6. Unpredictability

Unpredictability is a critical element of strategy whenever one side likes a coincidence of actions while the other wishes to avoid it.

- The Commissioner wants to audit tax evaders, those who have cheated hope to avoid an audit.
- The elder sister wants to rid herself of the younger brother, who wants to be included.
- An invading army wants its choice of place of attack to surprise, the defending army wants to concentrate its forces on the place of attack.
- The beautiful people want exclusivity, the hoi polloi want to be up with the latest trends.

While the taxman's or the attackers' decision on any occasion may be unpredictable, there are rules which govern the selection. *The correct amount of unpredictability should not be left to chance.* The odds of choosing one move over another can be precisely determined from the particulars of the game.

6.1 Anyone for Tennis?

The server, Stefan, wants to minimise the probability that the receiver, Rod, can return serve, and Rod wants to maximise this probability. It's a *zero-sum game*: Stefan's win is Rod's loss.

If Rod can anticipate Stefan's aim (to Rod's forehand or backhand) then Rod will move appropriately (forehand or backhand) to increase the probability of a successful return. Stefan will try to disguise or mislead Rod until the last second, hoping to catch Rod off guard and wrong-footed.

Tennis Serve & Return

Consider a 2×2 payoff matrix which sets out the percentages of Rod's successfully returning server:

Stefan: the *S*erver;
Rod: the *R*eceiver.

		Stefan's Aim	
		Forehand	Backhand
Rod's move	Forehand	90, 10	20, 80
	Backhand	30, 70	60, 40

TABLE 1. The percentage of times Rod successfully returns. A non-cooperative, zero-sum game. (Rod, Stefan).

No Nash equilibrium in pure strategies.

Stefan’s task

Stefan wants to keep the successful return percentage as low as possible; Rod has the exact opposite interest. If the two players decide on their strategies before the match, knowing the above probabilities, what should their strategies be?

To help answer this question, we now plot:

the percentage of times Rod returns serve against the probability of Stefan aiming to Rod’s forehand.

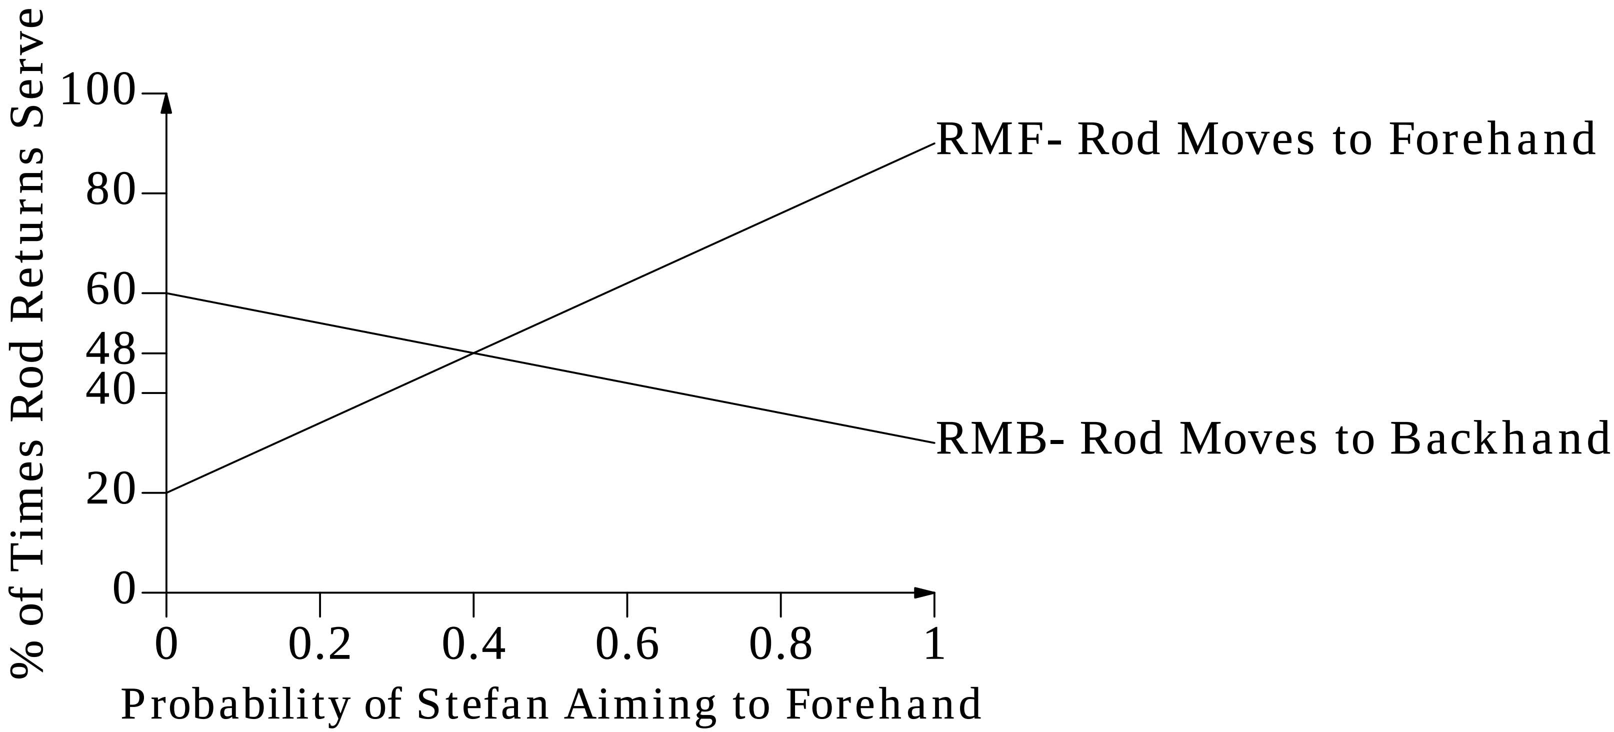
% of Times Rod Returns Serve
100
80
60
48
40
20
0
0
0.2
0.4
0.6
0.8
1
Probability of Stefan Aiming to Forehand
RMF- Rod Moves to Forehand
RMB- Rod Moves to Backhand

Mixing strategies

By plotting the two straight lines, we're considering the possibility that Stefan (and Rod) can *mix* their moves, using probability:

Stefan: "if I always serve to the forehand, then the serve will be returned 90% of the time, but if I always serve to the backhand, the percentage falls to 60%. In both cases, Rod learns to correctly anticipate what my (unchanging or pure) strategy is.

"What if I mix my shots and serve half to the forehand and half to the backhand at random? Then Rod will be kept guessing, and won't be able to anticipate correctly all the time."

- *If Rod anticipates forehand*, he will be right with probability half (and return 90% of the time) and will be wrong with probability half (and return only 20% of the time). The percentage of successful returns will be 90/2 + 20/2 = 55%.
- *If Rod anticipates backhand*, the percentage of success will fall to 60/2 + 30/2 = 45%.

The best mix

In a simultaneous-move game, which this is, Rod will be better off (55% success) if he moves towards Stefan's forehand.

For Stefan, a return percentage of 55% is better than the 90% or 60% of unchanging serving. (Remember: Stefan the server wants to *minimise* the percentage of successful returns by Rod the receiver.)

From the diagram, we readily see that Stefan's *best mix* is to serve to the forehand with probability of 0.4, which results in a successful rate of return of 48%, the best (lowest) Stefan can achieve. At this mix, Rod is indifferent between moving to forehand or moving to backhand: Rod cannot improve the success rate of 48%.

The exact proportions of the mix follow from the four outcomes of the basic interaction. If these numbers change, so will the best mixed strategy.

Rod’s task

From Rod’s point of view, we get a different chart:

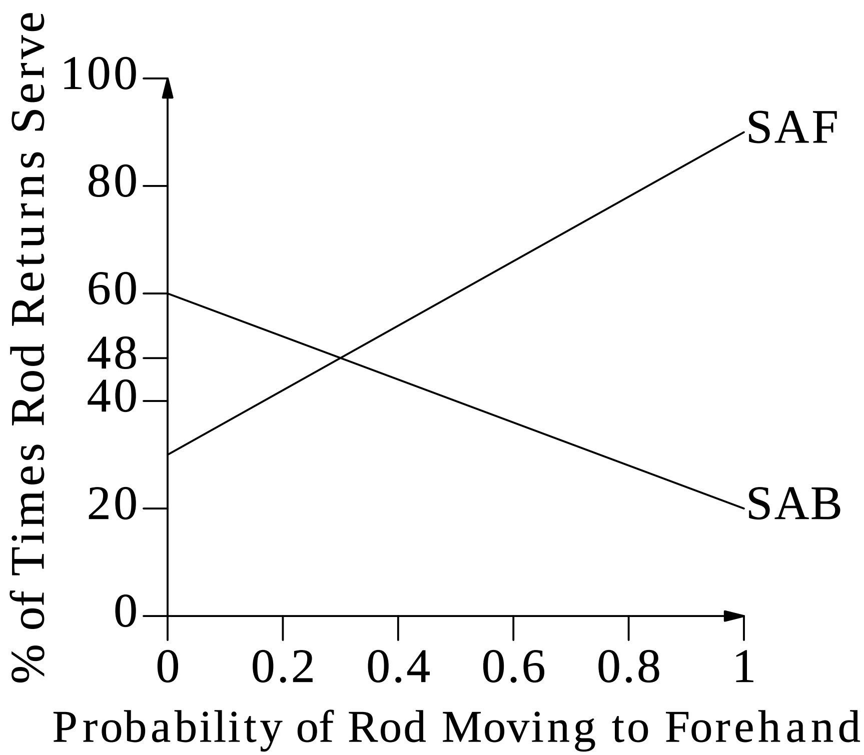

A Nash equilibrium at RMF: 0.3, SAF: 0.4.

SAF: Stefan aims at forehand
SAB: Stefan aims at backhand

A symmetry

There are two lines, one corresponding to Stefan aiming to forehand, one to backhand. The percentage of successful returns depends on both player's moves, from the payoff matrix.

We can readily see that that as Rod's probability of forehand returns increases, above 0.3, the rate of his success rises to 90%; below 0.3 forehand, the rate also increases to 60%. At 0.3 forehand, the rate of successful returns is 48%. Stefan responds appropriately.

Note that each player reaches the same rate of a successful return: 48%. Using his best mix Stefan is able to keep Rod down to this, the best Rod is able to achieve using his best mix.

The Min-Max Theorem

This property of zero-sum games is *the min-max theorem*:

> when, in zero-sum games, one player attempts to minimise her opponent's maximum payoff, while her opponent attempts to maximise her own minimum payoff, the surprising conclusion is that the minimum of the maximum payoffs equals the maximum of the minimum payoffs.

Neither player can improve her position, and so these (mixed) strategies form an (Nash) equilibrium.

An equilibrium

The server, Stefan, will act as if the receiver, Rod, has correctly anticipated his mixing strategy and has responded optimally. The minimum of Rod's maximum percentage occurs where the two payoff lines cross, at Stefan's probability of forehands of 0.4 and a success rate of 48%.

Rod is trying to maximise his minimum payoff. If he moves to forehand and backhand equally frequently, then his rate of successfully returning serve varies between (20+60)/2 = 40% (when Stefan aims to backhand) and (30+90)/2 = 60% (when Stefan aims to forehand).

Obviously Rod should anticipate backhand slightly more. If his probability of moving to the forehand falls to 0.3, then the rate of successful returns is 48% for any probability of Stefan's aiming for forehand.

Conditions apply

Min-max doesn’t work where the game is not zero-sum, or where there are more than two players, or more than two moves per player.

When mixing is necessary, the way to find your own equilibrium mixture is to act so as to make others indifferent about their actions: you want to prevent others from exploiting any systematic behaviour of yours. If they had a preference for a particular action, that would mean that they had chosen the worst course from your perspective.

Note that for the first time the payoffs must be *cardinal* (an interval scale) and not just ordinal: we’re now interested in how much more preferred one outcome is over another, not just that one is preferred to another.

6.2 Mixed Strategy Equilibrium

A *pure* strategy: calls for the selection of exactly one action at each decision node.

No Nash equilibrium (N.E.) in this non-zero-sum game (Ava, Rusty):

H Ava \ Trusty Rusty	Low	High
Low	\$100, \$50	\$75, \$100
High	\$50, \$220	\$200, \$200

Resolve with *mixed* strategy, in which players choose actions randomly.

Poker: fold, raise, see. Bluffing. Unpredictability important.

Tennis: passing, lob, volley, overhead smash.

Simultaneous advertising

Let's look at a market situation with two rivals, Honest Ava and Trusty Rusty, deciding whether to advertise their used cars as Low priced or High priced, when the customers can be influenced by this advertising.

It's a simultaneous-move game, and neither knows until the local paper comes out just what the other has done. By then, of course, it may be too late ...

6.2.1 Car Pricing (Honest Ava and Trusty Rusty):

- let p_A be the probability of Honest Ava advertising a low price
- let p_R be the probability of Trusty Rusty advertising a low price
- Ava will choose p_A to maximise her expected payoff $E(\pi_A)$:

$$= p_A \, [\; p_R 100 + (1 - p_R) 75 \;] + (1 - p_A) \, [\; p_R 50 + (1 - p_R) 200 \;]$$

$$= (\$200 - \$150\, p_R) - (\$125 - \$175\, p_R)\, p_A$$

- Similarly, Rusty will choose p_R to maximise his expected payoff:

$$E(\pi_R) = (\$200 - \$100\, p_A) + (\$20 - \$70\, p_A)\, p_R$$

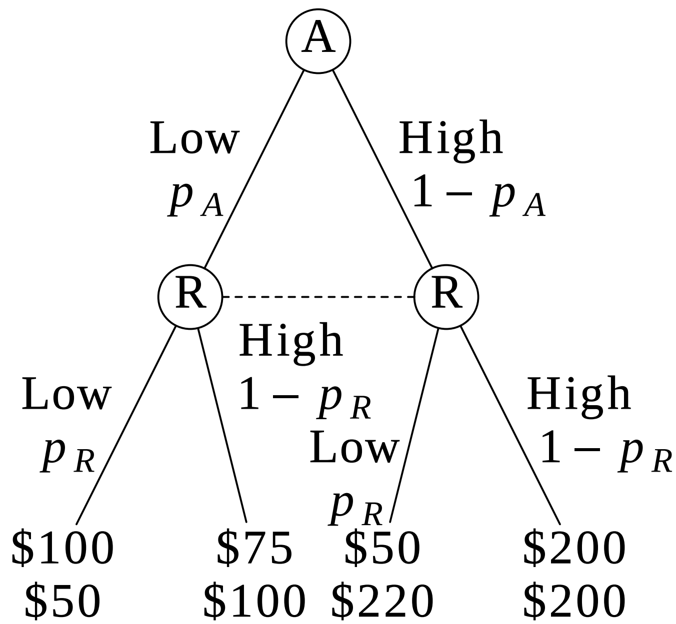

Forming beliefs

- *Honest Ava must form a belief about what Trusty Rusty believes she will do.* Not just a belief about what Rusty will do.
- Ava believes Rusty believes she (Ava) will choose a low price with p_A^e.
- If $p_A^e < 2/7$, then $\$20 - \$70\, p_A^e > 0$, and Rusty should set $p_R = 1$, and always price low.
- So Ava should also advertise a low price ($p_A = 1$).

  But this is inconsistent with Ava's conjecture of what Rusty believes she will do. ($p_A^e < 2/7$).
- If $p_A^e > 2/7$, then $\$20 - \$70\, p_A^e < 0$, and Rusty should set $p_R = 0$, and never price low.
- So: Ava should also never price low ($p_A = 0$), again inconsistent with the conjecture of $p_A^e > 2/7$. Only when $p_A^e = 2/7$ is Rusty indifferent between advertising low and high, — *unpredictable*.

An equilibrium

- And only when Rusty is unpredictable is it optimal for Ava to be unpredictable too.
- So: rational for Ava to believe Rusty finds Ava unpredictable only if Rusty's belief about Ava makes *him* unpredictable.
- And Rusty unpredictable only if he is *exactly indifferent between advertising low and high*, which
- happens iff $0 = \$20 - \$70 p_A$, or $p_A = 2/7$
- Similarly for Rusty, who forms beliefs about Ava's conjectures of his behaviour: Ava will be unpredictable iff $p_R^e = 5/7$.
- Ava's expected payoff $E(\pi_A)$ will be \$92.86 per week.
  Rusty's expected payoff $E(\pi_R)$ will be \$171.43 per week.

What about something such as Tit for Tat?

Such as Ava plays High and Rusty alternates between High and Low?

Then the profits alternate between ($50, $220) and ($200, $200).

But this is not an equilibrium. Why not? (Does either have an incentive to change?)

It relies on Ava keeping her price High.
But if it's Rusty's turn to price Low (and receive $220), then Ava will price Low too, with payoffs now of ($100, $50). She has doubled her payoff to $100 by lowering her price.

Absent some enforceable contract (with sufficient penalties), this proposal cannot be supported — it isn't an equilibrium.

Is there always a Nash equilibrium?

NB: The only pair of mixed strategies that is N.E. is Ava low with $p_A = 2/7$, and Rusty low with $p_R = 5/7$.

- Given what each believes the other player will do, and what each believes its rival believes it will do, neither has incentive to alter its beliefs, and so each is unpredictable: a N.E.
- Not necessary to randomise, only to appear unpredictable
- Mixed strategies are necessary for N.E. and sufficient:
- **Nash Existence Theorem:** *Every game with a finite number of players, each of whom has a finite number of pure strategies, possesses at least one Nash equilibrium, possibly in mixed strategies.*

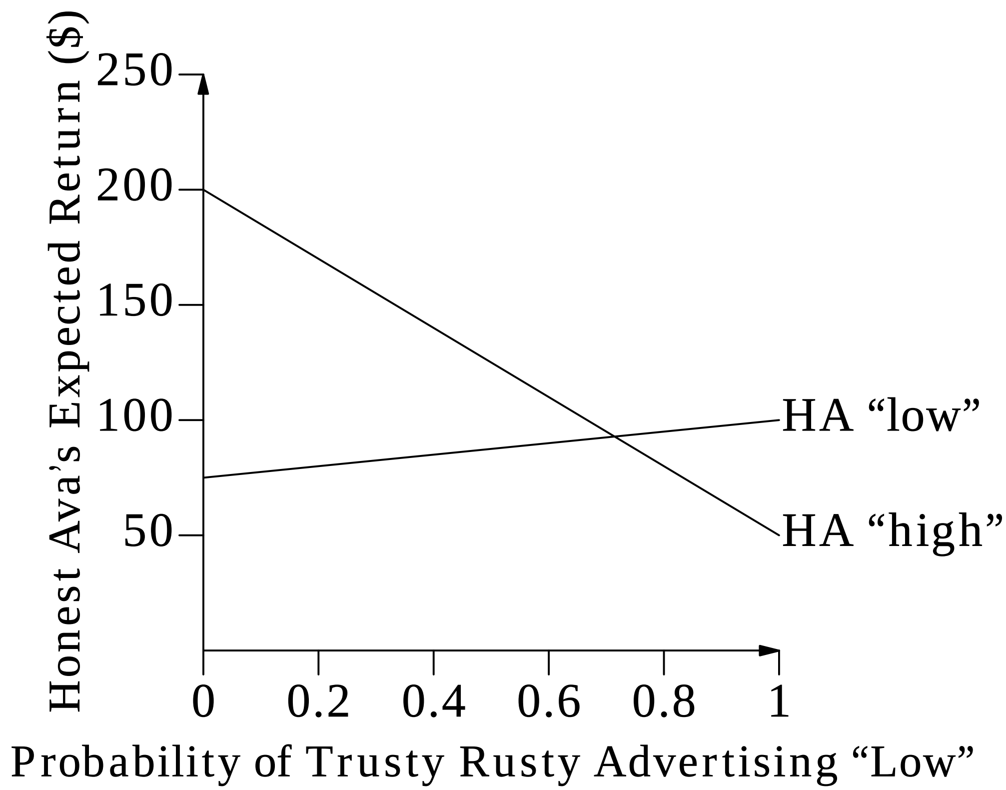
Honest Ava's Expected Return ($)
250
200
150
100
50
HA "low"
HA "high"
0
0.2
0.4
0.6
0.8
1
Probability of Trusty Rusty Advertising "Low"

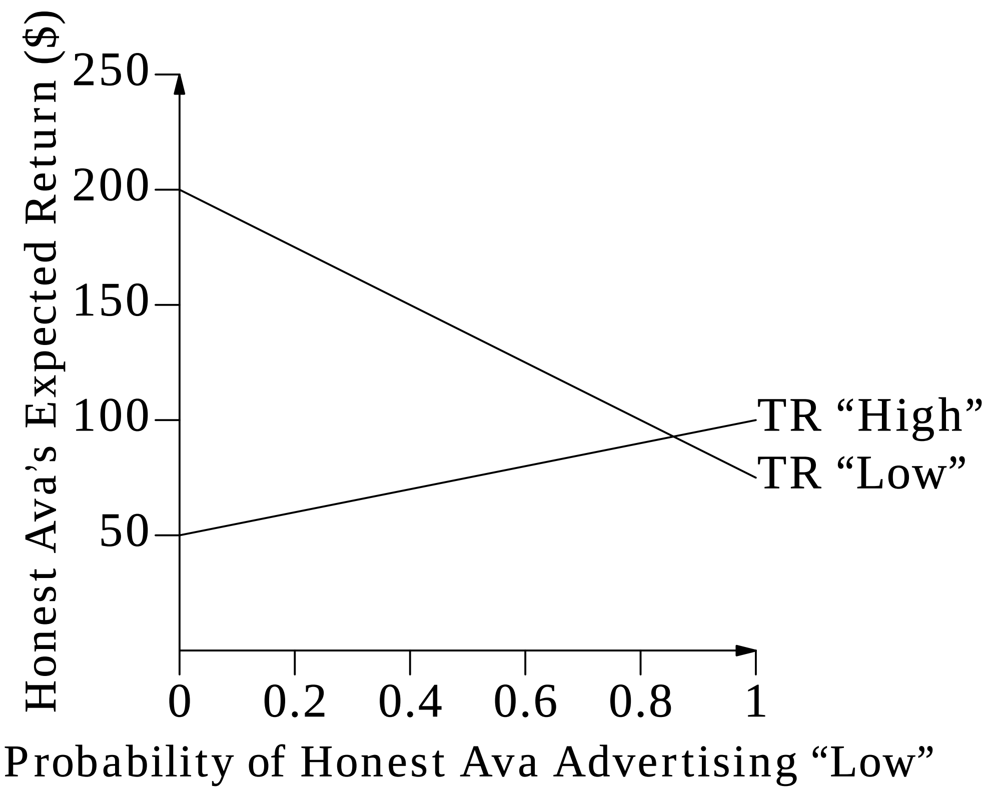
Honest Ava's Expected Return ($)
250
200
150
100
50
0
0.2
0.4
0.6
0.8
1
TR "High"
TR "Low"
Probability of Honest Ava Advertising "Low"

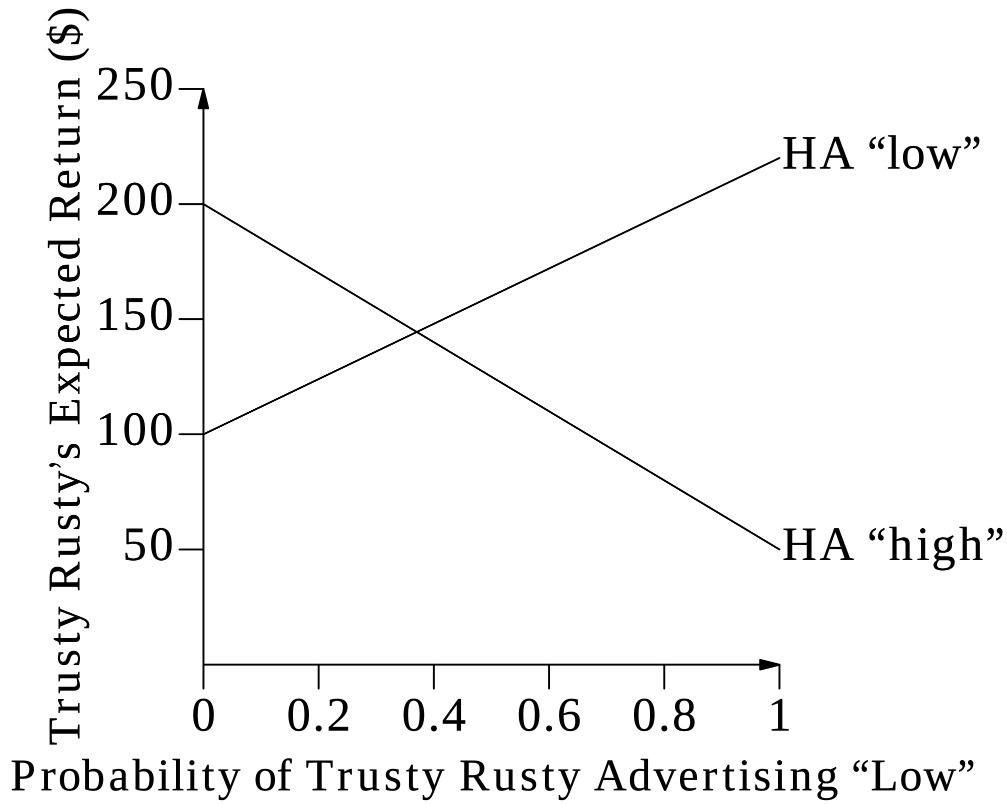
Trusty Rusty's Expected Return ($)
250
200
150
100
50
HA "low"
HA "high"
0
0.2
0.4
0.6
0.8
1
Probability of Trusty Rusty Advertising "Low"

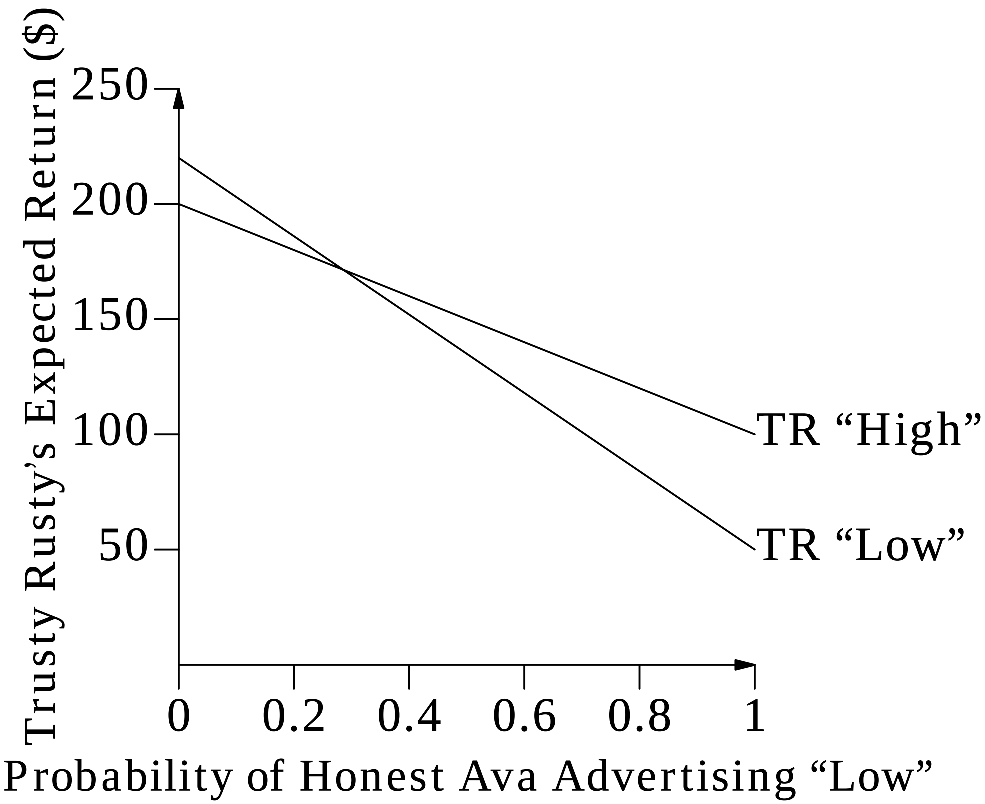
Trusty Rusty's Expected Return ($)
250
200
150
100
50
0
0.2
0.4
0.6
0.8
1
TR "High"
TR "Low"
Probability of Honest Ava Advertising "Low"

6.3 Choose the Right Mix

In the tennis example, if one player is not pursuing his equilibrium mix, then the other player can *exploit* this to his advantage. The receiver, Rod, could do better than a success rate of 48% *if* the server, Stefan, used any mix of strategy other than the equilibrium mix of 0.4 forehands and 0.6 backhands.

In general, if Rod knows Stefan's patterns and foibles, then he can react accordingly. Beware the hustling server, who uses poor strategies in unimportant matches to deceive the receiver when it matters: once the receiver deviates from her equilibrium mixture to take advantage of the server's "perceived" deviation, the receiver can be exploited by the server — a possible set up. Only by playing one's equilibrium mix is this danger avoided.

Each serve must be unpredictable: the nature of the randomness matters, lest the receiver take advantage of any patterns.

6.4 Why Not Rely on the Other’s Randomisation?

The reason why you should use your best mix — even if in equilibrium you are indifferent between moving to your forehand or your backhand as receiver — is to keep the other player using hers.

6.5 How Your Best Mix Changes as Your Skills Change

What if Rod's backhand return improves so that his rate of successful returns on that side increases from 60% to 65%? From the revised chart we see that Rod's best mix rises from 0.3 to 0.333 of moving towards forehand, and the overall probability of successful returns goes up from 48% to 50%.

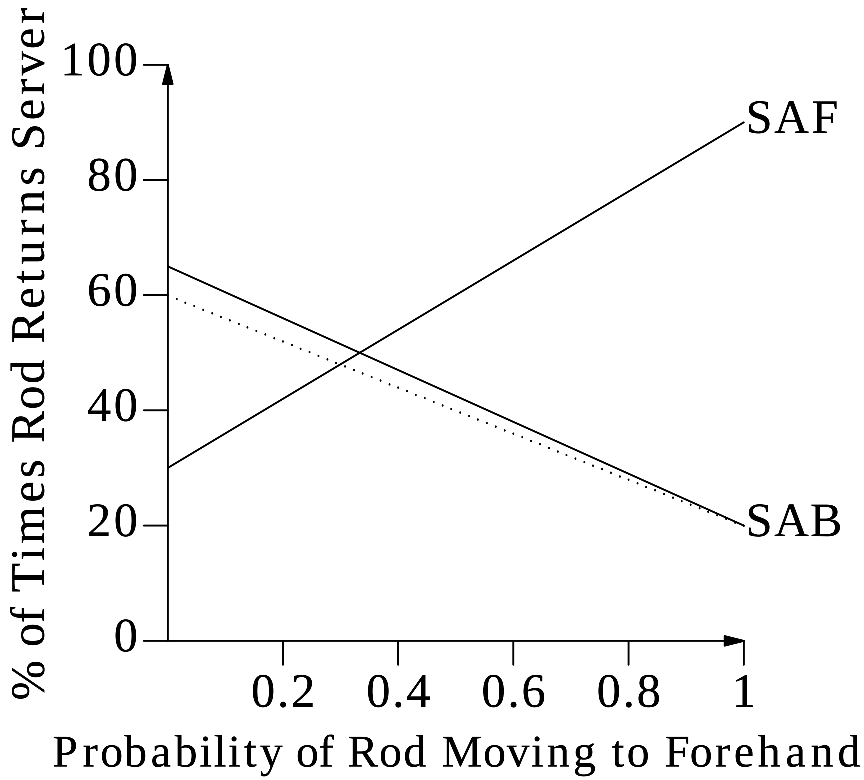

Effect of changed payoffs

The improved backhand is used less often, not more. Because of the interaction of the two players' strategies. When Rod is better at returning backhands, Stefan goes to the forehand more often (0.43 instead of 0.40). In response, Rod moves to his forehand more often, too. *A better backhand unlocks the power of your forehand.*

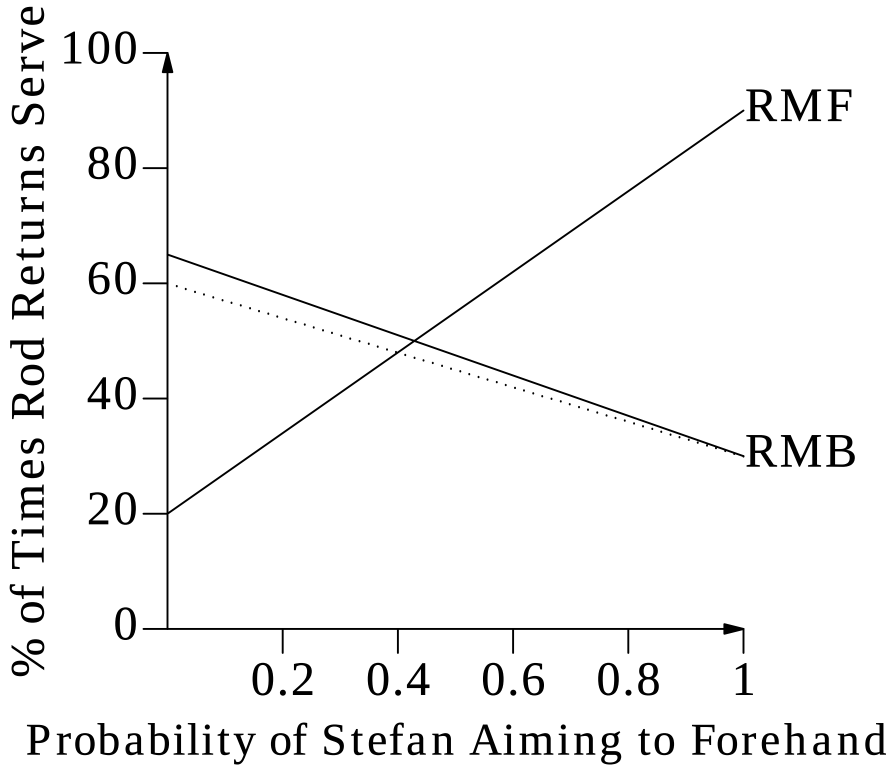

6.6 How to Act Randomly

To avoid putting order into your randomness, you need an objective or independent mechanism. Such as the second hand on your (analogue) watch: to act one way 40% of the time, do so if the second hand is between 1 and 24.

6.7 Unique Situations

Above is OK when we're in a repeating situation. What about unique, once-off situations?

To surprise the other side, the best way is to surprise yourself: keep your options open as long as possible, and then at the last moment choose between them using an unpredictable method. The relative proportions of the device should be such that the if the other side discovered them, they wouldn't be able to turn the knowledge to their advantage. But that is just the best mix as calculated above.

Even when using your best mix, you won't always have a good outcome. In games against nature (decision analysis) this is stated as the distinction between *good decisions* and *good outcomes*. Prudent decisions will on average result in better outcomes.

6.8 How Vulnerable?

If you are playing your best mix, then it doesn’t matter if the other player discovers this fact so long as he does not find out in advance the particular course of action indicated by your random device in a particular instance. The equilibrium strategy is chosen to avoid being exploited, so he can do nothing to take advantage of his knowledge.

But if you’re doing something other than your best mix, then *secrecy* is vital.

If the other side acquired this knowledge, they could use it against you. By the same token, you can gain by *misleading* the other side about your plans, especially in a non-zero-sum game.

Play the percentages

When playing mixed or random strategies, you can’t fool the opposition every time or on any one particular occasion. The best you can hope for is to keep them guessing and fool them some of the time.

e.g. When you know that the person you’re communicating with has some interest to mislead you, it may be best to ignore any statements she makes rather than take them on face value or inferring that exactly the opposite must be the truth.

“*Actions speak louder than words.*”
The right proportions to mix one’s equilibrium play critically depend on one’s payoffs. Thus observing a player’s move gives you some information about the mixing being used and is valuable evidence to help you infer your rival’s payoffs.

This is similar to tree flipping in games against nature.

6.9 Catch as Catch Can

Why so few examples of business using randomised behaviour away from the sports stadium?

Control over outcomes may militate against the idea of leaving the outcome to chance. Especially when things go wrong: it's not that mixing will always work, but rather that it avoids the dangers of the predictable and humdrum.

Companies using price discount coupons — similar to Della and Jim's coordination problem in the Battle of the Sexes.

Airlines and discount/stand-by tickets. If last-minute ticket availability were more predictable, then there would be a much greater possibility of exploiting the system, and the airlines would lose more of their otherwise regular paying passengers.

Most widespread use: to motivate compliance at lower monitoring cost — tax audits, drug testing, parking meters, etc. Explains why the punishment shouldn't necessarily fit the crime.

Appropriate incentives

If a parking meter costs \$1 per hour, then a fine of \$25 will keep you honest on average if you believe the probability of a fine is 1 in 25 or higher. (Risk neutral.) Which results in lower administrative costs and a better bottom line.

- No enforcement would result in misuse of scarce parking places;
- 100% enforcement would be too expensive.
- But the authorities don't want a completely random enforcement strategy: the expected fine should be high enough to induce compliance.

Other activities (random drug testing, tax audits) also require a sufficiently high expected penalty.

Those hoping to defeat enforcement can use random strategies to their benefit: they can hide the true crime amongst many false alarms or red herrings, so that the enforcer's resources are spread too thin to be effective.

6.10 Algebraïc Derivation of Optimal Mix

Consider a generalised payoff matrix:

Trusty

		Trusty: Low	Trusty: High
Honest	Ava: Low	A	C
	Ava: High	D	B

TABLE 2. The payoff matrix (Honest Ava's Payoffs)

With $D < C < A < B$.

Trusty chooses a probability P_R of playing Low so that Honest is indifferent between Low and High. That is:

$$P_R \times A + (1 - P_R) \times C = P_R \times D + (1 - P_R) \times B,$$

which implies

$$\frac{P_R}{1 - P_R} = \frac{B - C}{A - D}.$$

For Honest Ava, A = 100, B = 200, C = 75, D = 50, so

$$\frac{P_R}{1 - P_R} = \frac{200 - 75}{100 - 50} = \frac{125}{50} = \frac{5}{2},$$

which gives us Trusty's probability of playing Low: $P_R = \frac{5}{7}$.

Ava's mix can similarly be calculated as $P_A = \frac{2}{7}$.

Note that we derived P_R and P_A by looking for an equilibrium in which neither player had any incentive to alter their mix, given that the other was playing their best mix.

The reader is left to complete this exercise for Rod & Stefan. \&