

MacCready Theory with Uncertain Lift and Limited Altitude

John H. Cochrane¹

May 4, 1999

¹Graduate School of Business, University of Chicago, 1101 E 58th St. Chicago IL 60637, 773 702 3059, john.cochrane@gsb.uchicago.edu. I thank John Good and especially John Leibacher for many helpful comments. Programs and updated copies of this paper can be found at <http://www-gsb.uchicago.edu/fac/john.cochrane/>

1 Introduction

Given a MacCready ring setting, every textbook tells you how to fly. Take thermals greater than this setting, leave thermals when they get below the MacCready setting, and adapt to lift and sink via the usual MacCready rules. Much of the mystery, challenge and art of cross-country thermal flying comes down to a judgement what that setting should be. Of course, if we knew the strength of the next thermal, that would be the MacCready setting. But what setting should you use, given the fact that the strength and position of thermals is uncertain, and you may not have enough altitude to reach them before running into the ground?

This paper presents a mathematical solution to this long-outstanding problem. I solve for the optimal MacCready settings as a function of altitude and distance to go, given a statistical model of the thermals you are likely to find, trading the probability of a landout against a faster finish according to current U.S. contest rules for distance and speed points.

Unsurprisingly, the results describe philosophies already advocated by noted contest pilots, as pilots already had figured out to fly faster on strong days and in sink before MacCready published his original article. On the other hand, a quantitative and explicit solution helps to ground this advice and may help to refine it. In particular, the analysis helps one to make sense of conflicting advice whether to start final glides low and aggressively, hoping to “porpoise in” or find one last booming thermal, or whether to fly final glides conservatively so as not to lose disastrous points by landing a mile or two short. Also, as MacCready’s theory led to a useful instrument, the MacCready speed ring, an explicit quantitative solution for an extended MacCready theory may be useful to instrument designers. Modern flight computers could easily suggest and automatically implement the MacCready settings described below. Finally, the results give an explicit and quantitative account of why lower performance sailplanes must be flown more conservatively in the same weather, which may help the handicapping process.

MacCready rules are commonly disparaged, as flying 5 or 10 kts too fast or too slow, or using a setting 1 or even 2 knots too high or low, gives a very small change in overall cross-country speed. However, the MacCready value also determines which thermals one should take and when one should leave a thermal. This is perhaps the most important decision for cross-country and contest flying, and a mistake of 1-2 kts in the MacCready value used to make decision has a major impact on cross-country speed.

2 Results

The mathematical analysis supports the following rules:

1. The MacCready setting represents the time value of altitude. For example, a MacCready setting of 4 kts means “if I were 400 feet higher, I could, on average, finish the task one minute sooner.”
2. Given a MacCready setting, adapt speed to lift and sink via usual MacCready speed to fly rules.
3. Take any thermal greater than the current MacCready setting. Leave any thermal when its value drops below the current MacCready setting.
4. Steadily reduce the MacCready setting as you get lower – fly more slowly and take weaker thermals. Steadily increase the MacCready setting as you get higher – fly faster and only stop in or stay in stronger thermals.
5. The MacCready setting now is (roughly speaking) equal to the average value of the MacCready setting in the future. Obviously, there is no point in bombing along at a 8 kt MacCready setting, passing up 7 kt thermals, if you are very likely to end up low and desperate with a 2 kt setting later on. This principle is a little more subtle when lift is uncertain. For example, suppose you are at 4000'. If you glide 5 miles and do not find a thermal, you will want to set your ring to 3 kts, and take any 3 kt or better thermal. If you do find a thermal, you expect it to be about 5 kts. If those events are equally probable, set your ring to 4 kts now. That way, the setting now is the average value of the setting in the future.
6. You will often fly with MacCready setting well below the best thermals of the day. This gives you a better chance of reaching the good thermals without having to take bad thermals or landing out. You still should take any thermals stronger than the current setting.
7. Experts sometimes advise starting final glides low and aggressively. You can often achieve a glide angle better than that in still air by porpoising, and you give yourself altitude to use a good thermal if you find one. (Doug Jacobs was once quoted in *Soaring* magazine advocating this strategy.) However, other experts sometimes advise conservative final glides, reasoning that it is worth giving up a few speed points in return for making sure you don't land out. The results support both pieces of advice, and give guidance on which one to follow, depending on the weather conditions and the pilot's objective. Broadly speaking, you start final glides aggressively but finish them conservatively.

As per point 1, the MacCready setting tells you the rate at which to trade altitude for time. The basic principle underlying points 2-5 is to spend precious altitude evenly, in climb vs. cruise and now vs. farther down the course.

The contribution of this paper is in mathematically deriving what the optimal MacCready setting *is* for a given situation, when lift is uncertain and altitude is limited. The answers reduce to the classic special cases: if you know that you can reach a (say) 4 knot thermal, then the MacCready value is 4 kts. This is the classic MacCready speed to fly computation. If you know that a MacCready setting of 6 will just exhaust your altitude at the finish, then the MacCready value is 6 kts. This is the classic final glide computation. If you know that the next thermal will be 6 kts but you cannot reach it with a MacCready setting greater than 2, then the MacCready value is 2 kts.

Case 1: 1/10 chance of 4 kts to 5000'

I start with a case that is somewhat typical of conditions in Northern Europe and the Eastern U.S. This case is very simplified and will give somewhat extreme results, but it allows a clear exposition of the forces at work and a clear comparison with classical MacCready theory.

You fly a Discus on a 150 nautical mile competition task. Your objective is to maximize the expected value of your contest points, including speed points if you finish and distance points if you land out.

Thermals extend from 500 feet to 5000 feet. In each nautical mile there is a 1/10 probability of finding a 4 kt thermal. In addition, you may gain or lose 50 feet per mile due to lift and sink along the way. (More precisely, I add or subtract a random amount to your height each mile, with a standard deviation of 50 feet. This represents some of the uncertainty of glide angles, and is a useful device for smoothing out kinks in the numerical solution described below.)

Table 1 summarizes this and the following thermal models. It is useful to think not just of the probability per mile, but the chance of finding a thermal of given strength in 10 or 20 miles, which controls the probability of landing out. In the simple model, there is still a $100 - 88 = 12\%$ chance of finding no thermal over 20 miles.

Model	Miles	Thermal strength (kts)				
		1	2	4	6	8
Simple	1	10				
	10	65				
	20	88				
Realistic	1	20	10	5	2	
	10	99	84	52	18	
	20	99.99	96	77	33	
Strong	1		15	10	5	2.5
	10		98	85	54	22
	20		99.96	98	80	40

Table 1. Thermal models. In each case, the row marked 1 gives the percent probability of finding a thermal of the indicated strength in each nautical mile. The rows marked 10 and 20 give the percent probability of finding a thermal of the given strength *or better* in 10 and 20 nautical miles respectively.

Figure 1 presents the optimal MacCready settings for this environment.

Start with the lines for 100 and 150 miles out. These tell you how to fly on course, when final glide considerations don't enter the picture yet. We see right away that you fly with a MacCready setting substantially below the best thermals of the day. Even at the top of the thermals, the ring setting is about 3.0, less than the 4.0 thermals of this day. The slightly lower setting gives more range and thus a greater chance of finding a thermal. Earlier in the flight, the chance of landing out is greater, so you fly slightly more conservatively.

You also systematically lower the ring setting as you lose altitude (i.e. if you don't find any thermals). As you get lower, the importance of reaching the next thermal increases, so you fly at a more conservative setting. You do this gradually, rather than bomb along at 4 kts until you hit 1000' and then slow down to 1 kt in desperation. The principle that the current MacCready setting equals the expected setting in the future smooths out the curves by averaging over all the possible future events.

Neither feature means that you should cruise at a lower setting than the weakest thermal you would take. For example, suppose you are 100 miles out and at 2000', so the optimal MacCready setting is about 1.2. If you found a 2 kt thermal, you should take it. However, you should not take this thermal to cloudbase. You should climb to about 3000', since the 100 mile out MacCready value has risen to 2 at 3000'. At this point, you have enough altitude to make it worth leaving the weak thermal and going off in search of something better. (In this and the rest of the analysis of the paper, I ignore the fact that it may take time to center thermals and that thermal strength may vary with altitude. Both facts modify the advice in fairly obvious ways.)




Figure 1: MacCready setting vs. altitude and distance to go for a dry Discus in the simple thermal model, with 1/10 chance of finding a 4 kt thermal each nautical mile. The dashed lines give conventional final glide calculations: the height required to finish from 10 and 20 nm out in still air at the indicated MacCready settings.

Now, let's examine the initially weird-looking behavior close to the finish, in the 10 and 20 miles out lines. The solution for this thermal model says to stay above the glideslope, even if this means flying at a much lower MacCready setting than you would have used out on course for the same altitude.

For example, suppose you found a 4 kt thermal 20 miles out. As shown in Figure 1, the optimal MacCready value for 20 miles out is 4 kts at about 4000', which is also about the same altitude as given by the conventional calculation that you could just glide home at the 4 kts MacCready speed in still air. Thus, you leave the thermal at 4000' and a MacCready setting of 4.0. Suppose that good fortune put you 20 miles out at a somewhat higher altitude, 4500'. The 20 mile out MacCready value for 4500' shown in Figure 1 is higher, about 5. You raise the ring value to 5 and finish faster.

Conversely, if bad luck or sink leaves you 20 miles out at about 3100', on a glideslope of MacCready 1.0, the calculation says to stay on this glideslope, slowing down to MacCready 1 glide and taking any thermal over 1 kt for a few turns on the way home. It says this, even though you would be flying faster (MacCready about 2) at this altitude out on course. Furthermore, (and this is the really strange part) if

you really had bad luck and end up at 2000' 20 miles out, Figure 1 says you should now speed up (relative to 3000', 20 miles out), flying at a MacCready of 2 just as you would have farther out on course.

To make sense out of this advice, remember that the objective is to maximize expected contest points. A landout one mile from the finish is very costly – you get 650 distance points instead of 1000 speed points. Hence, if you were one foot below glideslope with little chance of finding a thermal, an extra foot of altitude is worth 350 points. An extra minute is worth only about 5 points. The MacCready value – the ratio of the value of time to the value of altitude – is thus very low.

However, if you are out on course, or well below glideslope, you will have to find a thermal for sure in order to finish. If you don't find one, an extra foot of altitude only gets you $L/D_{\max}$ extra feet of distance and thus a few extra distance points. Therefore, the value of altitude much less than it is right near glideslope. If you are well above glideslope, extra altitude just means finishing a bit quicker. Thus, the value of extra altitude is lower, and the MacCready setting is higher both well above glideslope *and* well below glideslope than it is right near glideslope, when there is any chance of landing out.

You can see the advice for cautious final glides directly. In this thermal model, Table 1 shows that there is a 12% chance of finding no thermal from 20 miles out. Examine the situation from 3100'. You can either glide home at a MacCready of 1.0, or you can press on at a MacCready of 2-3, as you would on course, hoping to find another thermal. If you glide home at 1.0 and do find that extra thermal, you will have lost maybe 10-20 contest points relative to someone who pressed on aggressively and did find a thermal. If you don't find that extra thermal, you lose 350 points relative to someone who played it safe. A 12% chance of losing 350 points is a worse deal than an 88% chance of losing 20 points, given that your objective is to maximize the expected (average) number of points.

Obviously, the nature of the final glide (and the whole flight for that matter) depends a lot on the thermal model and on the pilot's objective. I look at both issues in more detail below.

Case 2: A more realistic thermal model

Next, I pursue a somewhat more complex but more realistic thermal model that captures typical conditions in Northern Illinois (where I fly), the Eastern U.S. and northern Europe. There is in fact a range of thermals; many 2 kt thermals, fewer 4 kt thermals, and rare 6 kt thermals. In addition, there are frequent instances of weak lift that one can exploit for a mile or two in porpoise mode. The second model in Table 1 gives the distribution of thermal strengths that I assume for the next calculation, roughly calibrated from my experience on decent flying days in Northern Illinois. (I hope for 4 kts, I have to stop and take 2 kts here and there to stay up, and find 6 kts

maybe once per flight. In a 5 mile range there is almost always a 1-2 kt thermal that I can use in desperation to stay up.) In addition, I specify that one only gets 50% of the indicated lift while porpoising over a mile. Thermals where I fly are, alas, not a mile across.

Figure 2 presents the optimal MacCready settings for this case. The combination of the ability to porpoise and the existence of frequent weak lift that one can use to save a flight make one able to exploit higher MacCready settings on course. At most altitudes, the MacCready settings are 0.5-1 kt higher than in Figure 1.

The MacCready setting still declines with altitude, though less strongly so. At 2000', the MacCready value is about 2.3 rather than about 1.0. The MacCready settings decline with altitude mostly because you may be forced to take weak lift when down low than from the fear of landing out after finding no lift at all. Finally, since you are almost sure to stay up by taking weak enough lift when down low, you don't have to fly more cautiously at the beginning of the flight, since pure distance to fly doesn't make you more likely to land out. The 150 mile line and the 100 mile line are nearly the same.




Figure 2: Optimal MacCready settings for the “realistic” thermal model and a dry Discus. Thermals extend from 500 to 5000 feet with the strengths given in the “realistic” row of Table 1. The dashed lines give altitudes to finish from 10 and 20 miles out at the indicated MacCready settings

Bill Bartell was quoted in *Soaring* magazine that early in his career he “didn’t know what to do” with his MacCready ring, so he set it at 2 for 2000’, 3 for 3000’, 4 for 4000’ and so on. This is almost exactly the optimal policy predicted by the model.

The greater opportunities to save a flight in weak lift and to porpoise allow one to fly more aggressive final glides. At 20 miles out, the MacCready 4 value occurs at about 400 feet below the conventional still-air MacCready 4.0 glideslope. Thus, you should leave a 4 kt thermal, 20 miles out, about 400 feet below the conventionally calculated glideslope, and fly a fast MacCready 4.0 glide. If you leave here, you are likely to gain 400’ by porpoising up in the weak lift, and you may run into one of those rare 6 kt thermals and really beat the competition. The chances of landing out are low, since 1 and 2 kt thermals to save the flight abound, and you can stretch the glide by lowering the MacCready setting. The effect that a foot is really valuable one foot below glideslope is still there – at 2500’ 20 miles out you are flying a good deal more conservatively than you would on course, but you are still flying below glideslope for your MacCready setting.

However, 10 miles out, the chances of not finding anything and the huge penalty for a one-mile out landout start to loom larger, so you plan to be on glideslope by this time, and you will turn down the MacCready and glide home slow rather than search for lift 10 miles out if you are not on glideslope. The optimal final glide hopes to gain the 400’ on a MacCready 4.0 glideslope between miles 25 and mile 10, and then glide in comfortably.

Thus, in this thermal model and with the objective to maximize expected contest points, you start final glides aggressively, but end them conservatively.

Case 3: The effects of glider performance

Figure 3 presents the same calculation for a Schweizer 1-26, to give an extreme example of the effects of glider performance. I increased the air mass speed of thermals to compensate for the Schweizer’s higher sink rate, and its ability to circle tighter. Thus, the Schweizer and the Discus both climb at the same rate.

As Figure 3 shows, the 1-26 must be flown much more conservatively, in the exact same weather! 100 miles out, the best MacCready settings are barely over 2.0, while the Discus can use a 3.5 setting. Again, this means that not only will the 1-26 pilot cruise at a lower MacCready setting, but he should take weaker thermals more often and use them higher up, where the Discus pilot can ignore weak lift or leave it low in search for better lift. In addition, the 1-26 driver is more cautious early in the flight, as the chances of landing out are greater. The 1-26 pilot also does not get to use the aggressive final glide strategy. Conditions are still “weak” enough for the 1-26 pilot that he should fly conservative final glides.




Figure 3: Optimal MacCready settings for Schweitzer 1-26, in “realistic” thermal model.

This plot is also consistent with the common observation that lower performance gliders do not do as well, and why higher performance gliders do better, than pure MacCready theory suggests. Lower performance gliders must be flown more cautiously to avoid landing out. Equivalently, the plot reveals why good handicaps give more weight to lower performance gliders than pure MacCready theory suggests. It also suggests why pilots must take some time to adjust techniques to gliders of different performance.

Case 4: Strong conditions

The third group of rows in Table 1 gives a thermal model for strong conditions. I based the probabilities on typical conditions in Uvalde Texas. Roughly speaking, I moved the strength of the thermals up by 2 kts. 4 kt thermals are common, 6 kts are frequent, and you see 8 kts occasionally. In addition, thermals rise from 500 to 9000 feet. I allow half the thermal strength for porpoising.

Figure 4 presents the optimal MacCready settings for a wet (9.0 pounds per square foot) Discus in these conditions. Again, we see the pattern of smoothly decreasing MacCready settings with altitude. You fly a MacCready of 6 at the thermal tops, since the chance of finding a 6 kt thermal from 9000' is pretty good, despite their low

probability per mile. The MacCready setting declines to about 4 at 2000’.




Figure 4: Optimal MacCready settings for a wet Discus in the “strong” thermal model of Table 1. The dashed lines give altitude to finish from 20 and 35 nm out in still air at the indicated MacCready settings.

Figure 4 now shows quite aggressive final glides. At 35 miles out, you increase MacCready settings to stay almost 2000’ below glideslope. Again the chance of porpoising in or finding a 6-8 kt thermal is worth the chance of having to make a few turns in a 2-4 kt thermal if it doesn’t work out. Again, the final glide starts aggressive and gets gradually more conservative. At 20 miles out, you plan to be only about 500 feet below glideslope. One still flies a bit more conservatively down low on final glides: at 5000’ the 35 mile out MacCready line is one knot below the value one would fly on course.

These calculations do not (yet) include the fact that thermals are stronger in the middle band, say 3000-7000’ at Uvalde. This consideration would lower MacCready settings in the mid and lower altitudes, recognizing in MacCready settings the old adage to “get high and stay high.” They also ignore the fact that lift may be weaker and farther apart late in the day. (Or absent and raining!) This consideration would also result in less aggressive final glides.

3 Extensions

The basic method is easily adapted to a number of extensions that make the calculations more complex but simultaneously more realistic.

Objectives, and winning by taking small chances of a disaster

The objective of maximizing the expected or average number of contest points gets the highest score over the “long run” but may not be the way to win contests. One can win contests by taking “bad bets” with large chances of small gains and small chances of disastrous losses. In particular, this consideration may explain why the computed final glide strategy seems a bit more conservative than that advocated by some contest pilots.

Suppose that by starting an aggressive final glide you take a $1/20$ chance of landing out, at a cost of 400 points, in return for a $19/20$ chance of gaining 10 speed points. This is a poor decision when the objective is to maximize expected contest points – 20 speed points would be an even bet. But if it does work, you win the contest, so many pilots would take it anyway.

More precisely, suppose you do this every day in a 10 day national contest. If you never land out, you get 100 extra points, which is often enough to win. If you land out one of those 10 days, you lose. The chance of landing out at least once in 10 days is $1 - (1 - 1/20)^{10} = 0.40$. Even though it lowers their *average* score, or their chance of winning an infinitely long contest, many contest pilots would take a 40% chance of coming in (say) 10th place as a result of a landout in return for a 60% chance of winning.

It gets worse. Suppose there are 51 pilots flying in a contest. 50 of them lower their average scores by taking this bet every day – a $1/20$ chance of losing 400 points in exchange for a $19/20$ chance of gaining 10 points a day. You are the lone pilot who maximizes his average score, so that you are *certain* to win a sufficiently long contest – one in which all the others eventually lose their bad bets and land out. However, the chance of *all* the other pilots landing out once in 10 days is $(0.40)^{50} = 1.27 \times 10^{-20}$. Thus it is virtually certain that one of the betters will win!

In sum, to win contests, you have to take bets that lower your average score, but give a large chance of a small score gain in return for a small chance of a large loss. This, in the end, I think is the motivation behind the advice by many contest pilots to start final glides lower and earlier than specified by the above calculations.

This consideration reveals that many contest pilots don’t in fact want to maximize their *average* score, they want to win. They will trade all the chance of finishing second instead of tenth in return for a very slightly greater chance of finishing first.

The method can accommodate this kind of objective, at a cost in numerical and conceptual complexity. It can handle any objective that is a function of speed or (equivalently) time to complete the task, any function $E(U(T))$. For example if one were flying to set a record or win a contest one might want to maximize the probability of speed exceeding a critical value – the previous record or the target pilot’s speed. That can be achieved by substituting the objective

$$\max \Pr(V > V_{\min}) = E(\mathbf{1}_{T \leq T_{\min}})$$

where $\mathbf{1}_{T \leq T_{\min}}$ represents a function that takes on the value 1 if $T < T_{\min}$ and 0 otherwise. Then, we work back as before. It will lead to a much greater probability of landing out, of course. Especially if you are a bit behind, this objective will lead you to “win or land out,” taking the insane chances that we often see on the last few days of a contest. Alas, this variation will require the greater numerical effort of keeping track of time as a state variable, and hence value function iteration may be the preferable numerical method.

Thermal models

I used simple thermal models informally calibrated to my own experience. Obviously, one can do a much more careful statistical analysis of GPS traces to get more accurate thermal models for a site and season.

I left out the fact that thermals are generally weaker down low and near the top. This leads to more conservative flying, to “get high and stay high.” I also left out the fact that thermals are stronger in the middle of the day, and weaker at the beginning and especially at the end of the day when most final glides are taking place. I left out the fact that it may take a minute or two to center a thermal. This lowers the average thermal strength, and means that the MacCready setting is the lower of the average thermal strength and its strength in the first minute after you’ve centered it. It leads you to prefer longer climbs. All of these effects can be included at the cost of some complexity, and are likely to modify the theory in obvious ways. I kept it simple here in order to make the method clear, or at least a little less obscure.

Additional strategic issues

Most decisions in cross-country and contest soaring can be determined once one knows the appropriate MacCready value. Given this insight and the dynamic programming technique, the calculations can be extended in many directions. For example, a low MacCready value means that one can take larger course deviations – course deviations are another way of trading time for altitude, and the MacCready value tells you how much altitude you should trade for time. In return, the ability to make course deviations to trade altitude for time affects the MacCready value – good opportunities for course deviations are like good thermals. Also, one could add the effects of wind and up/downwind turnpoints. The standard advice is to “go to

upwind turnpoints low,” but how low? If one extends the calculations to include an upwind turnpoint, the MacCready values are likely to move to the right downwind of the turnpoint. You don’t “arrive low,” you just become more choosy and only stop for unusually good thermals.

4 Mathematical analysis

This section details the calculation of the optimal policies. The mathematical technique is *dynamic programming*. The basic insight is to figure out a set of *state variables* – here, height and distance to go – that summarize all the information a pilot needs to make his decision. (The thermal model and glider type does not change over the course of the flight, so these are parameters rather than state variables.) Then, if we know how what the expected speed to finish is from every altitude one mile farther out, we can figure out how to trade off altitude now for speed in getting there, and thus decide whether to climb and how fast to cruise.

I use a variant on this technique, working with the values of extra height and time directly. The basic idea is simple: If we know the appropriate MacCready value, for each height, at the next mile, $\lambda(h, x)$, then we can work backward to find the MacCready value for each height at this mile, using the principle that the MacCready value now should be the expected value of the MacCready value in the future. Since we know MacCready values at the finish and on the ground (boundary conditions), we can work back to find MacCready values everywhere.

Objective

I study the objective of maximizing the expected number of contest points, i.e. the expected speed relative to the winner.

$$\max E \left(\frac{V}{V_{\text{win}}} \right) = E \left(\frac{T_{\text{win}}}{T} \right).$$

Lift. Each mile, I draw a new value of lift or sink, denoted l , independently from a distribution $f(l)$. This lift/sink is valid for the next mile. The sequence of events is, 1) arrive at mile x , 2) find out the lift or sink l that will apply between mile x and mile $x + \Delta x$, 3) decide whether to climb or cruise, 4) if you climb, keep climbing until you either decide to cruise (MacCready value grows higher than thermal strength) or you hit the top of the lift, 5) cruise to mile $x + \Delta x$ 6) repeat. The expectation operator $E(\cdot)$ applies over the random variation induced by the fact that lift isn’t known ahead of time.

Value function

The *value function* W denotes the expected speed you will achieve, given that

you are at altitude h , distance x , time t and once you have found out the strength of the lift l ,

$$W(h, x, t, l) = \max E \left(\frac{T_{\text{win}}}{T} | h, x, t, l \right)$$

The vertical bar denotes a conditional expectation, i.e. this is the best guess about your final speed given the information h, x, t, l . T_{win} is known rather than random; as long as T_{win} is independent of actions you take, maximizing this objective leads to the same optimal flight path. Before you find out the lift at mile x , the value function is

$$W(h, x, t) = \max E \left(\frac{T_{\text{win}}}{T} | h, x, t \right) = E(W(h, x, t, l) | h, x, t)$$

Speed to fly

If you decide to cruise, you must decide how fast to cruise. You do this by picking the speed over the next mile (Δx) that maximizes the value function – expected speed as of the end of the next mile (or at the end of the task if this is the last mile). The answer to this question must then be the best expected speed at the beginning of this mile.

$$W(h, x, t, l) = \max_{\{v\}} W(h', x', t') \quad (1)$$

$$h' = h + (l - s(v)) \frac{\Delta x}{v} \quad (2)$$

$$x' = x + \Delta x \quad (3)$$

$$t' = t + \Delta x / v \quad (4)$$

$$s(v) = \text{sink rate at speed } v, \text{ the polar.}$$

To find the optimum speed to fly, we set to zero the derivative with respect to the decision variable v , resulting in

$$Mc(v) = v s'(v) - s(v) = - \frac{W_t(h', x', t')}{W_h(h', x', t')} - l = \lambda - l \quad (5)$$

I use subscripts to denote partial derivatives, $W_t = \partial W / \partial t$; $W_h = \partial W / \partial h$. Equation (5) is the familiar MacCready speed to fly rule, with MacCready value $\lambda = -W_t / W_h$. The function $Mc(v)$ defined by the left equality gives the MacCready setting corresponding to speed v . The right hand side represents the relative value of height and time, which is the MacCready setting. Thus, (5) says, fly at the speed v corresponding to MacCready setting λ and current lift/sink l .

Denoting by v^* (really $v^*(h, x, l, t)$) the optimal speed, we can now work backwards to find $W(h, x, t)$ from knowledge of $W(h', x', t')$ via

$$W(h, x, t, l) = W \left(h + (l - s(v^*)) \frac{\Delta x}{v^*}, x + \Delta x, t + \Delta x / v^* \right) \quad (6)$$

$$W(h, x, t) = \int W(h, x, t, l) f(l) dl$$

Value function iteration

We know $W(h, x, t)$ at the finish, because we know the achieved speed. We know $W(0, x, t)$ before the finish, as given by contest rules for distance points. At this point, we could solve the problem by working back from the finish to find $W(h, x, t)$ along the course, using (6), and then take derivatives W_h , W_t to find MacCready values. (This is called “value function iteration.”) This approach may be useful for some of the extensions that I described in the “Extensions” section. However, it is numerically intensive, since there are three state variables.

Updating the MacCready values

We are most interested in the MacCready values, and a little more analysis shows how to work backward to find the optimal MacCready values directly. Using this approach, I am able to make an approximation that reduces the state variables to two – height h and distance x – which makes a simple grid iteration numerically feasible, and captures the most important state variables for most actual flight decisions.

We want to find W_h, W_t at point x , given knowledge of their values at $x' = x + \Delta x$. Taking the derivative of (6) with respect to height, we find that W_h and W_t must equal their future values,

$$\begin{aligned} W_h(h, x, t, l) &= W_h(h', x', t') \\ W_t(h, x, t, l) &= W_t(h', x', t'). \end{aligned} \tag{7}$$

(Proof below.)

Taking expectations and iterating forward, W_h and W_t must equal their *expected* future values at any point farther on,

$$\begin{aligned} W_h(h, x, t) &= E[W_h(h', x', t') | h, x, t] \\ W_t(h, x, t) &= E[W_t(h', x', t') | h, x, t]. \end{aligned} \tag{8}$$

Rules (7) and (8) allow us to work back from the finish on the derivatives W_h and W_t directly. Start at a gridpoint h', x', t' where W_h and W_t are known. Find the point h, x, t, l that precedes h', x', t' knowing that the pilot will fly a MacCready setting of $\lambda = -W_t' / W_h'$ in lift l . Now we know W_h and W_t at the new point h, x, t, l . Averaging over l with (8), we find W_h and W_t at the new point x, h, t .

Proof of (7):

$$W_h(h, x, t, l) = W_h' \frac{\partial h'}{\partial h} + W_t' \frac{\partial t'}{\partial h} \tag{9}$$

Using (2),

$$\frac{\partial h'}{\partial h} = 1 - (l - s(v) + vs'(v)) \frac{\Delta x}{v^2} \frac{\partial v^*}{\partial h}$$

where v^* is the optimum speed to fly at current height, lift and distance to go. Using (5)

$$\frac{\partial h'}{\partial h} = 1 + \frac{W_t(')}{W_h(')} \frac{\Delta x}{v^2} \frac{\partial v^*}{\partial h}.$$

From (4),

$$\frac{\partial t'}{\partial h} = -\frac{\Delta x}{v^2} \frac{\partial v^*}{\partial h}.$$

Plugging everything in (9),

$$\begin{aligned} W_h(h, x, t) &= W_h(') \left(1 + \frac{W_t(')}{W_h(')} \frac{\Delta x}{v^2} \frac{\partial v^*}{\partial h} \right) + W_t(') \left(-\frac{\Delta x}{v^2} \frac{\partial v^*}{\partial h} \right) = W_h(') \\ W_t(h, x, t) &= W_t('). \end{aligned}$$

Thermals

So far, I have not included the possibility of stopping to thermal rather than cruising at the MacCready speed. If you choose to thermal, you go up and not forward. Therefore,

$$\begin{aligned} W(h, l, x, t) &= W(h', l, x, t') \\ h' &= h + (l - s_{\min}) \Delta t \\ t' &= t + \Delta t. \end{aligned} \tag{10}$$

Hence, we can work downwards to find W_h, W_l from their values at greater altitudes, rather than backwards from their values farther down the course.

$$\begin{aligned} W_h(h, l, x, t) &= W_h(h', l, x, t') \\ W_t(h, l, x, t) &= W_t(h', l, x, t') \end{aligned}$$

Sensibly, when thermaling, the MacCready value (how much sooner you could finish if you were 100 feet higher) is given by the thermal strength. Taking the derivative of (10),

$$\begin{aligned} 0 &= W_h(l - s_{\min}) + W_t \\ \lambda &= -\frac{W_h}{W_t} = l - s_{\min} \end{aligned}$$

You thermal if this value is greater than the MacCready setting you would use to cruise, and you stop thermaling when the cruise MacCready setting exceeds the thermal strength.

The one difficulty occurs if you reach the top of the thermal (cloudbase) in lift stronger than the MacCready value appropriate for cruising. To see what happens in this case, suppose lift died out gradually, so lift at altitude h is given by $l g(h)$. Then

$$W(h, l, x, t) = W(h + (g(h)l - s_{\min}) \Delta t, l, x, t + \Delta t)$$

In this case you can see

$$\begin{aligned} W_h(h, l, x, t) &= W_h(h', l, x, t')(1 + g'(h)l\Delta t) \\ W_t(h, l, x, t) &= W_t(h', l, x, t') \end{aligned}$$

You can see that W_h is modified and W_t is unaffected. Sensibly, the thermal top affects the value of altitude but does not affect the value of time. Hence with a sharp thermal top ($g'(h) \rightarrow -\infty$) W_h jumps to accommodate the higher thermal value, while W_t is given by the value corresponding to the next point in cruise.

Averaging over lift, and a little approximation.

Next, we have to average over lift l as specified by (8).

$$\begin{aligned} W_h(h, x, t) &= EW_h(h, x, t, l) = \int W_h(h, x, t, l) f(l) dl \\ W_t(h, x, t) &= EW_h(h, x, t, l). \end{aligned}$$

It is annoying to have to keep track of W_t and W_h separately, rather than just to keep track of MacCready values λ . However,

$$E\left(\frac{W_t}{W_h}\right) \neq \frac{E(W_t)}{E(W_h)} \quad (11)$$

so we do have to keep track separately of the value of time W_t and altitude W_h and then divide them to find the MacCready setting. My previous statements that the MacCready setting equals the expected future MacCready setting were not quite true. This holds for each component of the MacCready setting separately. However, for most of the rules of thumb to remember from the analysis, and especially for altitudes above 2000 feet or so, an equality in (11) is a very good approximation.

Boundary condition at zero altitude (landouts)

We might try to specify no landouts, or minimize the probability of landing out. However, this leads to infinitely conservative flying. With this objective in the current setup (I have not modeled the fact that the day will end), you minimize the probability of landing out by flying at 0 MacCready setting always. One must accept a higher probability of landing out in order to go fast.

I specify the value of landing out from a simplified version of US contest rules. Landing out at x gives you the same score as finishing at a speed $0.65 \times x/X \times V_{\text{win}}$.

(I ignore the day devaluation rules by which you lose less if lots of other pilots land out.) If you were Δh higher, you could land out at $x + \Delta h \times L/D_{\max}$ further, so this is the same as finishing with extra speed given by

$$W_h(0, x, t) = 0.65 \times \frac{L/D_{\max}}{X} \quad (12)$$

An extra mile is worth less on a long task.

If you're going to land out, it doesn't matter when you do it. Time has become worthless, so

$$W_t(0, x, t) = 0. \quad (13)$$

Sensibly, the resulting MacCready setting at zero altitude is $\lambda(0, x, t) = -W_t/W_h = 0$. (In reality, time is a bit valuable when it is clear that you will land out, because the day will end. I ignore this complication. It leads to MacCready settings slightly higher than zero even when a landout seems inevitable.)

Boundary condition at the finish

At the finish $x = X$, we have

$$W(h, X, T) = \frac{T_{\text{win}}}{T}$$

One is tempted to start with

$$\begin{aligned} W_h(h, X, T) &= 0 \\ W_t(h, x, T) &= -\frac{T_{\text{win}}}{T^2} \end{aligned}$$

These values are not a good place to start a numerical procedure however, because everything becomes singular at the finish. W is not differentiable: $W_h = 0$ for positive excursions but $W_h = \infty$ for negative excursions, as you lose 350 points for being an inch short.

Therefore, I start the program one mile out. The pilot simply glides home through still air so as to exactly use up his altitude, or, if altitude is insufficient, lands out. The MacCready setting is that setting which leads to just using up whatever altitude the pilot has. Formally, if you can make it to the finish, speed and finish times are determined from

$$\begin{aligned} \frac{s(v^*)}{v^*} &= \frac{h}{\Delta x} \\ T &= t + \frac{\Delta x}{v^*}. \end{aligned} \quad (14)$$

The speed v^* depends on (h, x, t) , but I again suppress that dependence to keep the notation simple. Therefore,

$$W(h, X - x, t) = \frac{T_{\text{win}}}{T} = \frac{T_{\text{win}}}{\left(t + \frac{\Delta x}{v^*}\right)}.$$

Taking derivatives,

$$W_t(h, X - x, t) = -\frac{T_{\text{win}}}{\left(t + \frac{\Delta x}{v^*}\right)^2} \quad (15)$$

$$W_h(h, X - x, t) = \frac{T_{\text{win}}}{T^2} \frac{\Delta x}{v^{*2}} \frac{\partial v^*}{\partial h} \quad (16)$$

$$= \frac{T_{\text{win}}}{\left(t + \frac{\Delta x}{v^*}\right)^2} \frac{1}{Mc(v^*)}. \quad (17)$$

The last equality comes from differentiating (14),

$$\frac{dh}{dv^*} = \frac{v^* s'(v^*) - s(v^*)}{v^{*2}} \Delta x = \frac{Mc(v^*)}{v^{*2}} \Delta x.$$

If you *can't* make it back, then, you get distance points, and W_h, W_t are determined by (12), (13). The cutoff between making it and not is determined from the max glide angle in lift/ sink, i.e. when you fly at zero MacCready setting.

Now we face a small technical nuisance. Right at the cutoff where you can make it back at best L/D , W_h goes to infinity, as $Mc(v^*)$ goes to zero. This occurs because one inch less and you only get distance points, so you lose 350 points in an inch— W is not differentiable right at the MacCready zero glideslope. However, this is a spurious result of the discrete approximation. In reality, the value of one inch, one mile out does not become infinite because there is still some uncertainty, not recognized by the discrete approximation. Technically, though $Mc(v^*)$ may go to zero as altitude declines to the cutoff for making it back, it is a function with no linear term in h (its first term is of order h^2). Therefore $W_h = E(W_h(l)) = \int (Mc(v^*(h)) - l)^{-1} f(l) dl$ is still finite for any continuous distribution on lift l . This means that the W_h function is in fact differentiable for the “real” problem we want to solve that allows uncertainty about lift all the way to the finish.

I avoid this problem by linearly interpolating $W(h, x, t)$ between zero (650 points) and the next altitude grid above MacCready zero. With this interpolation, the integral of W_h is still W , and expected values over W_h are therefore still quite accurately represented.

Getting rid of time as a state variable

It is a great inconvenience to have the extra state variable t . If we keep this variable, it means that we have to solve the problem numerically on a 3 dimensional grid which adds greatly to the numerical difficulty.

Time does belong as a state variable, because our objective is nonlinear in time. Intuitively, if you have been really slow on course so far, the value in terms of achieved speed of saving an extra minute is lower than it would be if you had been really fast

on the course. Also, if you have flown the course so far 65% slower than the winner, landing out one mile short has no effect on your score where it is disastrous if you fly as fast as the winner. Hence, somewhat paradoxically, you fly final glides more cautiously when you are winning than when you are losing. However, it is clear that this is a small effect for most flying (above 900 points), and keeping an extra state variable around to capture it is not worth the effort.

If the objective were to minimize $E(T)$, then the time spent on course so far would have no effect on how you minimize overall expected time. That objective cannot be used, however because it is always infinite: If there is any chance of landing out, there is a positive probability of $T = \infty$, and the expected value of any variable that can be infinite with positive probability is infinite. This approach leads to zero MacCready settings always. Also, contest rules do give points for speed not time.

To simplify the solution, then, I avoid time effects and calculate MacCready values assuming you're proceeding along course at the winning speed. Formally, I approximate (15) and (17) by

$$\begin{aligned} W_t(h, x) &\approx -\frac{1}{T_{\text{win}}} \\ W_h(h, x) &\approx \frac{1}{T_{\text{win}}} \frac{1}{Mc(v^*)}. \end{aligned}$$

Since time t disappears from these functions, it disappears from all W_h, W_t at earlier points in the flight when we iterate backwards.

Now we have $W_t(h, X - \Delta x)$ and $W_h(h, X - \Delta x)$ one mile out, for each value of lift l . Notice that the MacCready setting is

$$\lambda(h, X - \Delta x, t) = -W_h/W_t = Mc(v^*) \quad (18)$$

just as you would expect.

Moving back additional steps

At this point, we are ready to move recursively back to the start. It works as follows.

1. We have $W_t(h', x')$, $W_h(h', x')$ and hence $\lambda(h', x')$, the MacCready value at each height h' at mile x' .
2. For each value of lift l and altitude h' , work backwards to find the height h at which you must start, to glide through lift/sink l at MacCready setting $\lambda(h', l')$, and arrive at height h' . This gives have $W_h(h, x, l)$, $W_t(h, x, l)$ and hence the MacCready setting $\lambda(h, x, l)$ at $x = x' - \Delta x$ and h , by property (7), if you chose to cruise.

3. For values h below that reached by working back from $h' = 0$, you get the landout values of W_h , W_l as given by (12) (13)
4. Now check for thermaling. If lift $l - s_{\min}$ is greater than the MacCready value $\lambda(h, l, x)$ computed by cruising in step 2-3, then you want to climb. You will continue climbing until the cruise MacCready calculated in 2-3 is higher than the lift, or the top of the thermal, whichever comes first.
5. You now have $W_h(h, x, l)$, $W_h(h, x, l)$ one mile farther back. Find $W_h(h, x)$, $W_t(h, x)$ by averaging over possible lift strengths. Report the MacCready values $\lambda(h, x) = -W_h(h, x)/W_t(h, x)$. Go back to step 1.