GITC REVIEW
VOL.XIII NO.1 ~ DECEMBER 2014

THE MAGIC BEHIND MTIC (STATISTICAL REASONING IN TAX CASES)

by Michael Firth

We've all experienced it: the childlike wonder whilst the magician performs an impossible feat, followed by a few moments trying to work out how it was done, and then an acceptance that the clue was in the name: it was magic.

The aim of a magician is to produce the apparently impossible or hugely improbable at will. Whilst some tricks operate by way of slight of hand (or, on a grander scale, smoke and mirrors), other tricks play with perceptions of probability: the difference between what the audience perceive as the probability of the effect and the actual probability, as known to the magician.

For example, a magician hands you an ordinary, brand new, deck of cards. He (or she) invites you to rifle shuffle it three times, pick out a card, look at it, remember it and replace it anywhere in the deck. You hand the deck back and the magician tells you what your card was. How did he do it?[1]

To rule out what you are perhaps thinking, he did not look at your card or mark the cards in any way. In fact, until the deck is handed back the magician had no idea what card you chose.

All the information needed to work out how the magic happened is given above – there are no hidden extras or slights of hand. Instead, this is a perception of probability trick.

The starting point is that you were given an ordinary, brand new, deck of cards. That is significant because a new deck of cards comes in a specific order (typically Ace of Spades through to Ace of Hearts). Then you shuffled them – but you didn't just shuffle them in any old way, it was a rifle shuffle and the significance of a rifle shuffle is that it divides the original order of the deck into two and then interweaves those sequences. The basic sequence of each half of the deck remains the same, however, within the combined deck.

Thus, looking at a single suite, if the original order is A, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, J, Q, K, a perfect rifle shuffle will produce A, 7, 1, 8, 2, 9, 3, 10, 4, J, 5, Q, 6, K. An imperfect shuffle may produce A, 1, 7, 2, 8, 9, 3, 4, 10, 5, J, Q, 6, K. Repeating this with a second or third riffle shuffle will mean that there are potentially eight sequences in the deck.[2]

The magic works because if someone takes a card out of a deck that has been handled in this way, it is very unlikely that they will put the card back in exactly the same place. That card will be out of sequence which allows the magician to identify it. For example, continuing the example with one suite, if the sequence is as follows: A, 1, 7, 2, 4, 8, 9, 3, 10, 5, J, Q, 6, K, it is the four that is out of sequence and must be the one that was selected.

It can be seen, therefore, that what originally looked like a highly improbable feat to the audience – picking the right card out of the deck (1/52) – is actually a matter of high probability for the magician (subject only to the small chance that the card is put back in exactly the same place it was taken from). Once we know what the magician was looking for and why he was looking for it the magic become more of a trick.

The purpose of this article is to explore how differences between perceived probabilities and actual probabilities affect HMRC's arguments as well as First Tier Tribunal decisions and, potentially, lead them to reach erroneous conclusions. Whilst MTIC (missing trader intra-community fraud) itself is obviously criminal and without any justification, the cases that arise from it are some of the most fact-intensive around and thus provide a good opportunity to investigate how probabilities are treated. It will be seen that certain aspects of MTIC that appear to be magic, unless the taxpayer was participating in the fraud, are actually explicable on the basis of a difference between perceived probabilities and actual probabilities.

MTIC (in brief)

By way of brief background, simple carousel MTIC typically involves:

(1) An import of goods (e.g. mobile phones) into the UK.

(2) Onward sale of those mobile phones through a number of intermediary companies within the UK (typically referred to as "buffers").

(3) Sale and export of those mobile phones by a trader (typically referred to as a "broker") to a foreign customer.

The basic VAT analysis is that the importer has a liability to HMRC in respect of its onward supply to the first buffer, the buffers have only a small VAT liability to HMRC (because their input VAT cancels out most of their output VAT) and the broker is entitled to a repayment from HMRC because it paid input VAT to its supplier, but its onward supply is a zero-rated export.

To commit the fraud, the importer of the goods will be paid VAT by the buffer but will disappear without paying it over to HMRC. It becomes carousel fraud when the same goods are imported again to repeat the fraud. HMRC's response is typically to seek to deny the broker its repayment of input VAT on the basis that the broker either knew or ought to have known that, by its transaction, it was participating in fraud.[3]

Circular payments

In a large number of MTIC cases, HMRC produce evidence which they say demonstrates that the funds passing between traders buying and selling the relevant goods moved, ultimately, in a circular fashion – i.e. the money appears to originate in one company (typically a foreign company), flow through the various traders and end up back in that company (arrows show flow of money – goods flow in opposite direction):

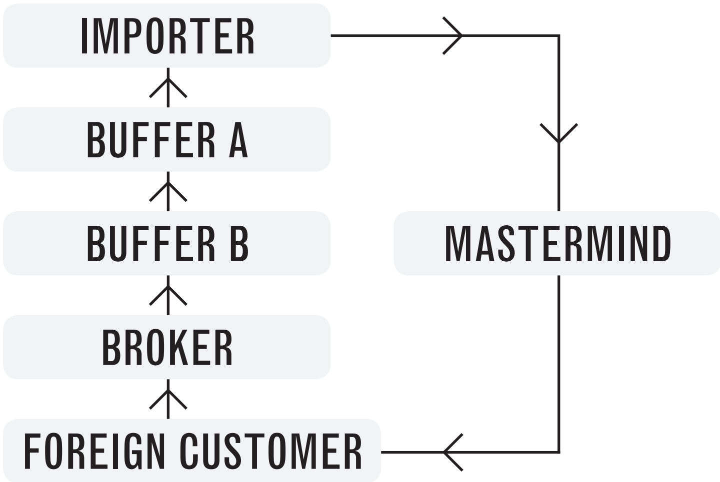

This, they typically say, demonstrates not only that there was an overall scheme to defraud the revenue, but also that the taxpayer (normally the broker) must have known of that scheme because if it did not know who it was supposed to buy from and sell to, the money would not be able to flow in a circle.

The First Tier Tribunal has accepted this reasoning. To take a recent case as an example, in *Honeytel Ltd v. HMRC* [2014] UKFTT 978 (TC), HMRC argued that the evidence showed that there was a mastermind behind the transactions, co-ordinating all of the deals. The Tribunal accepted:

> "Everything, in other words, was very plainly pre-arranged and it was clear that the money could not have completed its required circle had there been any chance that any of the parties might have purchased from an entity or sold to an entity, contrary to the planning expectations of the mastermind." (§52).
>
> "These further deals in accordance with the same pattern [including circular payment] diminish the chance of some incredible coincidence explaining the

> role of the Appellant and make it yet more obvious that the only conceivable explanation for the actions of the Appellant must be that the Appellant knew precisely what it was doing." (§125).

The taxpayer was self-represented and apparently baffled as to how the mastermind could have manipulated it in this way:

> "Declan Mundy [director of the taxpayer] periodically referred towards the end of the hearing to the fact that the Appellant had clearly been manipulated to do precisely as it had done, and that he could simply not work out how anybody had been able to achieve this result. He did not seek to advance the unarguable, namely the contention that the steps, including those either side of the Appellant, had been otherwise than pre-planned." (§54).

In the Tribunal's words, if the Appellant did not know that it was involved in MTIC, it would be an "incredible coincidence" for the Appellant to have sold to exactly the right customer with the result that the money went in a circle.

Somewhat contradicting its comment that the only conceivable explanation was knowing participation (§125), the Tribunal went on to consider the "conceivable explanation that the parties either side of the Appellant might have simultaneously approached the Appellant" (§132). In response to this, the Tribunal reasoned that such an explanation can only be used once in a deal chain and at other times the Appellant had been a "buffer":

> "Furthermore, with deal chains, the supposition that one particular participant (most obviously the exporter) might have participated by being duped by the parties either side of it can operate only once, and certainly cannot be advanced on behalf of every buffer company. Accordingly, once the Appellant had participated in a number of buffer deals, albeit that we were given no

> information about these deals other than that the profits were indeed minimal, this further reinforces the belief that the Appellant simply cannot have remained innocent and ignorant." (§32).

Before going further, it is only fair to point out that the Tribunal considered a lot of evidence besides the circularity of the payments and its reasoning in that respect will not be analysed here. Returning to the circularity of payments, the Tribunal's reasoning is relatively simple and prima facie attractive (with significant paraphrasing):

(1) Money moved in a circle, time and time again.

(2) If the Appellant had not purchased from a particular supplier and sold to a particular customer, the money would not have moved in a circle.

(3) It is possible for a "mastermind" to arrange for one innocent dupe in the circle by having companies approach that innocent dupe as supplier and customer respectively.

(4) It is not possible for a "mastermind" to arrange for two or more consecutive, innocent dupes in the circle because there can be no guarantee that the first innocent dupe will sell to the second innocent dupe (or, conversely, that the second innocent dupe will seek to buy from the first innocent dupe).

The first thing to note about this reasoning is that it is a statistical argument, based on probabilities. That is, essentially, what the Tribunal was saying when it referred to an "incredible coincidence" at §125. To demonstrate the underlying thinking, consider the following simplified trading environment:

At the bottom are nine offers of goods, say mobile phones, all of those offers are made to Buffer 1, who has three contacts and thus passes them on to Buffers 2(a), 2(b) and 2(c). In turn, those buffers have three contacts to each of whom the offers are passed (referred to as brokers, but they may or may not be exporters).

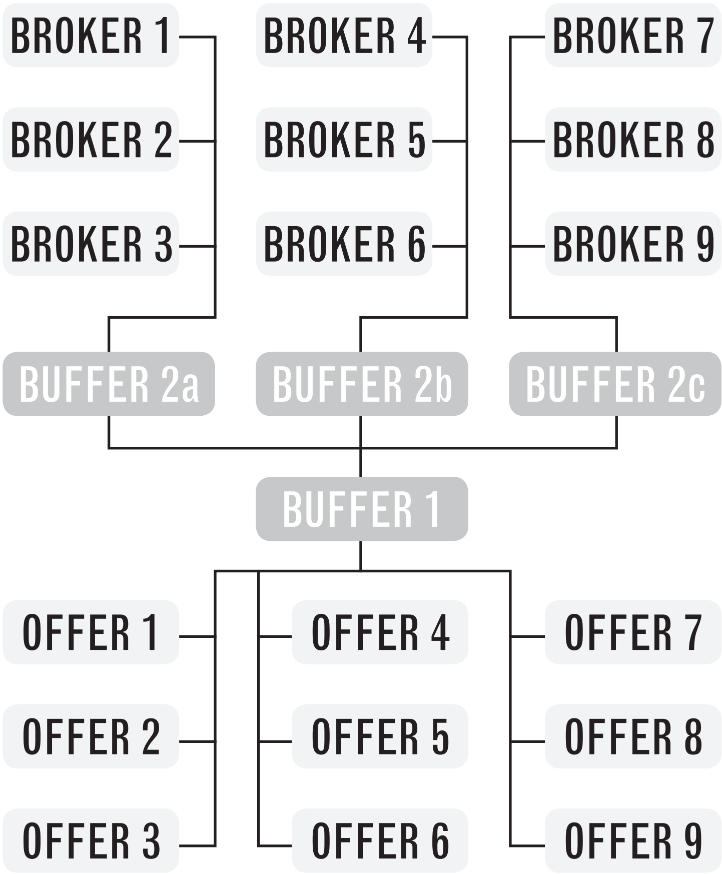

In the situation under consideration, we know that there is a "mastermind" looking to commit MTIC fraud, and he has arranged it such that an offer of goods is made to Buffer 1, say Offer 1 (i.e. assume Offer 1 is the MTIC offer). The other offers relate to other persons, unconnected with the mastermind.

If we assume that all of the buffers are innocent traders, not involved in the mastermind's MTIC, then if the mastermind is to involve three layers of innocent traders in his fraud, he (through his foreign company) must correctly choose the one

MTIC deal from amongst the nine being offered by the brokers, apparently without knowing from whom each offer originated (three stages earlier).[4]

Basic probability reasoning suggest that his chances of getting it right are 1/9. As a one-off occurrence, such a probability might raise an eyebrow, but it is not an “incredible coincidence” and certainly not inconceivable. However, if time and time again the mastermind appears to be able to correctly choose the right offer, the probability of being able to do that by chance alone drops rapidly. Indeed, performing it twice in a row has a probability of 1/81 (1/9 x 1/9) and five times in a row would have a probability of 1/59,049 (1/95). Further, when one takes into account the fact that in the real world there are many more offers and traders, the conclusion that the only conceivable explanation is that the buffers and brokers are in on the fraud starts to look fair.

Unless, of course, there is some magic going on here. After all, it is precisely such a statistically improbable feat that we would expect a magician to perform: correctly telling an audience member which card they chose by chance alone has a probability of 1/52; relatively unlikely, but not impossible. Correctly telling three audience members in row which card they chose has a probability of 1/140,608 – magic.

It will be recalled that some magic tricks rely on a difference between the perceived probability of an outcome and the actual probability. Exactly the same magical reasoning applies to MTIC: what appears improbable (correctly choosing the right offer, time and time again) is in fact highly probable, when understood properly.

The key point is to focus on how “information” about the MTIC deal can be communicated up the offer chain, without the intermediate traders being any the wiser that they are communicating information about an MTIC deal. In other words, how can the mastermind “mark” his deal, such that he will recognise it when it pops out of the other end of the legitimate market?

If we assume that all offers are the same, then no such information is communicated and the 1/9 probability in the simplified trading environment is correct. But there is a lot more to an offer than its mere existence, there is:

(a) the type of good (mobile phone, CPU etc.);
(b) the manufacturer of the good (Intel, AMD, etc.);
(c) the model of the good (each manufacturer makes a number of different models, with different speeds, etc.);
(d) the quantity offered;
(e) the timing of the offer.

What follows from this, is that the mastermind can insert an offer into the legitimate market via an innocent, unconnected party and can be relatively confident of identifying that offer popping out of the legitimate market at some other point because the chances of someone else offering:

(a) exactly that type of good;
(b) by exactly that manufacturer;
(c) with exactly the same model number;
(d) in exactly the same quantity;
(e) at or around the same time

as his original offer, is very small. Indeed, the probability may well become negligible once quantity is taken into account, given the numerous different quantities of good that can be traded in bulk (although some goods come in standard box sizes, partial box sizes are usually possible).

Furthermore, it is a mistake to think that the mastermind has to perform the same trick time and time again from scratch (as the magician does). Once it has been established that, for example, an offer made to Buffer 1 will be passed on to Buffer 2(c) who will pass it on to Broker 8, the mastermind can considerably reduce later searches for his original offer by going straight to Broker 8. Chances are, if the goods have been offered in that sequence once they will be offered in the same sequence again. In statistical terms, the taxpayer's participation in the second transaction chain is not independent of his participation in the first transaction chain so it is not correct to multiply the probabilities (see below, in relation to Sally Clark, for more information).

To increase the certainty beyond reasonable doubt, the mastermind can use a "tracer" deal. That would involve using a less common product (for example, a CPU manufactured by someone other than Intel or AMD) for the first offer. This will, by reason of it being less common, make it easier to identify the offer when it pops out the other side of the legitimate market. Once the mastermind has established that goods inserted through Buffer 1 will be offered by Broker 8 (for example), he can switch future offers to more common goods (e.g. Intel), remaining confident that the deal information will still allow him to identify the MTIC offer if/when it is made by Broker 8. For this reason it can be useful to try and identify the first deals in which the taxpayer was involved that were orchestrated by that mastermind (irrespective of which companies were inserted to do the purchasing and selling etc.): if it used a less common good, there was probably a good reason for that, namely, that it was a tracer deal.

It can be seen therefore, that once the magic behind MTIC is revealed, what appeared to be almost conclusive proof of the taxpayer's knowing involvement in fact becomes nothing more than a lesson in identifying and framing probabilities correctly. Nor is there actually any need to provide evidence that this is how the MTIC mastermind was operating. Aside from that being (presumably) impossible, the premise of the argument was wholly statistical and so it can be rebutted by showing that the statistical premise is wrong.

One final point is worth noting. HMRC typically produce evidence that a very high proportion of brokerage trading in the particular good at the particular time was MTIC trading. They use this to support the conclusion that there was an overall scheme to defraud the revenue (i.e. at least part of the chain is fraudulent), but usually there is no basis for saying that the taxpayer was aware of the proportion of fraud in the market. If one assumes, for the sake of argument, that HMRC are right that, say, 90% of brokerage trading in that particular good at the time was connected to MTIC, then in one fell swoop HMRC's argument relating to circularity has been considerably cut down.

The reason for this is that, if HMRC are right that 90% of trading is driven by MTIC fraudsters, then irrespective of the broker's knowledge, it is in the region of 90% likely that his customer will be a MTIC fraudster. Once it is almost certain that the customer would be an MTIC fraudster (irrespective of who the broker sold to), the question becomes: how many separate, non-communicating "gangs" of MTIC fraudsters are there operating in that environment? Without evidence on the point no assumptions can be made, and if the correct answer is a small number, then by that logic alone it becomes likely that money will move in a circle: 90% of all trading in this environment is controlled by only a few MTIC fraudsters.

Prosecutor's fallacy

Another statistical trap that lingers in relation to MTIC cases (and, in fact, many cases involving disputed factual evidence) is the prosecutor's fallacy. Such reasoning is as follows: X has happened; explanation Y (for X) is inherently very improbable; therefore alternative explanation Z is the probable explanation.

The famous example is the case of Sally Clark, who was convicted of the murder of her first two children in 1999. When her first child died, it was treated as arising by natural causes, probably "Sudden Infant Death Syndrome" (SIDS); when her second child also died she was arrested and charged.

There being no witnesses to either child's death, the prosecution's evidence consisted principally of expert medical evidence. One expert, Roy Meadows, gave evidence to the effect that the chance of one child in a family dying of SIDS was 1/8543 so the chance of two children in the same family dying of SIDS was about 1/73m (1/8543 x 1/8543). Professor Meadows also tried to give some context to this statistic:

> "it's the chance of backing that long odds outsider at the Grand National, you know; let's say it's a 80 to 1 chance, you back the winner last year, then next year there's another horse at 80 to 1 and it is still 80 to 1 and you back it again and it wins. Now here we're in a situation that, you know, to get to these odds of 73 million you've got to back that 1 in 80 chance four years running, so yes, you might be very, very lucky because each time it's just been a 1 in 80 chance and you know, you've happened to have won it, but the chance of it happening four years running we all know is extraordinarily unlikely. So it's the same with these deaths. You have to say two unlikely events have happened and together it's very, very, very unlikely." (*R v. Clark* [2003] EWCA Crim 1020 §99)

His first mistake was similar to that discussed above: multiplying the probability of one SIDS death by itself to find the probability of two SIDS deaths. It is only appropriate to multiply the probabilities of two events to establish the probability of them both happening if the two events are independent. Two events will not be independent if, for example, there is an underlying cause which causes them both. In relation to Sally Clark the cause of SIDS was unknown and thus, for example, it could have been a genetic defect that was being passed on to both children. In relation to circular payments in multiple chains, the events are not independent, because once it is established that offers flow in a particular way through the legitimate market, they are likely to do the same the next time.

The second mistake is the prosecutor's fallacy. Essentially, the expert's reasoning was as follows:

- Two children in the same family died in separate incidents.
- There are two possible explanations - an innocent explanation (two SIDS deaths in the same family) and a guilty explanation (murder).
- The innocent explanation is highly improbable, therefore the guilty explanation must be correct.

The error is to think that the probability of the guilty explanation is the inverse of the ex ante probability of the innocent explanation (which, in the circumstances, would be understood as 72,999,999/73m, i.e. certainty for most practical purposes).

To understand this, it is necessary to understand what the probability relates to. The 1/73m relates to the probability before any deaths occur that in a particular family, there will be two deaths caused by SIDS. Similarly, before any deaths occur, one could consider the probability that the mother would murder her first two children on separate occasions. Research is not required to say that that too is an unlikely event. Let us assume it is equally improbable (1/73m).

Ignoring any other causes, we can conclude that any given family, without any additional information, has a total probability of 2/73m of experiencing two infant deaths.

Occasionally, however, it will happen. After it has happened, one is essentially considering two highly improbable causes for the event that was, itself, highly improbable; but one of them must be true. If one wishes to use statistical reasoning, the correct approach is not to look at the ex-ante probability of the innocent explanation, see that it is highly unlikely and conclude that the other explanation, whatever it is, must be right. Instead, it is to compare the relative likelihood of all the possible ex-ante causes. On the premises adopted here, the innocent and guilty explanations have equal ex-ante probability, so on the available information one cannot conclude that either explanation is more likely than the other.

In more formal terms, the analysis needed is Bayes's theorem, which allows one to separate how likely alternative explanations are for an event that has happened from how likely it was that that event should have happened in the first place:

The equation is easier to understand than may first appear. Essentially, we are trying to work out the probability of our hypothesis (H) being correct in light of some new piece of evidence (B), i.e. Prob H given B.

As you might expect, we start with the probability of our hypothesis (H) being correct ignoring piece of evidence B (i.e. initial Prob of H). That is our base point – where we would be if we did not have evidence B – then we apply an adjustment to that initial probability based on piece of evidence B.

The adjustment is contained in the fraction. A probability of 1 is certainty. So if we assume that it is certain that we would find piece of evidence B if our hypothesis is correct (i.e. Prob of B given H = 1), then our adjustment is inversely linked to the general probability of B (ex ante Prob of B). In other words, the rarer piece of evidence B is, generally, the more likely our hypothesis becomes as a result of finding evidence B.

Thus, for Sally Clark, the hypothesis is that she committed double murder of her children, and the piece of evidence is the deaths of her two children. The initial probability of Sally Clark having committed the double murder of her children, without knowing whether her children are dead or alive is, on the assumed premises, 1/73m – very unlikely.

Then we adjust for piece of evidence B, namely, that her two children are dead. The probability of the two deaths occurring if our hypothesis (that she committed double murder) is correct is certainty, i.e. 1 (there is no more to this than appears – if she committed the double murder of her children then we would always expect to find that her two children are dead).

So it all turns on how generally prevalent the death of two children in the same family is. If the probability of two children dying is equal to the probability of double murder, then we would conclude that it is certain that our hypothesis is correct – double murder explains all the double deaths we see. Further, given that we are certain to find two deaths if our double murder hypothesis is correct, we can never have a smaller probability of double deaths than our initial probability of double murder.

In fact, we know that there is another cause for such double deaths: SIDS; so the probability of two deaths is higher than the probability of double murder (some double deaths will be caused by SIDS and not murder). Assuming, as we are, that SIDS and double murder have the same initial probability (1/73m), the general probability of double deaths is 2/73m (1/73m + 1/73m).[5]

Pulling this all together, the probability of our hypothesis (double murder) being correct in light of there having been two deaths is:

In other words, if the only fact we have is that two children died, we can only say that it is 50% likely that it was due to double murder.

There is a very good example of the prosecutor's fallacy in tax cases (aside from MTIC, on which see below); the case of *Joseph Okolo v. HMRC* [2012] UKUT 416 (TCC). Essentially, the taxpayer submitted self-assessment returns declaring self-employment income from a business of property development. Turnover disclosed was high, but expenditure meant that only a small taxable profit was left. HMRC investigated and issued closure notices on the basis that there was no evidence to support the expenditure (leaving the turnover intact). The taxpayer's explanation was that he had submitted entirely fictional tax returns (i.e. there had never been a property development business) as part of a scheme to create a false impression of a substantial trading history in order to improve his ability to obtain loans.

At the FTT, Mr Okolo lost because the FTT found it:

> "…wholly improbable that the appellant would have made up such an elaborate lie for the first reason that he has given [to obtain a loan]…
>
> "We find it beyond credence that the appellant would have overstated his income knowing that that would result in him having to pay tax on sums which, according to him, he did not earn." (§§16 – 17).

In other words, the (semi-)innocent explanation was highly improbable (that he had lied in order to boost his creditworthiness) so the guilty explanation should be accepted (or, at least, the Appellant had failed to discharge his burden).

This was an impermissible inversion of probabilities. The low probability that someone would submit false tax returns in the hope of getting a loan does not provide any grounds, of itself, for rejecting that explanation or attaching a high probability to the turnover being real.

Fortunately, on appeal, the Upper Tribunal corrected this error:

> "I agree with the tribunal that, at first blush, it appears implausible; but I agree with counsel for Mr Okolo that the alternative is even more implausible". (§33).

The alternative was, inter alia, that Mr Okolo, a person with no apparent experience of the building industry and employed full-time in a completely unrelated sector, should have carried on a substantial and highly profitable contractor's business in his spare time; that the turnover of that business should have been generated entirely in cash and the profits hidden in some unexplained manner (§32).

Returning to MTIC, the way HMRC typically present the argument is that the mastermind would not be able to cause the money to flow in a circle without the taxpayer's knowing involvement; money moved in a circle, therefore we can infer knowing involvement. If it was an actual impossibility for money to flow in a circle without the taxpayer's knowing involvement, the logic would be sound: there is no other possible cause. In fact, this is not true (one possibility is knowing involvement, the other is that the mastermind correctly identifies the MTIC deal by chance) and the highest HMRC can really put their argument is that it is highly unlikely ("an incredible coincidence").

Once that is recognised, we can see that their argument inverts the probabilities (i.e. makes use of the prosecutor's fallacy): the probability of circular money flows without knowing involvement is very low, therefore the probability of knowing involvement is correspondingly high. What is missing is a consideration of the initial probability of the hypothesis, namely that this person has knowingly participated in MTIC fraud.

Picking up the Bayes way of thinking (i.e. the correct way of thinking) we need to first work out what the initial likelihood of our hypothesis is, i.e. the taxpayer being knowingly involved in MTIC fraud (initial Prob of H). Depending upon the other evidence available this may be higher or lower than the general probability that a person caught up in MTIC was knowingly involved.

Next, to take account of our new piece of evidence (circular payments), we multiply by

We can assume, for present purposes, that the probability of circular payments if T is knowingly involved is 1.

So what we find is that the effect of circular payments on our initial confidence in our hypothesis (knowing involvement) depends on the general prevalence of circularity in MTIC deals. If, as explained above, there is a mechanism whereby the mastermind can correctly identify his MTIC deal with or without the taxpayer's knowing involvement, then the general probability of circular payments in MTIC could be expected to be 1 and the existence of circular payments has no effect on our confidence in our hypothesis.[6]

If the general probability of circular payments is less than 1, our confidence in our hypothesis increases inversely in proportion to that general probability. Thus, if the general probability of circular payments is 10/11, then we increase our confidence in our hypothesis by 10%.[7]

Whilst the above is specifically in relation to circular payments, the same way of thinking applies to all the evidence that is presented to the Tribunal: identify initial confidence in hypothesis, adjust to take account of new evidence.

Furthermore, a vital point is to avoid the prosecutor's fallacy in relation to the overall conclusion. One sometimes sees reasoning that looks suspiciously like: "pieces of evidence A, B, C and D would, taken together, be extremely unlikely if T was not knowingly involved, therefore it is very likely T was knowingly involved". In such reasoning there is no apparent consideration of the initial likelihood of the conclusion, before taking account of such evidence, and a proper conclusion must take into account the countervailing evidence – the evidence that makes the hypothesis less likely.[8] There is no probability, barring certainty, that renders it unnecessary to at least consider the evidence pointing in the opposite direction.

Conclusion

The purpose of this article has not been, in any sense, to encourage the use of complex mathematical calculations in tax cases. Sometimes such calculations may be appropriate, often they will not.[9]

Instead, the purpose has been to encourage critical reflection on the way we think about and assess probabilities when factual issues are disputed. Thus:

(a) We should be resistant to simply accepting assertions that something is very unlikely – it may be very unlikely, but we need to consider what the underlying mechanism that makes it unlikely really is, and whether there might be some complexity we are missing.

(b) We should always be suspicious of attempts to argue "explanation X is unlikely, therefore, alternative explanation Y is likely". It has an intuitive appeal, but as a general proposition it is wrong.

More often than not we do follow these rules without deliberately thinking about it, but we cannot and should not conclude from this that we always do. Magicians are a constant reminder that the probable can turn up dressed as the improbable:

> "One of the best-kept secrets we have as magicians is that laymen would never imagine we would work so hard to fool them."[10]

MTIC fraudsters are criminals, not magicians (although the two are not mutually exclusive), but HMRC, the Tribunals and innocent taxpayers should not underestimate the lengths they went to to fool them and to achieve their purpose.[11]

Endnotes

1 See further Fooling Houdini by Alex Stone, in particular at page 253.

2 Indeed, a pack of cards only becomes substantially random after seven rifle shuffles.

3 *Kittel*, ECJ, C-439/04.

4 Note that the risk that someone other than the intended foreign customer will buy the MTIC goods is not really a risk to the mastermind (assuming the goods are genuine): that legitimate person's money flows up to the importer, who defaults and the fraud is complete as normal. The problem with committing MTIC fraud in this way (i.e. relying on demand from the legitimate market) is that once legitimate demand is saturated, no more MTIC can be carried out. By posing as a foreign customer, the mastermind generates artificial demand (but does not have to pay VAT itself, because the export to the foreign customer is zero-rated).

What the mastermind does not want to do, however, is to buy someone else's goods for export, as that would mean money does not flow up to the mastermind's importer and there is no opportunity to default. The problem to be solved is thus one of the mastermind avoiding purchases of non-MTIC goods.

5 There is a more complicated way of reaching this conclusion, which is contained in some representations of Bayes theorem. The general probability of double death ("B") is:

The first stage presents no difficulty: the probability of a double death if our double murder hypothesis is correct is 1, and the probability of our hypothesis being correct without being aware of B is 1/73m.

The second stage requires care. The probability of not H is the inverse of the probability of H (i.e. 72,9999,999/73m), however the probability of B given not H is not 1/73m, it is 1/72,999,999.

This is again slightly counterintuitive – why does assuming there has been no double murder (i.e. not H) increase the probability of a double death due to SIDS? The answer is that within a sample of 73m mothers we would expect to see one double murder and one double SIDS. If we exclude the instance of double murder we exclude one instance where there has been no SIDS and our sample size decreases by one.

By way of analogy, consider the chances of rolling a dice and obtaining on even number if the dice does not show a 4. There are three causes of an even number: 2, 4 and 6. Initially, each has a probability of 1/6. By excluding the possibility of a 4, however, the probability of each cause increases (to 1/5 – there are now only five possibilities 1, 2, 3, 5, 6) even though the overall probability of the even outcome decreases.

6 The mastermind may or may not want circular payments in a particular case. However, once it is established that the mastermind can choose to have circular money flows even if T is not knowingly involved, then there is no reason to think that he would abstain from circular payments more often in cases where T is not knowingly involved as compared to where T is knowingly involved. Thus, one would have to revise the assumption in the numerator (that the probability of circular payments if T is knowingly involved = 1) by the same amount, with no overall effect.

7 1/(10/11) = 11/10 which is the same as multiplying our initial probability by 110%.

8 As long as correct reasoning is followed, it does not matter in which order one takes account of evidence: an increase in confidence in the hypothesis of 10% followed by a decrease of 40% is the same as a decrease of 40% followed by an increase of 10%.

9 In fact, the Court of Appeal has rejected the very concept of using probabilities to refer to past events as "intrinsically unsound":

"The chances of something happening in the future may be expressed in terms of percentage. Epidemiological evidence may enable doctors to say that on average smokers increase their risk of lung cancer by X%. But you cannot properly say that there is a 25 per cent chance that something has happened: *Hotson v. East Berkshire Health Authority* [1987] AC 750. Either it has or it has not. In deciding a question of past fact the court will, of course, give the answer which it believes is more likely to be (more probably) the right answer than the wrong answer, but it arrives at its conclusion by considering, on an overall assessment of the evidence (i.e. on a preponderance of the evidence), whether the case for believing that the suggested event happened is more compelling than the case for not reaching that belief (which is not necessarily the same as believing positively that it did not happen)." (*Nulty v. Milton Keynes BC* [2013] EWCA Civ 15 at §37).

This went far further than was necessary to decide the case in front of it, apparently amounts to a rejection of Bayes theorem and is inconsistent with basic notion that something being "more likely than not" is expressing a view on the probability of a past event, albeit not in specific percentage terms

10 Jamy Ian Swiss quoted in Fooling Houdini by Alex Stone at p.272.

11 We know that HMRC are perfectly willing to rely on other ways in which the fraudsters tried to fool HMRC, in particular, contra-trading.