Chapter 14

Interpreting the Odds: New Zealand's Sports Bookies as Expert Forecasters

John Fountain

Introduction

Risk taking is a fact of Kiwi life. Christchurch cyclists take incredible risks in head to head bets with automobiles on crowded suburban streets. Wellington primary schoolers learn to play pedestrian Russian roulette every morning at their local zebra-crossing. Auckland households faithfully lay down $20-$30 a week in bets with insurance companies on the suburban trifecta: adverse risks to the family home, the family car(s), and family health. Sex – protected or not - is another all time favourite Kiwi gamble. And don't forget about education, small businesses or financial market transactions.

But the fact that we face, and take, risks in every facet of our lives doesn't mean that we perceive or manage those risks effectively or efficiently. History and theory, as well as field and experimental social science research document all too clearly the limited abilities of people to deal effectively with risk (Kleindorfer and Kunreuther, 1993).

This is where the sports betting bookies of the NZ Totalisator Agency Board (TAB) come in. In a limited field (sports activities) this small group of people (6 bookies plus 1 sports betting manager) in Wellington have developed significant in-house expertise in essential tasks of risk management: inference, prediction, valuation and choice. The simple but profitable methods they use to communicate and analyse their uncertainties need to be known and appreciated more widely in management and in science.

This essay has two parts. The first part develops a simple “demand and supply” perspective on pricing in sports betting markets. The next part of the paper analyses a distinctive feature of the TAB bookies' expectations: their use of subjective probability ranges for assessing uncertain costs and the adaptation of those probabilities in light of new information.

Along with a dose of theory this essay is based on 9 hours of interviews with the TAB bookies in Wellington. I undertook these interviews hoping to discover information on subjective probability distributions for sporting outcomes based on explicit quantitative models, optimal information search and optimising decision making procedures. Instead I found that sophisticated quantitative modelling plays only a minor, secondary role, to on-the-spot informed expert judgement. No R-squares, no hypothesis tests, no significance levels, no p-values, no regressions, no likelihood functions, no prior/posterior probability calculations, no scoring rules, no calibration charts, no expected utility functions. Just odds, bets and profits. The conclusion to the essay suggests that perhaps this is all we (academics, scientists and policy makers) need to improve our reasoning under uncertainty.

Odds, percentages, and price setting: an economist’s perspective

Imagine a sporting contest between two teams A and B. A friendly bet between two sports fans is a simple exchange of promises. For example, Alpha, backing A to win, promises she will pay Beta $10 if A loses in return for Beta promising to pay Alpha $5, if A wins.

Bets with the TAB bookies differ in one important respect from friendly bets: the only promise exchanged is the TAB's. When Alpha bets with the bookies on A to win she actually pays the money up front, say $10. In return she receives the bookies' promise to pay pre-specified ***odds*** on A to win O_a, say $1.50, for every dollar spent if A wins and nothing otherwise. This change has no effect on the financial aspects of the exchange: Alpha still gets $5 if A wins ($15 from the bookies less the $10 she paid) and the bookies get $10 from Alpha if A doesn't win. But it does affect the risk of default. How many times have you made a bet with a friend and never been paid, or, for shame, reneged on a bet yourself? Friendly bets, especially the proverbial one with a stranger in the smoke filled backroom of a dingy pub, typically lack credibility. By contrast, bets with the bookies harness powerful mechanisms of reputations and outside enforcement of contracts to achieve credibility (Dixit and Skeath, 1999: 310-312).

During any one year the TAB bookies accept bets worth millions of dollars on hundreds of sporting contests with thousands of bettors just like Alpha. The natural question for an economist to ask is: how can one explain the terms of trade – quoted odds and the amount of contingent real wealth exchanged - in these transactions? The natural way for an economist to answer such question is in terms of demand and supply.

But just what is being demanded, and what supplied, and at what prices? To clarify these points, consider a simple win/lose contest between team A and team B (ignore the possibility of a draw). Define a *ticket* on whether A wins so that one A-ticket pays $1 if A wins and $0 otherwise. One ticket on A is really just an obligation of the TAB to pay out $1 to the bearer. Put another way, one event ticket specifies a claim to $1 of TAB wealth if the specified event occurs and nothing otherwise. Within limits (discussed later) the number of tickets on A winning that a bettor buys Q_a at any given price P_a is up to her. When the TAB offers odds O_a on A winning at $1.51, they announce their willingness to sell 1.51 tickets on A for $1. Hence tickets on A must have a *unit-price* P_a of 1/$1.51=66¢, the reciprocal of the odds O_a on A winning. Unit prices for tickets are known as "percentages" in betting circles.

Why do we distinguish odds on $1 spent from unit prices of tickets? In one respect it is only a matter of language. After all, any price is simply a rate of exchange. Just as oranges that sell 4 for $1 are 25¢ each, so too, A-tickets that sell at odds of $1.51 for $1 are 66¢ each. In another respect, since language shapes and limits the concepts we reason with (Deacon, 1977), some linguistic formulations offer insights that others do not. The economist's language here may be helpful because it connects us to the rich conceptual world of economic theory that can (maybe) improve our understanding of the terms of trade in sports betting markets.

Since the TAB has exclusive statutory rights within NZ to sell bets on sporting events, monopoly theory might help explain the terms of trade for TAB bets. Monopolists, whether they sell apples, oranges, or sports bets, will set their unit-prices (P) at a predictable mark-up factor $k/(k-1)$ above their unit costs (C), as in equation 1 below:

(1) $$P = \frac{k}{k-1} \bullet C$$

While no one expects this formula to predict prices to the last decimal point, it does highlight expected relationships between market prices (*P)* and a key supply side factor, unit costs (*C)*, as well as a key demand side factor, price elasticity (*k)*, when there is only one seller in a market.

Unit costs appear straightforward enough, but what is a price elasticity *k*? Elasticity of demand for a firm's product is a measure of how responsive (or not) customer demand is to price *changes* the firm might make. *k* is a *number,* like 1.5 or 3.7, indicating that if prices decrease by a small amount, say 1%, total quantities of the product demanded by consumers will increase by *k*%, 1.5% or 3.7% in our example. Having the mark-up factor *k/(k-1)* based on an elasticity measure captures our intuitive notion that monopolists can "rip off" their customers more when customers' price sensitivity (elasticity *k*) is low compared to when their customer's price sensitivity (elasticity *k*) is high.

Elasticities *(k)* are hard to measure because you have to be able to figure out what customers *would do* if the firm changes its prices, and only its prices, a little bit. Either an experiment has to be undertaken, which is costly and may not generate the data you need, or some sophisticated statistical inferences need to be made, which are also costly and may require data you don't have. These measurement difficulties mean it won't be possible in this paper to check out whether equation 1 is the best demand and supply theory to use, but it still points us in the right direction.

Equation 1 looks deceptively simple, perhaps too simple. For one thing, the TAB bookies are not monopolists, so why is monopoly theory being used? Second, TAB bookies sell many types of bets, not just one commodity with one price as equation 1 presumes. Does this matter? Third, the products the TAB bookies sell involve uncertainty for themselves and for their customers. How do these uncertainties enter equation 1? We will consider the first two issues, monopoly and multiple products, now, since the rest of the essay analyses the uncertainty question in detail.

The TAB bookies are *not* monopolists. NZ bettors can easily, and legally, place bets on a wide variety of professional sports on the websites of reputable sports betting agencies in Australia, England, the US, and many other countries. If TAB prices get too far out of line with the rest of the world their sophisticated customers (with elastic demands) will desert them and bet with other international agencies. The proportion of sophisticated bettors varies from match to match and sport to sport, but the bookies put it in a range of 25% to 50% of total turnover. The lower number is alleged to prevail on the well publicised matches involving important NZ teams (e.g., test rugby, cricket, netball). But don't underestimate the degree of market power (relatively inelastic, loyal demand) the Tab bookies have. For example, the recent Tua/Lennox match saw the Tab bookies selling large volumes of bets from loyal Kiwi fans on Tua to win around odds of $2.50 while the same bet was available from other bookies over the Internet at odds of between $4-$5.

However, the presence of competition for the TAB bookies doesn't alter the fundamental logic of equation 1. Profit maximisation implies there will still be a positive price/ cost mark-up $k/(k-1)$. But now the elasticity of demand for bets k with the TAB depends on strategic issues like prices charged by competitors and the willingness and ability of customers to switch between suppliers.

The logical argument behind equation 1 assumes a single product firm. Yet the TAB bookies offer hundreds of different types of bets across a wide variety of sports, from bets on winners to bets on exact point scores, on first goal scorers, on overall competition winners, and on combinations of these sorts of events. Surely elasticity of demand isn't the only factor governing pricing in these situations? Technical details of optimal multi product pricing can become quite complex due to spillover effects between demands and costs for different commodities (Wilson and Electric Power Research Institute, 1993). For example prices below unit costs appear to make losses, but such pricing may be profitable overall if positive demand for one commodity favourably influences the demands for, and profits on, other commodities. But as a general rule, the idea of larger positive price/marginal cost margins for less elastic demands is a good first approximation to profit maximising pricing behaviour, and we will stick with that for this essay.

We need to know something about the unit production costs (C) of supplying bets in order to use equation 1. As in any business, production costs are not the same thing as sales revenues. While there may be some fixed costs and scale economies associated with ticket sales and customer servicing, the largest costs facing TAB bookies are simply to meet the financial obligations on the tickets they sell. To keep things simple, ignore any costs other than the cost of meeting these contingent obligations.

At the time that the TAB bookies are selling tickets the cost per ticket sold is unknown. They can, however, arrive at an *expectation* about the unit cost of supplying an A or B ticket by assessing *subjective probabilities* π_a and π_b for the uncertain events "A wins" and "B wins" respectively. If the TAB bookies assess such subjective probabilities then their expected unit costs are $\pi_a \bullet \$1 + \pi_b \bullet \$0 = \pi_a$ per A-ticket and $\pi_b \bullet \$1 + \pi_a \bullet \$0 = \pi_b$ per B-ticket respectively. That is, the TAB bookies' *expected unit cost* of supplying an event ticket is numerically equal to their subjective probability for the event. It is these expected unit costs that are the relevant costs for *expected* profit maximising price setting using equation 1.

Don't be put off by the algebra or the language of probability here. This isn't higher mathematics - just good bookkeeping. Suppose the bookies think there is an 80% chance that a hot team like Manchester United will win any game it plays. Then if a punter buys exactly one ticket on Manchester United to win, the bookies expect to pay out one dollar 8 out of 10 times and nothing the other the other 2 times, an average of $0.80 per ticket sold. Naturally they will try to price tickets on Manchester United above their expected costs of 80¢ per ticket so as to make at least a small profit (ie set odds on an United win at something below 1/0.8=$1.25).

The general idea should now be clear. An economist using equation 1 as her demand and supply model will explain sports bets' prices in terms of **supplier subjective probabilities of outcomes of sporting events and elasticities of demand, or both**. Other things equal (i.e., elasticities of demand), when TAB bookies assess higher probabilities for events they will set higher unit-prices (lower odds) on those events. When the NPC division leaders Canterbury play last place Southland, the TAB bookies quoted low odds of $1.01, ie high unit prices around 99¢ per ticket on a Canterbury win. Remember a ticket on a Canterbury win pays its owner $1 for such a win and nothing otherwise. One suspects (using equation 1) that the key factor determining these odds is expected cost (high probability for a Canterbury win) rather than any price inelasticity of demand for Canterbury tickets. On the other hand, when the TAB bookies sold tickets on Tua to defeat Lennox at prices of 40¢ per ticket (odds of $2.50), one suspects that it is inelasticity of demand that is driving the TAB's prices (odds), not their expected costs. The TAB bookies I talked to in Wellington claimed that personally they didn't give Tua much of a chance to defeat Lennox, indicating low expected unit costs. But there is more to profitable pricing than low unit costs. A large number of price insensitive local bettors appeared to be backing the hometown boy with their hearts not their heads.

So far we (via equation 1) have treated pricing in sports betting markets separately. This makes some sense because fans loyal to team A may have different demand elasticities than fans loyal to team B, and so should be charged different prices. But there is a supply side factor to consider too, namely expected costs of supplying bets on both sides of the market. Suppose the TAB bookies manage to sell balanced packages of A and B tickets: 1 ticket on A to win at price P_a and 1 ticket on B to win at price P_b. No matter what happens they will have to pay out $1 (remember we have assumed away the possibility of a draw or postponement). But in return they have received revenues of P_a+P_b no matter what happens. Their net gain on such a balanced transaction is of P_a+P_b-1.

The TAB bookies refer to the sum of the unit-prices $P_a + P_b$ on a balanced transaction as the "*total percentage*" and to the difference $P_a + P_b - 1$ as the "*total percentage margin*". They *always* set their odds (reciprocals of the unit-prices) to ensure that their total percentage margins are positive. In fact the TAB bookies have a set of "rules of thumb" specifying uniform total percentage margins across a variety of sports, such as rugby, soccer, tennis, speed racing, and golf. For example, when there are only 2 possible outcomes, either side winning, total percentage margins are about 8%. Add in the possibility of a draw, and the total percentage margin rises to approximately 12%. On exact score bets in soccer games, or first try scorers in rugby games, where there may be 10, 20, 30 or more eventualities, total percentage margins can run as high as 40-50%. An interesting avenue for further research is how these rules of thumb on total percentages relate to optimal pricing rules like equation 1.

Why should total percentages exceed 1? If not, the bookies could be made into *sure losers*. For example, suppose that unit-prices are 0.66 (66¢) on A tickets and 0.30 (30¢) on B tickets, for a total percentage of 0.90 (90¢). A bettor could simultaneously buy 1 A-ticket and 1 B-ticket, *guaranteeing* herself $1 of TAB wealth while paying out only 96¢, for a risk-less profit of 4¢ per ticket. Holding 1 A-ticket and 1 B-ticket a bettor has a perfect hedge - the TAB has to pay out $1 either way, no matter who wins. To obtain this perfect hedge the bettor only had to pay 96¢ in total, 66¢ for the A-ticket and 30¢ for the B-ticket. Naturally TAB bookies don't like entering into like these where they are sure to lose. To avoid these situations they bookies set their total percentage above 1.

Of course these kinds of percentage calculations work in reverse too. Imagine the bookies have positive percentage margins, so that $P_a+P_b > 1$. If the bookies could find a bettor willing to buy one ticket on A to win and one ticket on B to win at the same time, that bettor could be made into a sure loser - and the bookies a sure winner. Yes, the TAB would have to pay such a bettor $1 no matter which team won or lost, but in return the bettor would have paid the TAB P_a+P_b which by assumption is more than $1.

But the TAB bookies don't need to find the *same person* willing to play both sides of the market to be sure winners. They only need to find *enough different people* betting in both A-ticket markets and B-ticket markets. For example the TAB may have sold 10,000 A-tickets at 66¢ a ticket (odds of $1.51=1/0.66 on A) and 8000 B-tickets at 42¢ a ticket (odds of $2.38=1/0.42 on B). It might appear that the TAB is facing a great deal of risk, but in fact their liability on the 8000 B-tickets is balanced by 8000 of the A-tickets. On these 8000 pairs of tickets they will have to pay out $8,000 for sure. After all, either A or B has to win the game. But on every one of these balanced 8000 A and B tickets they have a total percentage of 0.66+0.42=1.08. That is an 8% margin, an 8¢ margin, on each balanced pair for a guarantied profit of $640 (8% times $8000). To be sure they still have some risk should team A win, a possible gross liability of $2000 on the 2000 A-tickets that haven't been balanced by B-tickets. But remember these unbalanced 2000 tickets have pulled in $1320 in revenues ($1320=2000 tickets @66¢ a ticket), which the TAB keeps if A loses. If A wins the TAB's net liability on these 2000 tickets is only $680, $2000 to pay out less $1320 in revenues. Their overall wealth position is $640 for sure on the balanced portion of their transactions, plus either an additional $1320 if A loses or a loss of $680 if A wins.

In fact it may happen that the number of B-tickets sold exactly matches the number of A-tickets sold. In this case the bookies are said to have *balanced the book*. Balanced books at positive percentage margins are "nice" commodities for the TAB to have, since they have a risk-less profit rate equal to the total percentage margin. However, my conjecture is that in sports betting, balanced books are the exception, not the rule. TAB bookies are not content to merely balance books – they actively seek out bets which they may, but do not expect to, lose.

It may come as a surprise to you, but the TAB bookies <u>can</u>, and do, lose money betting head to head with their customers. While in Wellington I examined the frequency distribution of losses and profits to the TAB on head to head betting for an entire competition in a popular sport in a recent year. On a simple count, the number of matches where the TAB lost money was slightly greater than the number of matches where they made money. Practically, this means bettors were as likely to win *something* from the TAB as to lose *something* over the course of this competition. However, when the TAB bookies total up their losses and their wins, the overall result is that they made positive profits, despite taking losses on bets on more than half the matches.

Odds quotes as probability ranges

My original research plan was to learn how the TAB bookies assessed their expected costs of supplying tickets. I wanted to persuade them that they may be able to improve their forecasting efforts with sequential scoring rule methods used by meteorologists and statisticians (Kleindorfer and Kunreuther, 1993:98; Lad 1996, Ch. 6). Like most plans based on academic theories, it ignored the complexities of the TAB bookies life! There were no records of the on how they formed their expectations, only good and bad memories stored in their heads about the expectations setting process. Yet it turns out that there is a great deal of information revealed about bookies' subjective expectation formation from two indirect sources:

- their pricing behaviour and
- the ongoing dialogue the bookies have with one another around their computer terminals as they discuss appropriate adjustments to make to quoted odds on open books and planned odds for upcoming books.

Take a recent Monday night NFL football game in Denver between Denver Broncos/Oakland Raiders as a specific (hypothetical) example. (The choice of a sports contest far from Kiwi culture is deliberate, so Kiwi readers will have some empathy with the discovery process required by TAB bookies to formulate expectations about possibly profitable contracts.) We have already argued that at a conceptual level, the bookies' subjective probabilities π_r and π_b for the uncertain events R=(the Raiders win) and B=(the Broncos win) are the relevant marginal expected costs (remember we are assuming no draws for simplicity). It turns out we can work backwards and infer these subjective probabilities from the odds. For example, if a bookie suggest odds of \$1.54 on an R-win and \$2.31 on a B-win, these odds tell us that the bookie's subjective probability π_r on a Raiders win must be less than 65¢ but greater than 57¢, that is, π_r must lie within the range 0.57 to 0.65.

Where did these numbers come from? Consider first the upper limit of 0.65. Imagine a bettor with $10 to bet. At odds of $1.54 on a Raider win the bookie stands to lose $5.40 ($15.40 in payouts net of the $10 received) if the Raiders win and gain $10 otherwise. Her expected payoff from this gamble is her probability *weighted average* of the outcomes: $\pi_r \bullet (-5.40) + (1-\pi_r) \bullet 10$. What is an "expected" payoff for the TAB? An expected payoff is really just a kind of mental accounting for what *might* happen, before anything does happen, to reflect how likely one thinks the various possibilities are. The TAB is either going to lose $5.40 or win $10. The expected payoff concept averages out these two numbers to arrive at a figure somewhere "in between" that reflects the TAB bookie's assessment of which payoff number is going to occur. The term $\pi_r \bullet (-5.40)$ weights the negative number –$5.40 (a loss) by π_r, the chance the TAB assesses for a Raider win. The term $(1-\pi_r) \bullet 10$ weights the positive number $10 (a gain) by $1-\pi_r$, the chance the TAB assesses for a Raider loss. Each of these weights, π_r and $1-\pi_r$, is some fraction between 0 and 1, which is why we get something "in between" the two numbers, the –$5.40 loss and the $10 gain.

Now back to the original question, where does the upper limit of 0.65 on a probability assessment come from? Basically bookies don't want to enter into transactions they expect to lose money on. Of course *on the day* the bookie might win or might lose, but *before the day* it doesn't make a lot of sense to enter in to transactions where one expects to lose. Intuitively, for the TAB's expected payoff to be positive the chances she assesses for a Raiders win, where she loses $5.40, must not be too high. How high is too high? Well we actually have a formula for the TAB bookie's expected payoff: $\pi_r \bullet (-5.40) + (1-\pi_r) \bullet 10$. Using this formula we can write the requirement that the expected payoff be positive as an inequality $\pi_r \bullet (-5.40) + (1-\pi_r) \bullet 10 \geq 0$. A little algebra that rearranges the inequality shows that $\pi_r \leq 1/1.54=0.65$. That is, *the bookie's subjective probability on a Raiders win, π_r, must not exceed the quoted unit-price on R tickets, 65¢*. Otherwise she would be accepting a small gamble with a negative *expected* return.

Of course a bettor could also front up with a $10 bet on the Broncos at odds of $2.31. The bookie's expected payoff on this bet is: $\pi_r \bullet (10) + (1-\pi_r) \bullet (-10.31)$. The TAB's loss of $10.31 arises when the Broncos win, coming from a payout of $23.10 less the $10 received. Remember, since we've ruled out ties, the Bronco's win only when the Raider's lose, and $(1-\pi_r)$ is the probability assessed by the TAB for that event. So the weight applied to the $10.31 loss for the TAB is $(1-\pi_r)$. Similarly, the TAB wins $10 on this bet only when the Bronco's lose, that is the Raider's win. The TAB has assessed a probability of π_r for this event, so $10 gets weighted by π_r in the expected payoff expression $\pi_r \bullet (10) + (1-\pi_r) \bullet (-10.31)$. Reasoning as above, for the expected payoff on this bet on the Bronco's to be positive the bookie's probability for the Broncos winning, $(1-\pi_r)$, can't be too high either. Writing down the inequality that expresses the requirement that an expected payoff be positive, $\pi_r \bullet (10) + (1-\pi_r) \bullet (-10.31) \geq 0$. Following the line of reasoning in the previous paragraph, or just doing the algebra to rearrange the terms in the inequality, her probability on the Broncos winning $(1-\pi_r)$ must be less than the unit-price on B-tickets 43%=1/2.31. But if $(1-\pi_r)<0.43$, then $\pi_r \geq 0.57$. That is, the bookie's probability on the Raiders winning, π_r, must be at least 57%. Taking both inequalities $\pi_r \leq 0.65$ and $\pi_r \geq 0.57$ together, odds of $1.54 on an Raiders and $2.31 on a Broncos win imply that the bookie's subjective probability π_r on a Raiders win will lie within the range 0.57 to 0.65.QED.

Notice that we have not specified a particular number indicating what the bookies subjective probability π_r for an R-win is. Bookies, like the rest of us, have difficulty in pinning down their subjective probabilities for events to precise numbers. They are more comfortable with "ballpark" assessments expressed as ranges. The odds range, \$1.54/\$2.31, via the associated probability range 0.57 to 0.65, defines the size of the ballpark. Another way to interpret this range is as a bid-ask spread, the way prices are quoted in the stock market.

The example we just worked through had specific unit prices 0.57 and 0.65 to work with, and we crunched out the calculations for a specific Bronco's/Raiders matchup. But the line of reasoning employed doesn't depend on the specific teams or numbers involved. Imagine a head to head contest between teams A and B (assume no ties for simplicity) where the bookies quote unit prices P_a for A-tickets and P_b for B-tickets with a total percentage $P_a + P_b$ exceeding 1. Then just as the range 0.57 to 0.65 characterised the TAB's subjective probabilities for a Raider win in the numeric example above, the range $1-P_b$ to P_a will characterise their subjective probabilities for A to win in this general case. For 2 outcome bets where the total percentage $P_a + P_b=1.08$ the TAB's subjective probabilities for A winning will lie within a narrow range of width 8%, $P_a–0.08$ to P_a since $1- P_b=1- (1.08–P_a)= P_a–0.08$. The percentage margin for the TAB, the difference between $P_a + P_b$ and 1, here 0.08, tells us the width of the range that the TAB's subjective probabilities lie in. Returning to our ballpark assessment metaphor, we can infer the bookies' "ballpark" subjective probabilities for the outcomes of sporting contests from the odds or unit prices they quote, with the size of the ballpark determined by the total percentages expressed in their quoted odds.

The TAB bookies do not pull their subjective probabilities π_r π_b out of a hat. Rather, they search for and gather specific bits of information believed to be informative for an assessment of the chances that the Raiders will win. Probabilities for one event that are dependent on information on other events are called conditional probabilities. The conditioning event is called a *cue* (Kleindorfer and Kunreuther, 1993: 71). A list of relevant cues for a Raider win against the Broncos NFL game would include: whether the game is home or away, which team has the better recent form, whether any key players have injuries, whether the Raiders who currently lead their conference will rest key players for upcoming playoffs, whether the referee has a known home team bias, etc. No cue list is complete, a fact well recognised by the TAB bookies. The incompleteness of any cue list is one of the key reasons explicit quantitative modelling takes a back seat to on the spot expert assessment in their decision making. But cue lists do provide a useful checklist, and are used as such by teams of sports bookies.

However, keeping a cue list does *not* tell us how the TAB bookies actually process cue information into specific numerical subjective probabilities, or ranges of probabilities. During my short interviews with the bookies I listened to the bookies describe the cue lists they have in mind when arriving at specific odds, but I was unable to observe specifically how they process cue information into odds. Eliciting information form the bookies about this process and then articulating it into an "expert system" model would be an extremely interesting avenue for future research. The next few paragraphs uses our hypothetical NFL example to illustrate the complexities involved.

A cue list will be used by bookies to generate a *sequence* of odds (subjective probability ranges) starting from some rough, tentative assessment, such as an historic average of wins and losses between the competitors over the years. Suppose those historic statistics say the Raiders have won 65% of their matches against Denver. Take the first cue: Is it a home game to Denver? Yes, so tend to favour the Broncos by reducing the probability of an R-win below 65%. How much below? Well suppose publicly available statistics show the Raiders have not won a regular game in Denver against the Broncos in 3 years. On this information alone one might drop the probability for a Raiders win down to something like 7% to 15%, expressed as odds of $6.70 (1/0.15) on an R-win and $1.08 (1/(1-0.07)) for a Bronco win. But who is in better recent form? Recent statistics might show that the Raiders have 8 wins and 1 loss, with the only loss at the beginning of the season by a small margin to a top ranked team, whereas the Broncos are 4 and 4 with some losses by large margins against poorly ranked teams. Favour the Raiders and reduce the odds back to even, $1.85/$1.85, ie a unit price range 0.46 to 0.54, with a total percentage of 1.08.

The TAB bookies actively foster open dialogue about relevant known cues and how those cues might shape the odds each bookie would find acceptable. You can imagine the dialogue between all the bookies, especially the 2 NFL specialists, as they proceed through a further sequence of these specific, known cues. I was told that mistakes in judgement by TAB bookies are made more frequently by the bookies when only one bookie is staffing the office. The dialogue between experts appears to prevent them from either ignoring or over or under exaggerating the importance of relevant cue information – something we regular mortals are prone to do through a variety of judgement “biases” (Kleindorfer and Kunreuther, 1993: 91-94).

Not all cues may be known at the time opening odds are being discussed and analysed. However, this does not mean that such cue information is deemed irrelevant. Consider the question of whether the Raiders will rest some of their key players for the playoffs. This cue is uncertain, just as the outcome of the game itself is uncertain. In order to reduce their uncertainty about unknown but relevant cues the bookies specialising in NFL games will scour the news for up to date information on injuries, as well as credible public information on strategic plans to rest players for future playoff games.

After researching, suppose our bookie strongly believes that the Raiders will rest key players and field a weaker team. A check of the Internet reveals that two international betting agencies are quoting odds of $1.50 / $2.30 for the Raiders to win. These subjective probability ranges are 57% to 67% chances on the Raiders winning: the lower end of the range is 1-1/(2.30)=0.57 while the upper end of the range is 1/1.50= 0.67. Our TAB bookie puts a lot of weight on the Raiders resting some of their key players and is willing to transact at odds of around $2.20/$1.60, indicating a range of subjective probabilities for a Raiders win of 0.375 to 0.454 much less than other bookies are apparently assessing. The other bookies in the TAB team recognise the point being made but aren’t so sure about two key bits of reasoning. First, the chances of the Raiders making a strategic move like this may not be that great. After all whether they play home or away in the playoffs is still at risk. Second, even if the Raiders do play a second string, their depth of talent is so great they can beat anyone on the day. That is more cues are listed and integrated into the price setting process. The first bookie's experience in the field prevails but to accommodate the views of the other bookies opening odds are dropped back to $2.00/$1.70, a slightly higher subjective probability range of 0.42 to 0.50.

Being $0.50 above the big international bookies on a team like the Raiders, a bet for $10,000 from an overseas source quickly flashes across the computer screen at the office. Typically large bets like these are flagged for specific decision by the TAB bookies. Total betting is low and balanced on this Monday night football match in NZ, so by accepting this bet the TAB liability will be high, $10,000 if the Raiders win. A quick conference is held (the team has 1 minute to decide). The experienced bookie still wants to stick with his opening odds, but the computer flashes up the ID of the bettor. This guy is on a "hit-list" of 35 or so sophisticated big bettors to be wary of. Aha! The bookies make an inference that this bettor probably has private information about the game that they do not and they decline the bet, moving their odds back in with the rest of the world to $1.50/$2.30.

By the way, Denver won the game, even though the Raiders fielded a full strength team. It appears that the local TAB bookie's forecasting ability was better than that of the overseas investor, the overseas bookies, and the rest of the NZ TAB team. After all, the experienced TAB bookie expected a Raider win with lower probability ranges 0.375 to 0.454 than other overseas bookies whose probability range 0.57 is 0.67, and hence basically higher probabilities for the event of a Denver win. But these better expectations on the outcome were for the wrong reasons (poor expectations about unknown but relevant clues). Yet looking back at the sequence of odds and associated probability range assessments, without even considering that unknown cue event, the bookies would have forecast lower probabilities for a Raider win, 0.46 to 0.54, than the range revealed the rest of the world was forecasting, 0.57 to 0.67. So perhaps the bookies are still better forecasters than their international counterparts – at least on this particular NFL game.

Conclusion

The previous story is largely hypothetical, although events along the lines I described did in fact occur during my interview time. It highlights how the ongoing odds dialogue between TAB bookies reveals their conditional probability assessments in light of known, and expected, cue information. Their use of subjective probability ranges (implicitly revealed by their odds for and against a specific event) is unique amongst expert forecasters.

Why don't academics, managers, policy makers, scientists and other experts employ the TAB bookies form of predictive reasoning? The short answer is differences in incentives. That's probably also the long answer. Every forecast made by the TAB bookies in the form of odds can be "challenged" by bettors in the market. Bookies always have to back up their beliefs by "putting up or shutting up" in an ever changing world of sport. The methods they employ have been adapted to be successful at that task. But, business managers aside, many academics, policy makers, and scientists have developed and use a vast apparatus of classical and Bayesian statistics to support their endeavours – without ever having to bet money, or possibly their jobs, on their forecasting ability.

Of course prominent statisticians and econometricians have called for academics to change their ways about decision making and forecasting under uncertainty in the direction of a TAB model.

"Employ, until you have further experience, that expert [real or synthetic] whose past opinions, applied to your affairs, would have yielded you the largest average income" (Savage, 1971: 791).

"[A]uthors of scientific reports [should be] expected to select bid and asking prices for lotteries contingent on interesting and decideable events, and be required to back up those claims with currency.... The intent of this scheme is to provide a clear incentive for authors to report what they believe, but only what they believe" (Leamer, 1986: 220)

But what odds for this event do you think the TAB bookies would assess?