bailout, costly crisis**

Soon enough, Lehman Brothers had gone under and more and more banks were teetering on the edge of collapse. The rest is folklore: although the banks had been greedy and stupid, governments couldn't afford to let the financial system collapse so they began desperately to throw taxpayers' money at them. This incurred debts for which ordinary people are even now suffering through hefty tax rises and painful cuts in public services.

Or did it? There is truth in that analysis, but the story is not quite so simple.

The banking crisis has, indeed, cost citizens and taxpayers dearly. And in some countries – such as Ireland – the cost of bailing out the banks was

directly responsible for crippling levels of public debt. No wonder: the banks were large relative to the size of the economy and made astonishing losses. The official estimate in late 2010 was that Anglo Irish – ‘probably the world’s worst bank,’ says the hedge fund manager Theo Phanos – had managed to lose €30bn despite making only €72bn of loans. This is an almost inconceivable rate of losses. It is also almost a quarter of Irish national income – from just a single bank.

For comparison, the stricken British bank, Royal Bank of Scotland, lost similar sums of money: £24bn in 2008 alone, and £3.6bn in 2009. But RBS was a far bigger bank and the UK economy is well over ten times the size of Ireland’s. RBS’s losses were a few per cent of its loans; Anglo Irish’s losses were approaching 50 per cent.

While it’s easy to draw a straight line between the government bailout of Ireland’s banks and the Irish’s government’s debt problem, in other countries the story of the bailout is a subtler one.

In the UK, for instance, Gordon Brown’s Labour government ended up borrowing gigantic sums of money. The budget deficit reached around £150bn a year, the equivalent of £2500 of yearly borrowing on behalf of every man, woman and child in the country. But very little of this was to do with the banking bailout itself. While British taxpayers did indeed bail out British banks, the taxpayer is likely to get most of the money back. Some of the intervention took the form of guarantees for the banks, reassuring creditors that they could continue to lend money to any British bank. This was invaluable, but in the end the banks did not use the guarantee, so fortunately it turned out to cost nothing. The Treasury even turned a profit, because the banks paid fees for these guarantees. Other money was spent injecting capital into the banks; in

exchange for this, the British government received shares in the banks; this transaction may also make a profit.

It's difficult to tell exactly what the UK bailout will end up costing the taxpayer directly. The budget of June 2010 estimated that the likely cost would eventually be £2bn. That's a very uncertain figure. The final outcome could be a lot more, or taxpayers could make money on the deal: the cost was estimated at £20bn–£50bn in 2009 and £6bn in March 2010, so it is heading in the right direction. £2bn sounds a lot – it is a lot – but that is roughly what the government was borrowing every four days in the tax year 2010–11. 'The fact is that in most crises – this one and previous ones – the taxpayer is typically paid back,' says Andy Haldane, the Bank of England's executive director for financial stability. The bailout itself does not explain the huge government deficits that were accumulated. What does?

The real cost was the damage done to the British economy. As the banks teetered on the brink of collapse, they tried to suck up as much cash as possible to prevent their bankruptcy. That meant that every loan and every overdraft facility in the country was suddenly under scrutiny. Businesses suddenly found their access to credit drying up; consumers spent less, feeling the same pressures. Potentially healthy businesses were crushed by the financial turmoil. Andy Haldane reckons the crisis could have lowered the UK's national income by about 10 per cent – and put it on a permanently lower trajectory. The loss 'occurs year by year by year. We know from past financial crises that they leave a permanent scar.'

The total damage? Haldane reckons it could easily be £1700bn – and plausibly £8000bn–£9000bn. As he concludes, 'those are big sums'. It's because the economy has shrunk so much that taxes must rise and government

spending must fall if the government's books are to balance – not because the government threw money at bailing out the banks.

Indeed, it is worth contemplating how much worse things might have been if they hadn't.

When Lehman Brothers went bust, the money that it was holding immediately froze up. On the small scale it meant that the phone company, electricity company, and companies employing the cooks, cleaners and security guards were all suddenly not being paid until the bankruptcy experts figured out how much money there was to go around and who might get a share. That happens in any bankruptcy and it can take a long time to sort out, with much misery all round.

But what was different about Lehman – about any large bank – was the fact that it was also holding vast sums of money from other financial institutions, money it was investing on their behalf. And it also had billions of dollars of financial contracts open with other banks – promises to pay money (or be paid) depending on what happened to the price of oil, or gold, or the debt of the government of Japan, or whatever. Suddenly all these bets were off in a mess of confusion – and in a market where milliseconds can matter, investors were facing the prospect of not seeing their money for months or even years.

The clear danger was that other banks failed too, perhaps dragged down by their connections to Lehman Brothers, or perhaps because they were suffering similar problems. Given the interconnections between these institutions, there was a serious risk that if one failed – or maybe not one, but two, or three, or four – then they would all fail, and with them the infrastructure that modern everyday life depends on.

Under normal circumstances, the payment systems in most developed countries seem very robust. In the UK, for instance, as well as having the usual back-ups you might expect, additional security comes from the fact that the system that deposits your salary (BACS), and the system linking cash machines to banks (LINK) and the various systems that deal with debit and credit cards at shops, are all separate. Even if the cash machine network went down for a bit, you'd still be able to get cashback at the supermarket; or if the supermarket's system temporarily stopped accepting debit cards, you could always go to the ATM and buy your groceries with cash instead. But there is one scenario in which all these systems would fail simultaneously: if the banks that stood behind them collapsed.

What then? Imagine waiting for your salary to arrive in your bank account, and it never does – not because your employer is bankrupt, but because your employer's bank is. Imagine going to the cash machine and getting no cash, because the cash machine company doesn't believe your bank is good for the money. You go to the supermarket with your debit card and meet with exactly the same response. It's not much of an exaggeration to say we'd be looking at the end of western civilisation.

No wonder so many governments took the view that, although it was the banks that had got us into this mess, they effectively had no choice but to dump tens of billions of pounds of taxpayers' money on the banks to keep them running. If the banks had collapsed like a row of dominoes, disaster might just about have been averted – but it would have required some quick and decisive action. Someone – a regulator, perhaps – would have had to find a way of parking the unholy mess, all the arguments about who owed what to whom, somewhere out of the way. And while the mess was being sorted out, all the basics – the borrowing and the lending, paying salaries and getting cash out

of the hole in the wall – would have needed to be kept running smoothly somehow.

Such a trick would not be impossible – but it would take careful preparation and a solid legal footing, both of which were lacking in the crisis of 2008. It would make sense to be prepared for next time; whether we will be is open to question.\*

So by all means blame the banks for the behaviour that required a bailout – but don't blame governments for the bailouts. Though if it seems unsatisfying to be giving pats on the back to politicians for decisive action, fear not – there is still plenty of blame to go around. While governments probably had to bail out the banks, they didn't have to be so extraordinarily generous about the way they did it.

Pietro Veronesi and Luigi Zingales, two economists at the University of Chicago, studied the effect of the US bailout that saw Treasury Secretary Henry 'Hank' Paulson call bank bosses to a meeting on 13 October 2008, shortly after the collapse of Lehman Brothers. Paulson stumped up \$125bn of US taxpayers' money in exchange for shares in nine US banks. Some banks, such as Citigroup and JP Morgan, received as much as \$25bn each. The Treasury also guaranteed new issues of bank debt.

Veronesi and Zingales concluded that Paulson's move cost taxpayers \$20bn–\$45bn. That's a lot – although the US budget deficit that fiscal year was over \$1400bn, so as far as US government overspending or undertaxing was concerned, it was pretty much an average week. But they also concluded that shareholders and bondholders in the banks were about \$130bn better off as a

---

\* In my book., *Adapt: Why Success Always Begins with Failure*, I discuss some ways in which governments might be able to let banks fail while letting all those vital basics continue smoothly.

result of Paulson's gift. In short, every tax dollar spent created between \$3 and \$5 of wealth. There are very few government policies in history that can claim such a spectacular return on investment.

Why such a large effect? Veronesi and Zingales believe the gift prevented runs on the banks, the risk of which was depressing bond and share prices. A bank run is a situation where a bank can be bankrupted simply on the strength of the fear that it might happen. If a government guarantee can relieve such panic, it can create great value at little cost. And, indeed, the banks that gained most from Paulson's gift are also the banks that were in imminent danger of a run.

Sadly, that wealth accrued entirely to investors in banks – galling for taxpayers who had no desire to write a cheque to bank bondholders. Three weeks before Paulson's gift, the world's most successful investor, Warren Buffett, had injected capital into Goldman Sachs, and he had secured more generous terms. Veronesi and Zingales reckon that if Paulson had successfully demanded Buffett's terms, the taxpayer would have made more than \$40bn, a roughly even split with the banks' creditors.

The bailout was well worth doing. It's just a shame that Mr Paulson didn't drive a harder bargain.

## **The 'Greenspan doctrine' and efficient markets**

Once commentators took stock of the crisis, two high-profile tenets of faith in the investment world were widely deemed to be casualties – one justifiably, the other less so.

The first was the ‘Greenspan doctrine’, arguably *the* basic regulatory mistake that allowed the crisis to happen. This doctrine held that self-interest was the best guardian of banks’ balance sheets. Consumers would need protecting from mis-selling, and fraudsters had to be spotted and jailed; but if there was one thing the regulators didn’t need to lose much sleep over, it was the possibility that banks would simply make decisions so bad that they bankrupted themselves. So thought Alan Greenspan, Chairman of the Federal Reserve for nearly two decades in the build-up to the crisis.

A few weeks after the collapse of Lehman Brothers had become the defining moment of the crisis, Mr Greenspan admitted that he had been wrong. He told Congress, ‘Those of us who have looked to the self-interest of lending institutions to protect shareholder's equity -- myself especially -- are in a state of shocked disbelief.’

It may have been wrong to believe the banks would look after their own best interests, but it wasn’t absurd. It’s understandable that Greenspan was surprised by the self-destruction of the bankers. Many people were – it took months for it to dawn on most commentators that the crisis wasn’t about the banks selling junk to other people, but hoarding the junk for themselves. Greenspan’s error was not in overestimating the banks’ self-interest. It was in overestimating the quality of their management.

Many commentators also hastened to declare the collapse of the ‘efficient markets hypothesis’. David Wighton of *The Times* commented in January 2009, ‘The theory was officially declared dead yesterday at the World Economic Forum in Davos. There were no mourners.’ In June of that year, Roger Lowenstein wrote in *The Washington Post*, ‘The upside of the current Great Recession is that it could drive a stake through the heart of the academic nostrum known as the efficient-market hypothesis.’

But it's not at all clear why the efficient markets hypothesis should be blamed for the crisis – nor that the crisis told us anything we didn't know about the flaws in the idea. To see why, let's first explore what the hypothesis actually means.

The efficient markets hypothesis essentially states that anything which can be predicted by rational, well-informed investors will already be reflected in the price of assets. Rational people would buy shares today if it was obvious that those shares would go up tomorrow, and sell them if it was obvious that they would fall. But this means that any forecast that shares will obviously rise tomorrow will be wrong: shares will rise *today* instead, because people will buy them, and keep buying them until they are no longer so cheap that they will obviously rise tomorrow.\*

Paradoxically, this hypothesis predicts that if investors are all rational and well-informed, the behaviour of asset prices should be completely random – because when all the predictable trends are anticipated, the only thing left to drive market moves is unpredictable news. Stock prices should take what mathematicians call 'a random walk' – equally likely on any day to rise as to fall. (More correctly, the stock market should exhibit a 'random walk with a trend', meaning that it should on average edge up over the long term – even though, on any given day, this long-term trend will be dwarfed by random movements).

---

\* Strictly speaking, I have described the 'semi-strong form' of the efficient markets hypothesis. The 'weak form' says that investors cannot predict future share price movements by looking at past share price movements, but accepts that they might profit from a more fundamental approach, analysing corporate balance sheets and underlying trends. The 'strong form' of the hypothesis says that even insider traders can't make money, because their private information will leak out so quickly.

This is a paradox because if the behaviour of asset prices was truly unpredictable, there would be no incentive for investors to remain well-informed. They might as well buy shares completely at random. But then, if everyone bought shares at random, that would open up opportunities for investors to outperform the randomness by becoming more well-informed. Somewhere in the middle is a balancing point: a nearly random market with enough quirks to reward the informed investors who keep it nearly random.

You can see the same phenomenon at work in the more mundane setting of the supermarket checkout. Which queue is the quickest? The efficient markets hypothesis suggests it's not worth worrying about: if it was obvious which queue was the quickest, people would already have joined it, so you may as well just stand in the first queue you see. Anyone who has been in a hurry in a supermarket will know that's not quite right: some queues will obviously be slow, as they contain someone with a cart brimful of small items or clutching a sheaf of money-off coupons. Yet we also know it's right enough: you don't bother to analyse in detail which queue is likely to be fastest, because you're never surprised when the queue you stand in turns out to be agonisingly slower than expected.

There's plenty of evidence that the same is true of investing in the stock market. The efficient markets hypothesis says there are no obvious bargains, no easy forecasts, no get-rich-quick schemes. It tells us to ignore the blandishments of expert fund managers charging large commissions, ignore the prognostications of city economists, don't even bother trying to buy low and sell high – which in practice has always meant rushing in during booms and panic selling during busts. We should instead merely gently invest in a wide variety of shares, with no expectation of making a killing – we should diversify, keep charges low and avoid trying to be too clever. And in the past, investors

who have done just this have indeed outperformed those who overrated their own stock-picking skills.

So I do think the pundits who have written off the efficient markets hypothesis have done so too sweepingly. It remains a useful guide to humble, practical investing, even if it failed so badly to prevent the banking crisis. (The egg-carton products were being sold to investors with the claim that they were both cheap and almost risk-free. Any fan of efficient markets should simply have asked, ‘What’s the catch?’ It seems that the only investors who did were the ones betting *against* the subprime mortgage market. There’s a telling incident in Michael Lewis’s book, *The Big Short*, in which Greg Lippmann – a banker who effectively functions as a bookie, collecting bets for and against the subprime mortgage market – organises a dinner in which potential investors on opposite sides of the bet can meet each other. The smart, sceptical minority who’d been thinking of betting against subprime had held back because they thought there must be a catch: it was so obvious to them that the subprime egg cartons stank, they were sure they must have missed something. Then they met the guys buying the subprime egg cartons, and concluded they really were that stupid.)

I’m also puzzled by why the banking crisis should have caused anyone to lose a faith they’d previously held in the efficient markets hypothesis, given that we’d already had plenty of obvious and fresh evidence of its limitations: recall the dot-com bubble, a story told in the first edition of *The Undercover Economist*. Just as wiser commentators before the banking crisis called a bubble in the housing market and warned that subprime mortgages were storing up trouble for the banks, sage investment managers in the late 1990s clearly saw the dot-com bubble as a bubble. Investment managers such as the late Tony Dye called

it at an early stage and shunned dot-com shares. He was, eventually, proved absolutely right.

But not before he'd had to find himself a new job. By the time Tony Dye was proved right, his clients had missed out on handsome profits because the bubble had gone on inflating much further before it burst. Dye left the investment firm at which he was working, amid speculation that he'd been forced out. His problem was that while he was right that the stock market was overvalued, that wasn't enough to keep his clients happy while they watched others ride the dot-com bubble. Even the savviest investors, such as Tony Dye, cannot predict precisely when a critical mass of other investors will realise that shares are overvalued and rush to sell them.

Both the hi-tech shares on the Nasdaq stock exchange and the subprime mortgage-backed securities were badly mispriced. In the dot-com crash it was all too easy for investors simply to assume that the pioneers of new-fangled internet business models must have known what they were doing, and the same overconfidence afflicted investors when it came to the mind-bogglingly complex CDOs-squared and cubed and so forth. The difference is that the consequences of the banking crisis were far greater than that of the dot-com bubble – with a far longer and deeper recession and far greater need for government assistance. By the time it became clear that the banks' financial whizkids had got their sums catastrophically wrong, deep and lasting damage to the global economy was already inevitable.

## Extracted from *The Undercover Economist, 2<sup>nd</sup> Edition*, by Tim Harford

### References

There are now too many good books about the crisis to list, but for an accessible introduction I wholeheartedly recommend John Lanchester, *Whoops!*, Michael Lewis, *The Big Short*, Andrew Ross Sorkin, *Too Big to Fail* and Gillian Tett, *Fool's Gold*. Lewis's article in the March 2011 edition of *Vanity Fair*, 'When Irish Eyes Are Crying', is a must-read, as is Tett's prescient 'The Dream Machine' in the *FT Magazine*, 24 March 2006. My favourite more technical source is the *Squam Lake Report*, which consists of a series of brief notes that are available online. Readers may also be interested in my book, *Adapt: Why Success Always Starts with Failure*, which includes a chapter drawing some unexpected lessons from nuclear meltdowns and industrial accidents such as the Deepwater Horizon oil spill, and applying them to the financial system.

Fiona Exon's Eggs: The *Daily Telegraph* was among several newspapers to report 'Woman's Amazing Trillion to One Find' – by Andrew Hough, *Daily Telegraph*, 2 February 2010, <http://www.telegraph.co.uk/news/newsttopics/howaboutthat/7127206/Womans-amazing-a-trillion-to-one-find-all-eggs-in-box-had-double-yolks.html>. My little calculation about how often double-yolk cartons eggs would turn up assumes (simplistically) that the world population stays fixed at its current level. I am indebted to Alex Tabarrok for putting together a quick spreadsheet to perform similar calculations, which I modified for my own egg purposes. The key formula in Microsoft Excel is [1-BINOMDIST(x,6,y,TRUE)] where x is the egg's position out of six in a carton (usually 2 in my examples) and y is the probability that the egg is rotten. It's 0.05 initially and then the formula can be applied recursively to calculate the probability of default in the RMBS, CDO, CDO-squared and so on. Tabarrok's original spreadsheet can be downloaded from: <http://www.marginalrevolution.com/marginalrevolution/2010/05/the-dark-magic-of-structured-finance.html>

David Viniar's bad luck: Viniar was quoted by Peter Thal Larsen in the *Financial Times*, 13 August 2007, <http://www.ft.com/cms/s/0/d2121cb6-49cb-11dc-9ffe-0000779fd2ac.html>. The short academic paper that calculated just how unfortunate Mr Viniar was is: K. Dowd, J. Cotter, C. G. Humphrey and M. Woods, 'How Unlucky is 25-

Sigma?’ *Journal of Portfolio Management*, Volume 34 (Number 4), 2008, pp.76–80.

<http://www.nottingham.ac.uk/business/cris/papers/2008-3.pdf>

The costs of the crisis: I interviewed Andrew Haldane on Monday 20th December 2010 on behalf of BBC Radio 4’s *More or Less*. The Hank Paulson bailout is discussed in Tim Harford, ‘Why US Banks and Taxpayers Owe Big Thanks to Hank’, *Financial Times*, 23 January 2010, <http://timharford.com/2010/01/why-us-banks-and-taxpayers-owe-big-thanks-to-hank/> and Veronesi, Pietro and Zingales, Luigi, ‘Paulson’s Gift’ (1 December 2009). Chicago Booth Research Paper No. 09-42. Available at SSRN: <http://ssrn.com/abstract=1498548>. The UK government budgets contain estimates of the direct cost of the various interventions to support the banking system. In June 2010 they were estimated at £2bn, a figure that has shrunk dramatically over time. In January 2011, I interviewed Edwin Latta of LINK and Paul Smee of the UK Payments Council about the robustness of the UK payments system.

Alan Greenspan’s confession is reported by Mark Felsenthal, ‘Greenspan “shocked” at credit system breakdown’, *Reuters*, 23 October 2008.

<http://www.reuters.com/article/idUSTRE49M58W20081024>