Optimal Soaring
via Hamilton-Jacobi-Bellman Equations

Robert Almgren* and Agnès Tourin†

February 20, 2014

Abstract

Competition glider flying is a game of stochastic optimization, in which mathematics and quantitative strategies have historically played an important role. We address the problem of uncertain future atmospheric conditions by constructing a nonlinear Hamilton-Jacobi-Bellman equation for the optimal speed to fly, with a free boundary describing the climb/cruise decision. We consider two different forms of knowledge about future atmospheric conditions, the first in which the pilot has complete foreknowledge and the second in which the state of the atmosphere is a Markov process discovered by flying through it. We compute an accurate numerical solution by designing a robust monotone finite difference method. The results obtained are of direct applicability for glider flight.

1 Introduction

Competition glider flying, like other outdoor sports such as sailboat racing, is a game of probabilities, and of optimization in an uncertain environment. The pilot must make a continuous series of strategic decisions, with imperfect information about the likely conditions to be encountered further along the course and later in the day. In a competition, these decisions are made in order to maximize cross-country speed, and hence final score in the contest; even in non competition flying the pilot is under pressure to complete the chosen course before the end of the day.

Optimal control under uncertainty is, of course, an area of extremely broad application, not only of sport competition, but for many practical problems, in particular in the general area of finance and investments.

Another area of broad mathematical interest is that of free boundary problems. These are common in optimal control problems, where the boundary delineates a region of state space where it is optimal to take a different action: exercise an option, sell an asset, *etc.* The glider problem also has an optimal exercise boundary, corresponding to the decision whether to stop and climb in a thermal.

*Quantitative Brokers LLC and New York University Courant Institute of Mathematical Sciences, almgren@cims.nyu.edu

†New York University Polytechnic School of Engineering, Department of Finance and Risk Engineering, atourin@poly.nyu.edu

An entire area of mathematical research—viscosity solutions—has been developed in order to solve these problems with various forms of nonsmooth constraints. In this paper, we do not justify rigorously the applicability of these mathematical techniques. However, the theoretical results available provide important insight into this problem, which in turn gives a very concrete example to illustrate the techniques in practice. For an introduction to the theory and some of its applications, we recommend Fleming and Soner [1993] and Touzi [2013].

1.1 How a glider flies

A modern glider, or sailplane, is a remarkably sophisticated machine.[1] It has a streamlined fiberglass fuselage, and carefully shaped wings, to cruise at speeds in the range of 100–150 km/hr, with a glide ratio approaching 40: with 1000 meters of altitude the aircraft can glide 40 km in still air. Cross-country flights of 300–500 km are routinely attained in ordinary conditions. With this level of performance human intuition and perception are not adequate to make strategic decisions, and today most racing gliders carry sophisticated computer and instrument systems to help optimize in-flight decisions. But the inputs to these systems still come largely from the pilot's observations of weather conditions around him and conditions to be expected on the course ahead.

The glider is launched into the air by a tow from a powered airplane (most frequently—launching on a ground-based cable is also possible). Once aloft, the pilot needs to find rising air in order to stay up, and to progress substantial distances cross-country. In the most common type of soaring, this rising air is generated by thermal convection in the atmospheric boundary layer: no hills and no wind is necessary. (Soaring on ridges, and on mountain waves, is also possible but is not discussed here.)

Thermal lift happens when a cold air mass comes in over warm ground. Instability in the air is triggered by local hot spots, and once released, the air parcels continue to rise until they meet a thermal inversion layer. If the air contains moisture, the water vapor condenses at an altitude determined by the initial humidity and the temperature profile; this condensation forms the puffy cumulus clouds typical of pleasant spring and summer days. From the pilot's point of view, these clouds are useful as markers of the tops of thermals, and they also impose an upper altitude limit since flight in clouds is illegal. This altitude is typically 1–2000 m above ground in temperate regions, and as much as 3–4000 m in the desert. By flying in circles within the rising air, the glider can climb right up to this maximum altitude, although of course while climbing no cross-country progress is being made. Clearly, in order to complete a 300 km flight, the glider will need to climb in many different thermals, and one of the biggest determinants of cross-country speed is the proper choice of thermals to use: different individual thermals have different strengths.

In such unstable conditions, nearly the entire air mass is in vertical motion of one sign or another; the organized structures described as thermals are only the largest and most conspicuous features. The air is of course moving down at some points in between the rising masses, but in addition, there are vertical motions of smaller magnitude both

[1] More information about gliders and the sport of soaring is available from the Soaring Society of America, `http://www.ssa.org`.

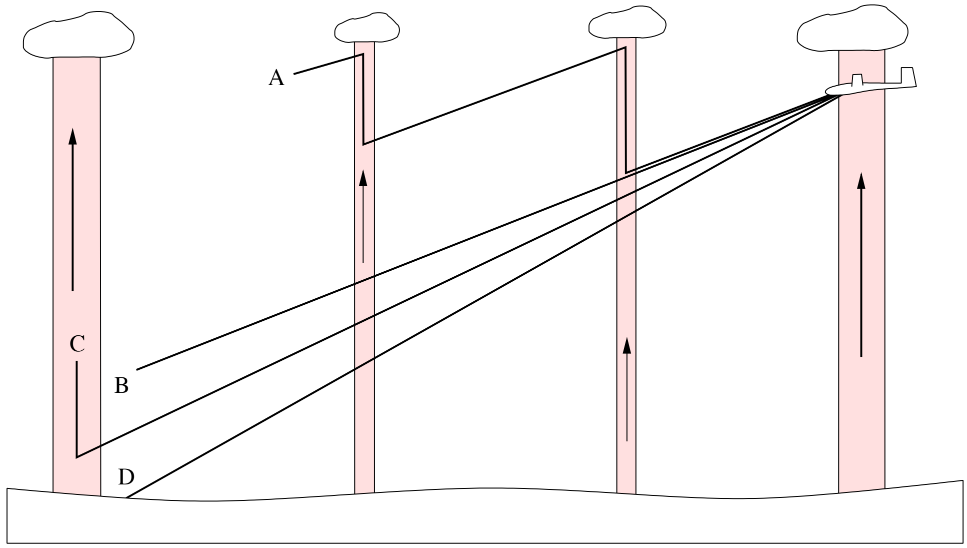

Figure 1: Gliding between thermals. Four gliders leave the thermal on the right at the same time; the trajectories shown cover the same length of elapsed time. Glider A stops in every weak thermal; he has low probability of landout but achieves a low total speed. Glider B heads toward the strong thermal at the speed of best glide angle, with minimum loss of height. Glider C flies faster: although he loses more height in the glide, the rapid climb more than compensates. Glider D flies even faster; unfortunately he is forced to land before he reaches the thermal. (Modeled on an original in Reichmann [1975].)

up and down. By exploiting these smaller motions during the "cruise" phase of flight, the pilot can greatly extend the distance of glide before it is again necessary to stop and climb. Generally speaking, the pilot slows down when passing through rising air and speeds up in sinking air; if the fluctuations are large enough then it can be possible to advance with almost no loss of altitude.

Since the air motion encountered in a particular flight is largely random, there is no guarantee that lift of the needed strength will be encountered soon enough. It is perfectly possible, and happens frequently, that the pilot will wind up landing in a farm field and will retrieve his aircraft by car. The probability of this happening is affected by the pilot's strategy: a cautious pilot can greatly reduce his risk of landout, but at the cost of a certain substantial reduction in cross-country speed.

Soaring competitions are organized around the general objective of flying large distances at the fastest possible speed. Although the details of the rules are bewilderingly complex, for the mathematics it will be enough to assume that the task is to complete a circuit around a given set of turn points, in the shortest possible time. Furthermore, a penalty for landout is specified, which we crudely approximate in Section 3.1 below.

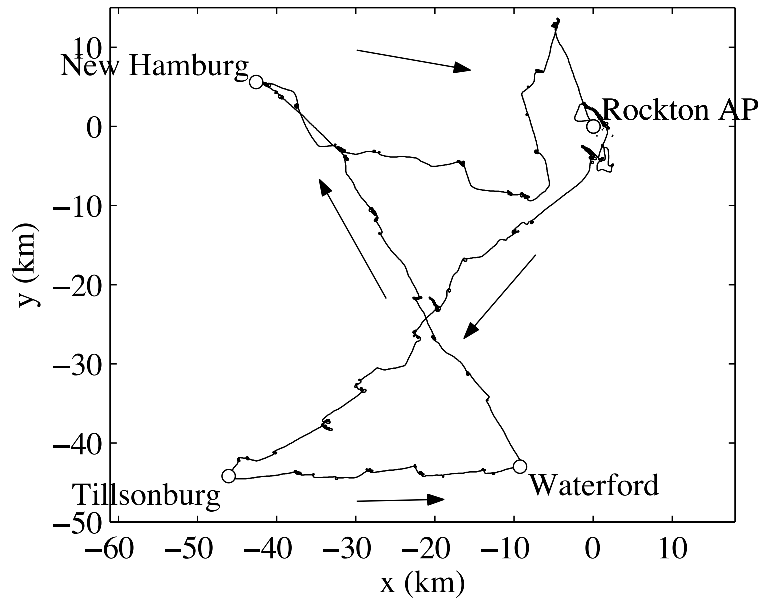

Figure 2: GPS trace from the 2001 Canadian Nationals competition, in xy projection. The course started at Rockton Airport, and rounded turnpoints at Tillsonburg, Waterford, and New Hamburg, before returning to Rockton. The glider does not follow precisely the straight lines of the course, but deviations are not too extreme (except on the last leg).

1.2 Mathematics in soaring

Soaring is an especially fruitful area for mathematical analysis, much more than other sports such as sailing or sports that directly involve the human body, for two reasons: First, as mentioned above, the space and time scales are beyond direct human perception, and as a consequence the participants already carry computing equipment.

Second, the physical aspects can be characterized unusually well. Unlike a sailboat, a sailplane operates in a single fluid whose properties are very well understood. The performance of the aircraft at different speeds can be very accurately measured and is very reproducible. The largest area of uncertainty is the atmosphere itself, and the way in which this is modeled determines the nature of the mathematical problem to be solved.

Indeed, in addition to vast amounts of research on traditional aeronautical subjects such as aerodynamics and structures, there have been a few important contributions to the mathematical study of cross-country soaring flight itself.

In the 1950's, Paul MacCready, Jr., solved a large portion of the problem described above: in what conditions to switch from cruise to climb, and how to adjust cruise speed to local lift/sink. In 1956, he used this theory to become the first American to win a world championship. His theory was presented in a magazine article [MacCready, 1958] and is extensively discussed and extended in many books [Reichmann, 1975].

This theory involves a single number, the "MacCready setting", which is to be set equal to the rate of climb anticipated in the *next* thermal along the course. By a simple construction described in Section 2.1 below, this gives the local speed to fly in cruise

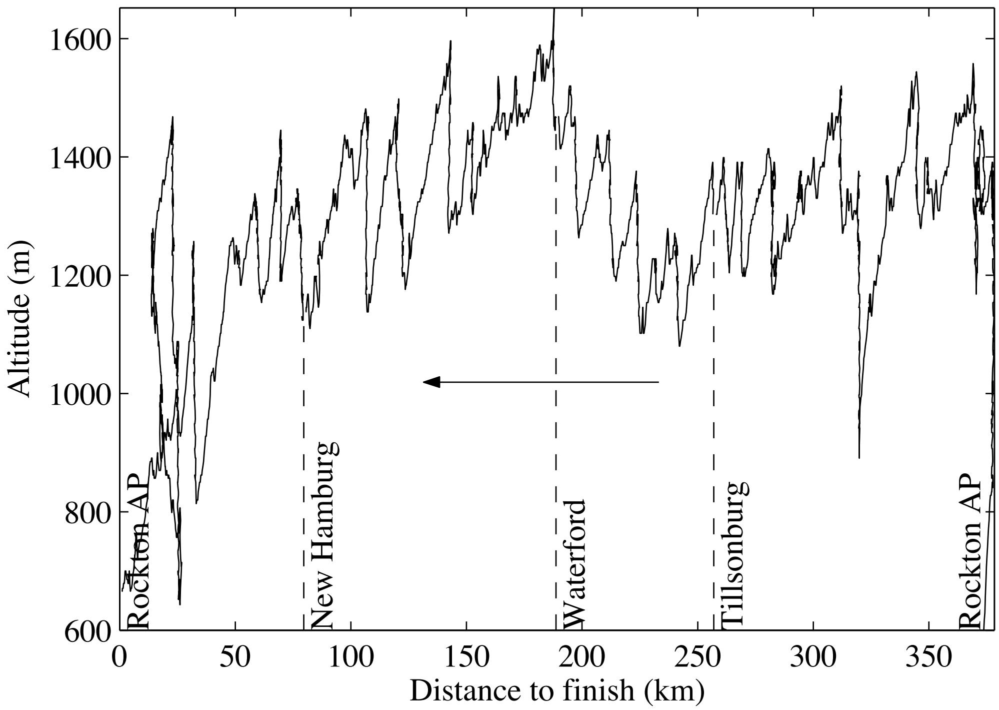

Figure 3: The same trajectory as in Figure 2, with altitude. The horizontal axis x is distance remaining on course, around the remaining turn points. Vertical segments represent climbs in thermals; not all climbs go to the maximum altitude. In the cruise phases of flight, the trajectories are far from smooth, representing response to local up/down air motion.

mode, and is simultaneously the minimum thermal strength that should be accepted for climb. MacCready showed how the necessary speed calculation could be performed by a simple ring attached to one of the instruments, and now of course it is done digitally. Indeed, modern flight computers have a knob for the pilot to explicitly set the MacCready value, and as a function of this setting and the local lift or sink a needle indicates the optimal speed.

The defect, of course, is that the strength of the next thermal is not known with certainty. Of course, one has an idea of the typical strength for the day, but it is not guaranteed that one will actually encounter such a thermal before reaching ground. As a result, the optimal setting depends on altitude: when the aircraft is high it is likely that it may meet a strong thermal before being forced to land. As altitude decreases, this probability decreases, and the MacCready setting should be decreased, both to lower the threshold of acceptable thermals and to fly more conservatively and slower in the cruise.

Various attempts have been made to incorporate probability into MacCready theory. Edwards [1963] constructed an "efficient frontier" of optimal soaring. In his model, thermals were all of the same strength, and were independently distributed in location; air was still in between. As a function of cruise speed, he could evaluate the realized average speed, as well as the probability of completing a flight of a specified length. His conclusion was that by slowing down slightly from the MacCready-optimal speed, a first-order reduction in landout probability could be obtained with only a second-order reduction in speed.

Cochrane [1999] has performed the most sophisticated analysis to date, which was the inspiration for the present paper. He performs a discrete-space dynamic programming algorithm: the horizontal axis is divided into one-mile blocks, and in each one the lift is chosen randomly and independently. For each mile, the optimal speed is chosen to minimize the expectation of time to complete the rest of the course.

In this model, the state variables are only distance remaining and altitude; since the lift value is chosen independently in each block, the current lift has no information content for future conditions. MacCready speed optimization is automatically built in. Solving this model numerically, he obtains specific quantitative profiles for the optimum MacCready value as a function of altitude and distance: as expected, the optimal value increases with altitude.

1.3 Outline of this paper

Our model may be viewed as a continuous version of Cochrane's. We model the lift as a continuous Markov process in the horizontal position variable: the lift profile is discovered as you fly through it. As a consequence, the local value of lift appears as an explicit state variable in our model, and the MacCready construction is independently tested. We find that it is substantially verified, with small corrections which are likely an artifact of our model.

The use of a Markov process represents an extremely strong assumption, representing almost complete ignorance about the patterns to be encountered ahead. In reality, the pilot has some form of partial information, such as characteristic thermal spacings, *etc.* But this is extremely difficult to capture in a mathematical model. Further, our model has only a single horizontal dimension: it does not permit the pilot to deviate from the course line in search of better conditions. This is likely our largest source of unrealism.

In Section 2 we present our mathematical model. We explain the basic features of the aircraft performance, and the MacCready construction, which are quite standard and quite reliable. Then we describe our Markov atmosphere model, which is chosen for simplicity and robustness.

In Section 3 we consider an important simplification of this model, in which the entire lift profile is known in advance. This case represents the opposite extreme of our true problem; it yields a first-order PDE which gives important insight into the behavior of solutions to the full problem.

In Section 4 we solve the full problem: we define the objective, which is the expectation of time to complete the course and derive, using the Bellman Dynamic Programming Principle, a degenerate second-order parabolic partial differential equation (PDE) that describes the optimal value function and the optimal control. Boundary conditions at the ground come from the landout penalty and at cloud base a state constraint is necessary. This nonlinear Hamilton-Jacobi-Bellman equation is surprisingly difficult to solve numerically, but numerical techniques based on the theory of viscosity solutions give very good methods [Barles and Souganidis, 1991]. We exhibit solutions and assess the extent to which MacCready theory is satisfied.

2 Model

In this section, the first and most straightforward part discusses the performance of the glider itself. Second, we present the MacCready optimization problem and its solution. Finally, we introduce our extremely simplified probabilistic model for the structure of the atmosphere; this model is the bare minimum that has any hope of capturing the partial information which is such an essential part of the problem.

2.1 The glider and its performance

The pilot can control the horizontal speed v of the glider by varying its pitch angle. For each speed v, the sink rate, or downwards vertical velocity, is given by a function $s(v)$ as in Figure 4. (In a powered aircraft, the rate of sink or climb is controlled by the engine throttle.)

The function s is positive and convex, with a minimum value $s_{\min}$ at speed $v_{\min}$. As v decreases below $v_{\min}$ the sink rate $s(v)$ increases, down to a minimum speed that the aircraft needs to sustain flight. In the rest of this paper we will consider v to be the net forward speed of the aircraft over the ground, which may be less than its instantaneous airspeed if it is not flying straight. We view forward speeds $0 < v < v_{\min}$ as being attained by an alternation of flying in a straight line at speed $v_{\min}$, with circling at airspeed $v_{\min}$ but net forward speed $v = 0$. Thus we set $s(v) = s_{\min}$ for $0 \le v \le v_{\min}$, although below we show that such speeds are never optimal. In practice the sink rate is slightly higher in circling flight than in straight, but we neglect this effect.

The *glide ratio*, or slope of the flight path, is $r(v) = v/s(v)$, which achieves its maximum value r_0 at speed $v_0 > v_{\min}$. As v increases beyond v_0, the glide ratio decreases, but in strong conditions it may be optimal to accept this poorer slope in exchange for the faster speed.

Beyond a maximum speed $v_{\max}$, flight becomes unsafe due to unstable oscillations of the control surfaces; the corresponding sink rate $s_{\max} = s(v_{\max})$ is finite. At any speed v, one can achieve sink rates larger than $s(v)$ by opening spoilers on the wings; this can be important near cloud base in strong lift, when flying at $v_{\max}$ does not yield enough sink to keep the glider from climbing into clouds. Thus the accessible region of horizontal and vertical speeds is the gray shaded area in Figure 4, rather than just its boundary curve.

For $v_{\min} \le v \le v_{\max}$, the function $s(v)$ is reasonably well approximated by a quadratic, and hence is fully specified by any three of the parameters $s_{\min}$ $v_{\min}$, r_0, and v_0:

$$\begin{aligned} s(v) &= \frac{v}{r_0} + \alpha(v - v_0)^2, & \alpha &= \frac{1}{4r_0^2(s_0 - s_{\min})} \\ &= s_{\min} + \alpha(v - v_{\min})^2, & v_{\min} &= r_0(2s_{\min} - s_0). \end{aligned}$$

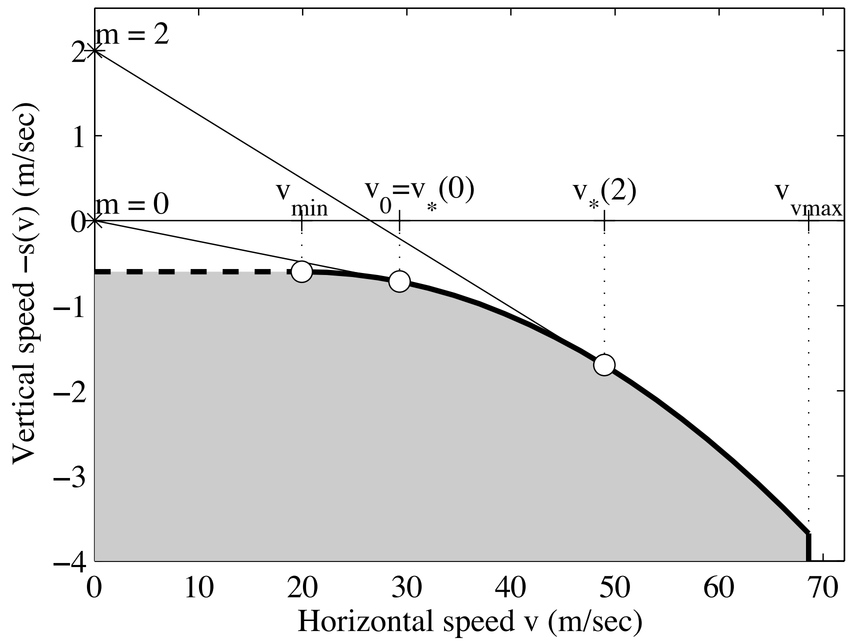

Figure 4: The sink rate, or "polar" function $s(v)$, with parameters as described in the text. Speed $v_{\min}$ is the speed at which sink rate is minimal; slower speeds correspond to circling flight. The tangent lines illustrate the MacCready construction: each value $m \geq -s_{\min}$ plotted on the vertical axis gives a MacCready speed $v_*(m) \geq v_{\min}$.

Reasonable values are those for a Pegasus (formerly owned by the first author): $s_{\min} = 1.16\,\mathrm{kt} = 0.60\,\mathrm{m/sec}$, $v_0 = 57\,\mathrm{kt} = 29.3\,\mathrm{m/sec}$, $r_0 = 41$, and $v_{\min} = 38\,\mathrm{kt} = 19.5\,\mathrm{m/sec}$.[2] Also, $v_{\max} = 133\,\mathrm{kt} = 68.6\,\mathrm{m/sec}$. See Figure 4.

2.2 MacCready Optimization

The speed v_0 that maximises the glide angle $v/s(v)$ gives the maximum glide distance for a given altitude loss, but it says nothing about time.

MacCready posed the following question: Suppose the aircraft is gliding through still air towards a thermal that will give realized climb rate m (upwards air velocity minus sink rate of glider). What glide speed v will give the largest final cross-country speed, after the glider has climbed back to its starting altitude?

The answer is given by the graphical construction in Figure 4. For any speed v, the final average cross country speed is given by the intersection with the horizontal axis of the line between $(0, m)$ and $(v, s(v))$. This speed is maximised by the tangent line.

We denote by $v_*(m)$ the MacCready value

$$v_*(m) = \operatorname*{arg\,min}_{v_{\min} \leq v \leq v_{\max}} \frac{m + s(v)}{v} \qquad \text{for } m \geq -s_{\min}. \tag{1}$$

[2]kt denotes "knot," one nautical mile per hour. Aviation in North America involves a confusing mix of statute, nautical, and metric units, but we shall consistently use SI units.

For the quadratic polar, this function is easily determined explicitly:

$$v_*(m) = \sqrt{v_0^2 + \frac{m}{\alpha}} \qquad \text{for } m \geq \alpha(v_{\min}^2 - v_0^2).$$

This $v_*(m)$ is an increasing function: the stronger conditions you expect in the future, the more willing you are to burn altitude to gain time now.

The function $v_*(m)$ also answers the following two related questions. First, Suppose the glider is flying through air that is rising at rate ℓ (sinking, if $\ell < 0$). What speed v gives the flattest glide angle, relative to the ground? The effect of the local lift or sink is to shift the polar curve up or down, and the optimal speed is $v_*(-\ell)$.

Second, combining the above two interpretations, suppose that the glider is flying toward a thermal that will give lift m, through air with varying local lift/sink $\ell(x)$. The optimal speed is $v_*(m - \ell(x))$. Each unit of time spent gliding in local lift $\ell(x)$ is separately matched with a unit of time spent climbing at rate m, so the optimization can be carried out independently at each x.

Thus, under the key assumption that future climb rate is known, the MacCready parameter m, and the associated function $v_*(m)$, give the complete optimal control. To emphasize the centrality of the parameter m, let us point out one more interpretation: m is the *time value of altitude*: each meter of additional altitude at the current location translates directly into $1/m$ seconds less of time on course. This is a direct consequence of the assumption of a final climb at rate m.

A consequence of this interpretation is that m is also the threshold for accepting a thermal. Suppose you are cruising towards an expected climb rate m, but you unexpectedly encounter a thermal giving climb rate m'. You should stop and climb in this thermal if $m' \geq m$.

The subject of this paper is how to determine optimal cruise speeds and climb decisions when future lift is uncertain. Thus the interpretation of m as a certain future climb rate will no longer be valid. But m will still be the time value of altitude, and local cruise speeds will still be determined by the MacCready function $v_*(m)$. The outcome of the analysis will be a rule for determining m in terms of local state variables.

As a final note, let us mention one extension: If the glider is cruising in a headwind of strength w, then the speed for flattest glide angle is

$$v_*(m, w) = \underset{v_{\min} \leq v \leq v_{\max}}{\arg\min} \frac{s(v) + m}{v - w}. \tag{2}$$

It is an increasing function of both m and w: you speed up when penetrating against a headwind. This function will become important when landout is inevitable, as a consequence of the time penalties assigned there.

2.3 Atmosphere

The model we propose is the simplest possible one that contains the essential features of fluctuating lift/sink, and of uncertainty about the conditions to be discovered ahead.

We assume that the glider moves in a two-dimensional vertical plane, whose horizontal coordinate x represents distance to the finish. Thus x decreases as the glider advances. The air is moving up and down with local vertical velocity (lift) $\ell(x)$, independent of altitude.

We take the lift to be a continuous Markov process, satisfying a stochastic differential equation (SDE) with x as the time-like variable, with the difference that x is decreasing rather than increasing. We choose the simplest possible Ornstein-Uhlenbeck model

$$d\ell_x \;=\; a(\ell_x)\,(-dx) \;+\; b\,dW_x, \tag{3}$$

where W is a time-reversed Wiener process, a, is the drift coefficient, and b is the standard deviation of the process. These are further specified by setting

$$a(\ell) \;=\; \frac{\ell}{\xi}, \qquad b^2 \;=\; \frac{2\bar{\ell}^2}{\xi}.$$

Here $\bar{\ell}$ is an average amplitude, or in other words, the standard deviation of the whole lift time history, and ξ is a correlation length (or alternately $1/\xi$ is often called the speed of mean reversion).

This model should not be taken too seriously. The true vertical motion of the atmosphere within its convective boundary layer has an extremely complicated structure that is not well understood. For example, rising air tends to be concentrated in narrow "thermals", often driven by condensation at the top (puffy fair-weather cumulus clouds) while sink is spread across broad areas, in contrast to our model which is symmetric between lift and sink. Also, the strength of lift generally increases with height, in contrast to our model in which ℓ depends only on x. Thus this model is only a crude approximation.

In principle, it might be possible to determine plausible estimates $\xi, \bar{\ell}$ empirically by analyzing flight data. In practice, noise in the data and other difficulties such as wind, three-dimensional effects, and the need to subtract the glider sink rate, make this extremely difficult. But based on experience, reasonable values are on the order of ξ = 1 km and $\bar{\ell}$ = 1.5 m/sec for good conditions.

The choice of a Markov process imposes very strict constraints on the nature of the information discovery. We shall consider two extreme forms of information.

First, in Section 3 below, we consider the following completely deterministic problem: we suppose that the entire profile $\ell(x)$ is known at the beginning; the randomness enters only in the initial determination of the profile, before the start of the flight and the pilot seeks to minimize his time to finish. If the lift is not strong enough, the pilot may not be able to complete the race. When an unavoidable landout occurs, we apply a penalty to the objective function. The model gives us a first-order Hamilton-Jacobi equation in the two variables (x, z), with $\ell(x)$ appearing as a parameter. Its solutions can be completely described in terms of characteristics, and it allows us to reproduce the classic MacCready theory [Cochrane, 1999].

Second, in Section 4, we consider the opposite extreme in which the pilot has information about what lies ahead. He discovers lift only by flying through it, though the continuity of the process (3) gives some information over distances less than ξ. Although in practice the pilot does have some ability to evaluate atmospheric conditions ahead, this version is more realistic than the fully deterministic problem, especially on the large length scales characteristic of competitions. It gives rise to a degenerate nonlinear second-order parabolic equation in two space-like variables (z, ℓ) and a time-like variable x.

3 Deterministic problem

3.1 Control problem

As described above, we suppose that the entire lift profile $\ell(x)$ is known to the pilot. The glider is in $x > 0$, flying inwards along a line toward a goal at $x = 0$; z denotes altitude which will vary between $z = 0$ and $z = z_{\max}$ (see below). The pilot's objective is to reach the goal as quickly as possible. In this version of the problem there is no uncertainty, so the pilot simply seeks to minimize the deterministic time.

Dynamics The control variables are the horizontal velocity $v \in [0, v_{\max}]$ and the sink rate $s \geq s(v)$, from which the vertical velocity is determined as $\ell(x) - s$. If the pilot advances at speed v and sink rate s, then his distance x and height z change by

$$dx = -v\,dt, \qquad dz = (\ell(x) - s)\,dt.$$

Away from the cloud base, there is no reason to sink faster than the minimum, so the preferred sink rate is $s = s(v) = s_{\min} + \alpha(v - v_{\min})^2$. The realized glide angle is then

$$\frac{dz}{dx} = \frac{s(v) - \ell(x)}{v}.$$

The pilot "cruises" if $v > 0$. As a special case of the above with $v = 0$, if $\ell \geq s_{\min}$, then he has the option to "climb" without advancing, and his position updates according to

$$dx = 0, \qquad dz = (\ell - s_{\min})\,dt.$$

State constraint As noted in the introduction, the cloud bases define a maximum height for safe flight. We let $z = z_{\max}$ denote the altitude of the cloud bases. At $z = z_{\max}$, even when the lift $\ell(x)$ is very strongly positive and the condition $v \leq v_{\max}$ becomes binding, the pilot can always avoid being drawn upwards into the clouds by opening the spoilers to increase the sink rate to a value greater or equal to $\ell(x)$. We thus impose the condition that the trajectories are forbidden to exit the domain even when it would otherwise be optimal to do so:

$$z(t) \leq z_{\max}$$

and the admissible controls at $z = z_{\max}$ must satisfy the extra condition $s \geq \ell(x)$.

Landout penalty Let $z = 0$ denote the height at which the pilot must stop trying to soar, and land in whatever field is available (typically $\sim 300\,\mathrm{m}$ above actual ground level). At $z = 0$, it is impossible to impose the conditions that trajectories may not exit; "landouts" may become inevitable when $\ell(x) < s_{\min}$, and when z is close to zero so that the pilot does not have room to advance toward a possible climb.

In this first-order problem, it would be possible to impose an infinite cost on trajectories that exit the domain, which would not affect nearby trajectories that do not exit. But since this would lead to an infinite value function for the second-order problem discussed in Section 4, and since it is incompatible with actual practice, we impose a finite value.

Contest rules impose a complicated set of penalties for failure to complete the course; these penalties are deliberately less than drastic, in order to encourage competitors to make safe landings with adequate preparation rather than to search for lift until the last possible moment. The penalty rules generally have two features: a discontinuity at the end, so landing one meter short of the finish is substantially worse than finishing; and a graduated reward for completing as much distance as possible.

Since we want a formula that can be simply expressed in terms of time, we suppose that when you land out, your aircraft is put onto a conveyer belt. This belt transports you to the finish in a time which depends only on the position on which you landed on the belt, not on the time at which you landed: the belt's speed is a function only of its own position, not of time. Thus the time at which you reach the goal is the time at which you landed, plus the time of transport; each minute in the air is always worth exactly one minute in final arrival time.

Thus we impose the penalty

$$\Psi(x) \;=\; T_{\text{pen}} \;+\; \frac{x}{V_{\text{pen}}}$$

where T_{pen} is the discontinuity at $x = 0$, and V_{pen} represents the distance penalty. In this paper, we use the more or less realistic values $T_{\text{pen}} = 600\,\text{sec}$, $V_{\text{pen}} = 17\,\text{m/sec}$.

Value function Let $u(x,z)$ denote the time needed to reach $x = 0$ from position (x,z). It is defined for $x \ge 0$ and $0 \le z \le z_{\max}$ by

$$u(x,z) = \inf_{(v,s)\in\mathcal{A}_{x,z}} \Big\{\tau + \phi\big(x(\tau), z(\tau)\big)\Big\}$$

where $\mathcal{A}_{x,z}$ is the set of admissible controls, and τ is the first exit time from the open set $\Omega = \{\, x > 0,\ z > 0 \,\}$ (recall that we forbid exit through $z = z_{\max}$):

$$\tau = \inf\Big\{\ t \ge 0 :\ \big(x(t), z(t)\big) \notin \Omega\ \Big\}.$$

Here ϕ is the discontinuous function defined on $\partial\Omega$ by

$$\phi(x,z) \;=\; \begin{cases} 0 & \text{if } x = 0 \\ \Psi(x) & \text{if } x > 0 \text{ and } z = 0 \end{cases} \tag{4}$$

where Ψ is the penalty defined above. In other words, the trajectories exit the domain, either when the glider reaches the finish line ($x = 0$), or else when it lands out before reaching the finish line ($x > 0, z = 0$); in the latter case, a strictly positive penalty applies.

The value function is nonnegative and grows linearly in the variable x. Since the glider can reach the goal from position x only by passing through smaller values of x at finite velocity, u is a strictly increasing function of x. Furthermore, u is a non increasing function of z. Consider indeed two different values of z; starting from the higher altitude, the pilot may always apply the strategy that would be optimal at the lower altitude and reach the finish line in the same amount of time. Furthermore, he may have access to a better strategy that would have led to a landout at the lower altitude but does not at the higher altitude. Finally, u may not be continuous when climbing is not an option ($\ell(x) \le s_{\min}$) because the penalty resulting from a landout causes a discontinuity.

3.2 Hamilton-Jacobi-Bellman equation

We do not expect the value function to be smooth and we present an informal derivation of the Variational Inequalities, assuming that the value function has continuous first derivatives. The theory of viscosity solutions tells us how to extend the formulation to non-smooth solutions.

Climb From height $z < z_{\max}$ in lift $\ell > s_{\min}$, the glider can access a new height $z' = z + (\ell - s_{\min})\Delta t > z$ by circling ($v = 0$) for time $\Delta t = (z' - z)/(\ell - s_{\min})$. Therefore, by the Dynamic Programming Principle,

$$u(x,z) \;\le\; u(x,z') + \frac{z' - z}{\ell - s_{\min}} \qquad \text{for all } 0 \le z \le z' \le z_{\max},$$

or

$$\frac{u(x,z') - u(x,z)}{z' - z} \;\ge\; -\frac{1}{\ell - s_{\min}}.$$

Letting Δt go to 0, we obtain the gradient constraint

$$u_z(x,z) \;\ge\; -\frac{1}{\ell(x) - s_{\min}} \qquad \text{wherever } \ell(x) > s_{\min}, \tag{5}$$

with $u_z = \partial u/\partial z$. As noted above, we have $u_z \le 0$ even when $\ell(x) \le s_{\min}$. Below we shall identify $m = -1/u_z$, and then this may be written

$$m \;\ge\; \max\{\, \ell - s_{\min},\, 0 \,\}. \tag{6}$$

Cruise Because the pilot has the option to fly forward at any speed v with $0 < v \le v_{\max}$,

$$u(x,z) \;\le\; \frac{\Delta x}{v} + u(x - \Delta x, z - \Delta z) \;\approx\; \frac{\Delta x}{v} + u(x,z) - \left(u_x + \frac{s(v) - \ell}{v}\, u_z \right) \Delta x,$$

with $\ell = \ell(x)$ and $\Delta z = \Delta x\ \bigl(s(v) - \ell(x)\bigr)/v$. Canceling the common $u(x,z)$ and retaining only the terms of size $\mathcal{O}(\Delta x)$, we have

$$0 \le -u_x + \min_{0<v\le v_{\max}} \left(\frac{1}{v} - \frac{s(v) - \ell}{v}\, u_z \right).$$

Next, we use the fact that $u_z \le 0$ and the MacCready index $m = -\frac{1}{u_z}$ to derive

$$0 \le -u_x - u_z \min_{0<v\le v_{\max}} \left(\frac{m - \ell + s(v)}{v} \right).$$

The minimum is attained at $v = v_*(m - \ell(x))$ from (1), so this is

$$0 \le -u_x - u_z \; \frac{m - \ell + s(v_*)}{v_*}.$$

The minimum in (1) is indeed well defined thanks to (6), which guarantees that whether $\ell(x) \le s_{\min}$ or $\ell(x) > s_{\min}$, there is a minimum v with $v_{\min} \le v \le v_{\max}$. Values $0 < v < v_{\min}$ are never optimal: the glider either climbs or it cruises.

We thus obtain the partial differential inequality

$$u_x \;+\; H(\ell(x), u_z) \le \; 0 \tag{7}$$

where

$$\begin{aligned} H(\ell, p) &= -\min_{v\in[v_{\min}, v_{\max}]} \left(\frac{1}{v} - \frac{s(v) - \ell}{v}\, p \right) \\ &= \frac{m - \ell + s(v_*(m-\ell))}{v_*(m-\ell)}\, p \\ &= -\frac{m - \ell + s(v_*(m-\ell))}{m\, v_*(m-\ell)} \end{aligned}$$

with $m = -1/p$.

Hamilton-Jacobi Variational Inequalities Finally, let $x_{\max} > 0$ be the length of the race. The value function u of the deterministic control problem is (at least formally) a bounded and nonnegative viscosity solution of the following HJB equation in $(0, x_{\max}] \times (0, z_{\max})$

$$\max\left\{ u_x + H(\ell(x), u_z),\; -u_z - \frac{1}{\ell(x) - s_{\min}} \right\} = 0 \text{ in } \{ \ell(x) > s_{\min} \} \tag{8}$$

$$u_x + H(\ell(x), u_z) = 0 \text{ in } \{ \ell(x) \le s_{\min} \}$$

associated with the state constraint

$$u_x + H(\ell(x), u_z) = 0 \text{ on } \{ x > 0,\, z = z_{\max} \}, \tag{9}$$

the initial condition

$$u(0, z) = 0 \text{ on } \{ x = 0,\, 0 \le z \le z_{\max} \}, \tag{10}$$

and the weak Dirichlet condition

$$u(x, 0) = \Psi(x) \text{ on } \{ x > 0,\, z = 0 \}. \tag{11}$$

3.3 Explicit solution

Differentiating (7) with respect to z and denoting $p = u_z$, in regions of cruise we have the nonlinear conservation law

$$p_x + H(\ell(x), p)_z = 0$$

which shows that u_z and hence m are constant along trajectories $z(x)$ with

$$\frac{dz}{dx} = \left.\frac{\partial H}{\partial p}\right|_\ell = \frac{\partial H}{\partial m}\frac{\partial m}{\partial p} = \frac{s(v_*(m-\ell)) - \ell}{v_*(m-\ell)}.$$

This is the physical trajectory of a glider flying in lift ℓ, at the optimal speed $v_*(m-\ell)$ for the value of m determined by the value function $p = u_z$.

The inequality (7) is to be solved with the constraint (6), along with the condition that at each (x,z), at least one must be an equality. For this first-order problem, we can explicitly describe the solutions.

If the cruise trajectories for nearby values of z all terminate in a climb in constant lift $\ell_* > s_{\min}$, then the achieved climb rate will be $\ell_* - s_{\min} > 0$, so u_z is constant across as well as along trajectories, with value

$$u_z = -\frac{1}{\ell_* - s_{\min}},$$

giving a constant value of $m = \ell_* - s_{\min}$. This reproduces the MacCready rule: In regions where cruise will terminate in a climb at a known rate, the velocity changes in reaction to local lift $\ell(x)$ by the rule

$$v(x) = v_*(m_* - \ell(x)), \qquad m_* = \ell_* - s_{\min}.$$

Thus the partial differential inequalities (6,7) are to be solved for $u(x,z)$ on the domain $x > 0$ and $0 < z < z_{\max}$. At $x = 0$ the "initial" condition is $u(0,z) = 0$, and x plays the role of time. We must specify boundary conditions at $z = 0$ and $z_{\max}$.

Let us determine the optimal strategy near landout, for small height z. Let us assume you cruise forward until you land; the lift ℓ will remain roughly constant. At speed v, you will land in time $\Delta t = z/(s(v)-\ell)$, and you will advance a distance $\Delta x = -zv/(s(v)-\ell)$. The optimal time from (x,z,ℓ) therefore satisfies

$$\begin{aligned} u(x,z) &\le \min_v \left[\Delta t + T_{\text{pen}} + \frac{x + \Delta x}{V_{\text{pen}}} \right] \\ &= T_{\text{pen}} + \frac{x}{V_{\text{pen}}} + \frac{z}{V_{\text{pen}}} \min_v \frac{V_{\text{pen}} - v}{s(v) - \ell}. \end{aligned}$$

The solution is easily seen to be

$$v = v_*(-\ell, V_{\text{pen}}).$$

with $v_*(m,\ell)$ from (2). When landout is inevitable, you should fly at the MacCready speed for the local sink rate, and as though you had a headwind of speed V_{pen}. The final result for u is

$$u(x,z) \;\le\; T_{\text{pen}} + \frac{x}{V_{\text{pen}}} - \frac{v_* - V_{\text{pen}}}{s(v_*)-\ell}\,\frac{z}{V_{\text{pen}}} \qquad \text{for } z \text{ small.}$$

The inequality recognizes the possibility that this strategy may be superseded by the ability to climb away from the bottom boundary. But when z is small, this possibility is negligible, and we may treat this as an equality.

3.4 Numerical solutions

Here we simply exhibit example solutions; the scheme is a simplification of the second-order one introduced below.

Figure 5 shows the trajectories and MacCready values for one realization of the lift, with values as described above. The trajectory starts 200 km from the finish, at an altitude half of the maximum thermal height $z_{\max} = 1500$ m. The upper panel shows the actual trajectory in x and z. The black regions indicate the “maximum thermal height:” the height to which it is optimal to climb in lift. At the top of the black region, the optimal trajectories leave the lift and cruise. The trajectories go right down to $z = 0$, since the pilot has perfect knowledge that the lift will be there.

The lower panel shows the local climb rate $\ell(x) - s_{\min}$, and the MacCready value m. By virtue of the constraint (6), the former is a lower bound for the latter. The MacCready value is approximately constant in regions of cruise. In contrast to classic MacCready theory, the value of this coefficient is typically equal to the climb rate experienced in the *previous* thermal, when the previous one is weaker than the one ahead. This is because with perfect foreknowledge, the optimal strategy when approaching a strong thermal is to enter it exactly at ground level. There is a collection of paths that all converge at this point. The physical analogy would be a rarefaction fan in compressible gas dynamics.

The increase in MacCready number shown when the trajectories approach ground level, for example at around $x = 80$ km is due to a computational artifact known as “numerical viscosity,” well-known in the numerical solution of hyperbolic partial differential equations [LeVeque, 1992]. For each trajectory that reaches a thermal at ground level at maximum glide, there is a neighboring trajectory just below that misses the thermal and pays the landout penalty. For the landing trajectory, the optimal strategy is to increase the MacCready value as though flying into a headwind, as discussed above. The true solution is discontinuous along this division line, and the discretization on a finite grid causes the landout value to diffuse upwards and affect solutions that do not land.

A really accurate solution of this problem would require explicit tracking of characteristics and trajectories, and would be extremely complicated geometrically. Since our real purpose is to solve the second-order problem below, in which numerical viscosity is dominated by the true uncertainty, we accept these minor errors.

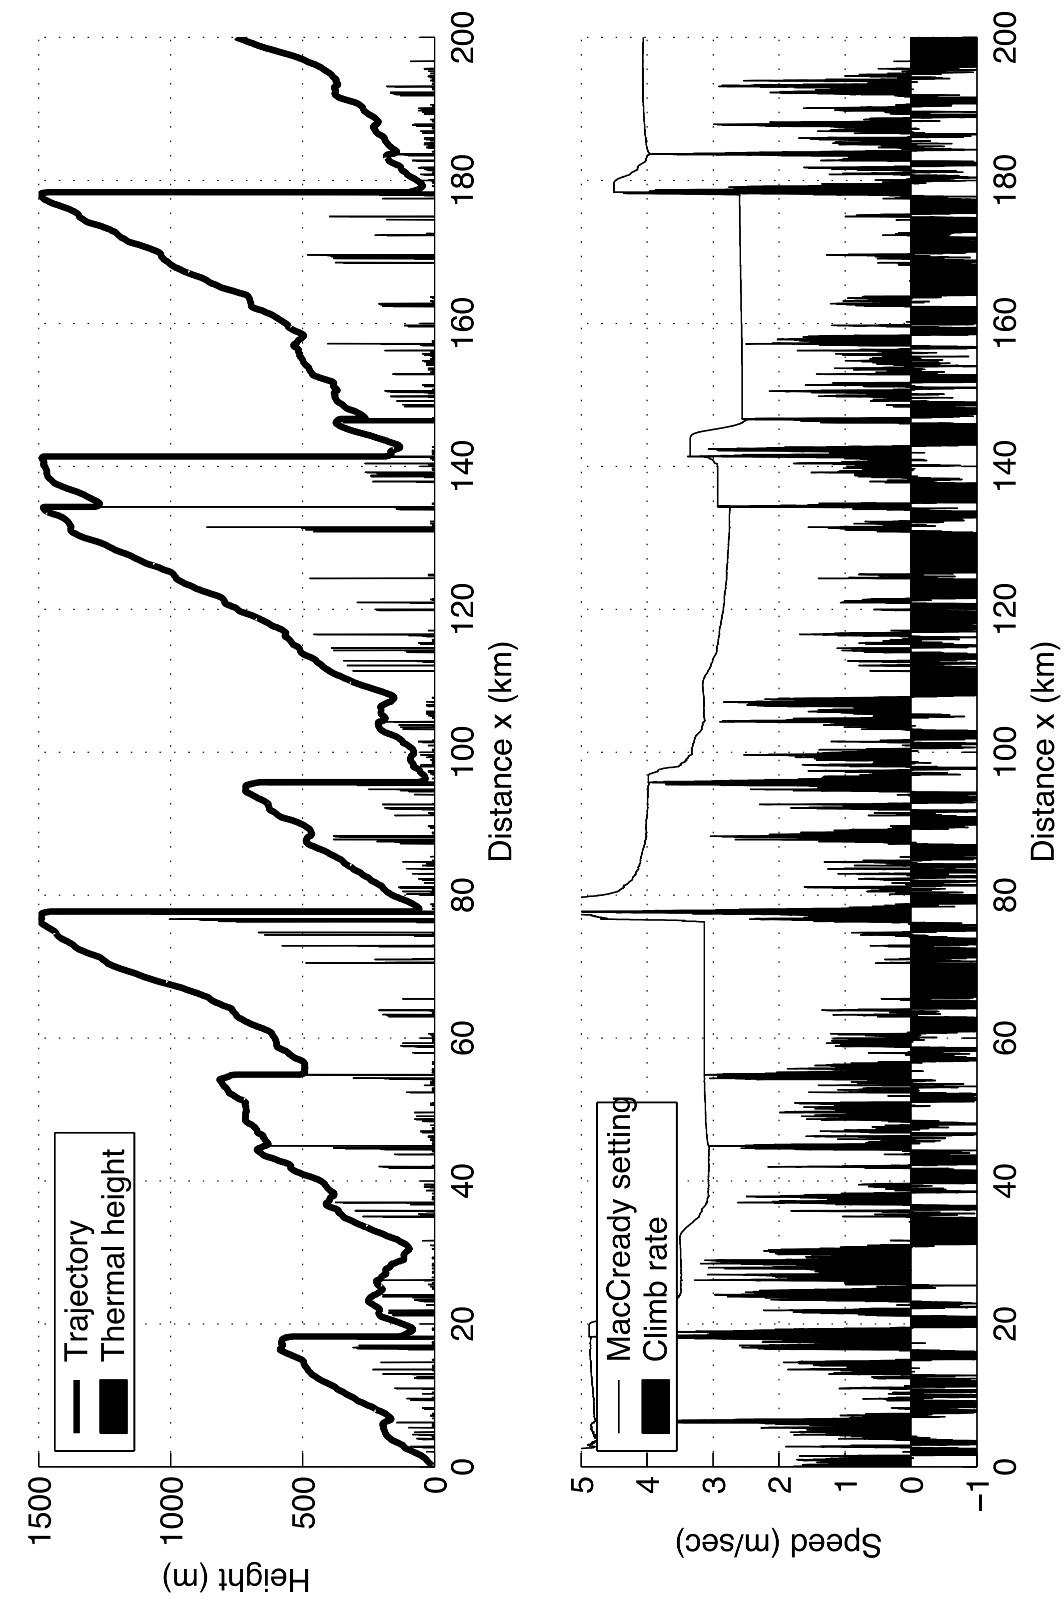

Figure 5: Optimal trajectory and corresponding MacCready values for one realization of lift for the first-order problem, with complete knowledge of future lift. In the upper panel, the thick line is the trajectory; vertical segments are where the glider stops to climb in lift. The vertical lines below the trajectory denote the height to which it is optimal to climb in each thermal; if one enters the thermal above that altitude it is optimal to continue without stopping. The lower panel shows the lift $\ell(x)$ (black region), which is fully known to the pilot at the beginning, and the optimal MacCready coefficient along the trajectory. Realized average speed is 139 km/hr.

4 Stochastic problem

We now address the much more realistic “stochastic” problem. The lift profile is again generated by the process (3); the only difference is that the pilot does not know the profile in advance. He or she learns it by flying through it, so his or her information structure is the filtration generated by the Brownian motion W_x in *decreasing* x.

Now the desire to finish “as quickly as possible” is more ambiguous. US and international contest rules give a bewilderingly complex definition of this criterion. Briefly, the pilot is awarded a number of points equal to his speed divided by the fastest speed flown by any competitor on that day.

A fully realistic model would take into account all of these effects, by introducing as many state variables as necessary. For example, elapsed time would need to be a state variable, since speed depends on that in addition to remaining time. Also, the uncertainty of the winner's speed would need to be incorporated, as an extrinsic random variable, with special consideration if this pilot believes himself to be the leader. Finally, a suitable objective function would need to be specified, taking proper account of the various ways to assess contest performance. Cochrane [1999] discusses these issues in detail.

To avoid these complexities, we define the objective function to be simply

$$u \;=\; \text{minimum expectation of time to reach } x = 0,$$

where the conditional expectation is taken over potentially uncertain future atmospheric conditions, and the minimum is taken over control strategies described below. This quantity depends on the starting position x, z, and the local atmospheric conditions as represented by the locally observed lift $\ell(x)$. Since $\ell(x)$ is Markovian, no past information is needed, and the value function depends on (x, z, ℓ).

4.1 Stochastic control problem

The state variables, x_t, z_t, ℓ_x are solutions of the controlled stochastic differential equations (3) and

$$dx_t \;=\; -v_t\,dt, \qquad dz_t \;=\; (\ell_{x_t} - s_t)\,dt, \tag{12}$$

with $s_t = s(v_t)$, where, as before, the sink rate $s(v) = s_{\min} + \alpha(v - v_{\min})^2$ for $v > v_{\min}$ and $s(0) = s_{\min}$.

We apply the same state constraint and impose the same landout penalty as for the first-order case. The value function is defined by

$$u(x, z, \ell) = \inf_{(v_t, s_t) \in \mathcal{A}_{x,z,\ell}} \mathbb{E}_{x,z,\ell}\Big[\tau + \phi(x_\tau, z_\tau)\Big]$$

where τ is the first exit time of the pair (x_t, z_t) from the open set $\Omega = (0, +\infty)^2$:

$$\tau = \inf\Big\{\, t \geq 0 \;:\; (x_t, z_t) \notin \Omega \,\Big\},$$

$\mathbb{E}_{x,z,\ell}$ denotes the expectation for the initial condition $x_0 = x, z_0 = z, \ell_{x_0} = \ell$, and ϕ is the function defined in (4).

As in the first-order problem, the value function is nonnegative and bounded on $[0, x_{\max}]$, is strictly increasing in x, nonincreasing in z; it is also nonincreasing in ℓ because stronger lift cannot be detrimental. As in the deterministic case, u may fail to be continuous when $\ell \le s_{\min}$.

4.2 Second-order Hamilton-Jacobi-Bellman equation

Constraint (5) applies just as before, since the lift does not change while the glider does not advance in x:

$$u_z \ge -\frac{1}{\ell - s_{\min}}. \tag{13}$$

The Dynamic Programming Principle and Itô's Lemma give the partial differential inequality

$$u_x \;+\; H(\ell, u_z) \;+\; a(\ell)\, u_\ell \;-\; \frac{1}{2} b^2\, u_{\ell\ell} \;\le\; 0 \tag{14}$$

with

$$H(\ell, p) \;=\; -\min_{v \in [v_{\min}, v_{\max}]} \left(\frac{1}{v} \;-\; \frac{s(v) - \ell}{v}\, p \right).$$

Boundary conditions The domain is now $0 \le x \le x_{\max}$, $0 \le z \le z_{\max}$, and $-\infty < \ell < \infty$. At $x = 0$ and at $z = 0, z_{\max}$, we have the same boundary conditions as we gave for the first-order problem in Section 3.2.

Furthermore, at every point, one and only one of the equations (13,14) holds, that is, we now are to solve the system of Variational Inequalities

$$\max\left\{ u_x + H(\ell, u_z) + a(\ell) - \frac{1}{2} b^2 u_{\ell\ell},\ -u_z - \frac{1}{\ell - s_{\min}} \right\} \;=\; 0 \quad \text{in} \quad \{\ell > s_{\min}\} \tag{15}$$

$$u_x + H(\ell, u_z) + a(\ell) - \frac{1}{2} b^2 u_{\ell\ell} \;=\; 0 \quad \text{in} \quad \{\ell \le s_{\min}\} \tag{16}$$

associated with the state constraint

$$u_x + H(\ell, u_z) + a(\ell) - \frac{1}{2} b^2 u_{\ell\ell} \;=\; 0 \quad \text{on} \quad \{z = z_{\max}\}, \tag{17}$$

the initial condition

$$u(0, z, \ell) \;=\; 0 \quad \text{on} \quad \{x = 0\} \tag{18}$$

and the weak Dirichlet condition

$$u(x, 0, \ell) \;=\; \Psi(x) \quad \text{on} \quad \{x > 0, z = 0\}. \tag{19}$$

We will denote the finite limits in the variable ℓ by $\pm\ell_{\max}$, which are an approximation to the infinitely distant boundaries in the PDE. In the computation, we impose some Neumann conditions at $\pm\ell_{\max}$. Specifically, we set them arbitrarily equal to 0:

$$u_\ell(\pm\ell_{\max}) = 0.$$

Because the characteristics of the first-order part flow strongly out of the domain (as x increases), errors at the boundary will have an effect confined to a thin boundary layer, and the exact numerical boundary conditions are not too important [LeVeque, 1992].

4.3 Numerical approximation

We continue with a detailed description of the finite difference numerical scheme we implemented. In 1991, Barles and Souganidis [1991] established the convergence of a class of numerical approximations toward the viscosity solution of a degenerate parabolic or elliptic fully nonlinear equation (see also Crandall and Lions [1984], Souganidis [1985] for earlier results). Although the second-order equation (15–19) is slightly different from the one considered by Barles and Souganidis [1991], it is appropriate to seek an approximation belonging to this class and we verify its convergence to a reasonable solution experimentally.

We apply two operator splitting techniques: the first one, which consists in splitting the parabolic operator into an operator in the z variable and an operator in the variable ℓ, is the nonlinear analogue of the well-known Alternate Directions method and is justified by Barles and Souganidis [1991]. The second one, which consists of treating separately the parabolic operator and the gradient constraint is specific to Variational Inequalities and has been applied in a variety of situations [Barles, 1997, Barles et al., 1995, Tourin and Zariphopoulou, 1994] .

4.4 Semi-discretization in "time"

Recall that distance x is our time-like variable. It is approximated by the mesh $\{n\,\Delta x\,|\,0 \le n \le N\}$, where $\Delta x > 0$ is the discretization step, $N\Delta x = x$, and the approximation of the value function at the point $n\Delta x$ is denoted by $U^n(z,\ell)$, i.e. $U^n(z,\ell) \approx u(n\Delta x, z, \ell)$. Here we want to compute $U^{n+1}(z,\ell)$ in terms of $U^n(z,\ell)$.

- Initialize $U^0 = 0$.
- Given U^n, solve for every ℓ,

$$u_x + H(\ell, u_z) = 0, \text{ for } z > 0, x \in (n\Delta x, (n+1)\Delta x], \tag{20}$$
$$u(n\Delta x, z, \ell) = U^n(z,\ell), \text{ for } z > 0, \tag{21}$$
$$u(x, 0, \ell) = \Psi(x), \tag{22}$$

  and denote its computed solution at $(n+1)\Delta x$ by U^{n*}.

- Given U^{n*}, solve for every $z > 0$,

$$u_x - \frac{1}{2} b^2 u_{\ell\ell} + a(\ell) u_\ell = 0, \text{ for } \ell \in \mathbb{R}, x \in (n\Delta x, (n+1)\Delta x], \tag{23}$$
$$u(n\Delta x, z, \ell) = U^{n*}(z,\ell), \tag{24}$$

  and denote its computed solution at $(n+1)\Delta x$ by U^{n**}.

- Form $C^{n+1} = \left\{(z,\ell) : \ell > s_{\min} \text{ and } U_z^{n**} < -\frac{1}{\ell - s_{\min}}\right\}$.

- Compute the solution v of

$$u_z = -\frac{1}{\ell - s_{\min}}, \text{ on } C^{n+1}, \tag{25}$$

$$u = U^{n**} \text{ on } \partial C^{n+1}. \tag{26}$$

- Set

$$U^{n+1}(z,\ell) = \begin{cases} U^{n**}(z,\ell) & \text{if } (z,\ell) \in [0,+\infty) \times (-\infty,+\infty) \backslash C^{n+1} \\ v(z,\ell) & \text{if } (z,\ell) \in C^{n+1} \end{cases}$$

4.5 Fully discretized Finite Difference scheme

We now turn to the fully discretized scheme. The approximation that will be applied when cruising in the interior of the domain is explicit in the variable z and implicit in the variable ℓ.

Specifically, let $U_{i,j}^n$ be the approximation of the value function at the grid point (x_n, z_j, ℓ_i) with $x_n = n\Delta x$, $z_j = j\Delta z$ and $\ell_i = -\ell_{\max} + i\Delta\ell$, where $0 \le i \le 2L$, $0 \le j \le M$, $M\Delta z = z_{\max}$ and $L\Delta\ell = \ell_{\max}$.

Step 1: explicit upwind difference in z For $z < z_{\max}$, we apply the following explicit upwind (monotone) Finite Difference scheme:

$$\begin{aligned} \frac{U_{i,j}^{n*} - U_{i,j}^n}{\Delta x} = \min_{v_{\min} \le v \le v_{\max}} \Bigg\{ \frac{1}{v} - \max\left\{\frac{s(v) - \ell_i}{v}, 0\right\} D_z^- U_{i,j}^n \\ - \min\left\{\frac{s(v) - \ell_i}{v}, 0\right\} D_z^+ U_{i,j}^n \Bigg\}, 0 < j < M. \end{aligned}$$

where $D_z^\pm$ denote the standard one-sided differences, with Δz in the denominator so that they approximate derivatives.

We also impose the Dirichlet boundary condition

$$U_{i,j}^{n*} = \Psi(x_{n+1}) = T_{\text{pen}} + \frac{x_{n+1}}{V_{\text{pen}}}, \text{ for all } i, j.$$

In practice, we solve in closed form the above minimization problem. We d not show this computation which is straightforward. We compute the optimal control v^* using the upwind difference construction

$v^* =$

	$s(v_-) < \ell_i$	$s(v_-) \ge \ell_i$
$s(v_+) \le \ell_i$	v_+	$\arg\min\left\{ \frac{1}{v_+} - \frac{s(v_+) - \ell_i}{v_+} D_z^+ U_{i,j}^n, \frac{1}{v_-} - \frac{s(v_-) - \ell_i}{v_-} D_z^- U_{i,j}^n \right\}$
$s(v_+) > \ell_i$	$v_{\min} + \sqrt{\frac{\ell_i - s_{\min}}{\alpha}}$	v_-

with

$$v_\pm = \min\left\{ \sqrt{v_0^2 - \frac{1}{\alpha D_z^\pm U_{i,j}^n} - \frac{1}{\alpha \ell_i}},\ v_{\max} \right\}.$$

Then we determine U^{n*} by the explicit difference formula

$$\frac{U_{i,j}^{n*} - U_{i,j}^n}{\Delta x} = \frac{1}{v^*} - \frac{s(v^*) - \ell_i}{v^*} \begin{cases} D_z^- U_{i,j}^n, & \text{if } s(v^*) \ge \ell_i \\ D_z^+ U_{i,j}^n, & \text{else.} \end{cases}$$

At the boundary $z = z_{\max}$, there is a state constraint and we have to modify the scheme accordingly: we do not assume any longer that $s = s(v)$ but we require $s \ge s(v)$ and $s \ge \ell$ instead. In addition, we drop the forward finite difference in the variable z at this boundary.

$$\frac{U_{i,M}^{n*} - U_{i,M}^n}{\Delta x} = \min_{v_{\min} \le v \le v_{\max}} \left\{ \frac{1}{v} - \frac{\max\{s(v), \ell_i\} - \ell_i}{v} D_z^- U_{i,M}^n \right\}$$

where the optimization problem above is solved explicitly in a similar manner as in the interior of the domain.

Step 2: Implicit upwind difference in ℓ For every $0 < j \le M$, we apply the Finite Difference scheme:

$$\frac{U_{i,j}^{n**} - U_{i,j}^{n*}}{\Delta x} = \frac{1}{2} b^2 D_\ell^2 U_{i,j}^{n**} - a(\ell_i) \begin{cases} D_\ell^- U_{i,j}^{n**} & \text{if } \ell_i > 0 \\ D_\ell^+ U_{i,j}^{n**} & \text{if } \ell_i < 0, \end{cases}$$

where $D_\ell^- U_{i,j}^{n**}, D_\ell^+ U_{i,j}^{n**}$ are the standard implicit one sided differences divided by $\Delta\ell$ and $D_\ell^2 U$ is the usual implicit difference approximation for the second derivative $U_{\ell\ell}^{n**}$

$$D_\ell^2 U = \frac{U_{i+1,j}^{n+1} + U_{i-1,j}^{n+1} - 2U_{i,j}^{n+1}}{\Delta\ell}.$$

The above scheme inside the domain is complemented by the Neuman boundary conditions at $\pm\ell_{\max}$

$$U_{1,j}^{n**} - U_{0,j}^{n**} = 0, U_{2L-1,j}^{n**} - U_{2L,j}^{n**} = 0, \quad 0 < j \le M.$$

Step 3: Climb (gradient constraint): For $\ell_i > s_{\min}$, we apply the gradient constraint at the points at which it is violated, from top to bottom in the variable j, using an implicit scheme. Let j_i^* be the smallest indice j for a given i, such that $U_{i,j}^{n**}$ satisfy the gradient constraint, *i.e.*,

$$D_z^+ U_{i,j}^{n**} \ge -\frac{1}{\ell_i - s_{\min}}, \qquad \text{for all } j \ge j_i^*.$$

Then

$$U_{i,j}^{n+1} = \begin{cases} U_{i,j}^{n**}, & \text{for } j \ge j_i^*, \\ U_{i,j+1}^{n+1} + \dfrac{\Delta z}{\ell_i - s_{\min}}, & \text{for } j = j_i^* - 1, \ldots, 1. \end{cases}$$

It is easy to verify that the scheme is *consistent*, *stable* and *monotone*, in the sense of Barles and Souganidis [1991], provided Δx is small enough. Checking its consistency is straightforward. A sufficient condition for the monotonicity is provided by the condition

$$\Delta x \;\le\; \Delta z \cdot \min_{v_{i,j}>0} \frac{v_{i,j}^n}{|s_{i,j}^n - \ell_i|} \tag{27}$$

where $v_{i,j}^n$and $s_{i,j}^n$ are respectively the approximations for the optimal speed and the sink rate at (x_n, z_j, ℓ_i). Under the above condition, the scheme is also stable.

Finally, we mentioned earlier the first-order problem where the lift is a known function of the distance to the target. An alternate method for solving this problem numerically is to use a modified version of the algorithm we described above. We just eliminate the partial derivatives with respect to ℓ and replace the variable ℓ by the given function $\ell(x)$. We can compare the results obtained with those coming from the explicit calculations carried out in Section 3, and this helps us validate our algorithm.

4.6 Numerical solutions

For our experiments below, we choose $\ell_{\max} = 10$ m/sec, $\Delta\ell = 0.1$ m/sec, $\Delta z = 7.5$ m, so that we have a 200×200 point mesh in the variables ℓ, z. The computed step Δx is equal approximately to 15.9 m. We cut off the boundary layer conditions due to the artificial Neumann boundary conditions and show the results only for $-5 \le \ell \le 5$.

Figure 6 shows the trajectories and MacCready values, for the same realization of lift used for Figure 5. The differences in strategy are apparent: First, in the second-order, uncertain, case, the trajectories stay higher above ground, because the pilot cannot rely on finding lift at a known location. This trajectory reaches its lowest point around 120 km, because of an extended region of sinking air.

The second difference is that the MacCready value is much less constant: it is continuously adjusted to respond to local expectation of future lift. Although it depends on local lift, its strongest dependence is on altitude; as the trajectory gets low around 110–120 km, the MacCready value is reduced: this indicates simultaneously that the pilot is flying slower to conserve altitude, and also that he is willing to accept a weaker thermal than when he is high.

In Figure 6, we show the MacCready function $m(\ell, z)$ for a fixed value of x. Classic theory would say that this function depends only on z (and x), but is constant over ℓ: this gives the response to local changes in lift. In fact, in this model, m has a slight dependence on ℓ, especially near the climb boundary at $m = \ell - s_{\min}$. To some extent, this is an artifact of our continuous-lift model, but it indicates that the deterministic MacCready theory cannot be carried blindly over to the stochastic case.

Large-x limit There is a close analog to the financial theory of the "perpetual option," in the limit of infinite time before expiration. As the distance x increases, the probability increases that the aircraft will encounter a large region of sinking air that is too broad to cross regardless of the strategy, because of the limited altitude band. Then an eventual landout becomes inevitable, and the value function is controlled by the penalty;

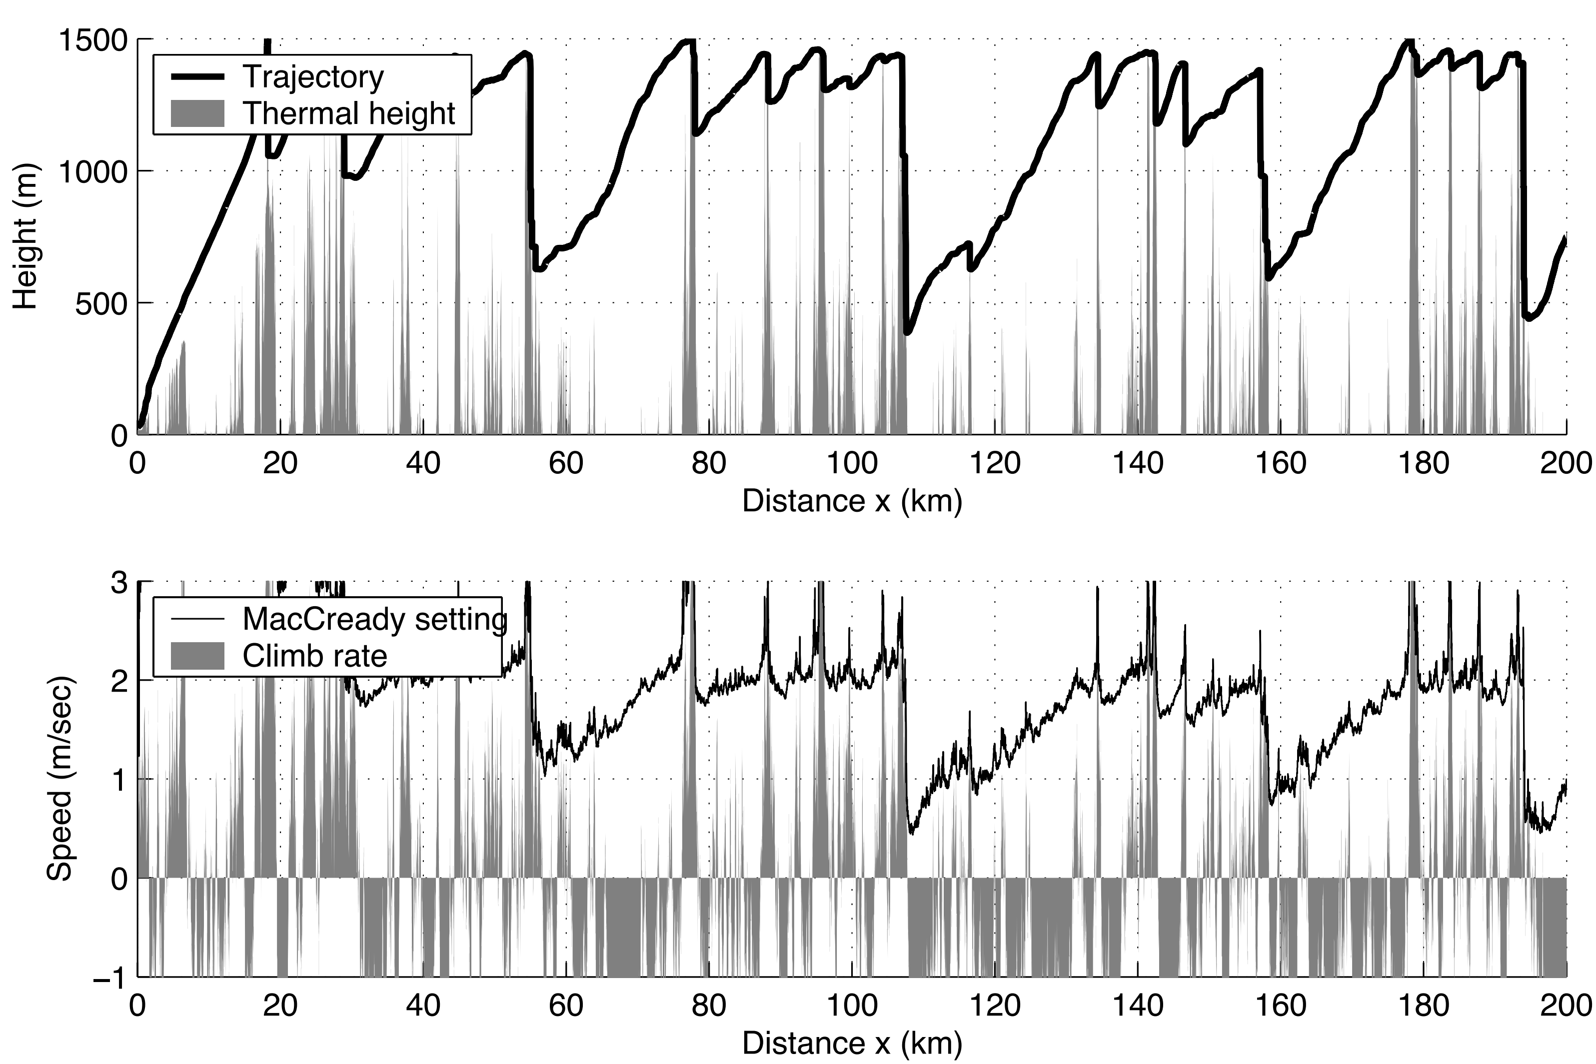

Figure 6: Optimal trajectory and corresponding MacCready values for one realization of lift for the second-order problem, with no knowledge of future lift. As in Figure 5, in the upper panel the thick line is the trajectory, including vertical segments where the glider stops to climb, and the vertical lines below show the maximum climb height. The lower panel shows the lift profile, which is the same as in Figure 5 but here discoverd by the pilot only when he flies through it. Realized average speed is 105 km/hr, less than with perfect knowledge.

mathematically this corresponds to an *Ansatz* of the form $u(x,z,\ell) = U(z,\ell) + x/V_{\text{pen}}$, where the function U is independent of x, and $m(x,z,\ell) = m(z,\ell)$. From its maximum height z_{max}, through still air the glider can travel a distance on the order of $r_0 z_{\text{max}}$, where r_0 is the maximum glide ratio from Section 2.1. The number of independent lift samples encountered is therefore $r_0 z_{\text{max}}/\xi \approx 60$, where ξ is the correlation length (Section 2.3). The sample distance required before encountering such a stretch of sinking air is thus on the order of $\sim \exp(r_0 z_{\text{max}}/\xi)$ which is extremely large, and hence the large-x limit is not relevant in practice.

Nonetheless, our computed profiles for the control variables do appear to be approximately stable after 50–100 km, and Figure 7 is thus a reasonable illustration of the "generic" case for long-distance cross-country flight.

5 Conclusions

Many real-world problems involve the choice of optimal strategy in uncertain conditions. Among these, problems arising from finance play a central role, because of their practical importance as well as their intrinsic interest. In financial settings, the principle of efficient markets suggests that price changes are pure Markov processes as well as martingales, which enables the use of a set of powerful mathematical techniques.

For problems in which the randomness comes from the physical world, there is no principle of efficient markets. Part of the modeling challenge is to correctly incorporate the appropriate degree of partial information. In this paper we have attempted to illuminate this difficulty by considering two extreme cases: one in which the agent has perfect foreknowledge, and the Markov setting in which he has none. The real situation is between these, and a true optimal strategy should incorporate elements of both solutions.

The sport of soaring provides a natural laboratory for this sort of modeling: the aircraft performance is very accurately known, the cockpits are equipped with ample instrumentation and recording capability, and the distances and angles are such that direct human perception is an inadequate guide for competitive performance. For these reasons, the sport has a history of incorporating mathematical optimizations, and we hope that these results may be useful in real competitions.

In the course of developing the mathematical model and in constructing the numerical solutions, we needed to draw on some common but relatively subtle aspects of the mathematical theory of Hamilton-Jacobi-Bellman equations. We hope that this problem can serve as a concrete example to stimulate further development of the theory.

5.1 Future work

Following suggestions by John Cochrane, we would like to propose two extensions of this model as possible subjects for future work.

The first extension is the role of wind. With a component w along the direction of the course, the first dynamic equation in (12) is modified to $dx = -(v + w)\,dt$; we would neglect the crosswind component. Description of the lift is more complicated since the thermals also drift with the wind as they rise, and their angle of climb depends on the local vertical velocity ℓ. This would cause the local lift to no longer be uniquely

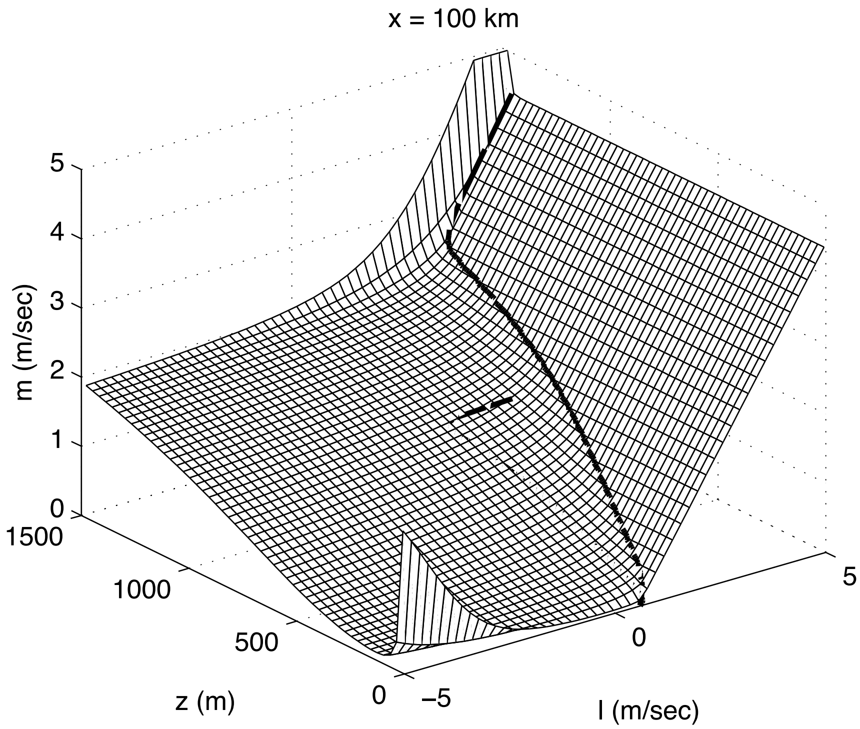

Figure 7: MacCready function for cruising flight, for the same parameters as in Figure 6, as a function of lift ℓ and height z at a fixed distance $x = 100\,\mathrm{km}$. The heavy line is the cruise/climb boundary, where $m = \ell - s_{\min}$. This function is used for smaller values of ℓ than this boundary (to the left); to for larger values of ℓ the glider switches to circling climbing flight. In classic MacCready theory this function would be constant over ℓ and z; in a heuristic modification it would depend on z but be constant over ℓ.

defined; for example, a slowly rising column might be overtaken from below by a rapidly rising column that had originated at a ground location further downwind. One possible formulation would be to use the present model (3) to describe the lift $\ell(x, 0)$ at ground level, and to extend $\ell(x, z)$ into $z > 0$ by the nonlinear convection equation $\ell\,\ell_z - w\,\ell_x = 0$ with an appropriate choice of weak solution. We would then obtain a problem of the form (8) or (14), with $H(\ell, p)$ depending on w.

At a turn point, where the course direction changes, the coefficient w would be discontinuous in x. First-order HJB equations with discontinuous coefficients are known in the area of image processing but theoretical results are incomplete. Careful analysis of the connection conditions across the discontinuity should reproduce the well-known soaring wisdom that one should enter upwind turn points low, and downwind turn points high, to benefit in both cases from the wind drift while climbing.

The second extension is to add a finite time cost in switching between cruise and climb modes: it can take one or two minutes for the pilot to properly “center” a thermal and attain the maximum rate of climb. To handle this effect, it is necessary to introduce a binary state variable describing whether the aircraft is currently in cruise mode or climbing. There may be locations in the state space (x, z, ℓ) where it is optimal to continue climbing if one is already in the thermal, but not to stop if one is cruising. It would be necessary to solve for two different value functions, one for the cruise mode and one for the climb mode.

5.2 Conclusion

In this paper, we propose a stochastic control approach to compute the optimal speed-to-fly and we test the classic MacCready theory. We find that it is verified to a large extent, with small corrections, namely some slight dependence on the value of the lift, which are likely an artifact of our model.

The remaining question is how to use these mathematical results to optimize soaring contest performance. Even with modern mobile computing, it is likely unrealistic to repeatedly solve a nonlinear PDE in two space-like variables, continually updating as conditions change. More plausibly, one could attempt to parameterize the computed solution shown in Figure 7 so as to rapidly evaluate it. This would likely be difficult because it is a function of three variables as well as several parameters. But more fundamentally the details of the mathematical solution depend on the exact model used for the atmosphere, which is realistic only in a qualitative sense. The penalization for landout is also not consistent with actual contest rules.

We expect the value of this work to be in the qualitative lessons one can draw from it. The surface in Figure 7 may be approximately summarized as a function of z, constant over ℓ; in practice this works out roughly to “MacCready equals altitude” (altitude in thousands of feet and speeds in knots). But understanding the theory and solutions presented here can give the pilot confidence and a rational basis for adjusting this strategy as conditions change.

References

G. Barles. Convergence of numerical schemes for degenerate parabolic equations arising in finance theory. In L. C. G. Rogers and D. Talay, editors, *Numerical Methods in Finance*, pages 1–21. Cambridge University Press, 1997.

G. Barles and P. E. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptot. Anal.*, 4:271–283, 1991.

G. Barles, C. Daher, and M. Romano. Convergence of numerical schemes for parabolic equations arising in finance theory. *Math. Models Methods Appl. Sci.*, 5:125–143, 1995.

J. H. Cochrane. MacCready theory with uncertain lift and limited altitude. *Technical Soaring*, 23:88–96, 1999. Available at `http://faculty.chicagobooth.edu/john.cochrane/soaring`.

M. G. Crandall and P. L. Lions. Two approximations of solutions of Hamilton-Jacobi equations. *Math. Comp.*, 43:1–19, 1984.

A. W. F. Edwards. A stochastic cross-country or festina lente. *Sailplane and Gliding*, 14: 12–14, 1963. Reprinted in *Stochastic Geometry*, E. F. Harding and D. G. Kendall, editors, John Wiley & Sons 1974.

W. H. Fleming and H. M. Soner. *Controlled Markov Processes and Viscosity Solutions*. Springer-Verlag, 1993.

R. J. LeVeque. *Numerical Methods for Conservation Laws*. Springer (Birkhäuser Basel), 2nd edition, 1992.

P. B. MacCready, Jr. Optimum airspeed selector. *Soaring*, pages 10–11, 1958.

H. Reichmann. *Streckensegelflug*. Motorbuch Verlag, 1975. Reprinted as *Cross-Country Soaring* by the Soaring Society of America, 1978.

P. E. Souganidis. Existence of viscosity solutions of Hamilton-Jacobi equations. *J. Differential Equations*, 56:345–390, 1985.

A. Tourin and T. Zariphopoulou. Numerical schemes for investment models with singular transactions. *Comput. Econom.*, 7:287–307, 1994.

N. Touzi. *Optimal Stochastic Control, Stochastic Target Problems and Backward SDE*, volume 29 of *Fields Institute Monographs*. Springer, 2013.