FLYING ON JUPITER: THE INTUITION OF THE MACCREADY SETTING

B. TAUB

Abstract. I develop intuition for the MacCready setting without elaborate stochastic control or numerical analysis. I express the speed to fly problem as a constrained optimization using a Lagrangian formulation; the Lagrange multiplier associated with the constraint then has an interpretation as a shadow price.

I then consider the effects of two types of uncertainty on the optimal speed to fly: thermals of random strength in known locations, and thermals of fixed strength in random locations.

Finally, I analyze the consequences of boundaries for the optimal speed to fly: the finite height of cloudbase, the ground, and the distance to the objective.

This paper began as a conversation with John Cochrane, whose detailed advice and suggestions I gratefully acknowledge. I also thank the editor, Arne Seitz, for several useful comments and suggestions.

[Filename: MacCready˙TS˙sub˙07122019]/July 14, 2019.

Notation

h	initial altitude (cloudbase)
v	airspeed while cruising between thermals
w	headwind
$v-w$	velocity over the ground while cruising at airspeed v with headwind w
$s(v)$	sink rate as a function of velocity in cruise mode
m	climb rate in a thermal, assumed net of sink rate at thermaling speed
$s(v)$	sink rate as a function of velocity in cruise mode (approximated by a concave quadratic)
x	distance to the next thermal with the required strength m (assumed fixed and known)
T	the elapsed time between setting off toward the thermal at altitude h and climbing back up to cloudbase h
X	remaining distance to the goal in a competition
λ	Lagrange multiplier for the polar curve constraint
ρ	parameter for the probability density characterizing the frequency of thermals when thermals are random
$E[\cdot]$	expectation operator
V	value function in Bellman equation expressing dynamic programming version of optimal speed problem

Introduction

One of the focuses of economics is optimization by individuals and by firms. A central feature of the analysis of the optimization problems individuals and firms face is the requirement that constraints be respected, and in the face of these constraints it is optimal to trade off choices: more wine, less cheese, or more labor, less capital. Where you finally settle for your picnic—one bottle of wine, half a kilo of cheese—depends on the price of wine in relation to the price of cheese. You can express the optimum with a simple formula that expresses a geometric fact: the tangency of the budget line to an indifference curve; in the case of a firm, the ratio of wages to rents is the slope of a tangency to the production frontier.

This notion that an optimum is expressed by a tangency carries over into other walks of life if the concept of price is generalized to the more encompassing notion of shadow prices. In this note I show as a starting point that the optimal strategy for flying a glider, whether it be a sailplane, hang glider or paraglider, can be expressed as a tangency.

Almgren and Tourin [1] provide the basic definition of the MacCready setting: it is the threshold value of net thermal strength (in meters per second or other velocity units) such that the time elapsed to a goal is minimized when flying between thermals at the optimum speed. I recapitulate and expand their intuition to demonstrate that the basic MacCready logic can be expressed as a shadow price of reducing flight duration, by expressing the optimization problem as a Lagrangian. The Lagrange multiplier then has the interpretation as a shadow price or marginal value, and has a surprisingly simple form that I believe to be previously unknown.[1]

Almgren and Tourin go on to analyze the potentially very complicated impact of stochastic thermals using a stochastic control approach. My second purpose here is to try to distill the basic intuition for the stochastic case. I do this by dividing the stochastic character of the thermals into two types. I show that the basic MacCready intuition is, surprisingly, extremely robust to the complications introduced by a stochastic environment. One's intuition is that due to risk aversion, one should slow down in a stochastic environment. In fact when parameterized realistically the MacCready optimal speed is unaffected.

Finally, in real flying situations there are boundaries that impose constraints on the optimization problem. In a competition task there is typically a distance to traverse in the minimum time; in cross-country flying there are limited hours before the sun drops and thermal activity ceases, and, finally, the cloudbase, which is the maximum attainable altitude for gliders flying legally, limits the distance a glider can fly before landing out.[2] I show that these boundaries can modify the basic MacCready velocity, but in simple and intuitive ways.

General characteristics of gliders. There are three basic types of gliders: sailplanes, hang gliders, and paragliders. Whilst not obvious from watching them flying from the ground, each has the means to

[1]The interpretation of the Lagrange multiplier as a marginal trade-off extends to physics, in which it is standard to pose problems of energy minimization with constraints; see Lemons [4], pp. 52 ff.

[2]These are the most elemental boundaries. There are other boundaries in practice, such as airspace restrictions, that can limit altitude.

control its forward speed through the air (keeping in mind that the air itself might be in motion overall due to wind). Sailplanes and hang gliders can increase their speed by tilting the nose of the craft down. Additionally, some sailplanes have flaps that can be deployed to slow the glider down. Paraglider pilots do not have the option to tilt the nose down, but they can increase the speed of the glider by pressing the "speed bar," which decreases the angle of attack, and so effects the same result as tilting the nose down, and then can also pull on the "brakes," which effectively act as flaps, to slow the glider.

In each case where the glider speeds up, a penalty is paid in terms of sink: higher speed, faster sink, due to the greater drag that develops at higher speed. Moreover the effect of drag increases with speed, so that the impact of increased speed on the rate of sink is nonlinear: sink is amplified at higher speed.[3] There is therefore a trade-off that sophisticated pilots interested in maximizing their distance, or minimizing the duration of their flight in a race, take account of in their decisions.

The intuition of the problem is as follows. Consider two extremes. In the first extreme you go as slowly as possible in order to minimize the possibility of sinking out. You reach the next thermal with high probability, and when you reach it you are very high. You climb back to cloudbase and continue.

[3]To elaborate a bit more, it might seem obvious that tilting the nose down increases sink, because now you are pointed toward the ground! But it isn't quite that simple: because you are going faster, *lift* also increases, and this increased lift offsets some of the downward sink. But the increased drag wins out.

The relationship between airspeed an sink rate is expressed by the so-called polar curve of a glider, which often is approximated by a parabola.

In the second extreme you fly as fast as possible—you want to win the competition. But because of the nonlinear effect of drag, you essentially "plunge" to the next thermal. When you arrive at the thermal you are near ground level, and so you need to spend a long time climbing back to cloudbase. The climb rate in the thermal is fixed: you can't improve it via any flying strategy beyond skillfully staying in the core.

These two extremes illustrate the main idea. In the first extreme you lose the competition because you are loafing along between the thermals. In the second your cruise speed is high, but you spend too much time climbing in thermals. There is an optimal cruising speed that trades off these two effects.

The basic MacCready logic

I begin by deriving Almgren and Tourin's objective in a more intuitive way. In this discussion I assume that the thermal strength m is net of the sink entailed by cruising in the thermal, as this sink rate is approximately constant due to the fixed circling speed. Because thermal cross-country entails repeated cruises and climbs, optimizing just one leg of this process captures the idea. Of course this approach does not yet reveal the optimal m—m will be held constant for now and I will return to this discussion later. (Moreover, the distance to the thermal and its strength m are treated as known—one could think of this as a "house thermal" that is a few kilometers ahead.)

The objective. The objective is to minimize the elapsed time T to traverse the distance to the thermal *and* subsequently climb back to cloudbase. There are two elements of this time, namely the time in cruise mode $\frac{x}{v}$, and the duration of the climb, $\frac{h}{m}$: see Figure 1. The control is the speed of travel in cruise, where the polar curve is relevant. There is also a constraint that the control v be physically feasible, that is, it is on the polar curve. Thus, the control problem is

$$\min_v \left\{ \frac{x}{v-w} + \frac{h}{m} \right\} \tag{1}$$

subject to the constraint

$$h = \left(\frac{x}{v-w} \times s(v) \right) \tag{2}$$

that is, the time spent sinking at rate $s(v)$ is in cruise mode, which lasts $\frac{x}{v-w}$ seconds.

A key aspect of the objective in equation (1) is that it *requires* that the subsequent climb back to cloudbase be undertaken after the cruise once the thermal is reached; this renders the optimization problem *stationary*.

A second key observation is that the height h is endogenous, that is, you don't automatically fly so fast that you get to the base of the thermal and then do a low save; in general you will arrive at the thermal at altitude.

I will express this problem as a Lagrangian in due course, because one can then interpret the multiplier as an appropriate shadow price, but before doing so it is illuminating to solve using brute force to establish the equivalence of the formulation with the Almgren-Tourin formulation. Substituting from the constraint, after cancellations we have

$$\min_v \left(\frac{m + s(v)}{v-w} \right) \tag{3}$$

which is Almgren and Tourin's objective. Notice that the distance x has factored out.

After some algebraic manipulation the first order condition is

$$(v^* - w)s'(v^*) = m + s(v^*) \tag{4}$$

With zero headwind this is the same as MacCready's (1954) original equation, $W + w_t = v\frac{df(v)}{dv}$ (see [5]), where w_t is the rate of climb in the next thermal, v is airspeed when cruising (the choice variable), W is the sink rate when cruising, and $f(v)$ is our $s(v)$.

We can interpret this first order condition geometrically: the right hand side is the vertical distance from the sink rate determined by the optimal cruise velocity v^* to the climb rate in the thermal, m, on the vertical axis of the polar diagram, and the left hand side is this same vertical distance as determined by the slope of a line $s'(v^*)$ times the horizontal distance to the cruise velocity v^* on the diagram, that is, there must be a tangency to the polar plot starting from the climb rate m.

The tangency is illustrated in Figure 2. With a positive headwind one simply shifts the polar curve to the left by the headwind; the tangency point moves to the right along the polar curve, so that the optimal airspeed (speed to fly) v^* increases.

The Lagrangian formulation. The way to think about the optimal speed to fly is to think (somewhat counterintuitively) in units of "seconds (gained) per unit of altitude," that is, in terms of $\frac{1}{v}$, not v. I now demonstrate that this perspective derives from the interpretation of the Lagrange multiplier associated with the constraint in the speed-to-fly optimization problem. The Lagrangian is

$$\min_v \left\{ \frac{x}{v-w} + \frac{h}{m} - \lambda \left(h - \left(\frac{x}{v-w} \times s(v) \right) \right) \right\} \tag{5}$$

where we present the usual formulation with the constraint term subtracted. Notice that the starting altitude h is the "income" term in the constraint, so it is already known that the multiplier λ will be the marginal increase in time per unit of altitude, that is, with units $\frac{1}{v}$. Also note that we could if we wanted re-phrase the problem with $\frac{1}{v}$ as the control, in which case it would become almost linear.

After algebraic manipulation the first order condition is

$$\lambda = -\frac{1}{s(v) - (v-w)s'(v)}$$

Because the condition (4) holds at the tangency point, we can write

$$\lambda = \frac{1}{m}$$

Because m is the climb rate of the thermal it is in units of velocity, and therefore the Lagrange multiplier is in units of the inverse of velocity, that is, seconds per meter. The interpretation is then straightforward: the multiplier is the shadow cost of a unit of duration in velocity, or, the marginal increase in altitude that is needed to justify a marginal reduction in the duration of the flight from speeding up by a marginal amount.

Notice that there is no effect of the headwind w on this shadow cost of airspeed!

Solving for v. Rearranging the geometric property yields

$$(v-w)s'(v) - s(v) = m = \frac{1}{\lambda}$$

Analyzing the left hand side geometrically reveals an inverted parabola shape, and the intersection with the constant m happens in two places; using only the left intersection, it is clear that increasing the climb rate m increases the optimal cruise *airspeed* v^*. Equivalently, increasing the net thermal strength m reduces the shadow cost λ of increasing speed, so you speed up.

Interpretation of the multiplier λ. Taking the derivative of (5) with respect to h yields

$$\frac{\partial(-T)}{\partial h} = \lambda$$

that is, λ is the marginal *decrease* in duration given a marginal increase in height, that is, it is the marginal value of height, which makes intuitive sense.

Allowing thermal strength m to vary. Now we can ask, what if there are two m's, one weak and one strong? The weak ones occur more often so the temptation is to take the weak thermals as well as the strong ones.

To think more analytically about this we would draw two tangent lines, resulting slow speed for the weak thermals (and smaller x for those thermals) and higher speed for the strong thermals.

But now it is obvious that you would slow yourself down by taking any of the weak thermals. As long as your objective is minimum time, you should cut out the weak thermals if it is physically possible to fly the strong thermals. The fact that they occur less frequently, so that the distance x is larger, is irrelevant to the strategy: x does not appear in the first order condition!

Stochastic thermals

Of course in real flying situations there is a degree of risk: if you don't find a thermal when you need it, you have to land out. If you are not in a race but want to maximize your cross-country distance (a common objective in paragliding) then you fall short.

There are two separate types of randomness to contemplate. The first random element is the strength of a given thermal. To envision this, one can posit that all thermals are "house" thermals, that is, thermals that are known to predictably develop at known locations: a knoll, a particular field, a rock formation, and so on. In that case, the pilot would be expected to adjust his MacCready threshold downward, given that he might encounter *two* weak thermals in a row: if he skips the first thermal because it is too weak by basic MacCready reasoning and then encounters a *second* weak thermal, he will land out. To prevent this he must slow down a bit on the glide/cruise phase; this reasoning here sounds a bit like risk aversion.

This reasoning is incorrect. In fact, the basic MacCready reasoning continues to apply, but with an average replacing the net thermal strength, m.

The second type of randomness is randomness in the geographical distribution of thermals. With this perspective, one could imagine that all thermals have the same strength, but that they are encountered randomly (that is, with Poisson arrivals in terms of geographical spacing). Again, the intuition is that you would slow down below the basic MacCready setting because you don't know if you will make it to the thermal after the current one. And again, there will be a "risk premium" driving the slowdown. In this case the risk premium could be expressed in terms of the arrival rate of the thermals, which behaves mathematically like an interest rate and is in this sense complementary with the idea of a risk premium.

First case: Evenly spaced thermals of random strength. This case turns out to be remarkably simple: it is identical to the basic deterministic MacCready model. The objective is

$$\min_v E\left[\left\{\frac{x}{v}+\frac{s(v)x}{mv}\right\}\right]=\min_v\left(\frac{\frac{1}{E\left[\frac{1}{m}\right]}+s(v)}{v}\right)xE\left[\frac{1}{m}\right] \tag{6}$$

But this is simply the original deterministic MacCready problem with $1/E[1/m]$ replacing m. The solution of the problem is then identical to the deterministic problem, but with the expected value of the inverse of net thermal strength replacing the fixed value. Because the expectation is of the inverse, there is however a bias toward *decreasing* speed. (That is, if there are two thermal strengths m_1 and m_2, each with 50-50 probability, then it is easy to see geometrically that $1/E[m] < E[1/m]$ due to Jensen's inequality, so the effective average thermal strength, $\frac{1}{E\left[\frac{1}{m}\right]}$, is less than the actual average $E[m]$. This shifts the MacCready tangency point to the left. But this slowdown effect has nothing to do with risk aversion!)

Second case: Randomly located thermals of fixed strength. If thermals are assumed to have equal strength but occur at random intervals then the probability density of the next thermal distance x' is $\rho e^{-\rho x'}$, that is, there is an arrival rate associated with *distance*, not time, and we need to integrate over this density.

Benchmark approximations. We can get some intuition about the random spacing case by ignoring the boundary formed by the ground, that is, we can assume that you are allowed to have negative altitude so there is no concern about landing out; this is what thermalling on Jupiter would be like! In that case the objective is

$$\begin{aligned}\min_v E\left[\frac{m+s(v)}{v}\frac{x}{m}\right] &= \min_v \int_0^\infty \rho e^{-\rho x'}\left[\frac{m+s(v)}{v}\frac{x'}{m}\right]dx' \\ &= \min_v \frac{m+s(v)}{v}\frac{1}{m}\int_0^\infty \left[\rho e^{-\rho x'}x'\right]dx' \\ &= \min_v \frac{m+s(v)}{v}\frac{1}{m}\frac{1}{\rho}\end{aligned} \tag{7}$$

that is, you integrate over the infinite line along which you are traveling.

It is immediately evident that the first order condition simply replicates the first order condition from the deterministic problem; the distance x to the next thermal is replaced by the expected value of this distance, $\frac{1}{\rho}$, but as with the deterministic case this expected distance has *no* effect on the optimal velocity; the MacCready optimum is not affected by the uncertainty! The only thing that can potentially influence the optimal velocity is the boundary.

What is going on here intuitively? It might seem that since you can have negative altitude you would just go as fast as possible. But remember that this formulation still requires you to climb in any thermal you reach, and this takes more time if your altitude is negative. So it is better to optimize the trade-off between cruising speed and the time it takes to climb in a thermal.[4]

[4]If you were flying on Jupiter, there would be no ground, and therefore you need never land out, but this therefore means that you never need to ascend in a thermal at all—so the requirement that you ascend in some thermals matters!

The effects of boundaries

We have seen that the MacCready logic is surprisingly robust to the presence of stochastic thermals, whether it is their strength or their geographic incidence that is random. However, to obtain these results we have ignored the effects of boundaries.

There are two main types of boundaries to consider. The first and most important is the cloudbase, which influences the possibility of landing out. Because gliders cannot legally (or safely) fly above the cloudbase, and because the cloudbase is roughly constant during a typical flying day[5], it is the maximum altitude in the practical version of MacCready analysis, and affects the possibility of landing out.

The second boundary is the physical distance to the goal of the task; the closer you are to the goal, the bigger the effect. Again, the desire to avoid landing out will affect the optimum. As both Almgren-Tourin and Cochrane [2] establish, these effects are potentially extremely complicated, in part due to the complicated scoring schemes used in competitions, as well as the dynamic programming effects of the boundaries. However, here as well there is some basic intuition that is easy to establish.

The impact of cloudbase when thermals are random across space. We will analyze the first boundary effect, the effect of the cloudbase altitude constraint, on the situation when thermals are randomly distributed geographically, but have equal strength.

[5] "Roughly" because the cloudbase tends to slowly increase during the course of the day. Also, whilst the cloudbase is fixed in the short run, traversing terrain of varying altitude changes the *effective* cloudbase.

To complete the formulation of the objective in the basic deterministic case, we needed to assume that the thermal height parameter h was exogenous, but now we will allow it to vary. When thermals are stochastically spaced, the altitude at which you arrive at a thermal depends on the arrival time: the longer you fly before you find a thermal, the more you have to climb back to cloudbase, that is, the bigger is h. We can express this using the previous constraint (2). If x is big, that is if you go for a very long distance without finding a thermal, you will land out. If we assume that the initial position is at the top of the current thermal, that is, at cloudbase, then if you sink more than the cloudbase altitude, you have by definition landed out. Call the cloudbase height $\overline{h}$. Then flight continues if

$$x \leq \frac{\overline{h}v}{s(v)}$$

otherwise you have landed out.

The conditional distribution and expected distance when there is landing out potential. Given that the possibility of landing limits the distance you can potentially fly starting from cloudbase, this imposes a restriction on the conditional distribution of the location of the next thermal. The truncated density is

$$\frac{\rho e^{-\rho x}}{1 - \int_{\frac{\overline{h}v}{s(v)}}^{\infty} \rho e^{-\rho \xi} d\xi} \tag{8}$$

Thus, the conditional density is normalized to the maximum possible distance traveled, and, importantly, this distance is endogenous to the velocity v. We can see now that this will have an impact on the optimization problem, as the distance x no longer will factor out of the objective.

There is a caveat here however: if one thinks of minimizing duration subject to not landing out, it is not appropriate to normalize the density. That is, if you speed up radically this will shorten your land-out distance and increase your potential to land out. Thus, the normalization should be removed. The expected distance to the next thermal is then

$$E[x] = \int_0^{\frac{\overline{h}v}{s(v)}} \rho e^{-\rho x'} x dx = \frac{1}{\rho}\left(1 - \left(1 + \rho\frac{\overline{h}v}{s(v)}\right) e^{-\rho\frac{\overline{h}v}{s(v)}}\right) \tag{9}$$

I next examine the impact of this expected value on the optimization problem.

Stationarity induced by the landing out boundary. One can now apply this conditional expected value to the cross country problem. The first observation is that if the flying day is potentially infinite, and the course has no geographical limits, then the optimization problem is stationary. Landing out is then the only concern. In that case the model is fully recursive, that is, you start over once you are at cloudbase. In that case, you can solve the non-recursive problem, which is to minimize the time to the next thermal *conditional on the requirement that you climb in the next thermal.*

Suppose you have the option of continuing to fly without the sun going down, as long as you don't land out/bomb out: you are on the infinite plane (or even line). But you want to go as far as possible. This then is equivalent to maximizing your average speed v, subject to the constraint that you have to fully ascend any thermal you encounter.

This problem is straightforward to write down as a non-dynamic-programming problem: you just integrate over the density of the thermal arrival, subject to not landing out. We already thought about this, except we didn't maximize speed, we minimized time. But they are equivalent problems. Thus, the stationary problem can still be expressed as the standard MacCready problem with uncertainty:

$$\min_v \left(\frac{1}{v} + \frac{\left(\frac{s(v)}{v}\right)}{m} \right) \frac{1}{\rho} \left(1 - (\rho \frac{\overline{h} v}{s(v)} + 1) e^{-\rho \frac{\overline{h} v}{s(v)}} \right) \tag{10}$$

If the arrival rate for the thermals is too low, or if the cloudbase is too low, or if the climb rate m is too low, then the expected value of x, and the objective function, become bell shaped, and thus non-convex, and so it is optimal to fly as fast as possible, not at the MacCready speed, simply because you have nothing to lose.[6]

However, with a high cloudbase, or strong thermals, or very frequent thermals so that ρ is large, the effect is to flatten the expected value of x, that is the truncation of the density becomes irrelevant, and the expected value is $\frac{1}{\rho}$, so when taking the derivative of the product of the original objective and the conditional expected value, then the derivative of the product is essentially that of the original MacCready objective, and the original MacCready velocity is more or less optimal.

[6]This effect shows up in models where the goal is to maximize survival, which is similar to the problem we face here; see [6].

That is,

$$0 = \frac{d}{dv}\left(\frac{1}{v} + \frac{\left(\frac{s(v)}{v}\right)}{m}\right)\frac{1}{\rho}\left(1 - (\rho\frac{\overline{h}v}{s(v)} + 1)e^{-\rho\frac{\overline{h}v}{s(v)}}\right) = \frac{d}{dv}\left(\frac{1}{v} + \frac{\left(\frac{s(v)}{v}\right)}{m}\right) + 0$$

(Just to emphasize: the "+0" term on the right is a *numerical* result that applies in reasonable calibrations.)

We can calibrate this model with a polar curve for a typical paraglider that has speeds on the order of 30 kilometers per hour, typical thermal strength of about 2 meters per second, and a polar curve with minimum sink on the order of one half to one meter per second, and typical flying conditions with cloudbase at 1,000 meters. Care must be taken to relate meters per second (m and $s(v)$) to kilometers per hour (v); the conversion ratio is

$$(1000 \text{ m/k})/(3600 \text{ sec/hour}) \times v \text{ k/hour} = 1/3.6 = .278,$$

that is, multiply velocity in km per hour times .278 to obtain meters per second. The remaining issue is the arrival rate of thermals; if there is a thermal every 500 meters on average, that is an arrival rate ρ of 2, and with this calibration the truncated expected value function is extremely flat in the speed range that matters. So the MacCready reasoning is basically unchanged.

I use this calibration, that is, expressed in meters and seconds, in the following simulations: $m = 2$, $\rho = .002$, and $\overline{h}$ described below. The standard approach seems to be to measure velocity v in kilometers per hour, whilst measuring the sink rate $s(v)$ and the thermal strength m in meters per second; this leaves the measure of the cloudbase height $\overline{h}$

ambiguous: if it is in meters then a reasonable calibration is 1000 meters. The result is shown in Figure 3. The basic MacCready construction is unaffected, because the expected value of x (orange line)—the distance to the next thermal—has a flat region around the minimum of the basic MacCready objective (blue line).[7]

Boundary effects in the competition task: close to goal. The Bellman equation has a flow element and a continuation value. The continuation value should reflect not only the influence of the optimization problem for the subsequent thermal, but for the possibility of reaching the goal and also of landing out. But the basic construction doesn't mesh with these boundary issues.

The intuition of the goal boundary is that if you are close to the goal, you speed up because the risk of landing out becomes nil. But there is another more subtle effect here. When we optimized the non-stochastic model, we built in the *requirement* that you re-ascend in the thermal you hit. This shapes the optimization significantly: even if you were at cloudbase only 50 meters from goal you would use the MacCready

[7]Notice that we have the ratio $\frac{h}{m}$ in the original objective; thus, if we are measuring m in meters per second, then h should be in meters, but the result is in seconds. The other term, $\frac{s(v)}{v}$, mixed meters per second ($s(v)$) and kilometers per hour (v).

In the Amgren-Tourin paper, which is oriented toward sailplanes, everything is measured in terms of meters per second; their polar plot figure has a MacCready speed of 50 meters per second. Dividing by .278 yields 180 kilometers per hour (the Pegasus has a maximum speed of 133 knots or 246 kilometers per hour—very fast, way faster than any paraglider!) A typical paraglider speed of 30 kilometers per hour converts to $.278 \times 30 = 8.34$ meters per second.

velocity, and this is obviously incorrect: it would be optimal to go at maximum speed. So the terminal value needs to reflect this. Moreover, this effect will have a recursive influence: it might be wise to speed up above the MacCready speed after the thermal *before* the last thermal because the risk is reduced.

We can rewrite the objective so that height is a function of the distance to goal, X, remaining:

(11)
$$V(X|m) = \min_v \int_0^X \rho e^{-\rho x} \left[\left(\frac{1}{v} + \frac{\left(\frac{s(v)}{v}\right)}{m} \right) x' + V(X - x \Big| m) \right] dx$$

Note that the probability measure that has zero support for distances beyond goal due to the normalization. It is therefore automatic that distances beyond the goal are not counted.

So, we can say that if your distance to goal is less than the land-out distance, then the optimum would be to speed up, choosing a v^* so that equality holds:

$$X = \frac{\overline{h} v^*}{s(v^*)} \tag{12}$$

Now we have a way to characterize the terminal value: it is the time it takes to fly using this rule:

$$V\left(X \Big| X \leq \frac{\overline{h}\overline{v}}{s(\overline{v})}\right) = \frac{X}{v^*}$$

where v^* solves equation (12), and where $\overline{v}$ is the maximum attainable speed (well above the MacCready speed).

We can take this further. Suppose there are N thermals before goal, labeled in dynamic programming countdown fashion as $M_N, M_{N-1}, \ldots, M_1$. Think about the penultimate thermal, M_2, with the final thermal, M_1, close to goal. We can anticipate reaching the final thermal from the penultimate thermal, and we can cook the altitude in the final thermal so that we only just squeeze into goal before landing. Should you speed up after thermal M_2? Yes! You can use the same reasoning as above: First of all, to squeeze out the the final run to goal at the maximum speed you only need to attain altitude $h^* = \frac{Xs(v^*)}{v^*}$. But now you can treat point X—the location of the final thermal M_1—as the goal. When you climb when you reach M_1, you only need to climb to h^*, and so the optimization problem can be modified to reflect this. Thus, there *is* a dynamic programming effect from the goal boundary, which iterates back through the earlier thermals.

This reasoning will also affect the stochastic model. In the stochastic model it is equally true that if you are close enough to goal you will speed up and behave deterministically.

Intuition for boundary effects in the cross-country problem. So far we have modeled landing out as a Poisson arrival problem, but indexed to distance rather than time. Intuitively, you are flying on some flats. And what we see is that once we put in the objective of not landing out between thermals, the convexity of the objective seems to be ruined: the problem becomes ill-posed.

The intuition for the ambiguity of the objective is easy to state: you can go far only if you don't land out. If you don't land out you want to go as fast as possible, as this maximizes your distance (again, because the day will end and you will then have to land). But to avoid landing out you might need to slow down. So the logic of the MacCready speed is vitiated. So the question becomes, how can a coherent objective be stated?

Let's simplify the problem as follows. Instead of thermals appearing randomly, suppose that they appear at fixed, discretely and evenly spaced locations, but not reliably. Thus, there might be a house thermal but the thermal turns on and off at random.

The simplest case is where you have two such intermittent-thermal locations ahead of you; call these M_1 and (counting down beyond M_1) M_0. If you catch both of them you can continue and go far. If you catch M_1 then you can also catch M_0 even if you fly fast (i.e. at the MacCready speed), although if it isn't active then you land out, but at least you made it that far (something like making it over the pass). But if you fly fast before M_1, and M_1 isn't active when you arrive, then you have to land out. On the other hand, if you fly slow, then if M_1 isn't active when you arrive, you still have enough altitude to make it to M_0 and at least have another chance.

This logic shows that if distance is the objective, you want to slow down. So the objective needs to be reformulated with distance at the objective. But this leads to another issue: the solution to this problem is trivial: just slow down to the minimum sink speed. To get beyond this trivial approach we need to add an additional ingredient.

The additional ingredient is that the sun goes down after a fixed number of hours, limiting the distance. So if the objective is to maximize distance given the time constraint, and also given the not-landing-out constraint, we can express this as an objective.

Boundary effects in the XC problem: you speed up at the end. With the intuition that you want to solve a constrained maximization problem rather than a minimization problem, we can see intuitively that as the end of the day grows near, you speed up, as there is now nothing to lose once the day has terminated. Thus, just as in the minimum time to goal problem, you speed up at the end!

Conclusions

I developed intuition for the MacCready setting without elaborate stochastic control or numerical analysis. (i) I motivated more clearly the objective function in the basic MacCready setting. (ii) I expressed the optimization problem as a constrained optimization using a Lagrangian formulation; the Lagrange multiplier associated with the constraint then has an interpretation as a shadow price. I demonstrated that this shadow price, which is the shadow price of reducing the duration of the flight marginally (thus reaching the goal sooner and potentially winning a competition) requires a marginal increase in altitude. When flying optimally this price—expressed as seconds per meter—is exactly the inverse of the net thermal strength; I believe this result to be previously unknown. (iii) I then considered the effects of two types of uncertainty. For the first type of uncertainty—thermals of random strength in known locations—I demonstrated that the MacCready logic is entirely preserved, and that it might be desirable to slow down, but this result is in no way driven by risk aversion. In the second case, thermals of fixed strength in random locations, I show that for very plausible parameters for real gliders that the MacCready velocity is entirely unaffected. (iv) Finally, I examined the effect of boundaries on the speed to fly, establishing that under some circumstances it is optimal to speed up beyond the MacCready velocity.

References

[1] Almgren, Robert, and Agnes Tourin (2014), "Optimal soaring via Hamilton-Jacobi-Bellman equations," working paper, New York University.

[2] Cochrane, John (1999), "MacCready theory with uncertain lift and limited altitude," *Technical Soaring* 23(3),88-96.

[3] Cochrane, John (2011), "Deviations," working paper, University of Chicago.

[4] Lemons, Don S. (1997), *Perfect Form.* Princeton: Princeton University Press.

[5] MacCready, Paul (1954), "Optimum airspeed selector," *Soaring*, March-April, pp. 8-9.

[6] Radner, Roy and Mukul Majumdar (1991), "Linear models of economic survival under production uncertainty," *Economic Theory*1, 13-32.

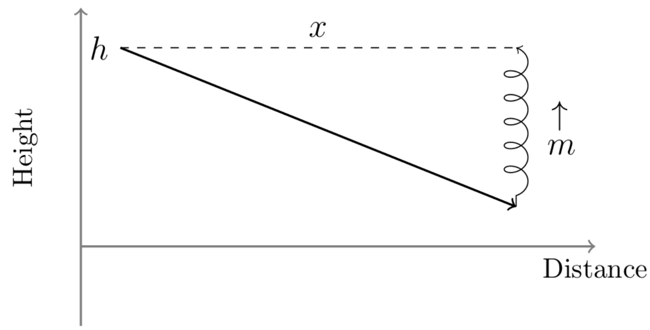

FIGURE 1. A cruise and climb leg. Repeated flight pattern of sink during cruise from starting height, followed by ascent in thermal to starting height.

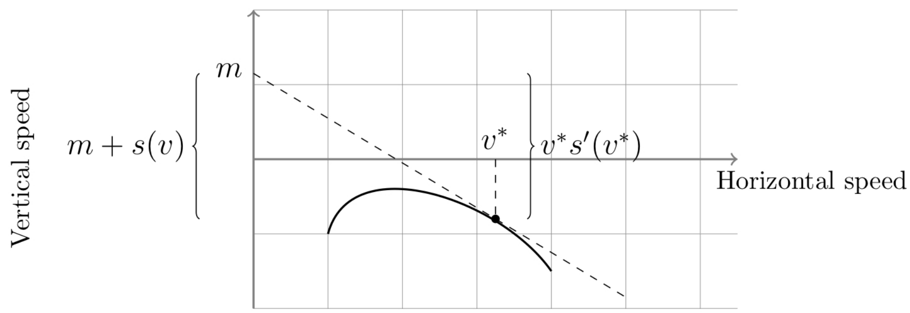

FIGURE 2. The tangency to the polar with a negative polar plot representation.

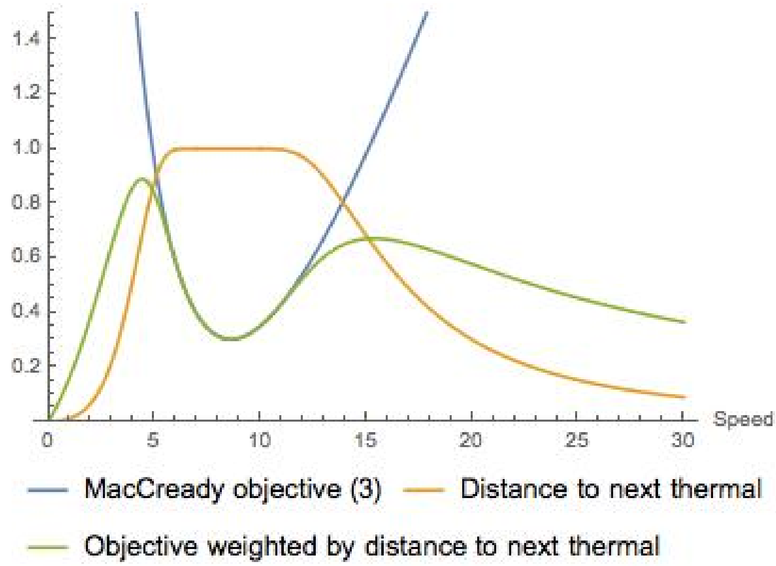

FIGURE 3. Geometry of the first order condition. The minimization problem with stochastic distance to the next thermal, with $m = 2$, $h = 1000$, $\rho = .002$ (units are in meters and seconds). The orange line is the expected value of the distance to the next thermal; the blue line is the original MacCready objective; the green line is the objective weighted by the expected distance to the next thermal. The green line and orange lines are multiplied by ρ to maintain scaling with the blue line. It is evident that the expected-value-weighted objective coincides with the original MacCready objective in the relevant region.