

Bounded Rationality and Risk Strategy in Thermal Soaring

John J. Bird*

The Pennsylvania State University

Daniel Sazhin†

New York University

Jack W. Langelaa‡

The Pennsylvania State University

Awareness and management of the risk of failing to encounter lift is fundamental to thermal soaring. When the weather changes or a thermal is missed the pilot may be exposed to a greater risk of landing out. In these situations the pilot may need to alter strategies in order to minimize risk exposure at the expense of speed, often referred to as “gear shifting.” In this work, we explore several models to explain why small changes in the environment can cause large changes in risk exposure, requiring this shifting. We also examine several flight strategies in simulation to define the relative risk and reward for adopting various levels of risk tolerance and for failing to “shift gears” when the risk of landing out increases.

I. Introduction

Thermal soaring is defined by uncertainty. Even if a pilot can see markers indicating the presence of thermals ahead, it is not certain that a thermal will still be working when the pilot arrives. Managing the risk of failing to find a thermal is an essential component in decision making. In high level competitions, failing to complete a task is often disastrous to a pilot’s overall standing in a contest. As a result, pilots must balance their goals of maximizing speed on each glide while minimizing the risk of an outlanding.

While managing risk is key to success in thermal soaring, flight planning and optimization has largely focused on maximizing speed and has not addressed risk explicitly. Since its initial development in the 1930s[1], speed-to-fly theory has been the dominant approach to soaring flight optimization. Paul MacCready’s development of the speed ring made the best speed to fly easier to compute in flight[2], and the MacCready setting has since been used both for speed optimization and as a proxy for risk[3, 4]. While most authors examine risk implicitly through the MacCready setting, Fukada finds the risk tolerance which achieves the highest average speed and which scores the most expected points[5]. Fukada only peripherally examines landing out however, and does not consider the approach a pilot would take to achieving a desired risk level[5]. While speed-to-fly theory and adaptations to it are very powerful flight optimization

*PhD Candidate, Department of Aerospace Engineering. jbird@psu.edu

†Master’s Student, General Psychology. ds5498@nyu.edu

‡Associate Professor, Department of Aerospace Engineering. jlangelaan@enr.psu.edu

tools, they do not provide a pilot a means to manage risk over the course of a flight or competition.

One of the challenges to addressing risk management in soaring is that human decision making is complicated and limited by human capabilities. Thermal soaring is cognitively taxing: there are an inordinate number of possible clouds or thermal sources to sample and every thermal opens up a branching tree of possible choices. Like a chess game, it is nearly impossible to compute all of the possibilities from a given position to determine the right move. Attempting to evaluate all options and make an optimal decision would simply overwhelm the pilot with information[6]. Instead, pilots engage in a number of strategies to omit suboptimal choices which makes the decision-making process manageable. These strategies employ heuristics, mental shortcuts that minimize cognitive workload[7].

In this paper we examine cross-country soaring from the perspective of risk. For our purposes “risk” represents sporting risk – the probability of landing out and no longer having the chance to finish the contest with a good score. We are explicitly not concerned here with risk from the perspective of flight safety: a “failure” ends the flight at a location and situation from which the pilot can make a safe pattern and landing. In this light, failure could also be interpreted as falling out of the lift band and having to “dig out” in a weak thermal, slowing a pilot down enough to preclude a competitive finish.

We seek to understand why risk management is challenging, how sensitive success is to risk, and to define a risk threshold which makes success in competitions likely. We then formulate a model for how humans address risk management, drawing on piloting experience as well as cognitive theory. This leads us to believe that humans bifurcate risk management into two dominant strategies: “racing” and “risk minimization.” Selection of a strategy is determined by the reliability and frequency of lift the pilot expects to encounter. Models of these strategies are implemented in numerical simulations to explore the utility of “gear shifting,” and the sensitivity of speed and task completion percentage to environmental conditions, risk tolerance, and pilot strategy. This leads us to a model of risk management in thermal soaring which employs simple heuristics in a systematic process that pilots can use to aid their decision making in the cockpit.

II. Assessing Risk Exposure

Before risk can be managed, it must be defined and an appropriate level of risk determined. As we discuss it in this paper, risk is the likelihood of landing out on a glide. In order to study the effect of risk on performance, we are explicitly neglecting variables which are present in reality but which could confound this study. In our analysis, we consider only homogeneous environments with consistent thermal strengths. This permits us to isolate the task of evaluating the risk a pilot is taking and the level of risk that is acceptable to succeed in a contest.




Fig. 1 Probability of completing a task given the number of glides required. The completion probability is shown for several levels of risk tolerance on each glide. For long tasks, the probability of completion is very sensitive to the risk tolerance, and even seemingly low risk tolerances can result in a significant probability of landing out.

A. Strategic Risk


Considering the risk of landing out, one can think of each glide as an independent event; a gamble with a probability of finding a climb (success) or landing out (failure). Looking at a contest day, we can consider the sequence of glides required to complete the task and compute the cumulative probability of success. Similarly, a contest is a continuation of such sequences. As such, the risk one accepts on each glide is compounded by the number of glides taken over a contest.

This raises the question: “what level of risk should a pilot try to maintain in order to succeed in a contest?” To answer this question, we assume that most competitors complete each task. This is commonly the case at competition sites with strong and consistent weather; such as in the western USA, Australia, or South Africa. In such contests, it is genuinely possible to tune risk preferences and maintain them over the course of a contest. In places with highly volatile weather, the immediate tactical situations predominate in a pilot’s decision making; when a pilot is simply concerned with staying airborne, strategic concerns become less relevant.


To determine an appropriate risk baseline, we first compute the likelihood of completing a contest day without landing out. We assign an accepted risk of landing out on each glide, $P_{tol} = \{0.01, 0.005, 0.001, 0.0001\}$. Figure 1 depicts the cumulative risk of landing out in a single contest day for several task lengths. This risk can be computed using Equation 1 where n is the number of glides required.

$$P_{success} = (1 - P_{tol})^n \quad (1)$$

It is apparent that very small changes in risk tolerance, perhaps imperceptible other than through long-term feedback,



(a) Probability of completing a contest without landing out assuming each contest day requires 15 glides to complete.



(b) Probability of completing a contest without landing out assuming each contest day requires 25 glides to complete.

Fig. 2 Probability of completing a contest without landing out for two different task lengths. As the contest and task length grows the level of acceptable risk shrinks.

have a very large impact on the likelihood of completing a task. A pilot that flies 15 glides on a contest day and has a risk tolerance of 1% only has an 86 percent chance of completing a task!

B. Cumulative Effect of Risk Over a Contest

To win a contest requires consistent performance over multiple contest days, further compounding the risk of failure. Assuming pilots maintain a consistent risk profile and fly 15 glides per contest day, the probability of completing the contest without landing out can be computed. This relationship is depicted in Figure 2a; in a five day competition a pilot flying at $P_{tol} = 0.01$ would have less than a 50 percent chance of completing the contest without landing out.

For longer tasks, such as in World Championships and competitive National competitions, the effect is even more striking. If we increase the number of glides flown per day to 25, even less risk is acceptable, as depicted in Figure 2b. For pilots flying at $P_{tol} = 0.01$, the likelihood of completing without landing out over a five-day competition falls to 30 percent.

This motivates the concept of a “strategic baseline” risk – a level which provides a good chance of completing every day of a competition without landing out. Figure 2 indicates that for this simple analysis, the strategic baseline risk per glide is about 0.001.

It is important to distinguish at this point between the strategic baseline and risk tolerance, P_{tol} . The strategic baseline represents a risk level which is likely to provide good results in a contest, while the risk tolerance represents the level of risk a pilot actually accepts when planning a glide. While they are nominally the same, there are instances where a higher or lower tolerance is preferable. For instance, Figure 2 can also be examined from the perspective of the number of days remaining in a contest. On the last two days of the competition, a pilot may actually be prudent shifting

to $P_{tol} = 0.005$ or even $P_{tol} = 0.01$ as the likelihood of finishing without landing out at this point is 80 to 90 percent.

When discussing risk tolerance, it is important to recall the gambler's fallacy; gambles have no memory! A pilot who "survives" a series of unlikely gambles on a given day should not be extraordinarily risk averse on the following days to help "replenish his luck." However, if the risks taken give the pilot a clear edge in points, it may be sensible to adopt a low risk tolerance to help protect the pilot's gains.

C. Tactical Risk

How can we estimate how much risk to accept while in the cockpit? Let us consider a pilot who is at the top of the lift band, assessing the thermal options ahead. The pilot picks a line and counts the number of potential thermal sources that can be sampled before running out of altitude and landing out. While a thermal source can be either a cloud or a particularly promising ground feature, we will refer to all options as "clouds" as they are simpler to visualize. Days with cumulus clouds are also useful since a cloud field often will give a fairly good picture of the amount of thermals one can possibly contact. Despite the fact that as the pilot gets lower the thermals are less likely to be connected to the clouds, the number and quality of clouds can still provide feedback as to the reliability of thermals in an area.

Each thermal option can be thought of as either a "hit" or a "miss," just like "heads" or "tails" when flipping a coin. This assumes that each thermal sampled is independent of the rest. There are circumstances that violate this assumption, such as on days with cloud streets, cirrus bands, or convergence lines. However, on days with "popcorn" cumulus and little wind, we believe it is reasonable to assume that thermals are largely independent of each other.

Continuing our coin toss model, a pilot unlucky enough to flip tails for each cloud sampled will land out. We can calculate the probability of a completely failed sequence and then compare it against a strategic baseline of $P_{tol} = 0.001$. So long as the pilot consistently keeps the likelihood of flipping all tails lower than $P = 0.001$, he or she is likely to complete a competition without landing out.

Without any experience other than occasionally encountering thermals underneath clouds, a pilot may expect that finding a thermal is really like a coin toss – 50/50. We can determine how many options are required to achieve a chosen risk tolerance, depicted in Figure 3. In order to maintain $P_{tol} \leq 0.001$, the pilot would need to keep at least ten clouds in range at all times. Note how little the risk changes for a pilot flying "aggressively" with only four clouds to sample as opposed to ten. In the short-term, a pilot who chooses this strategy may even be successful.

However, very small changes in risk exposure have a massive cumulative impact in the long run. The difference in risk accepted from having only seven clouds to ten clouds in a sequence with a fair coin is only one tenth of one percent. On a given contest day, this cannot markedly feel all that different. However, over an eight-day competition with 25 glides per day, this amounts to a 69 percent difference in the probability of completing without landing out!

Experienced pilots know that the likelihood of hitting a thermal under a cloud can be more predictable than a simple coin toss. If the pilot was routinely hitting thermals under clouds, she can reasonably believe that most of the clouds




Fig. 3 Number of thermal options required to achieve a desired risk tolerance assuming that each potential thermal source sampled has a 50% chance of working




Fig. 4 Number of clouds required to maintain a specified risk tolerance as a function of the probability that each potential thermal works.

ahead are “working.” Finding lift is not a certainty however, there is the possibility that a promising cloud dissipates or that the pilot misjudges the thermal location and misses a climb. Furthermore, there are days when the clouds are “dishonest” and the prudent pilot realizes that she must sample many more clouds before contacting a thermal. To incorporate the expected “honesty” of the clouds, we can use a weighted coin model:

$$n_{options} = \frac{\log(P_{tol})}{\log(1 - P_{option\ works})} \quad (2)$$

The number of clouds required to maintain a safe strategic risk profile is depicted in Figure 4. Since the honesty of clouds controls the number of options required, there is a strong emphasis on the degree to which clouds are working. Once the probability of contacting a thermal under a cloud is less than 50%, it becomes nearly impossible to maintain the strategic risk baseline. Once the probability that thermals work exceeds 70%, the number of options the pilot must




Fig. 5 Number of clouds required to maintain a desired risk tolerance as a function of the probability that each potential thermal works, assuming that one thermal with a 95% chance of working is known (e.g. a power plant, or a gaggle-marked thermal).

maintain becomes considerably more manageable. The risk is non-linear; when the days are “consistent” and “reliable,” the pilot can afford to have few options available and still have a very low probability of landing out. On the other hand, as the reliability of the lift diminishes, the pilot must maintain many more options in order to maintain an acceptable strategic risk exposure.

On a tactical level, maintaining even one “very likely” thermal option can significantly reduce a pilot’s risk exposure. If the pilot looks ahead and realizes that in his sequence of clouds to sample, there is one source that has a 95 percent chance of working; the likelihood of landing out on the whole sequence is much lower. Pilots who are especially good at reading ground sources or clouds can factor this in their tactical choices. Equation 2 can be modified to take into account one very likely cloud:

$$n_{options} = \frac{\log(P_{tol}) - \log(1 - P_{likely\ option\ works})}{\log(1 - P_{option\ works})} \quad (3)$$

The number of clouds required to achieve a desired risk tolerance is considerably reduced, depicted in Figure 5.

III. Modeling Pilot Decision Making

Now that we have a broad sense of how the quantity and reliability of thermals affect risk both in the long run of a whole competition and in the short run of a glide, the challenge is to model this decision making in terms that can be applied in the cockpit. Recognizing that it is *humans*, not computers which make decisions in soaring competitions, we must consider how people manage risk and make decisions under uncertainty when we develop risk management strategies.

A. The Brain as a Computer

While the brain is not a computer, it shares some basic characteristics with them: it processes inputs from the environment through the body's senses, integrates this data into perceptions, and generates a motor-driven output (i.e. moving the stick). The brain as an information processing unit has extraordinary capabilities but also significant limitations. Cognitively, one of the greatest limitations is working memory, limited to approximately 30 bits[8]. On the other hand, the brain has nearly endless capacity for long-term memory[9].

As a result, the brain is very effective at using long-term memory as a work-around for the limitations of working memory. Over time, the results of favorable computations become encoded and are retrieved given the right pattern of inputs. When a pilot identifies a cloud as particularly favorable, it is the result of having flown under many similar clouds with good outcomes. Thus, the brain offloads most cognitive tasks to programs or schemas of action; when there is a set of stimuli, to generate a certain output. It is through this process that many tasks become largely automatic or intuitive; the brain no longer needs to engage effortful cognitive processing in order to generate a good output.

Heuristics, or rules of thumb, are processes the brain uses to simplify the decision-making process[10]. For instance, a simple risk-related heuristic is to take every thermal on a cross country task. By taking every thermal, decision making is drastically simplified and the pilot is unlikely to land out. With experience, pilots refine and expand their heuristics, encompassing more and more variations in the environment.

Often, the goal is to make a good decision, not necessarily the best decision. Doing so is referred to as "satisficing"[11]. This permits acceptable outcomes while reserving cognitive capacity for other tasks. A sophisticated form of satisficing which pilots likely use is "elimination-by-aspects," which employs successive heuristics to exclude clearly suboptimal solutions[12]. For instance, in choosing the next cloud, a pilot may use criteria such as: "don't deviate more than 30 degrees", "fly under the clouds", "fly MC 1 ($m s^{-1}$) and stay on the upwind side of the course." By engaging these heuristics in sequence, the pilot can very quickly narrow down a large field of potential thermals to several "lines," saving the trouble of processing every single cloud and its respective decision tree.

B. Decision-Making Frames

While the elimination-by-aspects strategy helps pilots rapidly make decisions, the sets of heuristics employed can vary by situation. The heuristics a pilot uses when struggling at low altitude trying to "minimize risk" on a blue day are distinct from the heuristics used when at altitude, cloud streets are plentiful, and the pilot is "racing". Pilots choosing to engage their "risk minimization" program will process their environment differently than pilots engaged in "racing." These programs and the heuristics associated with them are called "decision-making frames."

The way in which frames are managed depends on how a pilot appraises his tactical situation and what losses are most immediate in his mind. In gliding, pilots are conflicted between two kinds of losses: losses in speed relative to competing pilots, and the catastrophic loss of landing out. Since people are averse to losses[13], the manner in which

losses are processed will greatly affect decision making. When facing an uncertain gamble that is framed as a choice among losses, people tend to overweight the impact of a loss[14, 15]. This is known as loss aversion, a phenomenon described by prospect theory[16].

Depending upon whether losing efficiency or landing out weighs more heavily in a pilot's mind will determine whether that pilot will make decisions from a racing or risk minimization frame. This is because the pilot evaluates gains and losses relative to the most salient loss in her mind[9]. Framing refers to the manner in which costs/benefits or risks are presented and interpreted. For instance, the likelihood a patient accepts a treatment is affected by whether he is told by his doctor that a treatment has a 95 percent survival rate, or that 1 in 20 people die, despite both options being mathematically equivalent. In gliding, when a pilot adopts a risk minimization frame, any action that increases the pilot's risk of landing out is experienced as a greater loss. When a pilot is in a racing frame, any action that diminishes speed is experienced as a loss. This is what drives a pilot to leave a thermal when a competitor merely bumps through it and continues: it is painful to give up points!

When a pilot's losses are reframed, such as when a pilot gets lower and becomes more concerned with the prospect of landing out than maximizing efficiency, the heuristics that are used in decision making change. Under the same conditions, the same pilot can generate distinctly different outputs depending on the decision-making frame adopted. This frame shift is the core of gear shifting; when the availability and reliability of lift dictate that a pilot must shift into a risk minimizing frame, the pilot who shifts earlier is more likely to make it home. If conditions change from unreliable to reliable, the pilot who shifts into racing will fly faster.

For the present work, we define models of these two decision-making frames and their respective heuristics for further analysis. This will enable us to explore the effect that the frames have on the speed a pilot achieves and how effective frame switching can be in managing risk in competition soaring.

1. Racing

The first behavior we call the racing frame. In this mode the pilot seeks to maximize speed while attempting to maintain a chosen risk tolerance. Diverting the flight path and stopping to thermal both decrease average speed, so the pilot will avoid these actions when possible. In a racing frame, the pilot will reevaluate the flight plan if an anticipated thermal fails to work, but will continue on a path that minimizes deviations and maintains a high speed. The pilot will attempt to satisfy a risk tolerance, but will not make large deviations or slow down to do so.

When racing, the flight path is generated using an iterative method which is initialized with a direct flight to the next turnpoint. Thermals are added to the flight plan until the number of thermals is sufficient to satisfy the acceptable risk of landing out. At each iteration, the longest leg between potential thermals is rerouted to visit another potential thermal (following a heuristic that failing to find a thermal after the longest glide will leave the pilot at the lowest altitude and is thus the most likely to cause a landout). The closest thermal which is in range along that leg is added to the plan. This




Fig. 6 Path planning method for a pilot in the racing frame. The pilot starts with a direct glide to the next waypoint (dashed line leading to the red point) and at each iteration chooses the thermal (blue points) nearest the longest segment of the path. This process is repeated until there are no thermals remaining in range or until the desired risk tolerance is met. The pilot is assumed to be able to detect any possible thermal within gliding range. The thermal added to the path at each iteration is depicted by iteration number and intermediate paths are shown in gray.

is repeated until the plan satisfies the risk constraint or until there are no more thermals within range. The planning model is intended to identify “lines” of favorable conditions, as illustrated in Figure 6.

When planning, the pilot is assumed to be able to predict the approximate location of any thermal within range (with a standard error of 400 m). When nearing a thermal, the pilot can determine the precise location and whether or not the thermal is working when within 700 m of the thermal center. When the pilot encounters a thermal, the planning model is run again and if the risk tolerance can be met from the current altitude then the thermal is skipped. If the risk tolerance is exceeded by skipping the thermal then the pilot stops to exploit the thermal. In all cases, thermals are not exploited if the current altitude is more than 80% of the boundary layer depth.

If a thermal is encountered which does not work, the pilot will replan the flight path from the current location. While the pilot attempts to keep the probability of landing out below the desired risk tolerance via planning, no action is taken if the risk rises above this level (i.e. if there are not enough thermals in range to allow the tolerance to be satisfied).

2. Risk Minimization

A second behavior is implemented where the pilot is primarily concerned with remaining aloft, but also desires to complete the task. We call this the risk minimization frame. In this frame, the pilot will seek out any lift which brings him closer to the next turnpoint, and will exploit any thermal encountered below 80% of the maximum altitude.

Again, the pilot can detect any thermal in range, and can determine if the thermal is working when within 700 meters. The pilot chooses as a destination the nearest thermal which brings the pilot closer to the waypoint. The pilot exploits that thermal if it is working or repeats the thermal selection process if the thermal does not work. The path




Fig. 7 Path planning for a pilot in the risk minimization frame. The pilot will fly to the nearest thermal option (in blue) which brings the pilot nearer to the next turnpoint (red). The pilot will accept large deviations to minimize the distance which must be flown before encountering a potential thermal.




Fig. 8 Schematic describing the decision-making process for risk aware thermal soaring. Gear shifting occurs when the pilot receives new information which reveals a change in the risk situation and reevaluates the decision-making frame. This could happen for example when missing an expected climb or when reaching a “blue hole” with few clouds.

planning strategy in risk minimization mode is depicted in Figure 7.

C. A Summary of The Decision-Making Process

Figure 8 illustrates the decision-making process we present here. Selection of the risk tolerance is described in Section II.A. Assessment of the tactical situation is described in Section II.C. Section III.B describes an approach to managing tactical risk which is rooted in the psychology of human decision making. At each step in the decision-making process heuristics are employed which reduce the cognitive load on the pilot and allow this schematic to be traversed rapidly in flight. The heuristics employed in this work are summarized in the flowchart.

IV. Monte Carlo Simulations

The coin toss model motivates the need for risk management and the psychology perspective introduces the concept of decision-making frames. This provides insight into the why and how of risk management but these approaches do not lend themselves well to analysis of flight to turnpoints, in limited altitude bands, or in the vicinity of areas of inhibited

lift. To enable deeper exploration of risk strategy in thermal soaring, we used a Monte Carlo approach.

Pilot behaviors representing the racing and risk minimization frames are implemented in a numerical simulation. A configurable frame-switching logic is implemented which permits allowing or prohibiting switching between frames. These pilot behaviors are then simulated over several hundred competition tasks to evaluate the effect of risk tolerance and frame switching. This is conceptually similar to work done by previous authors[3, 4] except that we explicitly explore the pilot’s risk tolerance rather than using MacCready setting as a proxy.

A. Simulation Environment

Thermals are defined using a Gaussian model:

$$w = w_{scale} \exp\left(\frac{-r}{R}\right) \tag{4}$$

Where r is the distance between the aircraft and the thermal center, R is the characteristic scale of a thermal, and w_{scale} is the maximum updraft velocity. At altitudes above the boundary layer top the thermal updraft velocity is set to zero.

Candidate thermal locations are drawn from a uniform distribution $\mathcal{U}(x_{scale})$. A weighting function is then applied to prevent thermals from occurring extremely close to each other. The weight is defined:

$$\theta_{inhibit} = \frac{1}{1 + \exp(-w(X_{candidate}) + 0.1)} \tag{5}$$

Where $X_{candidate}$ is the candidate location for a new thermal and $w(X_{candidate})$ represents the thermal updraft velocity at the candidate location due to any thermals already accepted into the updraft field. The factor 0.1 is used to allow thermals to slightly overlap, forming multi-core thermals. A uniform random thermal acceptance probability is generated, if it exceeds the weight then the thermal is accepted and added to the field. The parameters of each thermal are summarized in table 1. At generation, each thermal is assigned a working or not working state with a configurable probability.

Table 1 Parameters of the thermals used in the Monte Carlo simulations. The thermal strength is kept constant to isolate the effect of risk management and to simplify computation of the appropriate MacCready number. Approximate conversions to common U.S. units are given in the second column.

R	$\mathcal{N}(600 \text{ m}, 10000 \text{ m}^2)$	$\mathcal{N}(2000 \text{ ft}, 100000 \text{ ft}^2)$
w_{scale}	3.0 m s^{-1}	6 knots
z_i	1000 m	3300 ft

The aircraft model is a simple kinematic model whose states are the east, north, up position of the aircraft and the heading angle. Inputs to the system model are turn rate and airspeed. Lateral dynamics are neglected (the commanded

turn rate is achieved instantly) and the longitudinal aircraft dynamics are simulated with a first order lowpass filter on the airspeed command with a time constant of 5.0 s. The aircraft dynamic equations are summarized in equation 6.

$$\frac{\partial}{\partial t} \begin{bmatrix} x_{east} \\ x_{north} \\ h \\ \psi \end{bmatrix} = \begin{bmatrix} v_{east} \\ v_{north} \\ \dot{h} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} v_{ias} \sqrt{\sigma} \sin \psi \\ v_{ias} \sqrt{\sigma} \cos \psi \\ w_s(v_{ias}) \sqrt{\sigma} + w_{wind} \\ \dot{\psi} \end{bmatrix} \quad (6)$$

Where $\dot{\psi}$ is the commanded turn rate and σ represents the ratio of sea level density to the density at the aircraft location, computed using the 1976 standard atmosphere model. A quadratic speed polar is used to represent the sailplane's aerodynamic performance, given in equation Equation 7. The polar approximates a Discus 2 at a wing loading of 35 kg m⁻² (airspeed and sink rate are both specified in m s⁻¹)

$$w_s(v_{ias}) = -0.00285 v_{ias}^2 + 0.146 v_{ias} - 2.51 \quad (7)$$

B. Pilot Model

A pilot model is implemented which incorporates both basic airmanship and risk-based decision making. The airmanship portion is responsible for controlling the sailplane. A higher-level model implements the behaviors described in Section III.B.

1. Airmanship: Airspeed Selection and Trajectory Tracking

It is necessary to simulate the pilot's behavior in controlling the speed and direction of the aircraft. Airspeed commands are generated using speed-to-fly theory, with the MacCready setting determined by the pilot's frame. In the racing frame, the MacCready value is set to the climb rate achieved given the thermals defined in Table 1. In the risk minimization frame the MacCready setting is zero to maximize range. To represent the response time of the pilot and aircraft, the airspeed command is filtered with a first order lowpass filter with a time constant of 5.0 seconds. The filtered airspeed command is directly used in the aircraft state equations.

Trajectories for the pilot are defined as a series of points to visit, with each point being either a turnpoint or a potential thermal. The trajectory generator depends on the pilot frame as described in Section III.B. To follow the path, Park's nonlinear trajectory following controller[17] is used to generate turn rate commands. When thermalling the pilot tracks a circular orbit around the thermal location.

2. Risk Management Strategy

In order to determine the effect that frame shifting can have, two approaches to risk management are implemented. The first attempts to optimize the flight at a given risk tolerance at all times. While the pilot plans a path which seeks to maintain a given strategic baseline risk, it is not guaranteed that this threshold can be met at all times. This pilot will remain in the racing frame regardless of risk, pressing ahead at all times. In the simulation results this is referred to as the “racing” strategy.

The second pilot behavior switches frames depending on the current risk. The pilot continuously monitors the risk of landing out. If the risk rises above the pilot’s risk tolerance, the pilot will switch frames into risk minimization mode in an attempt to mitigate the risk of landing out. This mixed approach is called the “gear-shifting” strategy.

C. Results


Pure racing and gear shifting strategies are simulated for 350 iterations of a triangular assigned task 220 km in length. The start and finish cylinders each have a radius of 3 km, while turnpoints have 500 meter radii. The thermal reliability is varied ($P_{thermal\ works} = \{0.4, 0.7\}$) and several risk tolerances ($P_{tol} = \{0.1, 0.05, 0.01, 0.001\}$) are studied.

Figure 9 compares the flight path and altitude profile for gear shifting and non gear shifting pilots at a risk tolerance of 0.001 and for a thermal reliability of 0.7. Over small segments the flight paths of the two pilots are similar. However, when one pilot switches into risk minimization mode significant differences arise. When entering a tricky area, especially after missing a thermal, the gear-shifting pilot will occasionally make large deviations to remain connected with lift. At times, gear shifting occurs immediately upon finishing a climb if a path cannot be found that satisfies the risk tolerance. Gear shifting can also be triggered when an expected thermal fails, for example in Figure 9d at $t \approx 3000\ s$ (approaching the first turnpoint in Figure 9b). In this case, the pilot seeks a climb, trying several potential thermals before locating one.


Figure 10 compares the effect of two different risk levels on the flight path and altitude band used when conditions are inconsistent (thermals have a probability of working of 0.4) with no gear shifting. The more risk tolerant pilot typically uses more of the altitude band and flies a relatively direct course. The risk averse pilot makes large deviations to try to maintain the chosen risk tolerance, especially when a planned thermal does not work.

V. Discussion


The Monte Carlo simulations provide a means to evaluate the effect of risk management techniques and risk tolerance while navigating a task with altitude constraints. In particular, the simulations can reveal the effect of flight strategy on the probability of landing out and on the speed achieved on course.




(a) Sample flight path for a pilot who remains exclusively in the racing frame.



(b) Sample flight path for a pilot who shifts between racing and risk minimization frames.




(c) Sample barogram for a pilot who remains exclusively in the racing frame.




(d) Sample barogram for a pilot who shifts between racing and risk minimization frames.


Fig. 9 Effect of gear shifting on flight path and altitude utilization for a pilot with a risk tolerance of 0.001 with reliable thermals (probability of working is 0.7). Thermalling is depicted in blue, racing in green, risk minimization in cyan, and final glide in red. Potential thermal sources are depicted as blue dots and turnpoints as red dots. The task starts at $n=10$ km, $e=10$ km and proceeds anti-clockwise.




(a) Sample flight path for a pilot flying with $P_{tol} = 0.1$



(b) Sample flight path for a pilot flying with $P_{tol} = 0.01$




(c) Sample barogram for a pilot flying with $P_{tol} = 0.1$



(d) Sample barogram for a pilot flying with $P_{tol} = 0.01$

Fig. 10 Effect of risk tolerance on flight path and altitude utilization with unreliable thermals (probability of working is 0.4) for pilots always in the racing frame. Colors and task are as in Figure 9.



(a) Probability of failing to complete a task when the probability that a given thermal works is 0.7, the horizontal axis is plotted on a log scale.

(b) Average cross-country speed distributions when the probability that a thermal works is 0.7


Fig. 11 Probability of task completion and mean cross-country speed when thermals have a 70% chance of working and the pilot remains in the racing frame at all times. In consistent conditions, the pilot can achieve small rates of landing out in a pure racing frame with relatively small risk tolerances. Further, small risk tolerances are associated with only moderate reductions in average speed.

A. Effect of Risk Tolerance on Mean Speed and Probability of Task Completion


When conditions are consistent, risk tolerance and mean speed are not strongly related over a broad band of risk tolerance. Figure 11b shows that similar speeds are attained for risk tolerances between 0.01 and 0.1 when thermals have a probability of working of 0.7. This is likely because in consistent conditions only a few more potential climbs can significantly decrease risk, minimizing the deviation required. Only at $P_{tol} = 0.001$ does risk tolerance significantly affect the shape of the average speed distribution. The 0.001 level is what we identified as an appropriate “strategic risk baseline.” This explains why sailplane racing is such a challenge: at the strategic baseline risk, both speed and probability of landing out are sensitive to the risk tolerance, so the pilot must walk a careful line between flying efficiently and landing out.

Figure 11a shows that even when the mean speed is unaffected, risk tolerance has a substantial impact on the probability of finishing the task. The risk of landing out is lower than suggested in Section II.A. This is because when a thermal is missed, the pilot creates a new plan which attempts to maintain the desired risk level.

When conditions are unreliable, risk tolerance controls speed much more strongly. Figure 12b shows that the speed distribution varies progressively as a function of risk tolerance when thermals have only a 40% chance of working. Reducing risk in unreliable conditions requires many thermals, keeping this many thermals in range can require large deviations. This is illustrated in the sample flight paths in Figure 10. The most striking result is that when thermals are unreliable, the risk of landing out is very large. Even for a per-glide risk of 0.001, the pilot lands out about 30% of the time. This indicates that it is often impossible to keep enough thermals in range to achieve this risk level.



(a) Probability of failing to complete a task when the probability that a given thermal works is 0.4. the horizontal axis is plotted on a log scale.




(b) Average cross-country speed distributions when the probability that a thermal works is 0.4

Fig. 12 Probability of task completion and mean cross-country speed when thermals have only a 40% chance of working and the pilot remains in the racing frame at all times. In unreliable conditions, the probability of landing out in a pure racing frame is unacceptably high, greater than 30%. The probability of landing out is considerably higher than Figure 1 suggests, this indicates that it is often impossible to maintain enough options to satisfy any risk tolerance.

B. Impact of Gear Shifting Strategy

Unsurprisingly, gear shifting reduces the risk of landing out considerably. Figure 13a shows that by changing frames, the chance of landing out is reduced for every risk tolerance. Under consistent conditions (potential thermals worked with 70 % probability) not a single pilot landed out in 350 task simulations for risk tolerances of 0.01 and 0.001. The risk reduction in gear shifting comes at the expense of speed however. Figure 13b shows that speeds are reduced by more than 5 km h⁻¹ for each risk tolerance. In fact, depending on the acceptable risk of landing out on a contest day, it may be advantageous to abandon gear shifting. For example if the acceptable risk of landing out is 5% (perhaps reasonable on the last day of a close contest), the pilot can achieve a higher speed by adopting a risk tolerance of 0.01 but staying in the racing frame than by flying very aggressively (risk tolerance of 0.1) and using gear shifting.


The effect of gear shifting in unreliable conditions is illustrated in Figure 14. The most obvious effect is that it reduces the probability of landing out by approximately a factor of 10, from greater than 30% to less than 3%. Comparing Figure 13a and Figure 14a we can see another interesting effect: the gear-shifting pilot has a lower risk of landing out in unreliable weather than in consistent conditions at high risk tolerances ($P_{tol} = 0.01$). The reason for this is likely three-fold. First, in unreliable weather the pilot will almost always take any thermal encountered, as skipping a thermal would violate her risk tolerance. Second, it takes very little to drive the pilot into a risk minimization frame, so when a thermal doesn't work the pilot is on average higher in energy and has more freedom of action to mitigate the risk. Third, in reliable weather the pilot will fly with fewer options at a given risk tolerance so the loss of one thermal can sharply increase the risk of landing out.



(a) Probability of failing to complete a task when the probability that a given thermal works is 0.7 and the pilot can shift gears. The horizontal axis is plotted on a log scale.

(b) Average cross-country speed distributions when the probability that a thermal works is 0.7 and the pilot can shift gears.

Fig. 13 Probability that the pilot fails to complete the task and the speed distribution for a pilot who can shift gears. Gear shifting can substantially reduce the likelihood of landing out, at the expense of some speed. For smaller risk tolerances not a single pilot in 350 tasks landed out.



(a) Probability of failing to complete a task when the probability that a given thermal works is 0.4 and the pilot can shift gears. The horizontal axis is plotted on a log scale.

(b) Average cross-country speed distributions when the probability that a thermal works is 0.4 and the pilot can shift gears.

Fig. 14 Probability that the pilot fails to complete the task and the speed distribution for a pilot who can shift gears. Gear shifting again substantially reduces the risk of landing out, in this case by more than a factor of 10. In some cases, the pilot is actually less likely to land out than when gear shifting in reliable conditions.

Interestingly, while the risk of landing out is not dramatically different across risk tolerances, the mean speed varies considerably with risk tolerance in unreliable conditions when the pilot can shift gears, depicted in Figure 14b. This is likely because the deviations required to keep a low risk level are considerable. The effect would indicate that in unreliable conditions the pilot should take bigger risks on glides but aggressively switch into a risk minimization frame if a planned thermal does not work out.


VI. Interaction of Risk and Reward

Throughout the paper so far, we have focused exclusively on risk as a driver of decision making in thermal soaring. From this perspective, behaviors bifurcate into two distinct frames. The existence of these frames is supported by risk-management psychology and the experience of many cross-country pilots who frequently discuss “switching gears”. Simplifying these behaviors to their cores permits us to determine the effect these behaviors can have on the risk of landing out and on speed, but in some cases these behaviors are not so distinct.

When the objective is to maximize speed while completing a contest (rather than exclusively keeping risk below a desired level), the pilot will no longer fly the “pure” version of these frames and will adjust her outputs accordingly. In the real world, even when pilots are in a racing frame, they are still somewhat concerned with the risk of landing out. Furthermore, even when a pilot is in a risk minimization frame, she will still consider how her choices will affect her speed. As such, we are proposing a model of decision-making in which a pilot chooses a frame (racing or risk minimization) and runs its respective heuristics. Once the output is generated, the pilot will adjust the result depending upon the risk/reward of the tactical situation.

As we demonstrate in Section II, the probability of landing out is controlled by risk tolerance. Section V.A demonstrates that speed is also sensitive to risk. We can imagine these relationships schematically in Figure 15. In reality, the pilot only has imprecise knowledge of the relationship between speed, land out probability, and risk. In order to optimize speed while respecting the strategic risk baseline, the pilot must adopt an iterative approach. For a pilot in a risk minimizing frame this means “tuning” to increase speed without taking too much risk. A pilot in a racing frame will tune to decrease risk in ways that have little effect on efficiency.

This is likely to lead to behaviors where the pilot chooses a frame (perhaps subconsciously) which prioritizes whether losing speed or minimizing risk is the dominant concern and subsequently adjusts the output based on the secondary concern. The pilot chooses a frame based on the *perceived* relative sensitivity of speed and probability of landing out to the risk tolerance. This explains why pilots rapidly “shift gears” when experiencing a change in weather or an unexpectedly high risk situation – the pilot is confronted with the fact that she does not know the shape of the curves depicted in Figure 15 or her true position on them. Shifting gears provides the opportunity to gather information and tune risk tolerance while avoiding a high risk of landing out. Similarly, in slowly deteriorating or improving conditions a pilot may not change gears until having to “dig out” or when a late starter joins them in a thermal.



(a) In an environment which lightly penalizes risk, the pilot can initially be much more aggressive. We expect that this permits an initial behavior closer to the optimal speed while not exposing the pilot to too great a risk of landing out.

(b) In an environment which progressively and heavily penalizes risk taking, it is likely better to begin in a risk minimization frame and slowly tune the risk tolerance to increase speed.

Fig. 15 Two examples of the behavior that could occur when a reward behavior is introduced into the decision-making process. Depending on the relative shapes of the speed and failure curves, a racing or risk minimization frame might be chosen. In either case the pilot will tune their behavior toward the optimum from their starting frame. The pure risk minimization frame can be imagined as the left and racing as right side of the figure.

In this model, the two frames are still highly relevant, but they become a statement of what a pilot initially experiences as a loss. Whether the pilot is primarily concerned with losing efficiency or landing out anchors how he appraises his situation and is the primary driver of decision making. Subsequently, the pilot will tune the output to satisfy the secondary objective of maximizing speed or minimizing risk accordingly. When racing, slowing down is experienced as a loss but the pilot tunes his outputs based on the possible paths to increase the number of thermals available. For instance, if the pilot has two paths available that are nearly equally optimal in speed, but one path has more potential thermals, it is natural that even a pilot in a racing frame would sacrifice a little bit of speed for a path that meaningfully minimizes his risk exposure. On the other hand, when minimizing risk, anything that increases the likelihood of landing out is experienced as a loss. However, if a pilot has two paths available that are nearly equal in risk, but one path is faster, it is natural that the pilot would seriously consider taking a little bit more risk to meaningfully increase speed. The place that tuning holds in the decision-making process is depicted schematically in Figure 16.

Sometimes, pilots can be in completely different frames and their adjustments to their initial outputs can essentially converge on the same decision. Consider a scenario: two pilots are side-by-side and there is only one cloud ahead of them which has a 70 percent chance of working. One pilot is in the risk minimization frame; he wisely recognizes the high risk of landing out in this situation. However, since there is no chance of finding another thermal, he realizes that flying optimal MC speed does not meaningfully lower the likelihood of finding that thermal and speeds up accordingly. The other pilot is in a highly aggressive “racing” mode and is driving hard toward that one thermal, figuring that the reward of this particular climb justifies the risk. Both pilots are flying in the exact same manner, despite being in different frames.

However, sometimes the risk minimization and racing frames can lead to very different outputs, even when pilots are tuning their outputs to consider the effects of both risk and reward. To continue the proposed scenario, once the




Fig. 16 Considering reward in decision making can be thought of as adding an additional step to the process described in Figure 8. “Tuning” is the process of invoking some elements of the secondary decision making frame to either increase efficiency (when in the risk minimization frame) or decrease risk (in the racing frame). The strength of the tuning effect will depend on the relative intensity of the loss and reward felt by the pilot.

pilots climb up to cloudbase, they must now consider how they will pursue their next glide. They see a blue hole ahead and have two options: accept a major deviation around it, or make a highly aggressive dash across the middle, with very few thermal options on the other side. The pilot in the “racing” frame charges across the blue hole whereas the pilot in the “risk minimization” frame chooses a very different course which increases the number of clouds available in order to limit his risk exposure.

In such a case, the “racing” pilot may tune her output by flying somewhat slower across the blue hole to arrive at the other side at a greater altitude and thus able to reach more potential thermals. The “risk minimization” pilot may tune his output by accepting a path with slightly fewer clouds and flying closer to the MacCready speed than in the “pure” version of his frame. However, the initial outputs from the racing and risk minimization frames could be so divergent that it may be impossible for the pilots to converge on the same decision. When confronted with a tactical situation where shifting from racing to risk minimization yields a very different output, choosing the right frame becomes especially consequential as choosing incorrectly can be extremely costly.

VII. Conclusion

We explored the decision-making process to manage sporting risk in thermal soaring. We began by determining the level of risk which is appropriate for success in contest flying. Because landing out in most competitions is extremely costly, pilots must avoid landing out even once in order to be competitive. We show that pilots generally cannot accept a risk tolerance greater 0.001 and expect to succeed in the long run. The probability of landing out in a contest is extremely sensitive to risk: a risk tolerance of 0.01 greatly increases the likelihood of landing out when assessed over several days. However, a pilot may be justified in increasing risk tolerance toward the end of the competition. The sensitivity of landing out to risk motivates the definition of what we call the “strategic baseline” – a level of risk which

provides an acceptable probability of landing out.

Next we consider how this translates into the cockpit. Starting from the strategic risk permissible in a contest, we define tactical risk: the risk one can accept on a given glide. We use a coin toss model to assess how thermal reliability and the number of options to sample affect a pilot's risk exposure. We find that when thermal reliability is low, it is almost impossible to have enough options to maintain a manageable risk threshold. On the other hand, when reliability is high, the pilot can maintain very few thermal options and have very low risk exposure.

The sensitivity of tactical risk to the number and reliability of thermal options motivates the existence of two modes: racing and risk minimization. These decision-making frames are rooted in cognitive science: they represent an expression of what a pilot experiences as a loss in a given situation. We assert that "gear shifting" often discussed in soaring represents a transition from one frame to the other.

By using Monte Carlo simulations we are able to demonstrate that gear shifting can be used to significantly reduce the risk a pilot is exposed to if the pilot shifts frames when their risk tolerance is violated. We note that shifting gears carries an efficiency penalty as it requires greater deviations and slower speeds to maximize the number of options which can be sampled.

The steps outlined in this paper and illustrated in Figure 16 constitute a cognitive model for managing risk in thermal soaring. Assessing the level of risk, choosing a decision-making frame, tuning the dominant frame, and looking for new information while carrying out a flight plan forms a loop similar to the famous "OODA" loop for decision making[18]. While individual pilots may use different heuristics as they progress through this loop, the structure provides a systematic approach to evaluating and managing risk in thermal soaring. From the results of this investigation, the authors recommend several heuristics that can be applied by pilots in the cockpit to improve thermal soaring performance:

- The risk of landing out in a competition is extremely sensitive to the risk taken on each glide. The acceptable risk of landing out on each glide must be very small to complete a competition successfully.
- If more than half of the clouds are working, a pilot can become more selective about thermal choices. If fewer than half of the clouds work, conservatism is required to avoid landing out.
- Improvements in thermal "hit" probability can dramatically improve speed and reduce risk.
- Having strong confidence in at least one lift source ahead greatly diminishes risk exposure on the current glide.
- When assessing a tactical situation, ask yourself "should I be more concerned with speed or landing out?"
- In reliable conditions, a racing frame can be maintained with relatively few options. A low risk tolerance costs little in speed while reducing the risk of landing out.
- In unreliable conditions, a low risk tolerance reduces speed more than it reduces the risk of landing out. Shifting to risk minimization early can significantly reduce the risk of landing out, permitting a higher risk tolerance.
- Consider an immediate shift to risk minimization when a good "line" through the cloud field ahead is not clear.

References

- [1] Hirth, W., *The Art of Soaring Flight*, 1938.
- [2] MacCready, P. B. J., “Optimum Airspeed Selector,” *Soaring*, 1958, pp. 10–11.
- [3] Cochrane, J., “MacCready Theory with Uncertain Lift and Limited Altitude,” *Technical Soaring*, Vol. 23, No. 3, 1999, pp. 88–96.
- [4] Teter, M. P., “A Monte Carlo Approach to Competition Strategy,” *Science and Technology of Low Speed and Motorless Flight*, 1979, pp. 389–397.
- [5] Fukada, Y., “Speed to Fly with Management of the Risk of Landing Out,” *Technical Soaring*, Vol. 24, No. 3, 2000.
- [6] Janis, I. L., and Mann, L., *Decision making: A psychological analysis of conflict, choice, and commitment*, Free Press, New York, NY, 1977.
- [7] Kahneman, D., and Tversky, A., “Choices, values, and frames,” *American psychologist*, Vol. 39, No. 4, 1984.
- [8] Kaczmarzyk, M., Francikowsk, J., Łozowski, B., Rozpedek, M., Sawczyn, T., and Sułowicz, S., “The bit value of working memory,” *Psychology and Neuroscience*, Vol. 6, No. 3, 2013, pp. 345–349. doi:10.3922/j.psns.2013.3.11.
- [9] Bartol, T. M., Bromer, C., Kinney, J., Chirillo, M. A., Bourne, J. N., Harris, K. M., and Sejnowski, T. J., “Nanconnectomic upper bound on the variability of synaptic plasticity,” *eLife*, Vol. 4, 2015, pp. 1–18. doi:10.7554/eLife.10778.
- [10] Tversky, A., and Kahneman, D., “Judgment under Uncertainty: Heuristics and Biases,” *Science*, Vol. 185, No. 4157, 1974, pp. 1124–1131. doi:10.1126/science.185.4157.1124.
- [11] Simon, H. A., “Models of bounded rationality: Behavioral economics and business organization, vol. 2,” *The Massachusetts Institute of Technology*, 1982.
- [12] Tversky, A., “Elimination by aspects: A theory of choice.” *Psychological review*, Vol. 79, No. 4, 1972, p. 281.
- [13] Tversky, A., and Kahneman, D., “Loss aversion in riskless choice: A reference-dependent model,” *Quarterly Journal of Economics*, Vol. 106, No. 4, 1991, pp. 1039–1061.
- [14] Tom, S. M., Fox, C. R., Trepel, C., and Poldrack, R. a., “The neural basis of loss aversion in decision-making under risk-Supporting Material,” *Science*, Vol. 315, No. 5811, 2007, pp. 515–8. doi:10.1126/science.1134239.
- [15] Camerer, C., “Three cheers—psychological, theoretical, empirical—for loss aversion,” *Journal of Marketing Research*, Vol. 42, No. 2, 2005, pp. 129–133.
- [16] Kahneman, D., “Prospect theory: An analysis of decisions under risk,” *Econometrica*, Vol. 47, 1979, p. 278.
- [17] Park, S., Deyst, J., and How, J. P., “Performance and Lyapunov Stability of a Nonlinear Path Following Guidance Method,” *Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 6, 2007, pp. 1718–1728. doi:10.2514/1.28957.
- [18] Boyd, J., *A discourse on winning and losing*, Air University Press, 2017.