G13GAM — Game Theory (Solutions)

Note: These solutions should (of course) be read in conjunction with the questions. They form an integral part of the module, and should give you a lot of additional information about the content of the module, specifically about recommended ways of attacking problems — including the exam!

1 (*a*) Successively we have $e = -\min(0,-10) = 10$, $f = 19$, $g = -\min(-12,0,12) = 12$, $b = -\min(10,19,12) = -10$, $h = 50$, $i = 74$, $j = 60$, $c = -\min(50,74,60) = -50$, $k = -30$, $l = 0$, $d = 30$, $a = -\min(-10,-50,30) = 50$.

(*b*) Start at a with bounds $-\infty$ and $+\infty$.
Visit b with bounds $-\infty$ and $+\infty$.
Visit e with bounds $-\infty$ and $+\infty$.
$e1$ returns 0, $e2$ returns -10, so e returns 10 to b.
$\alpha < -10 < \beta$, so reset b's bounds to -10 and $+\infty$.
Visit f with bounds $-\infty$ and 10.
$f1$ returns -19, so f returns 19 to b.
$-19 < -10$, so visit g with bounds $-\infty$ and 10.
$g1$ returns -12, and $12 > 10$, so return 12 to b; $g2$, $g3$ not visited.
b returns -10 to a; so reset a's bounds to 10 and $+\infty$.
Visit c with bounds $-\infty$ and -10.
Visit h with bounds 10 and $+\infty$.
$h1$ returns -50 to h, so
h returns 50 to c.
$-50 > -\infty$, so reset c's bounds to -50 and -10.
Visit i with bounds 10 and 50.
$i1$ returns -74, and $74 > 50$, so return 74 to c; $i2$ is not visited.
-74 is below c's floor, so visit j with bounds 10 and 50.
$j1$ returns -60, so $j2$ isn't visited and j returns 60 to c.
c returns -50 to a, so reset a's bounds to 50 and $+\infty$.
Visit d with bounds $-\infty$ and -50.
Visit k with bounds 50 and $+\infty$.
$k1$, $k2$, $k3$ return 60, 30, 50, so k returns -30 to d.
$30 > -50$, so d returns 30 to a; l is not visited.
a returns 50 as its value.

(*c*) Note that the $\alpha-\beta$ bounds will always be 40 at the top and third levels of the tree, -40 at second and bottom. We visit b, e, $e1$ and $e2$, evaluating e as 10, as before. Since $-10 > -40$, we get a β-cutoff at b, and f and g are not visited; the value of b is at least -10, so is not interesting to a. The returned value of -10 from b to a negates to 10, which is below a's floor, so we visit c, h and $h1$, evaluating h as 50, as before. Since -50 is below c's floor, we visit i and $i1$; the returned value of -74 causes a cutoff at i, the value of i is at least 74 and will not be interesting to c. The returned value of -60 when we visit $j1$ from j similarly causes a cutoff, so c returns -50, meaning that (since this is below the floor) the value at c is at best -50. This value negates to 50 at a, which causes a cutoff, and d is not visited. The returned value from a is 50, meaning 'at least 50', as it is above the ceiling.

OK, now you understand all this, here is your follow-up question. Re-arrange the search order so as to minimise the number of nodes that need to be searched; that is, how many nodes absolutely must be searched in order to establish the value of the root?

2 Cases (a), (c) and (e) are where the result is correct, and no re-search will be necessary. Cases (b) and (d) are where you've been caught out, and the result γ is related to the wrong levels. In case (a), as $\gamma < \alpha$, the value returned is an upper bound [which probably resulted from a β-cutoff, though with the wrong β]. No action needs to be taken.

In case (b), the value is again an upper bound, but because this falls between the true α and the 'wrong' α', we don't know how γ relates to α. So we need to do a re-search with bounds α and γ to find an exact result [if the second search comes in between α and γ] or a useful upper bound [if it comes in below α. With luck, this will be fast, as there should be a smaller window and as many positions will have useful values in the transposition table. If it does come in above α, then we can proceed as in (c), if below then as in (a).

In case (c), γ is a true return from the search. We can raise α [and α'] to γ for the remaining moves at this node, and if this remains the best move, we can safely return γ as the value of this node. No other action is necessary.

In case (d), there is a β-cutoff, but with the wrong value of β. In other words, our child's value came back below its floor, presumably as a result of β-cutoffs in the grand-children, but using β'. We know the true value is at least γ, but not how it relates to β, so a re-search is necessary with bounds γ and β. This will return with either a true value [below β], and we can then proceed as in (c), or a lower bound [above β], and we can proceed as in (e).

Finally, in case (e), there is a β-cutoff, and although this is using the wrong β, as it was a lower bound we can still rely on the result. We can terminate the analysis and return γ.

This process is pretty much what an aspiration search or a minimum-window search does. If we're 'lucky', then we get good results much faster, because of the small window. If we're caught out, then we have to try again, and it's a matter of pragmatism and experiment whether we gain or lose.

3 There are many possible trees, including the one where all static evaluations are zero. The one shown is constructed as follows: we write down 0 at the root, then for each node which is the left-most child, we write down the negative of the value of its parent, and for each other node we write down a larger integer than that. For the particular tree, I simply wrote the next unused positive integer.

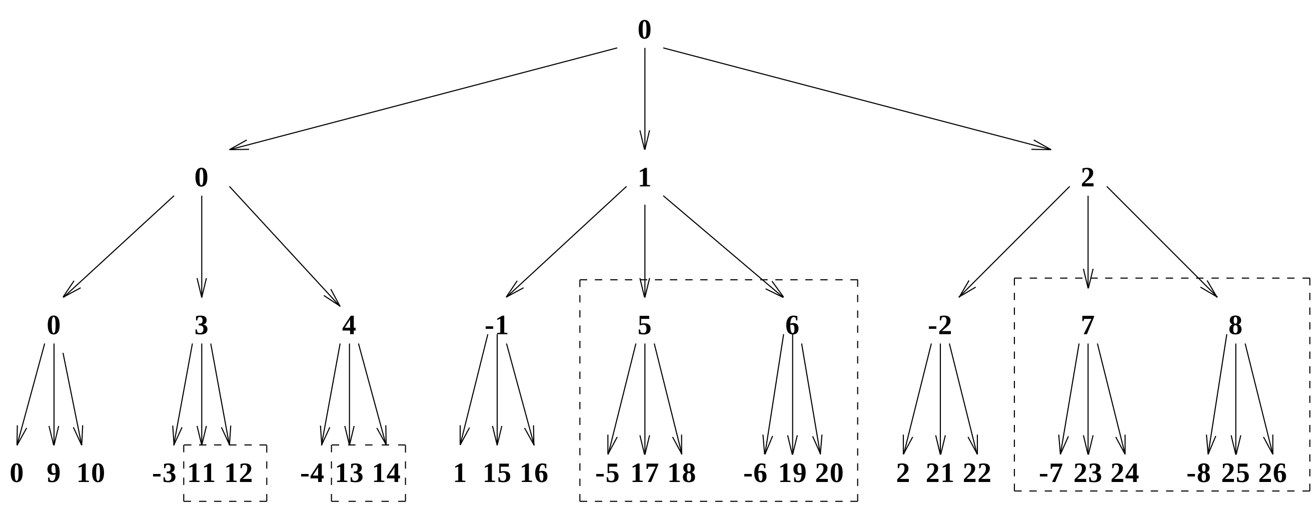

The nodes inside the dotted boxes are pruned, assuming the obvious left to right order of evaluation.

Note that because the $\alpha - \beta$ bounds are $\pm \frac{1}{2}$ and all the evaluations are integers, any node that is established as having a positive value will then have its analysis pruned. For example, at the node labelled 1 in the diagram, the value -1 of its left-most child establishes that its value is at least 1 and so we don't need to look at nodes 5 and 6; we already know that this value is going to be at most -1 when returned to the root of the tree, and therefore the corresponding move will be worse than that to the left branch.

Note further that if a different search order is used, then the pruning will be different. For example, again in the above diagram, if we had searched node 5 before the node -1, then the first value, -5 would have established that node 5 was worth at least 5, causing a cut-off and saving the search of nodes 17 and 18, and this would have returned -5 to node 1, below the floor, and we would then still have needed to look at node -1. If, perchance, the very leftmost node at the bottom had been -1 instead of 0, then the value backed up to the root would already have been 1, a β-cutoff, and we would not have needed to search nodes 1 and 2 at all.

4 [All such questions are somewhat bogus! Real games would be much too complicated to analyse, so we have to make some fairly drastic assumptions. In exam questions of this sort, the actual numbers are not very interesting — you would get most, if not all, of the marks by showing that you understand how pruning and move ordering affect which nodes have to be looked at — except as order-of-magnitude demonstrations of how important these techniques are. We also have to ignore effects such as the use of transposition tables (for faster results when the same position arises from different move sequences — this is a very important speed-up in games like Chess or Go where this often happens, less important in Othello, though still useful).]

(*a*) With no pruning, every continuation to 10-ply [5 moves by each side] must be examined, so static analysis is required on (about) $40\times20\times40\times20\times40\times20\times40\times20\times40\times20 \approx 3.28\times10^{14}$ nodes, taking roughly 3×10^{8} seconds, 10 years. [This is bogus, because after any White move other than the best, it is possible that White is now losing, and the move counts, based on the assumptions given, are therefore the other way round. But in real life, it is not too unrealistic — the number of available moves to each side mostly varies relatively little.]

(*b*) We come back to this after (*c*)!

(*c*) With pruning and move-ordering, all *Black* moves must still be looked at, for Black is losing, but for *White*, the first choice of move must be analysed completely, but other moves will turn out not to be as good. For each not-so-good move, Black has at least one reply which ‘refutes’ it, and causes a β-cutoff. With perfect ordering, this move will be tried first, so the search will be faster by a factor 20 from this cause alone; there will be similar improvements further down the tree, so we can ignore the time taken by analysing the not-so-good White moves, and the time taken is roughly that for $1\times20\times1\times20\times1\times20\times1\times20\times1\times20 = 3200000$ static evaluations, or 3.2 seconds; 10^{8} times faster than case (*a*)! Now back to case (*b*).

(*b*) With pruning but no move ordering, the news is not quite so good. We won’t normally find refutations quite so quickly — instead of finding one at the first move, if there is only one it will be found on average halfway through the analysis, but if there are several, one should be found quite quickly. Since this will apply at every level in the tree, even a factor of 2 speed-up will be a factor of 8 after three moves, and we can expect to do better than this. So let us assume again that the time taken to analyse moves that can be refuted can be ignored. Then the effect is that White moves will be pruned away unless they are the ‘best so far’ in each position. Assuming all the values to be different, and that the move ordering is random, the probability of the n-th move tried being the best so far is $1/n$, so the expected number of best-so-far moves for White is $1+\frac{1}{2}+\frac{1}{3}+\frac{1}{4}+\ldots+\frac{1}{40} \approx 4.3$, so the time taken will be $\approx 4.3^{5}\times20^{5}/10^{6} \approx 4500$ seconds, or 1¼ hours. [Note, much better than (*a*), much worse than (*c*).] [*Any* reasonable argument accepted here. For example, if we are merely looking to see whether White wins or loses, and if we assume that roughly half of White’s moves are winning, then White’s branching factor is 2 (instead of 4.3), and the time is around 100 seconds. There are lots of possible models.]

Rest of question — bookwork. In real life, the way to find out how much difference these ideas make is to try them in a real program.

5 (*a*) We have to search $12\times11\times\ldots\times1 = 12! \approx 479$ million sequences of moves. each terminating in a full board which takes 100 microseconds to evaluate. So the time taken is $\approx$ 47 900 seconds, somewhat over 13 hours. [This result, like the others here, ignores the possibility of transpositions—compare the comment in the answer to Q4.]

(*b*) In this case, we look at $1{\times}11{\times}1{\times}9\ldots{\times}1 = 10395$ sequences to establish that the computer's first move in each position wins. If the computer is satisfied to win, that terminates the analysis, in therefore just over a second. [The following argument was not expected, but would have earned bonus marks:] If it wants to be sure that it has found the very best move, then it has to look at its other moves, but each of these will be β-cutoff by the first reply, so there will be an extra $11{\times}1{\times}10{\times}1{\times}8{\times}\ldots = 42240$ such lines, so the time will become around 5.3 seconds.

(*c*) Here we have to look at all our moves, but the first move by the other side which is looked at will suffice to establish the loss. So we analyse $12{\times}1{\times}10{\times}1{\times}\ldots = 46080$ lines, taking around 4.6 seconds. Extending this, as in *(b)*, to establish exact values left as an exercise.

(*d*) Any reasonable argument accepted. The simplest says that on average, we have to look at two lines in a won position before we find one that wins [compare the average number of coin tosses that we need before the coin shows heads], so we must look at $2{\times}11{\times}2{\times}9{\times}2\ldots = 665280$ lines, taking just over a minute. A more sophisticated argument would establish and solve recurrence relationships for W_n and L_n, the number of lines that must be looked at in won and lost positions, respectively, when there are n moves to go. In a won position, we [on average] look at about one losing line before we find a winning line, so $W_n = L_{n-1} + W_{n-1}$. ['About one' breaks down when n is small; *e.g.* when $n = 2$, there is, by assumption, one winning move and one losing, and it's an evens chance whether you look at the losing line first, so you look at only half a losing line. Details left as an exercise.] In a lost position, we must look at all n moves, and these all are won for the other side, so $L_n = n{\times}W_{n-1}$. Now work back from $W_1 = L_1 = 1$ to find W_{12}. [Two minutes with pencil and paper!] Note that with perfect ordering, as in *(b)* and *(c)*, the equation for W_n becomes [simply] $W_n = L_{n-1}$, from which the numerical expressions obtained previously are easily found.

The important point in all of this is that to show you are winning, you need to find only one winning move; whereas a loss can be established only by looking at everything. Again, see the note at the top of the answer to Q4; these questions are *amazingly* easy, and a gift of lots of marks, if you take a pragmatic ['engineering' or 'common sense'] view of games, and amazingly difficult—indeed, ongoing research topics—if you expect exact answers.

6 (*a*) An m_{n+1} has a set of Left options and a set of Right options. Each option must be one of the m_n games of length no more than n; each such game may be either in or out of each set of options. So there are 2^{m_n} ways of constructing the left set and the same number for the right set, and 2^{2m_n} ways of combining them. Clearly $m_0 = 1$, so $m_1 = 2^2 = 4$ and $m_2 = 2^8 = 256$ [and $m_3 = 2^{512}$, is getting rather large]. So the number of games of length exactly 2 is $m_2 - m_1 = 252$.

In fact, they must be made up from combinations of the four games of length 1 or less, *viz.* −1, 0, * and 1. As Left options, −1 is dominated by the other three, and 1 dominates all the others, so only the combinations −1, 0, *, both 0 and *, 1 and empty do not contain dominated options. There are six such combinations, similarly for Right, so 36 games [of length $\leqslant 2$] without dominated options; each of the 252 games of length 2 must be equal to one of these 36.

If the Left options do not include either 0 or 1, then Left to play loses. There are four such sets of options, *viz.* the set containing just −1, that containing just *, that containing both and that containing neither. Similarly for Right, so there are 16 games of the 256 that are lost for the player to move, and so have value 0. One of these is 0 itself, the other 15 have length 2.

(*b*) Induction. There are, by inductive assumption, $1 + 2 + 4 + \ldots + 2^{n-1} = 2^n - 1$ numbers of length less than n, comprising the 2^{n-1} numbers of length $n-1$ separated by the earlier numbers. So the numbers of length n are produced one each side of those of length $n-1$, and there are 2^n of them. Clearly 1 has length 1, 2 has length 2, $1\frac{1}{2}$ has length 3, $1\frac{1}{4}$ has length 4, $1\frac{1}{8}$ has length 5, …, $1\frac{1}{256}$

has length 10.

7 No detailed solution given. Note when you draw the tree that Cross cannot initially complete a line, so, by the ordering rules given in the question, the first move to consider is where Cross plays in the bottom right corner; similarly Nought must then reply in the right middle, and the game is soon drawn. Your tree should now move to other Nought replies to Cross's first move; to all of these, Cross's move to right middle is a 'killer', causing instant pruning. Similarly, when you consider Cross alternatives at the first move, Nought has a killer in the bottom right.

8 We first show that $+_G > 0$. Left to play plays to 0, so wins; Right to play must play to $\{0 \mid -G\}$, in which Left can reply to 0 and win. So Left wins anyway, and $+_G > 0$. Note that phrases such as '$G > 0$ if Left is to play' are meaningless; G either is or is not greater than 0, and if all you know is that Left to play wins, then G could be either positive or fuzzy.

Next, we show that $+_G \leqslant \frac{1}{2}$. This means that $+_G - \frac{1}{2} \leqslant 0$, *i.e.* that $+_G - \frac{1}{2}$ is a Right win if Left plays first. [If Right plays first, either player could win, consistently with 'less than or equal', so we do not need to consider this case; similarly, '$\geqslant 0$' means that Left wins if Right plays first.] Left playing first could play to 0 in $+_G$; but this leaves the game $-\frac{1}{2}$, which Right wins. Or Left could play to -1 in $-\frac{1}{2} = \{-1 \mid 0\}$; but then Right replies in $+_G$ to $\{0 \mid -G\} - 1$, in which Left's only move is to $0 - 1 < 0$, so Right wins. So $0 < +_G \leqslant \frac{1}{2}$, as required. Note that we have not had to contemplate what might happen if we actually had to move in G, luckily. For $+_G = \frac{1}{2}$, see below.

If G is a negative number, then $-G$ is a positive *number*, say A, and $\{0 \mid A\}$ is the simplest number between 0 and A, so is a positive number B, say. Hence $+_G = \{0 \mid B\}$ is the simplest number between 0 and B, so is itself a positive number. [Note that $B \leqslant 1$, so $+_G \leqslant ½$, confirming the above in this special case. It is important that A is a *number*; if, for example, G is $\downarrow$, then you will find it difficult to slip in a number, simple or otherwise, between 0 and $\uparrow$, and the Simplicity Theorem doesn't apply.]

In particular, if $G = -1$, then $A = 1$, $B = ½$, and $+_{-1} = ¼$; while if $G = -\pi$, then $A = \pi$, $B = 1$ and $+_{-\pi} = ½$.

When G is a positive number, we must show that Right wins $+_G - \uparrow = +_G + \downarrow$. The basic idea is that Right has stronger threats in $+_G$ than Left has in $\downarrow$; see also Q13. Left to play can kill $+_G$, leaving Right winning $\downarrow$; or Left can play to * in $\downarrow$, when Right replies to $\{0 \mid -G\}$ in $+_G$. Now Left can either kill $\{0 \mid -G\}$, leaving Right to move in *, or kill *, leaving Right to play to $-G < 0$; so in either case Right wins. Alternatively, Right to play first plays to $\{0 \mid -G\}$ in $+_G$. Now Left can either kill $\{0 \mid -G\}$, leaving Right winning $\downarrow$, or can play to * in $\downarrow$, when Right replies to $-G$* and wins. So again Right wins, and $+_G < \uparrow$, as required.

[These explanations in words always seem a lot more complicated than they really are. If the above confuses you, try with the example of $+_2$ and $\downarrow$ from Domineering in front of you. $+_2$ is a Left win; but if Right gets first move, and plays to -1 ± 1 in it, the threat of moving to -2 is so strong that Left cannot take time out to mess around with $\downarrow$, but must immediately attend to the threat. The hottest threats are the most urgent.]

9 Clearly, Left to play can only block the position to 0, and Right to play has only one move, to *LRL.RR*, with an obvious notation. In this position, Left can play only to *LR.LRR* [which we must show to be 0], and Right to *LRLR.R* [which we must show to be $-\frac{1}{4}$].

(*a*) In *LR.LRR*, Left to play must play to *.RLLRR*, leaving Right [only] a free move, while Right to play must play to *LRRL.R*, which is clearly a first-player win in three more moves. So *LR.LRR* is a first-player loss, *i.e.* is 0.

(*b*) In *LRLR.R*, Right to play must block to 0, while Left must play to $LR.RLR = \{ .RLRLR \mid LRR.LR \}$. Now *LRR.LR* is clearly a first-player loss in at most two moves, so $LRR.LR = 0$, while $.RLRLR = \{ \mid R.LRLR \}$. Finally, *R.LRLR* is clearly a first-player loss, as Left cannot move and Right can only play to *RRL.LR* in which Left to play can block to 0. So: $.RLRLR = \{ \mid 0 \} = -1$; and $LR.RLR = \{ -1 \mid 0 \} = -\frac{1}{2}$; and $LRLR.R = \{ -\frac{1}{2} \mid 0 \} = -\frac{1}{4}$, and the result follows.

10 (*a*) The game n, $n-1$, $n-2$, …, 2, 1 shows that the game can last for n moves, for any n, and so its length is unbounded.

Suppose that the first move is n, and at any later stage after this, suppose that the product of the coins other than n is m. Let the GCD of m and n be d. Then m/d and n/d are co-prime, so all integers greater than $(m-d)(n-d)/d^2$ can be written as $(pm+qn)/d$; in other words, all but a finite number of multiples of d are now 'killed'; within at most $(m-d)(n-d)/d^2$ steps, a coin must be nominated that does not have d as a factor, and thus the GCD of n with the product of the remaining coins must reduce. As this can happen no more than n times, the game must terminate.

(*b*) After playing 2, all even coins are 'killed'; the response 3 kills 3, 5 = 3+2, 7 = 3+2+2, …, *i.e.* all odd numbers other than 1, thus forcing the play of 1. Similarly, 2 is a winning response to 3. So the first play of either 2 or 3 loses.

Playing 4 and 5 kills 8 = 4+4, 9 = 4+5, 10 = 5+5, 12 = 4+4+4, 13 = 4+4+5, 14 = 4+5+5 and 15 = 5+5+5, and hence all greater coins, which are 4 more than a coin already killed. [Or Sylvester: 4 and 5 are co-prime, so all numbers $\geqslant 3\times4 = 12$ are killed.] So the only remaining coins are 1, 2, 3, 6, 7 and 11. 1, 2 and 3 are always bad, as just seen. If Left plays 6, this kills 11 = 6+5, so the reply 7 leaves only 1, 2, 3 for Left and Right wins; similarly, if Left plays 7, the reply 6 wins for Right. So Left should play 11, forcing Right to choose between 6 and 7, Left takes the other, and only 1, 2, 3 are left for Right. Left wins. [Actually, 6 is a winning reply to 4, and an initial 5 wins for Left; but these are quite deep results.]

11 No solution given—all parts are either bookwork or routine.

12 Note that an all-small game is either 0 [neither player can move] or else is of form $G = \{a, b, c, \ldots \mid d, e, f \ldots\}$ where all the options are themselves all-small and neither list is empty [both players can move]. We have to show that $G-1 < 0$ (that is, is a Right win). Informally, what happens is that Right wins by playing in G whenever possible; all-smallness means that this is possible until G is reduced to zero, when only the -1 is left.

More formally, assume the result to be true for all options of G. Note that in the limiting case when $G = 0$, there are no such options, so there is no need for an inductive base, beyond observing that the induction must terminate because we're always moving to simpler games. Left to play must play to some option for which the inductive hypothesis holds, so Right wins; or else Left cannot play at all [when $G = 0$], in which case Right has already won. With Right to play, if $G = 0$ then $G-1 = -1$, in which Right can move to 0 and win; if $G \neq 0$, then because G is all-small it has a Right option to which the inductive hypothesis applies, so again Right wins. So in all cases Right wins, *QED*.

For the next bit, you must exercise care. It is tempting to say that if $G \geqslant x$, then $G+G+\ldots+G \geqslant nx$, where n is the number of copies of G, that $G+G+\ldots+G$ is all-small, and that for sufficiently large n, $nx \geqslant 1$, which contradicts the previous result. All of this is true, and it shows that G is not $\geqslant x$. But that is not equivalent to $G < x$! G is not necessarily a number; it may be fuzzy with x. However, we can extend the previous idea. The only difference is that Left now may have extra moves in $G-x$, because there are moves in x. However, any such move, as x is a number, is to $G-y$, where $y > x$, and in particular no move by Left can run Right out of moves in $-x$—all it can do is give Right more spare moves. So, whatever Left does, Right simply plays in the all-small game until it vanishes, when Right is left with a negative number and a win. Formalising this is very similar to the above.

Note that $\uparrow$ is an all-small game which Left wins, so is strictly positive. [This result shows that it is, however, smaller than every positive number.] Q8 shows that $+_2$ is a positive game which is less than $\uparrow$, and so is also less than every positive number; since $+_2$ has a Right option which has a Right option of -2 [imagine Right playing in the two middle columns], which is non-zero but has no Left option, $+_2$ is not all-small. All impartial games are all-small [if either player can move, both can, as they both have identical moves], so all impartial games are infinitesimal.

13 We have to show that $+_y + -_x > 0$, where $-_x$ ('miny-x') is $- +_x = \{\{x \mid 0\} \mid 0\}$. The technique is for Left to use the threat of playing to x to overcome Right's threat of playing to y. So Left to play plays to $\{x \mid 0\}$ in $-_x$; Right can either play to 0 in this game, leaving Left winning $+_y$ (as in Q8), or to $\{0 \mid -y\}$ in $+_y$. In this latter case, Left plays to $x+\{0 \mid -y\}$, and whether Right replies to $x-y$ or to $p+\{0 \mid -y\}$ (where p is some Right option of x, and necessarily a positive number), in which Left can reply to p, Left is winning. If, on the other hand, Right plays first, then playing to 0 in $-_x$ leaves Left winning $+_y$, and the only other move is to $\{0 \mid -y\}$ in $+_y$. In this case, Left has a reply to $\{x \mid 0\}$ in $-_x$, and again, because $x > y$, the result will be a clear Left win after one more move on each side.

Informally, the same idea works when playing $+_y + n\cdot -_x$. Left can 'kill' all the $-_x$'s in turn, by using the threat to move to x in one of them. [Note that, following Q12, Q8 and the result in the first part, $+_x$ and $+_y$ are infinitesimal, so any positive number dominates any finite collection of them.] We can see this in action in Domineering—imagine a large collection of $+_2$s, or 2 by 4 rectangles, all won for Left, and one measly $\downarrow$ [the wedged-together *s, and note that $\downarrow$ is $-_0$]. Right wins by simply picking off all the rectangles in turn; the threat to move to -2 forces Left to keep 'fire-fighting', replying to 0 in each rectangle. Thus $+_x$ is not only less than $+_y$, though still positive, it is infinitesimal compared with it. Thus, game theory includes an infinite hierarchy of infinitesimals, since x can be any positive real number. Formalising this into a proof (by induction) is left to the interested reader.

14 $-A$ is the game obtained by interchanging the roles of Left and Right; $A+B$ is the disjunctive sum of A and B—the game wherein a move consists of moving in *either A or B*; $A > B$ means that Left (whether or not on move) wins the difference game $A-B = A+(-B)$; $A \parallel B$ means that whoever is to move wins the difference game.

(a) For example, $A = 2$, $B = \pm 1$.
(b) No such game; Left to play wins by playing first in B, replying in whichever game Right moves in.
(c) For example, $A = 1$, $B = \pm 2$.
(d) For example, $A = 2$, $B = -1 \pm 2$.
(e) For example, $A = 2 \pm 1$, $B = -1 \pm 2$.
[Many other solutions, possibly using games like *.]

Consider the game $H = G+n$, for some number n. First of all, we show that H cannot be fuzzy. Suppose it is; then Left to play wins, either by playing in n to some number $p < n$, or by playing in G to some left option L of G. But Right to play *also* wins H, by assumption, and so Right to play wins $G+p$ (which is better for him than $H = G+n$), and also wins $L+n$, which is better for him than $G-\varepsilon+n < H$. This is a contradiction. [Informally, playing in either G or n makes your position worse, so if Left wins with first move, Left must also win without.]

So one of $H > 0$, $H < 0$ and $H = 0$ holds. Clearly, if $H > 0$, then $H+q > 0$ for any positive number q, so there is some real number x such that $H > 0$ if $n > x$ and $H < 0$ if $n < x$. [Informally, we have defined a Dedekind section of the real line.] What happens if $n = x$? Well, Left to play can play, as before, either to $G+p$, where $p < x$, so $G+p < 0$, by the construction of x, or to $L+x < G-\varepsilon+x < 0$, as $x-\varepsilon$ is a number less than x; so Left to play loses. Similarly, Right to play loses, and $H = 0$, so $G = -x$ is a real number. [Note, things would get very murky if ε were, say, $+_2$, which is positive, but not a number; then $x-\varepsilon$ would not be a *number* less than x, so would not be in the left partition of x. Note also that things can get murky if G (and hence x) is infinite. If you are happy with infinite numbers, there is no problem!]

Each left option L of G is therefore less than $-x-\varepsilon$, and each right option R is greater than $-x+\varepsilon$. There must be a multiple of ε between these values, and so x must be a multiple of ε by the Simplicity Theorem.

15 Parts *(a) (d)* and *(e)* are bookwork. For *(b)*, suppose that $G = \{A, B, C, \ldots \mid D, E, F, \ldots\}$; then $-G = \{-D, -E, -F, \ldots \mid -A, -B, -C, \ldots\}$. If Left to play first plays to $A+G$, then Right has a reply to $A+(-A) = A-A = 0$ [by induction, since A is simpler than G]; or if Left plays to $G-D$ then Right has a reply to $D-D = 0$, and similarly for all Left moves, so Right wins. Similarly, if Right plays first, Left wins, and $G+(-G) = 0$. For *(c)*, if G is impartial, then $-G = G$, and we can use the result of *(b)*.

16 Explanation—bookwork. The given games are, respectively, 2, 2, 0, $-1*$, $-2\frac{1}{2} \pm 1\frac{1}{2}$, $-\frac{1}{2}$, 1 and $\frac{1}{2}$.

If p and q are not numbers, but $p < q$, then we cannot directly apply the simplicity theorem; we will have to do a case-by-case examination. Nevertheless, the conclusions of the simplicity theorem may hold; recall that the ‘proof by example’ in lectures depended on the result being intermediate between p and q but having its options outside that range. So $\{\uparrow \mid 1\} = \frac{1}{2}$, as $\uparrow < \frac{1}{2} < 1$, but the options [0 and 1] of $\frac{1}{2}$ do not satisfy those inequalities. In the second case, $\{0 \mid \uparrow\}$, there is no number between 0 and $\uparrow$, so the simplicity theorem cannot be applied at all; this is, in fact, the canonical form of this game. In $\{\pm 1 \mid 2\}$. there are numbers strictly between ± 1 and 2; the simplest such number is 1½. But this game is clearly in fact 0, as if Left plays to ± 1, then Right gets to move to -1 and wins, while if Right plays first it must be to 2, a Left win.

Exercise for the reader: What happens if you add 2 to both options, thus $\{2 \pm 1 \mid 4\}$? Is this 2? What if I add 1½? Is it then 1½? Can you generalise your findings?

17 (a) Part *(i)* is a matter of strategy. To show that $G \geqslant H$, we must show that $G-H \geqslant 0$, in other words that $G-H$ is a win for Left if Right moves first [we don’t care who wins if Left moves first—it will be Left if $G > H$ or Right if $G = H$]. But Right moving first must play to one of $D-H$, $E-H$, $F-H$, …, $G-B$, $G-C$, …, and in each of these Left has a reply to the zero game $D-D$, …, $B-B$, …. The only move by either side to which there is no symmetric response is Left’s move to A in G, and that cannot be played by Right. Part *(ii)* is bookwork.

Part *(iii)* is trivial if you spot the answer, otherwise you need to apply common sense. In the supposition, move A is the best move available. Does it make any difference if move A can’t be played? Well, it does if the move is good, and it doesn’t if the move is futile. Now it’s easy to construct examples where $G = H$. For example, the games $\{-2, -1 \mid 1\}$ and $\{-2 \mid 1\}$ are equal [both zero]. Summary—it can’t do you any harm to be given extra possible moves, and it may [but won’t necessarily] do you some real good [even if the extra moves are the best available].

(b) First part is bookwork. If p is not a number, then $p+*$ and $\{p \mid p\}$ will differ in general. One example is the case $p = \uparrow$; an even easier one to analyse is the case $p = \pm 1$ [for $\pm 1 + *$ is obviously a first-player win, while $\{\pm 1 \mid \pm 1\}$ is a first-player loss].

18 First part—bookwork. $G = 3 + 4\frac{1}{2} \pm 2\frac{1}{2} - 4 \pm 5 - 1 \pm 1 = 2\frac{1}{2} \pm 5 \pm 2\frac{1}{2} \pm 1$; so playing the ‘hottest first’ strategy, Left plays to $2\frac{1}{2} + 5 \pm 2\frac{1}{2} \pm 1$, Right replies to $7\frac{1}{2} - 2\frac{1}{2} \pm 1$, and Left reaches $5 + 1 = 6$, while if, on the other hand, Right plays first we get successively to $2\frac{1}{2} - 5 \pm 2\frac{1}{2} \pm 1$, $-2\frac{1}{2} + 2\frac{1}{2} \pm 1$, and -1, as required.

Note that $H = 2\frac{1}{2} \pm 3\frac{1}{2}$, so $G-H = 2\frac{1}{2} - 2\frac{1}{2} \pm 5 \pm 3\frac{1}{2} \pm 2\frac{1}{2} \pm 1$, in which Left to play reaches $5 - 3\frac{1}{2} + 2\frac{1}{2} - 1 = 3$ and Right to play reaches $-5 + 3\frac{1}{2} - 2\frac{1}{2} + 1 = -3$. So, $G-H$ is a first-player win, and in particular is not 0, so $G \neq H$. [Trivially!] If we take $K = -H = -2\frac{1}{2} \pm 3\frac{1}{2}$, then, as we have just seen, $G+K$ is a first-player win, while $H+K = 0$ is a second-player win. Since $\pm n + \pm n = 0$ for any n,

$G+G = 5 = H+H$.

19 First part: virtually bookwork – a number game is the disjunctive sum of numbers, stars and switches, so $2H$ is numbers plus an even number of stars [and $*+* = 0$] and pairs of switches [and $\pm n \pm n = 0$], so we are left only with numbers. In the given case,

$$H = \{3|6\}+\{3|1\}+\{8|1\}+\{-4|4\}+\{-4|-4\}+\{4|-4\} = 4+2\pm 1+4\tfrac{1}{2}\pm 3\tfrac{1}{2}+0-4+*\pm 4,$$

i.e. $H = 6\frac{1}{2}+*\pm 4\pm 3\frac{1}{2}\pm 1$ and $2H = 13$.

20 [Think of $G{:}H$ as H 'sitting on top of' G, so that when you play in G, H disappears, rather like Hackenbush, but playing in H is 'independent'.] So $1{:}-1 = \{0 \mid 1{:}0\} = \{0 \mid 1\} = \frac{1}{2}$. In $*{:}*n$, which is clearly impartial, either player can play to 0 [by playing in the left-hand $*$] or to $*{:}*k$ for any $k < n$; but, inductively, $*{:}*k = *(k+1)$, so either player can play to $*p$ for any $p \leqslant n$, and this is $*(n+1)$.

In $*{:}1$, Left can play to 0 or to $*{:}0 = *$; Right can play only to 0. So this game is $\{0, * \mid 0\}$; which you may recognise as our old friend $\uparrow+*$, or as the Hackenbush 'flower' with a green stem and a blue top, and making $\uparrow$ if you add a single blade of grass.

We have to show that $G{:}H-G{:}K \geqslant 0$, so that Left wins if Right plays first. If Right plays in G to some option A, then Left can play to $-A$ in $-G{:}K$, leaving a zero game with Right to play; similarly, if Right plays in $-G$. If Right plays in H or $-K$, then as $H-K \geqslant 0$, Left has a winning response, leaving a simpler game $G{:}H'-G{:}K'$ [where either one move has been played in each of H and $-K$ or two moves have been played in one of them] where it is Right to play and $H' \geqslant K'$, so Left wins, by induction. If $H{=}K$, then $H \geqslant K$ [Left wins $H-K$ if Right plays first], and so $G{:}H \geqslant G{:}K$; similarly, $G{:}H \leqslant G{:}K$; so whoever plays first in $G{:}H-G{:}K$, the other player wins, so $G{:}H-G{:}K = 0$ or $G{:}H = G{:}K$. [We have to be a little careful, you can't *just* use elementary algebra, as these aren't numbers but games, and some of them may well be fuzzy. Exercise: if $H > K$, does it follow that $G{:}H > G{:}K$?]

A Hackenbush tree is the corresponding bush on a stalk; so $T = *{:}B$. If $B = *n$, then our previous results show that $T = *{:}*n = *(n+1)$, which is the Tree Principle. [More generally, in any Hackenbush position, impartial or not, any sub-picture which is joined to the ground at only one node may be replaced by any other sub-picture of the same value; the Tree Principle is the special case where an impartial picture is replaced by its Grundy-equivalent snake.]

21 (*a*)
$$\uparrow+* = \{0|*\}+\{0|0\} = \{\uparrow+0, 0+* \mid \uparrow+0, *+*\} = \{\uparrow, * \mid \uparrow, 0\}.$$

Now, as a R move $\uparrow$ is dominated, as $\uparrow > 0$, and as a L move it is reversible, through its R option $*$ [as $* \leqslant \uparrow+*$], to the L option(s) of $*$, or 0. Thus,

$$\uparrow+* = \{0, * \mid 0\}.$$

(*b*) Let $H = \{\uparrow|\uparrow\}$, and let $G = \{\uparrow|\uparrow\}+\downarrow$. Then

$$G = \{H+*, \uparrow+\downarrow \mid H+0, \uparrow+\downarrow\} = \{H+*, 0 \mid H, 0\}.$$

Now, as a R option, H is dominated, as $H > 0$ [H is a game in which all options are left wins]; as a L option, 0 is dominated, as $H+* > 0$ [L can play to H, R must play either to H or to $\uparrow+*$ with L to play (which we have just seen is a left win)]. So $G = \{H+* \mid 0\}$. Now look for reversibilities. In $H+*$, R can move to $\uparrow+*$, so we need to check whether $\uparrow+* \leqslant G$, that is, whether $G-(\uparrow+*) \geqslant 0$, that is whether L wins $G-(\uparrow+*)$ if R plays first. Now, in $G-(\uparrow+*) = H+\downarrow+\downarrow+*$, R can play to $\uparrow+\downarrow+\downarrow+* = \downarrow+*$, or to $H+\downarrow+*$, or to $H+\downarrow+\downarrow$; in which L can reply to $*+* = 0$, or to $H > 0$, or to $H+\downarrow+* > 0$ respectively. So $H+*$ is reversible through $\uparrow+*$ to the L options of $\uparrow+*$, that is [as in part (*a*)] 0 and $*$. That is,

$$\{\uparrow|\uparrow\}+\downarrow = \{0, * \mid 0\}.$$

[With the hindsight of now spotting that $G = \uparrow+*$, we could derive this much more easily by showing that $G+\downarrow+* = 0$.]

(c) Although this game, call it K, is not as it stands impartial, it is composed of Nim heaps; so we suspect, using mex theory, it has value $*2$. Indeed, if Left plays first to $*3$, say, then Right has an option $*2$; and if we look at $*2-K$ with Left to play, then we see that Left must either play in $-K$ to a Nim heap of size different from 2 or in $*2$ to a Nim heap which is present as a Right option in $-K$, so Right wins in either case. That is, $*2-K \leqslant 0$, and the move to $*3$ is reversible through $*2$ to Left's options in $*2$—that is, to 0 and $*$. As these are already present in K, effectively $*3$ can be deleted. Similarly, so can the other larger stars, and K can be reduced to $*2$. [This is round the houses, to show how reversibility can be used instead of mex theory; but it is actually just as good to look at $K-*2$ and show that this is the zero game, so that $K = *2$.]

(d) $$2 + -_2 = \{1|\}+\{\{2|0\}|0\} = \{2+\{2|0\}, 1+-_2 \mid 2+0\} = \{\{4|2\}|2\} = \{3\pm1|2\}$$

[either of the last two forms acceptable, and in the previous form, one of the left options was dominated as $-_2$ is infinitesimal].

(e) $$\{\downarrow|\uparrow\} + \{4|0\} + \pm2 + 2* = 0 + (2\pm2) \pm2 + 2* = 4* = \{4|4\}.$$

[either of the last two forms acceptable].

22 The game is in five disjoint parts. Three of these are old friends: $*$, ±1 and $\uparrow$. The other two are the [upside-down] T-shaped region in the top middle, and the 3×2 rectangle at the bottom right. The T-region is clearly $-\frac{1}{2}$, either by noting that the extra square compared with an L-shaped region is no use to Left, or by observing that Left can only play to -1 and Right can play to 0. The rectangle is hot. Left can play to 2 or to ±1 [clearly dominated]; Right's moves are all equivalent, by symmetry, and are all moves to $-\frac{1}{2}$. So the rectangle is $\frac{3}{4}\pm\frac{5}{4}$, and the total position is $* + \uparrow \pm1 - \frac{1}{2} + \frac{3}{4}\pm\frac{5}{4} = \frac{1}{4}* \pm\frac{5}{4} \pm1 + \uparrow$.

(a) Left playing first should claim the two free moves in the bottom right, and Right should reply in the ±1 at top right, leaving the position $\frac{1}{4}* + \frac{5}{4} - 1 + \uparrow = \frac{1}{2}* + \uparrow > 0$, and Left wins. *(b)* Right playing first should equally play first in the bottom right, and Left should reply in ±1, leaving the position $\frac{1}{4}* - \frac{5}{4} + 1 + \uparrow = * + \uparrow$. This game is fuzzy, and it is Right to play, so Right wins by playing either move in the $\uparrow$ in the bottom middle. This will leave two $*$'s, a $+1$ and two $-\frac{1}{2}$'s, which all cancel nicely, leaving a zero game with Left to play.

23 Suppose without loss of generality that Left plays first. In general, after a few moves, there will be several types of region: type A—some empty fields bounded by two Left fields; type B—some empty fields bounded by two Right fields; type C—some empty fields bounded by one Left field and one Right; type D—some empty fields bounded by one Left field and the end of the farm; type E—some empty fields bounded by one Right field and the end of the farm.

In Col, we have to show a strategy by which Right wins, in other words a way of playing such that wherever Left plays, Right always has a reply. General idea—types B and E are clearly bad for Right, so we show how to avoid them. Left's initial move must create two Ds, of which one [but not both, as $n \geqslant 2$] can be empty. In a D, Right can play in the end field, creating a [possibly empty] C with Left to move. In a C, Left to move must create an A [which cannot be empty, as he can't play next to his own field] and a C, which could be empty. Equally, in a D, Left to play must create another D [perhaps empty] and an A [which cannot be empty]. So when it is Right's turn to play, after the first move, Left *must* have just created an A region, in which Right can play [anywhere] creating two [perhaps empty] C regions with Left now to move. So Right always has a move, and his strategy is to move in an end square on his first move, and thereafter always to move [anywhere] in the A region that Left has just created. As Right always has a move available by following this strategy. Left must lose [when all the C regions become too short to move in].

In Snort, we have to show that Left, playing first, can always win. Here, the simplest strategy uses symmetry. If n is odd, play first in the central field. If n is even, play first in one of the central fields, and note that this prevents Right from ever using the other central field. Now Left can always copy Right's moves symmetrically about the centre. By the strategy, the only difference between the two halves of the farm is that Left has occupied the centre, which reserves the adjacent fields for him and denies them to Right. So any move which Right plays has a symmetric counter, and Left cannot be run out of moves, so Left wins.

24 In Col, the LH region is reserved for Red, value -1. In the rest of the diagram, Left to play can play either in the triangle, reserving the remaining three regions for Red, or (clearly better) in the RH region, giving Red only two more moves. Right to play can play in the triangle, leaving Blue one move, or in the top or bottom region, allowing Blue to play in the RH region and allowing only one further Red move, or (clearly best) in the RH region, forcing Blue to the triangle. So the value of the game is

$$-1 + \{-2 \mid -1\} = -2\tfrac{1}{2}.$$

In Snort, Blue clearly plays to either the triangle or the RH region, reserving all four remaining regions. Red can play to the triangle, killing the top and bottom regions, and leaving one reserved region each, or (worse) to the RH region, allowing Blue to take the top and reserve the two remaining live regions. So the game is

$$\{4 \mid 0\} = 2 \pm 2.$$

25 There are four uncoloured regions; call them A, B, C and D from left to right. Then, in Col, the position is $\{A{:}, C{:}, D{:} \mid {:}A, {:}B, {:}C, {:}D\}$, where the notation, for example, $PQ{:}R$ means the position where the regions P and Q are coloured blue and R is coloured red. Now use brute force! (Note that any position where at least two of A, B, C and D have been coloured can be evaluated by inspection.)

$A{:} = \{AC{:}, AD{:} \mid A{:}B, A{:}C, A{:}D\} = \{-2, -1 \mid \tfrac12, 1, 0\} = \{-1 \mid 0\} = -\tfrac12;$
$C{:} = \{AC{:} \mid C{:}A, C{:}B, C{:}D\} = \{-2 \mid -1, 0, -\tfrac12\} = -1\tfrac12;$
$D{:} = \{AD{:} \mid D{:}A, D{:}B, D{:}C\} = \{-1 \mid -1, 1, *\} = -1*;$
${:}A = \{C{:}A, D{:}A \mid {:}AC, {:}AD\} = \{-1, -1 \mid 1, 1\} = 0;$
${:}B = \{A{:}B, C{:}B, D{:}B \mid {:}BD\} = \{\tfrac12, 0, 1 \mid 2\} = 1\tfrac12;$
${:}C = \{A{:}C, D{:}C \mid {:}AC\} = \{1, * \mid 1\} = 1*;$
${:}D = \{A{:}D, C{:}D \mid {:}AD, {:}BD\} = \{0, -\tfrac12 \mid 1, 2\} = \tfrac12.$

So the original position is $\{-\tfrac12 \mid 0\} = -\tfrac14$.

Similarly in Snort, the original position is now $\{A{:}, B{:}, C{:}, D{:} \mid {:}A, {:}C, {:}D\}$, which we can evaluate using

$A{:} = \{AB{:}, AC{:}, AD{:} \mid A{:}C, A{:}D\} = \{\{1|0\}, 2, 2 \mid -1, *\} = \{2 \mid -1\};$
$B{:} = \{AB{:}, BC{:}, BD{:} \mid B{:}D\} = \{\{1|0\}, 2, 2 \mid 1\} = \{2 \mid 1\};$
$C{:} = \{AC{:}, BC{:}, CD{:} \mid C{:}A\} = \{2, 2, \{1|0\} \mid 1\} = \{2 \mid 1\};$
$D{:} = \{AD{:}, BD{:}, CD{:} \mid D{:}A\} = \{2, 2, \{1|0\} \mid 1\} = \{2 \mid 1\};$
${:}A = \{C{:}A, D{:}A \mid {:}AC, {:}AD\} = \{1, 1 \mid -1, -1\} = \pm 1;$
${:}C = \{A{:}C \mid {:}AC, {:}CD\} = \{-1 \mid -1, *\} = -1*;$
${:}D = \{A{:}D, B{:}D \mid {:}AD, {:}CD\} = \{*, 1 \mid -1, *\} = \pm 1.$

So the original position is $\{\{2|1\} \mid -1*\}$, as required.

26 In Col, if Left plays first, then whatever the move, Right can always colour the symmetrically opposite region, so Left must lose. Similarly, Right to play loses, so this is the zero game.

In Snort, there is no way in which Left to play first can colour a region so as to prevent Right from moving, but Left can do the next-best thing, by playing in either rectangle. In the resulting position, Left could play in the other rectangle to +4 (and there is clearly nothing better), while Right can only play in the opposite curved region, killing the other rectangle and leaving the position as +3. So Left playing first plays to $\{4 \mid 3\} = 3½ \pm ½$. Similarly, Right playing first plays to $-3½ \pm ½$, and the starting position is therefore $\{3½ \pm ½ \mid -3½ \pm ½\}$.

27 [This is a ‘jumbo’ solution to all versions of this question as set in various years!]

(a) In $S_{n+1}+S_{n+1}-S_n$, Left to play can either *(i)* chop down an S_{n+1}, in which case Right can take the top edge in the remaining S_{n+1}, leaving $S_n - S_n = 0$ with Left to play, or *(ii)* chop at least one edge from $-S_n$, say to $-S_k$, where $0 \leqslant k < n$, in which case Right can chop the same number of edges from an S_{n+1}, which leaves $S_{n+1}+S_{k+1}-S_k \leqslant S_{k+1}+S_{k+1}-S_k = 0$ (inductively). On the other hand, Right to play first can *(i)* chop the $-S_n$, which is a trivial Left win, or *(ii)* chop one edge from an S_{n+1}, leaving $S_{n+1}+S_n-S_n = S_{n+1} > 0$, or *(iii)* chop more than one edge from an S_{n+1}, leaving (say) $S_{n+1}+S_k-S_n$, in which Left can chop $n-k$ edges from $-S_n$, leaving $S_{n+1}+S_k-S_k = S_{n+1} > 0$. So $S_{n+1}+S_{n+1}-S_n$ is a second-player win, hence zero. Since $S_0 = 1$, and the S_n are numbers, by assumption, we deduce that $S_n = 2^{-n}$.

(b) S_∞ is still a Left win, so $\varepsilon > 0$. On the other hand, S_∞ is better for Right than any (finite) S_n, as you can see by playing $S_\infty - S_n$, so $\varepsilon < 2^{-n}$ for every n; the result follows. On the third hand, if we play $S_\infty - \uparrow = S_\infty + \downarrow$, we see that Left wins, by following the prescription of Q12 (play in the all-small game whenever possible), and we deduce that $\varepsilon > \uparrow$.

You can interpret S_∞ as the game in which Left chops to zero, Right can select which S_n to play to. Similarly, there is an even-better-for-Right game $S_{\infty+1}$ in which Left chops to zero, but Right can ‘pass’ for one move before deciding which S_n to play to, and an $S_{\infty+2}$ in which Right can ‘pass’ twice, and an $S_{2\infty}$ in which Right chooses how often to pass, and an $S_{3\infty}$ in which Right chooses how often to pass before eventually playing to $S_{2\infty}$, and an S_{∞^2}, and so on. All of these games are Left wins, as Left can chop them down at any time; they are all different infinitesimals, as you can see by playing them against each other, and they are all large compared with the tiny games like $+_2$. In fact, you can interpret ε as $1/\infty$, as long as no real mathematician is looking, and then there are lots of (strictly) smaller infinitesimals like $1/(\infty+1)$ and ∞^{-2} and even $\infty^{-\infty}$, all corresponding to perfectly reasonable games, but all behaving just like numbers. They are small compared with ordinary real numbers, but large compared with all-small games and tiny games. For much more on this see *Winning Ways*, or Conway’s earlier book *On Numbers and Games*.

(c) If G has edges of only one colour, then the result is trivial, so we assume that there are both red and blue edges, and so both left and right options. Inductively, we can assume that all options of G are equivalent to numbers.

If A is a left option of G, it is obtained by removing from G a blue edge **e** and the [possibly empty] lumber of **e**. We obtain $-A$ by swapping blue and red edges in A. Now consider $G-A$, the Hackenbush picture obtained by drawing both G and $-A$ in the same picture.

Left to play wins $G-A$ trivially by playing to $A-A = 0$. Right to play may possibly take a red edge in the lumber of **e**; since this edge cannot be **e** itself, Left can still take **e**, removing what is left of its lumber, and again reaching $A-A = 0$. Every other Right move removes a red edge (and its lumber) which has a blue counterpart in the other component. In the resulting position after Blue removes this counterpart, call it $G'-A'$, following therefore exactly one move in G and one in $-A$, either the move from G to G' has removed **e** (and its lumber), in which case $G' = A'$ (and Right, now to move, loses), or it hasn’t, in which case A' is the option of G' obtained by taking **e**, and, inductively, $G'-A'$ is a left win.

So, $G-A$ is a left win; similarly, $B-G$ is a left win, where B is any right option of G, and so $(B-G)+(G-A) = B-A$ is a left win, or $B > A$. That is, G is a game all of whose options are numbers, by inductive assumption, and all the left options are less than all the right options; so the simplicity theorem applies, and G is the simplest number between these two sets of numbers. That establishes the induction, *QED*.

28 Clearly Abbreviations is a thin disguise for Red-Blue Hackenbush played on strings rather than more general pictures. Using previous results, every position is therefore a number; *banana*, *i.e. CVCVCV*, is worth $-1+\frac{1}{2}-\frac{1}{4}+\frac{1}{8}-\frac{1}{16}+\frac{1}{32} = -\frac{21}{32}$; and the given string *VVCCVC CVCV CV V CVC CVVC* has value $(2-\frac{1}{2}-\frac{1}{4}+\frac{1}{8}-\frac{1}{16}) + (-1+\frac{1}{2}-\frac{1}{4}+\frac{1}{8}) + (-1+\frac{1}{2}) + (1) + (-1+\frac{1}{2}-\frac{1}{4}) + (-1+\frac{1}{2}+\frac{1}{4}-\frac{1}{8})$, that is $1\frac{5}{16}-\frac{5}{8}-\frac{1}{2}+1-\frac{3}{4}-\frac{3}{8} = \frac{1}{16}$ [details left as an exercise!]. So Violet just wins, and as the only word with sixteenths in it is *Euclid*, the only winning move is to play the last vowel therein, as required.

If a phrase includes a y, then either player may play it, and so at least one Left option is equal to at least one Right option, This contradicts the definition of number. In the given phrase, *CVyVCC CyCVVCC VCV VCV*, the bits before any y's have total value 0 [exercise], so whoever first takes a y will lose to the other player taking the other y. So the game hinges on who wins *VCC CVVCC*, thereby forcing the other player either to take a y or to move in the bits before the y's. Clearly this is Connie, as *CVVCC* is the negative of *VCC* with two extra consonants; so Connie has the waiting move *tycoo*, making the rest of the game symmetric, while if Violet moves first, she must play a move which has a symmetric counter.

29 [My ‘model solution’ has pretty pictures, which I’m not going to attempt to reproduce here! Denote instead a ‘stalk’ by a sequence such as *WWBWBB*, meaning a white edge at the bottom, surmounted by another white edge, then a black, then a white, then two blacks at the top.]

House: Left’s moves are to $W+BB = -1$ and to *BBW*, in which Right can’t move and Left can move only to $BB = -2$. Right’s moves can be obtained by symmetry. So

$$House = \{-1, \{-2\,|\,\} \mid +1, \{\,|\,+2\}\} = \{-1, 0 \mid 1, 0\} = \{0\,|\,0\} = *.$$

Giraffe: Right to play has two possible moves: chopping the rear leg loses to chopping the front leg; so R must instead chop the back, leaving the back legs [clearly zero] and the stalk *WBBW*, which is clearly a loss for whoever goes first, so is zero. Left to play has three possible moves: chopping the head moves to -2; chopping the middle leg moves to [as it happens, see later] $\frac{1}{2}\pm\frac{1}{2} = \{0\,|\,-1\}$; chopping the front leg gives the back legs surmounted by the stalk *BBBW*. In this last position, chopping the head is still silly, and chopping either leg moves to a stalk which is clearly zero. What about the middle position? Left to play still obviously loses by chopping the head, but chopping the front leg gives the stalk *BBBBW*, clearly zero. Right to play can chop either the back or the back leg, leading in either case to a position where Left has no useful move, but Right still has one choppable edge to reach a zero stalk. Hence the result. So

$$Giraffe = \{-2, \{0\,|\,-1\}, \{0\,|\,0\} \mid 0\} = \{*\,|\,0\} = \downarrow.$$

[Note that $\{0\,|\,0\}$ dominates $\{0\,|\,-1\}$; a game cannot be improved ‘merely’ by making some moves worse.]

Chair: Left to play moves to -1 by taking the chair back or to $BBW = 0$ by taking the back leg. Right to play moves to $\{BW\,|\,2\} = \{0\,|\,2\} = 1$ by taking the front leg, or to $B+WW = -1+2 = 1$ by taking the seat. So

$$Chair = \{-1, 0 \mid 1, 1\} = \tfrac{1}{2}.$$

[These are all taken from *Winning Ways*, p54, but at least two of their values are wrong!]

30 Suppose the players are labelled *A*, *B*, *C*, are playing in that cyclic order, and that *A* and *B* form the coalition. From (2, 2), if *A* plays first, then *A* can take one heap, *B* the other and win. If *B* plays first, then he can remove one match/pebble/coin, leaving *C* with the position (2, 1); *C* cannot win immediately in this position, and must leave no more than two matches, so cannot get another chance to play, so must lose. If *C* plays first, he must leave either one heap or two; if one, then *A* can remove it and win;

if two, then A can remove one and B the other, and win. So in all cases, C is one of the losers.

If there are at least two heaps of size at least two, call this property P, then [firstly] if the position is exactly (2, 2) then C loses as just shown; [secondly] otherwise, there is at least one move that preserves property P. We can assume inductively that simpler positions with property P are won for the coalition, so A or B to move have at least one winning move. C to play must either play to a simpler position with property P and therefore lose inductively, or to a position with exactly one heap of size $\geqslant 2$. A to play can reduce this heap to either one or zero; when all heaps have size 1, the number of remaining moves is exactly the number of heaps, so the winner is determined exactly by the remainder when that number is divided by three; A can ensure that, after his move, that remainder is not 2 and therefore that C does not win.

If there is no heap of size > 1, then C wins or loses depending, as above, on the remainder when the number of heaps is divided by three; specifically, C wins if and only if that number is $3n+1$ for some integer n when C is to move. If there is one larger heap, then C can force a win only if it is his move, and then only if he can leave $3n$ singleton heaps after the move; that is, if the number of singleton heaps *before* his move is either $3n$ [so that removing the entire large heap wins] or $3n-1$ [so that removing all but one of that heap wins]. C cannot force a win in any other position.

31 The explanations are bookwork. In (a), we are initially playing a game equivalent to Nim on heaps 4, 1, 3 and 5 [alternate gaps!], Nim-sum 3, and you win by moving the marked coin as far as it will go; with the fixed coin, the heaps are 4, 1 and 2, Nim-sum 7, and you win by moving the rightmost coin left by one square. In (b), we are playing Nim on the gaps between opposite-coloured pawns, so with heaps of 6, 0, 5 and 2, Nim-sum 1, so White has to make the gaps 7, 1, 4 or 3 respectively, thus there are four winning moves, three involving moving away from the opponent, but the most direct win is to move the third pawn one square to the right. If the pawns don't necessarily alternate, then nothing much happens if only two pawns of the losing colour are ever adjacent [unless one has some room near the edge]; but if there are three with some space between, or two near the edge, then the otherwise-losing side can lose a move without changing the gaps by moving the 'protected' pawn in the middle. Fine details left as an exercise.

32 Cram is impartial, so any position is equivalent to a Nim heap.

For (a), we can construct the Grundy sequence very much as we did for Kayles [details left to the reader—indeed, Cram played on strips is the octal game **·07**, compared with Kayles **·77**]; it is 0, 1, 1, 2, 0, 3, 1, 1, mumble, 3, and a strip of length 10 is equivalent to a Nim heap of size 3. [We don't need the result for 9, as we can't move to it from strip 10.]

For (b), the simplest solution is to note that wherever the first player plays, the second player can play so as to block off a 2×2 corner, leaving a strip of five squares along the edge. This must leave one more move for each player [obvious, but the result from (a) also gives this], so the 3×3 board is a second-player win, equivalent to a Nim heap of size 0.

33 From a heap of size 13, the available moves are to split into heaps of sizes 1 and 12, or 2 and 11, or 3 and 10, or 4 and 9, or 5 and 8, or 6 and 7. Using the table, these moves result in positions of Nim values $0\oplus1 = 1$, $0\oplus2 = 2$, $1\oplus0 = 1$, $0\oplus1 = 1$, $2\oplus2 = 0$ and $1\oplus0 = 1$, respectively. So $G(13) = \text{mex}\,(1,2,1,1,0,1) = 3$. Similarly, from size 14, the moves are to heaps of sizes 1 and 13, or ... or 6 and 8 (but not 7 and 7), with corresponding Nim values of $0\oplus3 = 3$, $0\oplus1 = 1$, $1\oplus2 = 3$, $0\oplus0 = 0$, $2\oplus1 = 3$ and $1\oplus2 = 3$, so that $G(14) = \text{mex}\,(3,1,3,0,3,3) = 2$. A single heap of size 13 should be split 5 and 8, as that was the combination that produced the 0 in the mex.

Given heaps of sizes 13 and 14, the routine way to win is to note that the position is equivalent to Nim with heaps 3 and 2, which is won by taking 1 from the heap of size 3. The (only) combination that produced a 2 in the mex for $G(13)$ was 2 and 11, so a winning move is to split the 13 into 2 and 11. As it happens, there are also some 3s in the mex for $G(14)$, so splitting off 1, 3, 5 or 6 from the heap of 14 also wins.

From a heap of 27, there are lots of possible moves, so we have to hope that we find the win quickly. If 13+14 doesn't work, then the next to try is 12+15, for which we need $G(15) = \text{mex}(2, 3, 0, 2, 2, 0, 2) = 1$ [details left as an exercise, but once you've done 13 and 14 it should get routine]. Luckily, $G(12) \oplus G(15) = 1 \oplus 1 = 0$, so that's it, the opponent should have split 27 into 12 and 15. [For the curious, or persistent, $G(27) = 4$, and 1 and 26, 4 and 23 and 7 and 20 also win.]

34 A single sweet is 'dead', so has Grundy value zero. From two sweets, the only 'move' is to eat both sweets, which is 'dead', so the Grundy value is 1. From three sweets, you may move to two sweets by eating one, or you may 'kill' the pile by eating all three; as the corresponding Grundy values are 1 and 0 respectively, the Grundy value is $mex(1, 0) = 2$. From four sweets, the only move is to three, and $mex(2) = 0$. From five sweets, the only move is to four, and $mex(0) = 1$. From six sweets, the only move is to five, and $mex(1) = 0$, and so on. So the Grundy sequence is 0, 1, 2, 0, 1, 0, 1, 0, 1, ….

So three aniseed balls, four bon-bons, five comfits and six dragees have Nim-equivalent values of 2, 0, 1 and 0 respectively. The only winning move is to change the 2 to a 1 by eating one aniseed ball.

From there, the winning *strategy* is 'easy'. Because of the "She loves me, she loves me not" nature of the Grundy sequence, you can play anything you like *as long as* you never leave just three sweets of the same kind, as this is a guaranteed winner among 'heaps' of Grundy values 0 or 1 [like having one heap of size two and lots of singletons when playing Nim, you control the parity of the number of singletons]. When your opponent leaves such a threesome, as it came from one of size four, worth 0, you must immediately eat all three sweets to kill that sort.

[I found this problem in the Putnam competition, an annual 'Olympiad'-style test for students in the USA. As they didn't know about Grundy numbers, it was described as 'difficult'. *Our* students found it easy! Some students were worried about whether they could eat 'both' sweets from a pile of three; I don't think English allows such an interpretation, but if it did, it wouldn't matter as the one sweet remaining is 'dead'. Note also that 'strategy' means rather more than 'move'; it means explaining what you intend to do for the whole game, as opposed to just in the current position.]

35 We can easily draw up a table. Denote by P_n a Polite Nim heap of size n. Then the moves in P_n are to P_k where $\frac{1}{2}n \leqslant k < n$.

n	permitted k	corresponding Grundy numbers	mex
1			0
2	1	0	1
3	2	1	0
4	2, 3	1, 0	2
5	3, 4	0, 2	1
6	3, 4, 5	0, 2, 1	3
7	4, 5, 6	2, 1, 3	0
8	4, 5, 6, 7	2, 1, 3, 0	4
9	5, 6, 7, 8	1, 3, 0, 4	2
10	5, 6, 7, 8, 9	1, 3, 0, 4, 2	5
11	6, 7, 8, 9, 10	3, 0, 4, 2, 5	1
12	6, 7, 8, 9, 10, 11	3, 0, 4, 2, 5, 1	6
13	7, 8, 9, 10, 11, 12	0, 4, 2, 5, 1, 6	3
14	7, 8, 9, 10, 11, 12, 13	0, 4, 2, 5, 1, 6, 3	7

So the Grundy sequence is 0, 1, 0, 2, 1, 3, 0, 4, 2, 5, 1, 6, 3, 7, …, and we observe that $G(P_{2n+1}) = G(P_n)$ and $G(P_{2n}) = n$. We can now see why this pattern happens: For even n, the Grundy numbers [third column] are some permutation of the integers up to $\frac{1}{2}n - 1$, so the mex is $\frac{1}{2}n$. When we go to the next row, the first of these numbers drops out and is replaced by $\frac{1}{2}n$, so the mex is the number that has just dropped out, *i.e.* $G(P_{\frac{1}{2}n})$. When we go to the next row again, that mex comes back in, so we have restored the pattern. [You weren't asked for an explanation!]

Faced with $P_{12}+P_{13}+P_{14}$, the corresponding Grundy numbers are therefore 6, 3, 7, so the position is equivalent to playing Nim with heaps of those sizes. We first work out $6\oplus3\oplus7 = 2$, then $6\oplus2 = 4$, $3\oplus2 = 1$ and $7\oplus2 = 5$, so there are three winning moves, in which we play the heaps to smaller heaps of values 4, 1 and 5 respectively. Looking back at the table, the corresponding k values are 8, 11 and 10, so we have a choice of playing $P_{12} \rightarrow P_8$, $P_{13} \rightarrow P_{11}$ or $P_{14} \rightarrow P_{10}$, or in other words of taking 4 matches from the heap of size 12, 2 from the heap of size 13 or 4 from the heap of size 14.

36 (*a*) Note first that when one end 'burns' away, the game is exactly two-heap Nim, so is lost if and only if the two remaining heaps are equal. From a position where $p = r \neq q$, the player to move must reduce either p or r; by symmetry, we can assume p. So, after a move, the new position is p', q, r where $p' < r$. Three cases: if $p' = 0$, then as $q \neq r$ the second player wins by equalising the remaining heaps; if $p' = q$, then the second player wins by removing the r-heap; otherwise, the second player wins by removing the same number of matches from the r-heap, moving to a position $p' = r' \neq q$ which we can assume to be lost for the first player by induction. So in all case, the first player loses. If the condition $p = r \neq q$ does *not* hold in the initial position, then a very similar argument shows that the first player can move to bring it about, so winning.

(*b*) In the position 1, p, q, 1, the first move must take away one of the ends, leading to 1, p, q or p, q, 1. But 1, p, q is lost if and only if $q = 1$ and $p \neq 1$, while p, q, 1 is lost if and only if $p = 1$ and $q \neq 1$. So, 1, 1, 1, 1 is lost, 1, p, q, 1 is lost if $p \neq 1$ and $q \neq 1$, and all other positions with 1 at both ends are won. So, 1, p, q, r for $r > 1$ is won if p and q are both equal to 1 or both greater than 1, by reducing r to 1; and is won if $p > 1$ but $q = 1$ by reducing r to zero. This leaves the case 1, 1, q, r, with q and r both greater than 1. This is lost if $r = 2$, as every move reduces it to a previously analysed won postion; otherwise, it is won by reducing r to 2. Summary: If $r = 1$, then the position is lost unless exactly one of p and q is equal to 1; if $r > 1$, then all positions are won except 1, 1, q, 2 with $q > 1$.

(*c*) Faced with 2, 3, 4, 5 we see from the above that reducing either end to 0 or 1 leads to a position won for the other side. So the [only] winning move is to reduce the 5 to 2, leaving 2, 3, 4, 2.

37 The rule for calculating Grundy numbers given in the question says that we can ignore the rightmost coin if it is a tail, and that if it is a head, then the row is equivalent to the disjunctive sum of the row with that coin removed and a $*n$, where n is the length of the row. The first part follows from the rules of the game; a rightmost tail can never be flipped, so it plays no part at all in the game and might as well be removed [and then we can appeal to induction]. So we need only consider the case where the rightmost coin is a head.

Consider the game consisting of the given Turtle row, another identical Turtle row except that the rightmost coin is removed, and a Nim heap of size n. We need to show that this is a zero game, *i.e.* that to every move there is a winning counter. To every move in the shorter Turtle row there corresponds the same move in the first $n-1$ coins of the longer, and *vice versa*, resulting in a simpler position of the same sort which can be assumed to be zero by induction. To flipping the rightmost coin [only] there corresponds the removal of the entire Nim heap, and again *vice versa*, resulting in a zero position as there are two equivalent rows. To flipping the rightmost coin together with the kth, $1 \leqslant k < n$, there corresponds the move to $*k$ in $*n$, and again *vice versa*; in the resulting position, there are two shorter Turtle rows that differ only in the kth place, whose Grundy numbers are given [inductively] by the rule given in the question, so that they are together equivalent to the $*k$, so this position too is zero. That covers all moves, establishing the result.

In the given position, coins 1, 6, 7, 8 and 10 are heads, so its Grundy number is $1\oplus6\oplus7\oplus8\oplus10 = 2$, as required. If we flip coin 1, this would create a $*(1\oplus2) = *3$, but we can't flip coin 3. If we flip coin 6, this would create a $*(6\oplus2) = *4$, so flipping coins 6 and 4 wins. If we flip coin 7, this would create a $*(7\oplus 2) = *5$, so flipping coins 7 and 5 wins. If we flip coin 8, this would create a $*(8\oplus 2) = *10$, so flipping coins 8 and 10 wins. If we flip coin 10, this would create a $*(10\oplus 2) = *8$, so flipping coins 10 and 8 wins, as just seen. So there are three winning moves: 6 and 4, 7 and 5 or 10 and 8.

38 In Bond viewed as a Kayles-like game, the legal play is to knock down exactly three skittles from anywhere in a row of at least three. So, in a row of 11, the legal moves are to rows of 8, or 7 and 1, or 6 and 2, or 5 and 3, or 4 and 4, with Nim values 0, $2 \oplus 0 = 2$, $2 \oplus 0 = 2$, $1 \oplus 1 = 0$ and $1 \oplus 1 = 0$, respectively, so the mex is 1.

Given heaps (rows) of sizes 8, 9, 10 and 11, Grundy numbers 0, 3, 3, 1, the most obvious win is to turn the 1 into a 0, for example by taking the end three skittles from the row of 11. [There are several other ways of winning.]

From the 'equivalent game', we see that the Treblecross position consisting of a gap of n squares between two X's is equivalent to a Bond heap of size $n-2$. So the given position is equivalent to heaps of 4 and 9 in Bond, Nim equivalents 1 and 3, so a winning move should turn the 3 into a 1. In Bond, from a row of 9 we can move to 6, or 5 and 1, or 4 and 2, or 3 and 3, with Nim values 2, 1, 1 and 0 respectively, so we win by playing to 5 and 1 or 4 and 2 in Bond, that is to gaps of 7 and 3 or 6 and 4 between X's in Treblecross. So the winning moves in the given position are to place the next X in the 4th, 5th, 7th or 8th square past the middle X.

39 In **0.6**, you must take one man from a squad, leaving a non-empty squad behind, and you may split the squad into two if you like. [So this game is Officers.] From a squad of 11, the legal moves are to 10, to 9 and 1, to 8 and 2, to 7 and 3, to 6 and 4 and to 5 and 5. Using the table, the moves result in positions with Nim values 3, $2 \oplus 0 = 2$, $1 \oplus 1 = 0$, $3 \oplus 2 = 1$, $2 \oplus 0 = 2$ and $1 \oplus 1 = 0$, respectively. So $G(11) = \text{mex}\,(3, 2, 0, 1, 2, 0) = 4$. There are two 0s in that list, corresponding to 8 and 2 or 5 and 5; so the winning squads are 2, 5 or 8.

In Fish Kayles, you may take one skittle from the end of a row, corresponding to digit 3 (you may leave no heap or one heap, but not two); and you may take two skittles from anywhere, corresponding to digit 7 (leave no heap, one heap or two). So FK is **0.37**, first cousin to **0.6** and with Grundy sequence 1, 2, 0, 1, 2, 3, 1, 2, 3, 4, ... So an FK row of length 9 is equivalent to an Officers squad of size 10. From the table, we see that in Officers, the moves from 10 are to squads 9, 8 and 1, 7 and 2, 6 and 3 or 5 and 4 with respective Nim values 2, 1, 1, 0 and 1, so the only winning move in Officers is to 6 and 3. The corresponding move in FK is to 5 and 2; so faced with a row of 9, the (only) winning move is to knock down the 3rd and 4th skittles from (either) end.

40 The Grundy table is

7	8	6	9	0	1	4	5
6	7	3	1	9	10	3	4
5	3	4	0	6	8	10	1
4	5	3	2	7	6	9	0
3	4	5	6	2	0	1	9
2	0	1	5	3	4	8	6
1	2	0	4	5	3	7	8
0	1	2	3	4	5	6	7

[Work out from the bottom-left corner; each number is the mex of those left, down, or diagonally left and down from it. Eg, the value 6 just below the 8 in the main diagonal is the mex of (going left) 7, 2, 3, 5, 4, (going down) 0, 4, 3, 5, (going diagonally) 2, 5, 0, 1; in which we see 0, 1, 2, 3, 4, 5 but not 6.]

In the given initial position, the Grundy numbers of the initial counters are 8 and 4; a winning move therefore is to 'balance' the numbers by playing the 8-counter to a square of value 4, by sliding it three squares left or down. [That is the systematic win; there is also, as it happens, a 'sporadic' win by sliding the 4-counter five squares down to a square of value 8.]

The given initial position has Nim-equivalent value $8 \oplus 4 = 12$. As no square on the board has value 12, any third counter must leave the position unbalanced, so that the position has non-zero value and is a first-player win.

41 First part: bookwork. Note that **0.45** and **0.177** are first cousins. In **0.177**, the moves are to take one singleton match, or to take two or three matches from any heap with the option of splitting the result. So the moves from a heap of 11 are to 9, 8 and 1, 7 and 2, 6 and 3, 5 and 4, 8, 7 and 1, 6 and 2, 5 and 3 or 4 and 4, with Nim values 4, 5, 0, 3, 1, 4, 0, 0, 1 and 0 respectively, and a mex of 2, as required. [The Grundy sequence for **0.45** would be the same but with an extra 0 at the beginning.]

So, in **0.45**, heaps size 8, 9, 10 and 11 have respective Nim equivalents of 1, 4, 4 and 3 respectively, and we have to find a move to Nim value 1 from the heap of size 11, or equivalently from the heap of size 10 in **0.177**. In **0.177**, the moves from 10 are to 8, 7 and 1, 6 and 2, 5 and 3, 4 and 4, 7, 6 and 1, 5 and 2 or 4 and 3, with respective Nim values 4, 0, 0, 1, 0, 1, 0, 2, and 0. The moves that lead to a value of 1 are heaps 5 and 3 or a single heap of 7; back in **0.45**, the corresponding moves are (from the heap of 11) to take one match and split to 6 and 4 or to take two matches and split to 8 and 1. [There may accidentally be winning moves in the other heaps also.]

In **0.177**, the heaps have values 4, 4, 3 and 2, so a winning move would be to play from the heap of 10 to heaps 5 and 2. By accident, there was also a move in the heap of 11 to a Nim value of 3 by playing to heaps of 6 and 3; there may also be winning moves in the heaps of 8 and 9.

42 Explanations — bookwork [but a good test of your understanding!]. **·04**, **·042** and **·0421** all have identical sequences, as the extra moves are ineffectual [they amount to doing extra things using the skittles (or whatever) at the end of the rows, which would be dead anyway — try it!]; they are all first cousins to **·007**, which will have the same sequence but shifted left one place.

The sequence for **·04** starts 0, 0, 0, 1, 1, 1, 2, 2, 0, 3 [left as an exercise]. So rows of lengths 8, 9, 10 have combined Nim value $2 \oplus 0 \oplus 3 = 1$; the simplest win is to make the row of 10 have value $3 \oplus 1 = 2$, by taking away the second and third skittles.]

43 Take the components in turn:

tree: The RH branch is worth $(2 \oplus 2)+1+1 = 2$ as far as the main fork. The top branch is also worth 2 at that fork. The LH branch is worth 3 to the same fork. So, using the tree principle, the whole tree is worth $([(2 \oplus 2 \oplus 3)+1] \oplus 2)+1 = (4 \oplus 2)+1 = 7$.

house: Everything except the TV aerial is in loops, with a total of 11 edges, so can be fused into 11 petals, total value 1. The cross-bars on the TV aerial come in pairs, so can be ignored, and the rest is worth $(5 \oplus 1)+1 = 5$, using the tree principle, so the whole house is worth $5 \oplus 1 = 4$.

door: Fuses to 3 petals, worth 1.

car: The steering wheel and the aerial are worth 1 each and cancel. The other 21 edges can be fused into petals, total value 1.

So, the whole picture is worth $7 \oplus 4 \oplus 1 \oplus 1 = 3$. To win, we must chop the tree down to value $7 \oplus 3 = 4$, or the house up to value $4 \oplus 3 = 7$, or the door or the car to value $1 \oplus 3 = 2$. So, one winning move is to chop down either door jamb, and there may be something clever in the house or the car (in either case, you'd have to break up the loops, so a good place to try might be somewhere in the house roof, effectively lengthening the aerial, or somewhere like the car radiator, bonnet or boot — left as an exercise). The systematic win is to chop down part of the tree.

We work up from the root. If the whole tree is worth 4, the part above the lower fork is worth 3, but the current branches meeting there are worth 2 and 4, so we must either chop the short left branch to $3 \oplus 4 = 7$ (impossible), or the main branch to $3 \oplus 2 = 1$, and hence the part above the upper fork to 0. The branches meeting there are worth 3, 2 and 2, so we must chop the LH upper branch to 0 or either other branch to 1. This can be done by lopping off the whole upper LH branch, or the RH branch leaving just one component, or by plucking the top apple. Aesthetically, this last is clearly the best move!

44 Choosing a3 leaves just the top two rows. Suppose these have lengths p and q respectively. If $q = p-1$, then any move disrupts this relationship; any move in the second row makes q less without changing p, while any move in the top row makes p and q equal. Conversely, if $q \neq p-1$, then there is a move which creates this relationship; if $q < p-1$ then you can take $p-q-1$ squares from the top row, while if $q \geqslant p$ you can take $q-p+1$ squares from the second row. Since the (only) final position with $p = 1$, $q = 0$ is of this form, we see that a two-row position is lost for the player to move if and only if row 2 is one shorter than row 1, and hence the result.

Choosing b2 leaves row 1 and column a. Any further move is equivalent to playing Nim on heaps of their lengths, less one for the poisoned square; so the resulting position is lost if and only if they are balanced.

In the given position, there are seven possible moves. Now, c1, a3 and b2 should not be played, by the rules just given; b3 and d1 lose to c1 and a3 (or b2) respectively; b1 loses to a2 and a2 to b1. Hence, this position is lost for the player to move.

In a rectangle larger than 1×1, consider the effects of taking the bottom rightmost square. If this wins, then we have nothing more to prove. If it loses, then there must be a winning reply; but this reply was also available to the first player, and would have resulted in the same position (including the bottom rightmost square). So the first player has a winning move (which either is or isn't taking the bottom righmost square!).

As we have just seen, a winning move from the 4×3 rectangle is to take b3, leaving the reflexion about the long diagonal of the position we analysed as a loss.

45 (*a*) A stalk of height n has Nim value n. For sufficiently large n (for example, more then the number of edges in the rest of the picture), this stalk cannot be balanced by the Nim value of the rest of the picture, so the total picture must have non-zero Nim value and be won for the player to move.

(*b*) If chopping the blue edge wins, there is no more to prove. Otherwise, there is a good reply, say **e**, to chopping the blue edge; let Blue chop **e**, and continue to play as though the blue edge were not present. Red can either acquiesce in playing as though the blue edge were missing [as indeed it may be], or at some point must chop an edge which depends on the blue one [necessarily not the blue edge itself]. At this point, chopping the blue edge wins, by construction.

(*c*) For example, put all the blue edges on top of a green [purple] stalk.

46 Take the components in turn:

tree: Working down from the top, we have alternately to ordinary-add 1 to go down the trunk, and Nim-add 2 to incorporate the side-branches. So, we successively get 1, [Nim-add 2] 3, [add 1] 4, [$\oplus$2] 6, [+1] 7, 5, 6, 4, 5, 7, 8, 10, 11, 9, 10, 8, 9, 11, 12, 14, 15, 13, 14, 12, 13. The tree has value 13. [You might find this calculation easier in binary, where Nim-adding 2 means swapping the next-to-rightmost bit: 01, 11, 100, 110, 111, 101, 110, 100, 101, 111, 1000, 1010, 1011, 1001, 1010, 1000, 1001, 1011, 1100, 1110, 1111, 1101, 1110, 1100, 1101. At the risk of flogging the example to death, note that the bottom two bits repeat in a cycle of length 8, and that once round each cycle a one is carried to the next bit to the left, so each cycle has values 4 more than the previous cycle; this makes it easy to calculate the Nim value of a tree of this form of any height.]

dog: Can be fused to an odd number of petals apart from the head and neck (worth 2) and the tail (worth 1), so total value is $1\oplus2\oplus1 = 2$.

girl: The head fuses to two petals, and then there's the ponytail and the arm, so from the neck upwards she's worth $1\oplus1\oplus1\oplus2 = 3$. Treeing her chest makes 4; her skirt fuses to 4 petals, cancelling each other out; finally, the 4 trees through her leg to 5.

light: The shade and bulb fuse to 6 petals, together worthless. This is on a stalk of three components, so the light is altogether worth 3.

So the total picture is worth $13\oplus2\oplus5\oplus3 = 9$. Since we clearly can't get such large values from the dog, girl or light, we must chop down the tree to value $13\oplus9 = 4$. This is done by working up the tree from the bottom, alternately subtracting 1 to go up and Nim-adding 2 to get past a side-branch, until we reach zero. So, from 4 we get successively 3, 1, 0 [but this is uselessly at a fork, so there is no corresponding chop], 2, 1, 3, 2, 0 [now having gone past the 4th side-branch]. There is a unique winning move, which is to chop the tree at the 5th segment up the trunk. [Left as an exercise: check that the resulting stump indeed has value 4.]

47 First part is bookwork. To evaluate the given position, it makes life easier if you appeal ruthlessly to symmetry. The two lanterns are [supposed to be] the same, so cancel. Likewise the girls arms, the hair-ribbon, any two of the three blooms on the flower, and the topmost branch of the tree with the branch opposite to it. Then the girl's skirt and face fuse to an even number of edges, so can be deleted. So, the girl is *4 [two lengths of hair plus body plus leg(s)], the flower is *3, the bridge is an odd cycle which therefore reduces to *1, and the tree consists of a *3 on the left and a *2 on the right [easy exercise], combining to a *1 on a trunk making *2. Tree, flower and bridge therefore sum to zero, and the whole position is worth *4.

A simple winning move is therefore to zap the girl by chopping off her leg(s). I don't think there is any other winning move! [But chopping the top apple or a bridge-span next to the flower come close.]

48 First part—bookwork. From a row of 9, you can move to 7, to 6 and 1, to 5 and 2, or to 4 and 3 with respective Nim values of 1, 3, 1 and 3, and a mex of 0. From a row of 10, you can move to 8, to 7 and 1, to 6 and 2, to 5 and 3 or to 4 and 4, with respective Nim values of 1, 1, 2, 1 and 0, and a mex of 3. To win with three rows of lengths 8, 9 and 10, so with Nim values of 1, 0 and 3, you have to find a move to value 3 in the row of 8, to value 2 in the row of 9, or to 1 in the row of 10; there are several ways (left as an exercise) of doing this.

The Boxes position given is a Dawson's Vine equivalent to a Dawsons's Kayles row of length 7 (the number of internal tendrils). There are four winning moves (two symmetric pairs), corresponding to the lines $c56$, $e56$, $f56$ and $h56$ using the notation of Q49.

49 (*a*) The move $c34$ sacrifices two boxes by $b34$ and $cd3$, leaving three long chains, one down the *def* columns, one snaking round between b and e in the bottom two rows, and one along the top leading to a still open position in the left column. There is also a *1 in the bottom left corner. In the open position, there is really just one move, $a34$ ($ab3$ is equivalent impartially, but sacrifices a box needlessly), so this is also a *1, and the total is zero, so the player who is to move after $c34$ loses. It's nice to see the theory confirmed, so we note that in the diagrammed position 23 moves have been played, so the second player is to move, and wants to create an odd (as there are 30 dots) number of chains, such as three. After $a34$, the (originally) second player sacrifices the bottom left corner box by $a12$ (or $ab1$), then drops four more boxes in the first two long chains, winning by 13 boxes to 7. But ....

In the version of this question as actually set in 1989, the right-hand column of four boxes was missing, so that the winning margin was 9 to 7. Unfortunately, in the resulting diagram, instead of $a34$, the second player can sacrifice with $ab4$, and after the replies $b45$ and $cd5$, sacrifice the bottom left corner by a12. The first player must then play ab3 (after ab1) to throw the move, sacrificing two more boxes. So the first player loses the two boxes sacrificed initially, two more boxes in the a-column, and two boxes in each of the first two long chains, total score 8 all. No-one noticed this possibility for five years! So, with best play, the loser squeezes one more box than above, and the total score is 12 to 8.

(b) The reply to b34 is cd3, throwing the two middle boxes but switching the move. Everything above is then reversed, and the player to move in the diagram loses by 12 to 8 [or 13 to 7 by missing the ab4 idea].

(c) The move a34 is a blunder. The other player must attempt to keep the number of long chains down to two, so must prevent the split by c34. This can be achieved by the move bc2, sacrificing two boxes in order to make the central boxes part of the long chain running down to near the bottom right corner; d12 and (less obviously) ab3 or a23 have the same effect. After the replies bc1 and c23, the first player has nothing better than ab3 (or a23), sacrificing a box; the other player then gives up the bottom left box, and we are left with just the two long chains. The other player wins by 15 to 5. You might wonder whether the first player does better with b34 instead of ab3, forcibly creating the third chain, and hoping for some extra enforced sacrifices; but no, the other player takes all the five available boxes, then sacrifices the bottom left, and still wins by 15 to 5.

It's worth noting that after a34, the first player has converted the central position into a Kayles vine with two nearby tendrils, and so equivalent to K_2, Nim value *2. The other player wins by zapping one (but not both!) of the tendrils, giving a *1 to add to the *1 in the bottom left. Again, it's nice to see the theory confirmed!

50 First part—bookwork. In a position with n chains of 3, and no other moves, the player to move must commit suicide in one chain. Suppose the resulting position is worth a swing of $f(n)$ boxes to the player to move. This player has the choice between grabbing all three boxes, but then having to move first in the next chain (thus losing $f(n-1)$), and taking one box but sacrificing two to make the other player move first in the next chain. So

$$f(n) = \max(3 - f(n-1), -1 + f(n-1)).$$

Clearly, $f(1) = 3$, so $f(2) = \max(0, 2) = 2$, $f(3) = \max(1, 1) = 1$, $f(4) = \max(2, 0) = 2$. Also, if $f(n-1) = 1$, then $f(n) = 2$ and *vice versa*, so all further values of f are 1 or 2, according as n is odd or even.

51 In *(a)*, if we use the notation of Q49, there are 7 moves, ab2, ab3, bc2, bc3, a12, a23 and b12. Of these, bc3 and b12 are suicides, while a23 and ab3 are equivalent, so only one of them needs to be considered. Taking the remaining four moves in turn:

ab2 divides the board into a chain of length 3 and a box with two sides already drawn, worth *1;

after ab3, the non-suicidal replies are a12 (creating a chain of length 4) or ab2 or ab3 (both sacrificing a box to leave a chain of length 3), all therefore leading to *0, or bc2 (sacrificing two boxes to reach a position in which the only non-suicide is ab2), thus leading to *1, so ab3 leads to a position whose Nim value is the mex of 0 and 1; *i.e.* to *2;

bc2 sacrifices two boxes to reach a position in which ab2 and a12 sacrifice a box to reach a *1 while ab3 and a12 reach the *1 considered previously, so bc2 leads to *0 (and note that playing bc2 on move 1 is the way to ensure that no long chain is ever created in this position); and

after a12, the only non-suicidal moves are ab3 and a23, both giving a chain of length four, so a12 leads to *1.

Thus, the Nim value of (a) is the mex of 1, 2, 0 and 1; *i.e.* 3, as required.

In (b), ab1, de1, a12 and e12 all lead to a chain of three attached to a one-sided box, the same position as that reached in (a) after a12, so these positions are worth *1, while b12 and d12 are suicides, and c12 sacrifices two boxes to reach a *0 (by symmetry, or as *1 + *1). The mex of 0 and 1 is 2, so position (b) is worth *2.

The position shown in (c) has three components: a *3 in the top-left corner, a *2 down the right-hand side, and a long chain of length 8 down the middle. So, the good move is to play to *2 in the top-left corner, by filling in either of the edges from the extreme top-left corner dot (*ab*5 or *a*45).

Note that there are 25 dots, and 17 lines already drawn, so the second player is to play, and wants to create just one more long chain. Each 'live' component is now in a state where the next move will determine whether or not a long chain will result, so the other player will have to decide for one component, and the second player can take the opposite decision for the remaining component, thus ensuring exactly two long chains.

52 First part—bookwork. The position of *(a)* was discussed after the move *ab3* [with what is now our standard notation!] in part *(a)* of the previous question. So the position *(b)* consists of *1 [derived during the analysis of *(a)*] in the top left-hand corner, position *(a)* [rotated through 90°] worth *2 in the top right-hand corner, a long chain, and a *1 in the bottom right-hand corner. So the winning move is to convert the *2 to 0, by the move *cd5*. Sacrificing the very top right-hand box also wins, but loses a box unnecessarily; every other move loses. If the opponent instead of *cd5* sacrifices the bottom right-hand box, then we are left [after capturing that box] with a *2 and a *1; the winning move is to convert the *2 to a *1 by the two-box sacrifice *d34*. Again, *every* other move loses.

To compare with theory, note that there are 25 dots, and 19 lines have been drawn, with no captures, so the originally-second player is to move in position *(b)*. So the player to move should be trying to make the number of long chains even. There is one long chain already, and the only place to create a second is in the rotated version of *(a)* in the top right-hand corner. The player to move should thus create a long chain there, and the other player should, given the chance, prevent this.

53 The first part is bookwork.

In the left-hand position, there are 16 dots, 11 lines have been drawn, so it's the second player to move, wanting to force an odd number of chains. There is already one chain along the top, and no way of creating a second from those five boxes. So it is necessary to prevent a chain in the bottom right. You can either sacrifice two boxes by playing *c*23 with our standard notation, or create a four-box loop by playing *bc*1 or *cd*1. *Exercise:* show that these all win by five boxes to four.

In the right-hand position, there are 16 dots, 12 lines drawn, so the first player is to move, with an even number of chains wanted. But the second player has [presumably!] just suicided in the top left. There is no way of getting a second chain, so the good move is to force an extra double cross; you should decline the offered boxes, and play *ab*4 in the top row, winning by six boxes to three.

54 Note first that if there are no chains or loops, then C and L are both zero, $N = B$, and since $B > 3$ the condition

$$C+2L < \tfrac{1}{4}B-\tfrac{1}{2}N+1$$

fails, so there is nothing to prove. If there is at least one chain or loop, then the impartial game will be won by the usual double-cross theory. In the worst case, the impartial winner will lose all the N uncommitted boxes, two boxes for each of the chains except the last, and four boxes for each loop, so the impartial loser may score as many as $N+2(C-1)+4L$ [or two less than this if $C = 0$ but $L > 0$]. So the impartial winner can certainly [with best play!] win the partisan version provided this score is less than $\tfrac{1}{2}B$, that is, if

$$N+2(C-1)+4L < \tfrac{1}{2}B$$

which easily re-arranges to the above form.

To construct a case where the outcomes differ; try $B = 4$ [as we are given $B > 3$], then we need $N+2C+4L \geqslant 4$. However, any loop involves all four boxes, and the winner picks up all four; so we must have $L = 0$. There is no room for two chains, and if $C = 1$, then $N \leqslant 1$, so $N+2C \leqslant 3$, and there can be no differing outcome there. So there must be no loops or chains. The simplest position to analyse is the 2×2 square [9 dots] with the central ‘cross’ drawn in, so dividing the board into four separate regions each of one box. This is clearly a second-player win in the impartial version, but a two-all draw in the partisan version. Left to the reader to find cases where the other player actually wins—there are easy examples in a five-box strip [2×6 dots].

55 The ‘explain’ bit is bookwork. The *d5*, *d6*, *d7* complement problem starts with a count of 12 and finishes with a count of 16, so is insoluble.

In the *d7* complement problem, there is an actual slack of 12, but *d7* must be cleared between first and last move, costing at least 8, so only 4 is usable by other moves. Since the given 6-pack costs 4, all other moves must be free, which rules out *c2–e2* and *a3–a5* of those given.

In the *d6*, *d7* problem, the slack is 2. The first move must be *d4–d6*, and the only way of getting a peg to *d7* is by playing *d5–d7*, clearing *d6* again. Since any further move involving *d6* or *d7* other than *d4–d6* cost more than the slack, in particular all ways of clearing *d7*, the result follows. In addition, all moves clearing *a4* or *g4*, other than those given, cost 4 and so cannot be used.

56 (*a*) For example, start with $c4-c2$. Now clear out the *ab* wing with a 6-pack; this creates the catalysts to clear out the top and bottom wings also by 6-packs, and then the *fg* wing similarly. Now clear the 5th rank using a 3-pack, and finish off with $e3-c3$ and $e4-c4-c2$. Many other solutions.

(*b*) Start with an exercise in Reiss classification. You can count pegs; but it’s easier to imagine appropriate moves and packs. From Weeping Willow, chase the $a3$ peg as far as it will go: $a3-a5-c5-c3-e3-e5-g5$. This leaves 3-packs in the g file and the 1st rank, which can be ignored, and a single peg in $d2$. So, since the standard board is self-complementary, the single vacancy start must be in $d2$ or the equivalent. As we want to clear out the top of the board, the most sensible equivalent is probably $d5$. So, start with $d5$ empty. Clear out the top wing with a 6-pack, and the 5th rank with two 3-packs. This leaves the first four ranks fully occupied, and everywhere else empty (this is a good position to aim for; imagine unplaying moves like $b3-b5$). Clear $c-e4$ with a 3-pack, and now it’s easy: just play $b3-b5$, $c2-c4$, $e2-e4$ and $f3-f5$.

It’s also possible, though harder (unless I’ve missed a trick!), to reach WW from the $d2$-empty position. Start with $d4-d2$ and $b3-d3$. Next some ‘put and take’ routines: $c5-c3$, $c2-c4$ (clearing squares three apart provided the intervening squares are opposites), clear $b4$ and $e4$ similarly, and then $a5$ and $d5$. Now clear $e5-g5$ with a 3-pack, clear the top wing with a 6-pack, finishing with $f3-f5$ and $e2-e4$.

57 Reiss classification! The usual starting position is in the same class as its complement, and hence (making an unmove) as the position with pegs at (for example) *b4* and *c4*. Now unmove *c6–c4* and move instead *c5–c7* to get one peg to *c7*. The other peg has to be on the intersection of diagonals like those through *b4*, and is hence at one of the places given.

Many possible solutions. For example, start with a 3-pack on *e4–6*, then clear out the *f–g* wing and the *1–2* wing with 6-packs, and *d5–7*, then *c5–7*, and *c–e3* with 3-packs to reach the suggested position. [Note: you can’t clear out *c–d,5–7* with a 6-pack, as the catalyst is missing.] Now there is a neat triple-move: *a3–c3–c5*, *a4–c4–c6* and *a5–c5–c7* to finish off.

58 (*a*) For example, clear out the second row with a 3-pack, then the left, top and right wings with 6-packs, then the fifth row with a 3-pack. Now jump the fourth row down over the third row, clear out the second row again with a 3-pack, and jump $e1$ to $c1$. Many other solutions.

(b) Exercise in Reiss classification! We can remove all the pegs in the e-file without affecting the class; the resulting two-peg position has even occupancy in each of the NW to SE diagonals, and so is in a different class from any single-peg position [which must have just one odd occupancy]. Since the standard board is self-complementary, it is also in a different class from any single-*vacancy* position. So Letter J cannot be reached from any such position.

59 Suppose that neither m nor n is divisible by 3. Then all but the bottom right-hand $i\times j$ rectangle, where i and j are the remainders when m and n are divided by 3, can be tiled in an obvious way by 3×1 and 1×3 rectangles, which are in the same Reiss class as their complements. But each of the possible $i\times j$ rectangles (that is, 1×1, 1×2, 2×1 and 2×2) has Reiss diagonals of differing parity (for example, the 2×2 has α and γ diagonals of length 1, but a β diagonal of length 2), so that when the position is complemented some of the counts will change parity and others won't, and the Reiss class will change. Hence the result.

On the continental board, the standard start position is in the same Reiss class as the empty board (either by direct counting, or by systematic removal of 3-packs until the board is empty enough to make such counting very easy), and so in a different Reiss class from any single peg. For a two-peg solution, both pegs must lie on the same class of diagonal (to make all counts even), so if one is on *d6*, the other must be on *a3*, *d3* or *g3*. One solution is as follows: Start with 3-packs on *b–d3* and *b–d2*. Next an L-pack on *c–e1–4*, a 3-pack on *f2–4*, another L-pack on *g3–5–d*, another 3-pack on *d–f6*, and a final L-pack on *e–c7–4*. This leaves 6 pegs in *a3–5* and *b4–6*. Finish with *a5–c5*, *a4–c4–c6* and *b6–d6*, leaving pegs in *a3* and *d6*. [The key to this packaging is the way the 3-packs and L-packs can rotate neatly to fill in the wings; once that has been spotted, it's easy to get rid of most of the pegs, leaving a compact and relatively easy configuration to reduce. Of course, your first attempt is very likely to be in the wrong orientation, but that's life!]

60 By symmetry, we only need to consider three cases, for example $a1$, $a2$ and $a3$. There are many solutions; here are mine.

Firstly, consider the $a1$ problem [hole in the corner]. We can clear $b1-3$ with a 3-pack, and this gives us the space and catalysts to clear out four 6-packs; successively, $cd4-6$, $ef4-6$, $ef1-3$ and $ab4-6$. If we now jump $d1-3$ over $c1-3$ into $b1-3$, we have another 3-pack in $b1-3$ and a final move of $a3-a1$.

This same solution works equally well with the $a3$ problem; the only change is that the final move is $a1-a3$.

For the $a2$ problem. start with the hint: play $a4-a2$ and $c3-a3$. Now we can clear out $a-c12$ and $d-f12$ with 6-packs. There are no more obviously-useful 6-packs, but we can do the same thing by 'eating away' with 3-packs in $d-f3$ and $d-f4$. Now we again have 6-packs available, so can clear out $a-c56$ and $d-f56$, and finish off with $c4-a4-a2$.

61 In the complementary position, first construct an L-configuration by $c4-c6$, $c2-c4$, $b3-b5$, $a5-c5$. This needs a catalyst, so play $f3-f5-d5$. Clear the L-pack, and rescue the isolated g-peg by $d6-d4-f4$, $e2-e4-g4$ and clear to a single peg by $g5-g3$. [Note that the rescue is impossible if $e234$ is cleared by a 3-pack; and if $f3-f5$ followed by a 3-pack on $efg5$ is played, there is no way of getting a catalyst to $d5$.]

To solve the original position, start with a full board apart from a vacancy in $g3$, and play the above moves in reverse order: $g5-g3$, $e4-g4$, $e2-e4$, …, $c4-c6$.

If, initially, a peg is added in $d6$, the resulting position is in the same Reiss class as the empty [or full] board—seen by direct computation or by inspection from the single-peg position above. So it is in a different Reiss class from any single-vacancy position. Finding a two-vacancy start is left as an exercise!

62 Start by finding the Reiss class of ‘Arrow’. Lots of ways, including brute force; *e.g.* move [‘rule of 3’] $e2$ to $e5$ and $e6$ to $e3$, giving a 6-pack in the e- and f columns that can be removed; then move the $g4$ peg to $d4$ to cancel the peg there; then play $a5$ to $c5$ to $c3$ [jumping over and removing $b5$ and $c4$], and remove the 3-pack on the 3rd row, leaving a single peg in $b4$. So a single vacancy in the same class must [as the English board is self-complementary] also be in $b4$ or else in another place in the same class, hence [‘rule of 3’ again] in $e4$ or $e1$ or $e7$.

Now we look at the suggested resource count. ‘Arrow’ has resource [sum of values of occupied holes] $-1-1+1+1+1+1+2+1+1+0+1+1+1+1 = 10$, while the full board also has resource [sum of all values] 10. So, with the single vacancy at $b4$ or $e4$, the initial resource is $10-1 = 9$, so is insufficient. With the vacancy at $e7$, the initial resource is $10-(-1) = 11$, so there is an initial slack of 1. But a first move of $e5-e7$ costs 2, which is too much, and the only other move is $c7-e7$ which is free, but then the second move must be $c5-c7$ or $d5-d7$ which both cost 2. Similarly [by symmetry], an initial vacancy at $e1$ is no better. So in all cases, after no more than 2 moves the current resource is no more than 9, from which it is impossible to reach a position such as ‘Arrow’ which has resource 10.

63 (*a*) Bookwork—the ‘Desert Patrol’ or ‘Chinese Army’.

(*b*) Work backwards. To get to 0,4, we need pegs in 0,3 and 0,2. Defer 0,2, and consider 0,3. We’ll try to do this with pegs ‘to the right’ of the diagram, so that we can use the space ‘to the left’ for the deferred problem. To get to 0,3, we need pegs in 0,2 and 0,1. To get to 0,2, simplest is to have pegs in 0,0, 0,−1, 1,0 and 2,0; then you can jump ‘up’ and ‘round the corner’ in the most obvious way. We now have to get a peg to 0,1 starting without the four pegs already initially placed. If we place pegs in a block of four at 1,−1, 2,−1, 1,−2 and 2,−1, then we can jump these up and across to get a peg to 0,0; then add pegs at 0,−2 and 0,−3, and this last peg can double-jump up to 0,1 and thence to 0,3. [We could do this slightly more easily except that we’re trying to leave lots of space to the left for our ‘deferred’ problem.]

Now we have to get a second peg to 0,2 without using pegs initially ‘to the right’. So we need to get pegs to 0,0 and 0,1. We can get to 0,1 by having a block of four at −1,0, −2,0, −1,−1 and −2,−2, and jumping them up and across [or across and up]. Now to get to 0,0. Simplest is to observe that we can use the shape we used to get to 0,3, but turned through 90 degrees so as to get to 0,0 using only pegs from three columns or more to the left and no higher than the ‘x-axis’; if you do this, and then see how it works, you’ll find you can save a peg or two, but you weren’t asked for the solution with fewest pegs.

There are other ways of doing it, but the ‘Desert Patrol’ shows that it won’t be trivial. It may help if you draw a diagram or three.

(*c*) Resource counts! Simplest, since you have already used it in part *(a)*, is to use the array of σs; but a Fibonacci-style count can be used instead. With 7 pegs, the most you can get is $\sigma^0+3\sigma^1+3\sigma^2 = 1+3\times(\sigma+\sigma^2) = 1+3\times1 = 4$, whereas the target hole has resource $\sigma^{-3} \approx 4.23$, so you have insufficient resource. If you don’t quite trust your calculator, then $\sigma^{-3} = \sigma^{-2}+\sigma^{-1} = 2\sigma^{-1}+\sigma^0 = 3+2\sigma = 4+(2\sigma-1) = 4+(\sigma-\sigma^2) = 4+\sigma^3 > 4$, making ruthless use of the defining equation $\sigma+\sigma^2 = 1$.

64 Left, the attacker, can split the forces 4–0, 3–1 or 2–2; Right, the defender, can split 3–0 or 2–1. In each case, the choice of which road to attack or defend more strongly should clearly be taken at random with probability ½. The payoff matrix is easily found to be (taking strategies in the order given above):

$$\begin{pmatrix} \tfrac12 & \tfrac34 \\ 1 & \tfrac12 \\ 1 & \tfrac12 \end{pmatrix}$$

Turn Over

The row minima are all ½, so the lower value is ½. The column maxima are 1 and ¾, with minimum ¾, so the upper value is ¾. Since these differ, a mixed strategy is in order. Plainly, Left's strategies 2 and 3 are equivalent, so we can take them together. Assume Left plays strategy 1 with probability x, and Right plays strategy 1 with probability y. Then the expected payoff is

$$\tfrac{1}{2}xy + \tfrac{3}{4}x(1-y) + 1(1-x)y + \tfrac{1}{2}(1-x)(1-y) = \tfrac{1}{4}(2 + x + 2y - 3xy)$$

with critical value ⅔ when $x = ⅔$ or $y = ⅓$ (obtained by equating the partial derivatives to zero). [So, Left should usually attack in strength, and Right should usually split the forces. In reality, Left should presumably always send at least a small force along the 'other' road, if only as a diversion. and Right should probably keep one division in reserve.]

65 Holes B and C are indistinguishable, so may be taken together, with the dog or rabbit choosing between them by tossing a mental coin when appropriate. So the dog (Left) has four available strategies: wait near A, wait near B or C, wait between B and C and wait overlooking all three. The rabbit (Right) has two strategies: exit through A and exit through B or C. The corresponding payoff (probability of catching the rabbit) matrix is:

$$\begin{pmatrix} 1 & 0 \\ 0 & \tfrac{1}{2} \\ 0 & p \\ q & q \end{pmatrix}$$

The lower value is plainly q, and the upper value is the greatest of p, q and ½, so a mixed strategy is needed unless q is particularly large. Comparing the dog's second and third strategies, we see that strategy three is dominated if $p < ½$; this is intuitively obvious—if p is small, then he does better to choose a hole and stick to it, and if it is large he does better to hedge his bets. The fourth strategy, guarding all three holes, is inferior to choosing a hole at random if $q < ⅓$; again, this is intuitively obvious, but can be rationalised by comparing row four with choosing row one with probability ⅓ and row two with probability ⅔.

In the case $p = \frac{3}{4}$, $q = \frac{2}{5}$, suppose the rabbit chooses hole A with probability y. Then the dog's strategy two is dominated, and the other three strategies pay y, $p(1-y)$ and q.

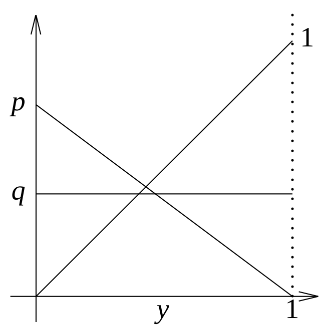

From the graph, and remembering that the rabbit is trying to minimise the maximum payoff, we see that strategy four is dominated, and the value of the game is $\frac{3}{7}$ when $y = \frac{3}{7}$. To achieve this payoff, the dog should wait by hole A with probability $\frac{3}{7}$ and between holes B and C with probability $\frac{4}{7}$; the rabbit should choose hole A with probability $\frac{3}{7}$, and holes B and C with probability $\frac{2}{7}$ each.

66 Note that you either win your opponent's stake or lose your own; construction of the payoff matrix and of lower and upper values left as an exercise. If Left stakes £1 with probability x and Right stakes £1 with probability y, then the expected payoff is

$$xy - x(1-y) - 2(1-x)y + 2(1-x)(1-y) = 2 - 3x - 4y + 6xy = (2 - 3x)(1 - 2y).$$

From this factorisation, we see that either Left or Right can force the payoff to be zero (hence a fair game) by choosing $x = \frac{2}{3}$ or $y = \frac{1}{2}$ respectively. (The same result follows by the more usual techniques.)

Since the amount you lose is your own stake, the amount you win is your opponent's stake, and the outcome depends only on the parities of the stakes, there is no point staking a higher number of pounds if there is available a lower number of pounds with the same parity. So only the lowest possible stake of each parity is relevant. In case (*a*), only stakes 1 and 2 should be considered, and we recover the previous game. In case (*b*), only stakes 2 and 9 are relevant; the analysis is left as an exercise (very similar to the above). In case (*c*), Right can only bet an even number (hence £2), and Left wins £2 by betting any even number.

67 Labelling strategies by the number of fingers shown, the payoff matrix is:

	RIGHT 1	2	3	minima
LEFT 1	−2	3	−4	−4
2	3	−4	5	−4
3	−4	5	−6	−6
maxima	3	5	5	

The lower and upper values, −4 and 3, differ so a mixed strategy is in order. There are no obvious dominations, so we try letting Left play strategies 1, 2, 3 with respective probabilities x, y and $1-x-y$, and Right similarly s, t and $1-s-t$. This leaves you with some tedious but simple algebra and calculus [left as an exercise]. You will find that $\partial V/\partial s = 0$ yields $y = \frac{1}{2}$ and then that $\partial V/\partial t = 0$ yields $x = \frac{1}{4}$, and similarly $t = \frac{1}{2}$ and $s = \frac{1}{4}$, and that for these values, the game has value zero [so is 'fair'].

Now we've seen the answer, the reason becomes obvious. Note that the payoff [for each player] of strategy 2 is the negative of the average of strategies 1 and 3. So if either player plays strategy 2 with probability $\frac{1}{2}$, amd the other two strategies each with probability $\frac{1}{4}$, then the expected payoff against any opposing strategy is zero. Since either player can enforce this, neither player can do better than this.

68 If $\alpha \geqslant 1$ then for each player strategy 2 dominates strategy 1, so each player should play strategy 2 and the value of the game is 1. If $\alpha < 1$, then the lower value of the game is α, the upper value is 1, and a mixed strategy is called for. If Left plays 1 with probability x, Right with probability y, then the expected payoff is

$$1xy + 0x(1-y) + \alpha(1-x)y + 1(1-x)(1-y) = 1 - x - (1-\alpha)y + (2-\alpha)xy$$

so the critical values are $x = (1-\alpha)/(2-\alpha)$, $y = 1/(2-\alpha)$, and the value of the game is $1/(2-\alpha)$.

If we take Left to be the attacker, with strategy 1 being 'sitting duck' and 2 being 'out of the sun', Right to be the defender playing 'stare at the sun' or 'look away', then we recover the previous matrix with $\alpha = \frac{19}{20}$. So, the value of the game is $\frac{20}{21}$, Left should play sitting duck with probability $\frac{1}{21}$ and Right should stare at the sun with probability $\frac{20}{21}$.

If Left refuses to play sitting duck, then Left strategy 1 disappears, Right strategy 2 is dominated, Right should always stare at the sun and the value of the game is $\alpha = \frac{19}{20}$. The difference is $(1-\alpha)^2/(2-\alpha) = \frac{1}{420}$; realistically, it will be very hard to persuade pilots to play sitting duck for such a small advantage.

69 Clearly, if the opponent has already thrown but missed, then you should wait until the third round, when the hit is certain. Otherwise, the available strategies are to throw on the first, second or third round. Assuming a payoff of 1 if Dum wins, −1 if Dee wins, and 0 for a draw, the payoff matrix is

	Dee 1	2	3	minima
Dum 1	$-\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{3}$
2	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{3}$

3	$\frac{1}{3}$	$-\frac{1}{2}$	1	$-\frac{1}{2}$
maxima	$\frac{1}{3}$	$\frac{1}{2}$	1	

Since the lower and upper values are both $\frac{1}{3}$, they should play the corresponding pure strategies:

Dum should pass on the first round, intending to throw on round 2, unless Dee throws and misses; Dee should throw at the first opportunity, unless Dum throws and misses, and thereby wins with probability $\frac{1}{3}$; and Dum wins with probability $\frac{2}{3}$.

With phaser stun guns, since you cannot tell whether your opponent has fired and missed or just not yet fired, the strategies are simply to fire [if you still can] on the first, second or third round. Now, for example, if Dum plays 2 and Dee plays 1, the probability of Dee winning is $\frac{1}{3}$, of Dum winning is therefore $\frac{2}{3}\times\frac{3}{4}$, and so the payoff is $\frac{2}{3}\times\frac{3}{4}-\frac{1}{3} = \frac{1}{6}$. The payoff matrix is, continuing similarly:

	Dee 1	2	3	minima
Dum 1	$\frac{1}{9}$	$-\frac{1}{6}$	$-\frac{1}{3}$	$-\frac{1}{3}$
Dum 2	$\frac{1}{6}$	$\frac{9}{16}$	$\frac{1}{2}$	$\frac{1}{6}$
Dum 3	$\frac{1}{3}$	$-\frac{1}{2}$	1	$-\frac{1}{2}$
maxima	$\frac{1}{3}$	$\frac{9}{16}$	1	

so the lower value is $\frac{1}{6}$ and the upper value $\frac{1}{3}$; as these differ, a mixed strategy is indicated. Dum1 is clearly dominated, and then Dee3 is also dominated (either by inspection or by graphical solution). If Dum plays Dum2 with probability x, and Dee plays Dee1 with probability y, then the payoff is easily found to be

$$(\tfrac{1}{3}-\tfrac{1}{6}x)y+(\tfrac{17}{16}x-\tfrac{1}{2})(1-y) = -\tfrac{1}{2}+\tfrac{17}{16}x+\tfrac{5}{6}y-\tfrac{59}{48}xy.$$

The stationary value is $\frac{13}{59}$, when $x = \frac{40}{59}$ and $y = \frac{51}{59}$.

Dum should pass on round 1, should fire on round 2 with probability, $\frac{40}{59}$, and on round 3 with probability $\frac{19}{59}$; Dee should fire on round 1 with probability $\frac{51}{59}$ and on round 2 with probability $\frac{8}{59}$; Dum wins with probability $\frac{1944}{3481}$, roughly 0.56, Dee with probability $\frac{1177}{3481}$, roughly 0.34, and there is a probability $\frac{360}{3481}$, roughly 0.10, that both miss.

70 Left has four pure strategies: LL, LH, HL and HH, where, for example, LH means ‘bet low on a picture, high on a non-picture’. Similarly, Right has the four strategies CC, CD, DC and DD, where CD means ‘concede a low bet, double a high’. So the payoff matrix is (with Left winnings positive):

	CC	CD	DC	DD
LL	1	1	$2p-2q$	$2p-2q$
LH	$p+5q$	$p-10q$	$2p+5q$	$2p-10q$
HL	$5p+q$	$10p-q$	$5p-2q$	$10p-2q$
HH	5	$10p-10q$	5	$10p-10q$

where $p =$ probability of picture $= \frac{3}{13}$, $q = 1-p = \frac{10}{13}$. By inspection, LL and LH are dominated by HL and HH respectively [‘always bet high on a picture’], and in the reduced matrix CC and CD are dominated by DC and DD [‘never concede a low bet’], so this is equivalent to the 2×2 payoff matrix in the bottom right-hand corner. As the lower value, $5p-2q = -\frac{5}{13}$, differs from the upper value, $10p-2q = \frac{10}{13}$, both players should play mixed strategies. Suppose Left plays HL with probability x, Right plays DC with probability y, then the payoff is

$$V = \tfrac{1}{13}(-150xy+80x+135y-70),$$

so that V is stationary when $y = \frac{8}{15}$ and $x = \frac{9}{10}$, and for these values of x and y, $V = \frac{2}{13}$.

So Left should bet £5 on any picture, £1 on any other card with probability $\frac{9}{10}$ (and £5 with probability $\frac{1}{10}$). Right should double any £1 bet, and concede a £5 bet with probability $\frac{8}{15}$ or double it with probability $\frac{7}{13}$. Left expects to win £$\frac{2}{13}$.

71 Len has three pure strategies: *(a)* play the real fly first; *(b)* play the imitation first, then (if necessary) the real fly; *(c)* play the imitation both times. So does Rick: *(a)* swat on the first go; *(b)* pass, then (if necessary) swat; *(c)* pass both times. So the payoff matrix is (with Len's winnings positive and scaled to 1):

L\R	a	b	c
a	–1	0	0
b	1	–1	0
c	1	1	–1

(or some linear scaling from this, depending on what you think L or R is winning).

As the row minima are all –1, while the column maxima are 1, 1 and 0 (minimum of 0), no pure strategy is optimal. Also, by inspection, no strategy is 'obviously' dominated. Assume therefore that L plays (a, b, c) with probabilities x, y, $1-x-y$ respectively, while R plays them with probabilities s, t, $1-s-t$. Then the payoff is

$$V = -sx+sy-ty+s(1-x-y)+t(1-x-y)-(1-s-t)(1-x-y).$$

For stationary values, $\partial V/\partial x = 0$, *etc.*, so we get the four simultaneous equations:

$$-s-s-t+1-s-t = 0,$$
$$s-t-s-t+1-s-t = 0,$$
$$-x+y+1-x-y+1-x-y = 0, \text{ and}$$
$$-y+1-x-y+1-x-y = 0;$$

whence (E&OE) $x = \frac{4}{7}$, $y = \frac{2}{7}$, $s = \frac{1}{7}$ and $t = \frac{2}{7}$ [details omitted, but not difficult!], and so $V = -\frac{1}{7}$. *Etc.*

72 Each player has two pure strategies, to go either to the Bower or to the Valley. If they choose the 'same' strategy, then the expected payoff for Galahad is half the number of sisters at the destination, if 'opposite', then he rescues all the sisters. Setting up the matrix and showing that a mixed strategy should be played is now absolutely routine; if G goes to the Bower with probability x, and M with probability y, then the expected payoff is

$$V = 4xy/2+4x(1-y)+3(1-x)y+3(1-x)(1-y)/2,$$

and [details omitted] this is stationary when $x = \frac{3}{7}$, $y = \frac{5}{7}$, $V = 2\frac{4}{7}$. as required.

If *(a)* neither G nor M has heard about the escape, then their strategies will be unchanged, but the expected payoff is now

$$V = 3xy/2+3x(1-y)+3(1-x)y+3(1-x)(1-y)/2,$$

so with the same x and y, the payoff is reduced to $V = 2\frac{29}{98}$. If *(b)* G knows but M doesn't, then M should follow the original strategy, so we can substitute $y = \frac{5}{7}$ in the new formula for V to obtain $V = \frac{9}{14}(4-x)$. This is maximised by taking $x = 0$, so G should always go to the Valley, and his expected payoff is $\frac{36}{14} = 2\frac{4}{7}$, as before. [This is *not* a coincidence!] If *(c)* M knows but G doesn't, then G should follow the original strategy, so we can substitute $x = \frac{3}{7}$ in the new formula for V to obtain $V = \frac{3}{14}(10+y)$. This is minimised by taking $y = 0$, so M should always go to the Valley, and the expected payoff is $\frac{30}{14} = 2\frac{1}{7}$. If *(d)* both knights know, then we have a new game. However, this is obviously symmetric between B and V, so each knight should toss a coin to determine his destination, and the expected payoff for G is $\frac{1}{2}(3+\frac{3}{2}) = 2\frac{1}{4}$, intermediate [as you should expect!] between the previous two values.

73 I have four pure strategies: be 'creative' on neither company [N], on chess alone [C], on bridge alone [B] or on both companies [CB]. The authorities can investigate the chess company [C], or the bridge company [B], but not both [by assumption]. As a matrix game, this gives rise to the game [with payoffs in tens of thousands of pounds, to simplify the use of percentages]:

	C	B
N	-400	-400

C	$-400-3p$	-100
B	-300	$-400-p$
CB	$-300-3p$	$-100-p$

[It would actually be quite sensible to work this problem from the point of view of the authorities, then all the payoffs are positive, and one potent source of errors is eliminated!] Assuming $p > 0$, the row minima are -400, $-400-3p$, $-400-p$ and $-300-3p$, so the lower value of the game is -400 if $p > 33\frac{1}{3}$, and $-300-3p$ otherwise [it pays to cheat if the maximum penalty is less than the minimus tax saved]. [If $p \leqslant 0$, then my strategy CB dominates all the others—you should definitely cheat if the authorities pay you to do so!] The column maxima are [again, assuming $p > 0$] -300 and -100, so the upper value is -300. As the upper value is greater than the lower value [for $p > 0$], a mixed strategy will be in order.

As there are only two strategies for the authorities, but four for me, none of which is obviously dominated, a graphical solution is easiest—a rough sketch will suffice. Suppose the authorities 'play' C with probability x; then the payoffs are roughly as shown [the diagram shows $p = 50$].

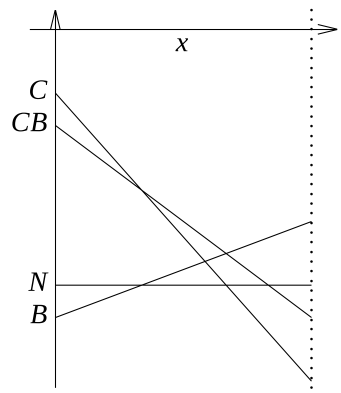

(a) In the case $p = 50$, we see from the diagram that the lowest point on the highest strategy occurs where B and CB cross, and at this point N and C are dominated. This occurs when $-150-300x = -450+150x$, or $x = \frac{2}{3}$, and the payoff is -350. [Check: at this value of x, both N and C pay -400, so are indeed dominated.] Against the strategy of checking the chess company, my strategy B pays -300 and CB pays -450, so I should play B with probability y such that $-300y-450(1-y) = -350$, or $y = \frac{2}{3}$. So, I should always 'cheat' with the bridge company, and should also 'cheat' with the chess company with probability $\frac{1}{3}$, while the authorities should check the chess company with probability $\frac{2}{3}$ and otherwise the bridge company; I expect to save £500 000 this way.

(b) As p increases, the three 'cheating' strategies move down the diagram until the point of intersection just obtained crosses the 'honest' strategy. The honest strategy is worthwhile compared with strategy C when $-100-(300+3p)x \leqslant -400$, or $x \geqslant 100/(100+p)$; with strategy B when $-400-p+(100+p)x \leqslant -400$, or $x \leqslant p/(100+p)$; and with strategy CB when $-100-p-(200+2p)x \leqslant -400$, or $x \geqslant \frac{1}{2}(300-p)/(100+p)$. [These can be read off from the diagram (generalised to arbitrary values of p).] The value of x can satisfy all three inequalities if $p \geqslant \max(100, \frac{1}{2}(300-p)$, so if $p \geqslant 100$. So honesty is worthwhile if cheating costs [at least] double tax. Note that in the case $p = 100$, the optimal value of x is $\frac{1}{2}$, so the authorities should toss a coin to decide which company to investigate, and against this strategy all four of my strategies result in a tax bill of £4 000 000.

74 Len has four possible pure strategies: *(L1)* always tell the truth; *(L2)* always lie; *(L3)* tell the truth on 'Heads', lie on 'Tails', *i.e.* always say 'Heads'; and *(L4)* always say 'Tails'. Similarly, Rick has four pure strategies: *(R1)* always believe; *(R2)* always call; *(R3)* believe 'Heads' but call 'Tails'; and *(R4)* believe 'Tails' but call 'Heads'. This gives rise to the following matrix:

	R1	*R2*	*R3*	*R4*

L1	$\frac{1}{2}$	2	$1\frac{1}{2}$	1
L2	$\frac{1}{2}$	-2	$-\frac{1}{2}$	-1
L3	1	0	1	0
L4	0	0	0	0

For example, in *L1 vs. R1*, Len wins £1 on 'Heads' and nothing on 'Tails', so his expected winnings are 50p; in *L3 vs. R4*, Len always says 'Heads' so is always called, and Len is equally likely to be telling the truth or lying, so his expected score is 0. The row minima are $\frac{1}{2}$, -2, 0 and 0 so the lower value is $\frac{1}{2}$; the column maxima are 1, 2, $1\frac{1}{2}$ and 1 so the upper value is 1; so a mixed strategy will be indicated.

We note that $L2$ and $L4$ are dominated by $L1$, so Len should choose between $L1$ and $L3$, and always tell the truth on 'Heads'. With $L2$ and $L4$ deleted, $R4$ dominates $R2$ and $R3$, so Rick should choose between $R1$ and $R4$, and always believe 'Tails'.

Suppose Len plays $L1$ with probability x and so $L3$ with probability $1-x$, while Rick plays $R1$ with probability y and so $R4$ with probability $1-y$. Then the expected payoff is

$$V = \tfrac{1}{2}xy + (1-x)y + x(1-y) + 0 = x + y - 1\tfrac{1}{2}xy,$$

so $\partial V/\partial x = 1 - 1\frac{1}{2}y = 0$ when $y = \frac{2}{3}$, and $\partial V/\partial y = 1 - 1\frac{1}{2}x = 0$ when $x = \frac{2}{3}$. So Len should always tell the truth on 'Heads' and bluff one-third of the time on 'Tails'; Rick should always believe 'Tails' and call one-third of the time on 'Heads'; and the expected outcome is a payoff to Len of £$\frac{2}{3}$.

75 First bit is bookwork. If $x = 1$, then the expected payoff is $\frac{1}{2}y(W-I)$, and if $I < W$ then this is maximised by $y = 1$; in other words, pure Hawks do better in a population of Hawks than do deer with any mixed strategy. If on the other hand $I > W$, then the payoff decreases with y, and mixed strategists do better than pure Hawks in a population of Hawks; that is, pure Hawk is not E.S.S. A population of pure Doves would be characterised by $x = 0$, with expected payoff therefore $\frac{1}{2}W(1+y)$. The payoff increases with y, so that mixed strategists do better than pure Doves ($y = 0$), and pure Dove is not E.S.S. either.

Bullies are indistinguishable from Doves except when a Bully fights a Dove, in which case the Bully always wins. So, if an E.S.S. is to play Hawk with probability p, Dove with probability q and Bully with probability r, where $p+q+r = 1$, then a male playing Bully whenever he 'should' play Dove will now expect to score W instead of $\frac{1}{2}W$ whenever he encounters a Dove, and $\frac{1}{2}W$ instead of 0 whenever he encounters a Bully; the expected gain of $\frac{1}{2}Wq(q+r)$ is strictly positive if $q > 0$. So, for an E.S.S., q must be zero, and there can be no Doves.

In such an E.S.S., Bullies are indistinguishable from Doves, so the expected payoff, from the formula given in the question, when a deer plays Hawk with probability y is

$$\tfrac{1}{2}W\left(1 - p + y - \frac{I}{W}py\right).$$

This is better than the E.S.S. strategy if

$$\tfrac{1}{2}W\left(1 - p + y - \frac{I}{W}py\right) > \tfrac{1}{2}W\left(1 - p + p - \frac{I}{W}p^2\right),$$

which simplifies (left as an exercise) to

$$(W - Ip)(y - p) > 0.$$

Since p can be neither zero nor one, from the first part of the question, either $y = 1$ or $y = 0$ will cause this inequality to be satisfied unless $W = Ip$, as required.

The payoff matrix with Retaliators is left as an exercise. If $I > W$ then [it is easily seen from the matrix that] Doves dominate Retaliators, so there can be no Retaliators in an E.S.S. If $I < W$, then Retaliators dominate Doves; and *in the absence of Doves* they *also* dominate Hawks. So, no E.S.S. can contain Doves, and therefore one cannot contain Hawks either. However, Retaliator is not an E.S.S., because in the absence of Hawks, Doves are just as fit. [So there is no E.S.S.]

76 This is actually a minor tweak on the 'bookwork' display behaviour covered in lectures. There, the cost of displaying for time t was αt; proceed the same way, replacing αt by $w(t)$, and replacing the value of winning by V, to get the expected payoff of

$$\int_0^p (V - w(q)) f(q)\,\mathrm{d}q + \int_p^\infty - w(p) f(q)\,\mathrm{d}q$$

As in lectures, for ESS, this must be independent of p; differentiating wrt p gives

$$(V - w(p)) f(p) + w(p) f(p) - \int_p^\infty w'(p) f(q)\,\mathrm{d}q = 0,$$

or

$$V f(p) = w'(p) \int_p^\infty f(q)\,\mathrm{d}q.$$

as required. In the case $w(p) = p^2$, this is

$$V f(p) = 2p \int_p^\infty f(q)\,\mathrm{d}q.$$

Differentiate again wrt p; then

$$V f'(p) = -2p f(p) + 2\int_p^\infty f(q)\,\mathrm{d}q = -2p f(p) + V f(p)/p,$$

whence $V f'/f = -2p + V/p$, which integrates directly to give $f(p) = Ap\mathrm{e}^{-p^2/V}$ for a suitable constant A, determined by the fact that $\int_0^\infty f(q)\,\mathrm{d}q = 1$. Thus, $\frac{1}{2}AV = 1$, and $f(p) = 2pV^{-1}\mathrm{e}^{-p^2/V}$. [More generally, $f = w'V^{-1}\mathrm{e}^{-w/V}$.]

77 Effectively, this is a simple application of the Principle of Indifference. At each moment, you must decide whether to throw or wait. If you throw too soon, you are likely to miss, and then you get killed. If you wait too long, your opponent is likely to throw and kill you. You throw when it doesn't matter; this is when the chance of me [say, Cain] missing, $1 - c(x)$, is the same as the chance of you hitting, $a(x)$, or when $c(x) + a(x) = 1$, as required. This strategy is optimal for both of you; if you throw earlier, then your chance of winning is less [because $c(x)$ is monotonic]; and if you throw later, then your oppponent can limit your winning chance to that by following the 'correct' strategy, but can also do better by waiting until just before you throw, in which case you have an increased chance of losing.

If Abel is optimistic, then Cain does not need to change strategy. The effect is that Abel can be expected to throw too soon, that is when $c(x) + a'(x) = 1$, and Cain's winning chance is therefore $1 - a(x) > 1 - a'(x)$ as $a'(x) > a(x)$. If Abel is pessimistic, then Abel can be expected to throw too late. Cain can retain the same winning chance as above by throwing at the correct time, but if Abel is known to be playing by the above strategy, then Cain can afford to wait until Abel is about to throw, so getting a winning chance of nearly $1 - a'(x) > 1 - a(x)$.

End