



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

## On risk attitude and optimal yacht racing tactics

**Citation for published version:**

Tagliaferri, F, Philpott, A, Viola, IM & Flay, RGJ 2014, 'On risk attitude and optimal yacht racing tactics', *Ocean Engineering*, vol. 90, pp. 149–154. <https://doi.org/10.1016/j.oceaneng.2014.07.020>

**Digital Object Identifier (DOI):**

[10.1016/j.oceaneng.2014.07.020](https://doi.org/10.1016/j.oceaneng.2014.07.020)

**Link:**

[Link to publication record in Edinburgh Research Explorer](#)

**Document Version:**

Publisher's PDF, also known as Version of record

**Published In:**

Ocean Engineering


**General rights**

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

**Take down policy**

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [openaccess@ed.ac.uk](mailto:openaccess@ed.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.





## On risk attitude and optimal yacht racing tactics



F. Tagliaferri <sup>a,\*</sup>, A.B. Philpott <sup>b</sup>, I.M. Viola <sup>a</sup>, R.G.J. Flay <sup>c</sup>

<sup>a</sup> Institute for Energy Systems, School of Engineering, The University of Edinburgh, United Kingdom

<sup>b</sup> Yacht Research Unit, Department of Engineering Science, The University of Auckland, New Zealand

<sup>c</sup> Yacht Research Unit, Department of Mechanical Engineering, The University of Auckland, New Zealand

### ARTICLE INFO

#### Article history:

Received 14 November 2013

Accepted 30 July 2014

Available online 10 September 2014

#### Keywords:

Yacht race

Tactics

Risk aversion

### ABSTRACT

When the future wind direction is uncertain, the tactical decisions of a yacht skipper involve a stochastic routing problem. The objective of this problem is to maximise the probability of reaching the next mark ahead of all the other competitors. This paper describes some numerical experiments that explore the effect of the skipper's risk attitude on their policy when match racing another boat. The tidal current at any location is assumed to be negligible, while the wind direction is modelled by a Markov chain. Boat performance in different wind conditions is defined by the output of a velocity prediction program, and we assume a known speed loss for tacking and gybing. We compare strategies that minimise the average time to sail the leg with those that seek to maximise the probability of winning, and show that by adopting different attitudes to risk when leading or trailing the competitor, a skipper can improve their chances of winning.

© 2014 Elsevier Ltd. All rights reserved.

### 1. Introduction

In this paper we model and analyse the problem faced by a skipper who wants to sail an upwind leg of a yacht race, rounding the mark before his opponent. This problem falls into the category of stochastic shortest-path problems, where the cost function to be minimised is the time needed to reach the mark, and it depends on stochastic quantities such as wind direction. Many problems fall into this category and involve routing for emergency response, both civil (Yamada, 1996) and military (Resch et al., 2003), and applications in logistics (Fleischmann et al., 2004) and transport (Shuxia, 2012). The aim is to find a path between two vertices of a graph such that the sum of its constituent edges, often representing a cost, is minimised. When cost depends on random quantities this becomes a stochastic problem, and the standard objective is to minimise expected costs (where costs include time) (Bertsekas and Tsitsiklis, 1991). For yacht races, models which minimise the expected time to finish, or to reach the next mark, have been studied in a number of papers (Philpott and Mason, 2001; Philpott, 2005). This might be appropriate in fleet races where corrected time over a number of races forms a basis for scoring points. Even so, such scoring systems assign rankings in each race and it is well known that rank-based scoring leads to different incentives than those from performance on average (Anderson, 2012).

As observed in Philpott (2005) rank-based scoring takes its most extreme form in match racing, where the objective is to maximise the probability of arriving before the competing yacht. Indeed the time difference between the two boats is not of interest, as opposed to its sign. In this context, the attitude towards risk of the skipper assumes a greater importance. The aim of this work is to show that by changing the skipper's attitude to risk, it is possible to define a strategy that performs better in match races than strategies aimed at minimising the expected time to finish.

Of course, in most forms of match racing, the interaction between the boats is important. A leading yacht will attempt to cover a trailing yacht, not only for tactical reasons, but also to spill turbulent air on the trailing yacht's sails to reduce their drive. Forcing another boat to tack to avoid a collision is also a tactical ploy to increase a yacht's advantage. In this paper we choose to ignore these effects, as well as assuming identical yachts and crew expertise. This is done for modelling convenience as well as simplicity. By focusing solely on risk attitude we can see to what extent this is important, other effects being equal.

The paper is laid out as follows. In the next section we describe the model of the yacht and basic sailing strategy for the upwind leg of a match race. We then review dynamic programming as an approach to finding the strategy that minimises the expected time to reach the next mark. The following section shows how this is implemented in a routing model that accounts for different risk attitudes of the skipper. We then present the results of some simulations of the strategies that emerge from the routing model.

\* Corresponding author.

E-mail addresses: [f.tagliaferri@ed.ac.uk](mailto:f.tagliaferri@ed.ac.uk) (F. Tagliaferri), [a.philpott@auckland.ac.nz](mailto:a.philpott@auckland.ac.nz) (A.B. Philpott), [i.m.viola@ed.ac.uk](mailto:i.m.viola@ed.ac.uk) (I.M. Viola), [r.flay@auckland.ac.nz](mailto:r.flay@auckland.ac.nz) (R.G.J. Flay).

### 1.1. Sailing strategy

The speed of a sailing yacht depends on the wind speed and on the angle between boat heading and wind direction. It is usually expressed as a polar diagram like the one shown in Fig. 1. The numbers around the semicircle represent different true wind angles, while the radial ones represent the boat speed. The red line corresponds to the plot of boat speed for a particular true wind speed. While no direct course is possible straight into the wind, it is possible to sail upwind with an angle between wind direction and sailed course which is usually between  $30^\circ$  and  $50^\circ$ . Sailing closer to the wind direction (lower angle) makes the course shorter, but when sailing at higher angles a boat is faster. Velocity made good (VMG) is the component of yacht velocity in the wind direction. With a constant wind direction from the top mark, an optimal policy maximises VMG. This is typically attained at a true wind angle of around  $40\text{--}45^\circ$  (as in this example). In a polar diagram like the one in Fig. 1, it is possible to find the maximum VMG for a given wind speed by finding the intersection between the polar corresponding to the wind speed and the line perpendicular to the upwind direction. For this reason the common route towards an upwind mark, or in general towards the direction from which the wind blows, is a zigzag route. Such a route requires changes of direction which are called *tacks*. When manoeuvring for a tack, a boat points for a few seconds directly into the wind, therefore causing a temporary decrease in boat speed. If the wind is constant during the race and all over the racing area, trying




Fig. 1. Example of a polar diagram (velocities in m/s and angles in degrees). (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)




Fig. 2. Example of upwind routes. (a) Constant wind and (b) left wind shift.

to do the minimum number of tacks is the best choice. Fig. 2(a) shows two possible routes. In a constant wind, the route on the left is faster because it involves just one tack. Fig. 2(b) shows a situation in which the wind shifts towards the left over the duration of the leg. The best policy in this case is to go to the left of the course (referred to as being on *starboard tack*), and then tack and point towards the mark, while a myopic policy that begins the race going to the right (referred to as being on *port tack*) turns out to be suboptimal.

In real races the evolution of the wind can be much more complicated than these examples, with temporary shifts or gusts that a sailor seeks to take advantage of. Moreover wind has a random component. While racing, it is difficult to know how the wind is behaving at another location, or to foresee how it will behave once that point is reached. In the presence of randomness the optimal course in Fig. 2(b) might turn out to be worse than a myopic policy that tacks on every wind shift. For this reason sailors tend to try and stay in the centre of the course to enable shifts in wind direction to be exploited by tacking, while avoiding the risk of overlaying the mark.

In the presence of a competitor, a policy that avoids the course boundaries while staying close to the competitor reduces the risk of being beaten, at least when the competitor is the trailing boat. On the other hand, when the competitor is leading, it can make sense for a skipper to take a risk and explore the corners of the course hoping for a favourable wind shift. This is the phenomenon that we seek to model in this paper.

### 1.2. Dynamic programming

Finding an optimal set of tacks when the wind varies randomly requires a *stochastic dynamic* optimisation model. In contrast to the deterministic case, a solution does not consist of a single optimal path for a specific wind realisation, but a *policy* that is optimal over a range of wind realisations. Policies can be computed a priori and respect the principle of optimality: an optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision (Bellman, 1957). A policy that respects this principle can be found with *dynamic programming* (Bertsekas, 1995). Dynamic programming has been successfully applied in sailing in both ocean races and short course racing (see Philpott and Mason, 2001; Philpott, 2005). In this work we adapt the short-course model described in Philpott and Mason (2001) and Philpott (2005) with the aim of incorporating the skipper's attitude towards risk in their actions.

The risk that a skipper is willing to take is usually influenced by his position with respect to the opponent. A common behavioural pattern is to be conservative, or *risk averse*, when in a leading position, while being *risk seeking* when losing. Here we interpret risk aversion as being pessimistic about wind shifts, believing that any shifts we observe will not be to our advantage. In contrast, a risk-seeking skipper will be optimistic about wind shifts and act as if these are more likely to be to his advantage. Such attitudes can be modelled by altering the transition probabilities of the process that defines wind shifts.

To understand the effect of risk-averse or risk-seeking skippers, we develop a race modelling program (RMP) for simulating races between two boats. The first RMP was developed in 1987 for the America's Cup syndicate Stars and Stripes and is described in Letcher et al. (1987). Since then, RMPs have been used mainly in America's Cup applications to compare different designs (see e.g. Philpott et al., 2004). In our case, since we are interested in comparing tactical choices, we model two identical boats (i.e. they have the same polar diagram).

## 2. Method

### 2.1. Dynamic programming

We consider an upwind leg of 6000 m (corresponding to 3.24 nautical miles, which approximates the length of the 2013 America’s Cup course), and 4000 m wide. In the coordinate system used the starting line is located on the  $x$ -axis, and centred around the origin, while the upwind mark is located on the  $y$ -axis. The racing area is discretised into a rectangular grid with  $N=20$  increments  $\Delta x$  across the course and  $M=400$  increments  $\Delta y$  in the direction of the course, as shown in Fig. 3. The  $N-1$  lines defining the grid that are perpendicular to the  $y$ -axis will be referred to in the following as “cross sections”. The dynamic program is at stage  $i$  when the yacht crosses the  $i$ th cross section.

The state variables are the yacht’s position  $x_i$ , the wind direction  $w_i$  observed at stage  $i$ , and the current tack  $z$  (where  $z=0$  denotes starboard tack and  $z=1$  denotes port tack). The wind direction  $w_i$  is random and satisfies the Markov property, namely that the probability distribution for the variable  $w_i$ , conditioned on all the previous values, is equal to the distribution for the variable  $w_i$  conditioned just on the last event:

$$\begin{aligned} \mathbb{P}(w_i = v | w_{i-1} = v_{i-1}, w_{i-2} = v_{i-2}, \dots, w_0 = v_0) \\ = \mathbb{P}(w_i = v | w_{i-1} = v_{i-1}) \end{aligned} \quad (1)$$

for every  $i > 0$  and for every  $w_i$  in the state space.

The actions at each stage are whether to tack the boat (i.e. change  $z$  to  $1-z$ ) or continue on the same tack. As mentioned in the Introduction, a tacking manoeuvre implies a time loss that will be denoted as  $\tau$ . Given a yacht’s polar and its location, we can compute  $t(i, x, x', w, z)$ , defined to be the time to sail from location  $(x, i\Delta y)$  to  $(x', (i+1)\Delta y)$  if it is on tack  $z$  and the observed wind direction is  $w$ .

We define the value function  $T_i(x_i, w_i, z)$  to be the minimum expected time to sail from location  $x_i$  on cross section  $i$  to the top mark given wind observation  $w_i$  and current tack  $z$ . Clearly  $T_M(x, w_i, z) = 0$  when location  $x$  is at the top mark, and we choose  $T_M(x, w_i, z) = \infty$  otherwise.

We compute  $T_0(x_0, w_0, z)$  for  $(x_0, w_0, z)$  corresponding to the boat’s position and tack on the start line, using a dynamic programming recursion. First define at stage  $i$  the function

$$F(i, x, x', w, z) = t(i, x, x', w, z) + \mathbb{E}_w [T_{i+1}(x', w', z) | w], \quad (2)$$




Fig. 3. Schematic representation of the course.

where  $w'$  is the wind direction that is observed at stage  $i + 1$ . Now we can define the recursion as follows:

$$T_i(x_i, w_i, z) = \min \begin{cases} \min_{x_{i+1} \in \mathbb{X}} F(i, x_i, x_{i+1}, w_i, z) \\ \tau + \min_{x_{i+1} \in \mathbb{X}} F(i, x_i, x_{i+1}, w_i, 1-z) \end{cases} \quad (3)$$

where  $\mathbb{X}$  is the set of  $x$  coordinates of positions  $(x_{i+1}, (i+1)\Delta y)$  that can be reached at stage  $i+1$  from position  $x_i$  at stage  $i$ . More details on the recursive procedure defined by Eqs. (2) and (3) can be found in Philpott and Mason (2001).

### 2.2. Wind modelling

We assume the wind speed to be constant during the race, focusing on the changes in wind direction. As discussed in the previous section the dynamic programming algorithm we use assumes that the wind direction satisfies the Markov property. Although more refined wind models are being developed (see for instance the recent reviews by Costa et al., 2008 and Bitner-Gregersen et al., 2014), Markov models are computationally very efficient and can still capture most of the statistical properties that are relevant in certain applications (Shamshad et al., 2005; Sahin and Sen, 2001).

For tactical purposes we are interested in changes in wind direction that significantly affect the racing time. We therefore define a finite number of wind direction states: namely  $-45^\circ, -40^\circ, \dots, 0^\circ, +5^\circ, \dots, +45^\circ$ , where  $0^\circ$  represents the wind direction at which the upwind mark is set, and the other states represent shifts of  $\pm 5^\circ$  from that direction.

For a system with a finite number of states the stochastic process is uniquely defined with an initial distribution for  $w_0$  and a transition matrix  $\mathbf{P}$ . The matrix elements  $P_{jk}$  represent the probability that the system at time step  $i$  is in state  $k$  conditioned on the fact that it was in state  $j$  at the previous time step  $i-1$ :

$$P_{jk} = \mathbb{P}(w_i = k | w_{i-1} = j)$$

In order to obtain a realistic transition matrix we considered a time series of wind measurements from a weather station installed on the Newcastle University research vessel, and then built the matrix  $\mathbf{P}$  using a maximum likelihood estimator. As we use for the model a grid with 15 m resolution in the upwind direction and the decisions are taken every time the yacht reaches a cross section, the wind is modelled using a time step of 3 s, which is the time spent on average to move between two consecutive cross sections. The recorded wind direction signal was sampled every three seconds, and the corresponding wind directions were placed in  $K$  bins of amplitude  $5^\circ$ . The number of jumps from bin  $j$  to bin  $k$  divided by the total number of jumps out of bin  $j$  defines the value  $P_{jk}, j, k = 1, 2, \dots, K$ , in the transition matrix.

Given a transition matrix  $\mathbf{P}$ , Eq. (2) becomes

$$F(i, x, x', w_j, z) = t(i, x, x', w_j, z) + \sum_{k=1}^{k=K} P_{jk} T_{i+1}(x', w_k, z). \quad (4)$$

### 2.3. Risk modelling

We now turn our attention to the risk attitude of the yacht skipper. There is an enormous literature on modelling risk (for a recent introduction see Anderson, 2013). To model risk aversion, we adopt an approach based on the theory of coherent risk measures (Artzner et al., 1999). As shown in Artzner et al. (1999) coherent risk measures can be expressed as the worst-case expectation over a convex set of probability distributions to give a risk-adjusted expectation. Given the current wind direction state,



the probability distribution that we work with is the corresponding row of the transition matrix. To model risk aversion we choose the worst possible transition probabilities from a convex set  $\mathcal{D}$  of transition matrices. In other words, (4) becomes

$$F(i, x, x', w_j, z) = t(i, x, x', w_j, z) + \max_{P \in \mathcal{D}} \sum_{k=1}^{k=K} P_{jk} T_{i+1}(x', w_k, z). \quad (5)$$

An interpretation of (5) is illuminating. A boat skipper who is winning will be risk averse. She will try to behave safely, trying to stay ahead and to minimise her losses in bad wind outcomes. Using (5) in a recursion is pessimistic about the next wind shift and assigns a higher probability to the worst outcomes (i.e. heading shifts). Being pessimistic about random outcomes reduces risk, at some loss in expected performance.

Risk seeking behaviour has been less well studied, although it is often given as an explanation for participation in lotteries and negative expectation gambles, where optimistic participants place greater weight on winning probabilities than their real values. In our context we model risk seeking by choosing the best possible transition probabilities from a convex set  $\mathcal{D}$  of transition matrices. In other words, (4) becomes

$$F(i, x, x', w_j, z) = t(i, x, x', w_j, z) + \min_{P \in \mathcal{D}} \sum_{k=1}^{k=K} P_{jk} T_{i+1}(x', w_k, z). \quad (6)$$

This has the following interpretation. A boat skipper who is losing will seek risk. If she adopts a minimum expected finish time strategy against another skipper who minimises his expected time to finish, then she will tend to make the same decisions (unless the boats see very different winds) and lose the race almost certainly. She will instead seek different wind conditions from the competitor. Using (6) in a recursion will be optimistic about the possible advantageous wind shifts and assign a higher probability to these outcomes (i.e. lifting shifts). Being optimistic about random outcomes increases risk, as well as incurring some loss in expected performance.

We implement (5) and (6) in the recursion by adding a transformation in the solver that post multiplies the transition matrix by another matrix which redistributes the probabilities. The resulting matrix has to be normalised in order to represent again a probability distribution.

### 3. Results

Fig. 4 shows a graphical representation of the transition matrix for the Markov model obtained with the maximum likelihood estimator as described in the previous section. With a notation that will be used throughout this paper, we use a grey scale to represent values in the interval  $[0, 1]$  where white represents 0 and black represents 1. It can be noticed that the diagonal is dominant, meaning that, in general, if the wind is in state  $i$ , the most probable state for the next step is to remain in state  $i$ .




Fig. 4. Representation of the transition matrix obtained for the wind model.

Moreover, when the wind has deviated from the mean, the event of a shift back towards the mean value is more likely than one in the same direction.

The wind for the simulations was generated as described in the previous section. The Markov chain defines a discrete wind direction. This can be made continuous by superimposing a mean-reversion noise process (see Philpott et al., 2004). However we did not do this as we found that the behaviour of the simulated wind signal, achieved with no additional noise component, was similar to the empirical one, as can be seen in Fig. 5, with close values of mean and variance on different sub-intervals. A wind history of 400 values was generated for each of the 4000 simulated races.

Fig. 6 shows a histogram of the time needed by a yacht following the policy generated to minimise the expected time of arrival, according to the wind distribution previously modelled. The distribution is asymmetric, and this is due to the fact that even with a very favourable evolution of the wind there is a minimum time needed to complete the course. On the other hand, even with a policy which is effective in the majority of the cases, it is possible to be very unlucky and need a much higher time.

This policy was generated using a risk-neutral transition matrix for wind direction as pictured in Fig. 4. When the skipper is risk seeking or risk averse we replace this with a modified transition matrix. A sailor who is losing will seek risk. This corresponds to increasing her confidence of a lifting wind shift while discounting the likelihood of a heading wind shift. The transition matrices we use to represent a risk-seeking skipper are shown in Fig. 7(a) and (b). As shown in the figures, advantageous shifts (cells below the diagonal when the skipper is to the left of the opposition, and cells above when on the right) happen with higher probability than in the risk-neutral case. The remaining probabilities in each row are reduced to add to one.

The transition matrices for a risk-averse skipper are constructed similarly. Here bad wind shifts (above the diagonal when the skipper is to the left of the opposition, and below the diagonal




Fig. 5. Sixty-minute example of artificially generated wind and sixty-minute example of recorded wind.






Fig. 6. Distribution of arrival time of boat following the optimum policy.



**Fig. 7.** Modified transition matrices for a risk-seeking skipper. Advantageous wind shifts occur with higher probability than disadvantageous ones. (a) Yacht on the left-hand side of competitor and (b) yacht on the right-hand side of competitor.



**Fig. 8.** Histogram of arrival time of B minus arrival time of A.

when on the right) happen with higher probability than in the risk-neutral case. In our experiments we have obtained transition matrices for a risk-averse policy by simply swapping the matrices in Fig. 7(a) and (b).

Simulations were carried out in order to verify the differences between a risk-neutral policy that minimises expected arrival time at the top mark, and a policy generated assuming either risk seeking or risk averse behaviour. Results showed that policies that minimise expected arrival time won more races than either being consistently risk seeking or risk averse.

However, combining the strategies together (to allow both risk-seeking and risk-averse behaviour at different times) can lead to a significant improvement in the chances of winning. We simulated races between two boats that are denoted as boat A and boat B. Both boats start the race at the same time, on two different (random) points along the starting line. Boat A experiences the simulated wind and always follows the risk-neutral policy (to minimise expected arrival time). Boat B experiences the same wind as A if their distance apart is less than  $d_{min} = 10$  m, an independent wind if their distance is greater than  $d_{max} = 100$  m, and a linear combination of A's wind and an independent sample if their distance is between  $d_{min}$  and  $d_{max}$ . At every step of the simulation, if B is more than 15 s behind A, she uses the risk-seeking policy depending on the side of the course; if B is more than 20 s ahead of A, she uses the risk-averse policy, while she uses the optimum risk-neutral policy otherwise. Results of those simulated races are shown in Fig. 8.

The x-axis shows the arrival time of boat B minus the arrival time of boat A at the top mark. The average time difference is positive (actually 16 s in this plot). This means that B arrives 16 s later on average than A, as one would expect, since A is using the optimum policy to minimise the average time. However about 63%

of the race outcomes are to the left of zero, meaning that B wins 63% of the time (always by a small margin). Of course sometimes B is hopelessly outclassed, losing by 400 s (just around 0.01% of the times, and those are extremely unfavourable events) but this is because B takes high risks when behind. If we consider  $p=0.5$  win probability as a null hypothesis, then the probability of winning more than 63% of 5000 races by chance is the probability that a binomial random variable with mean  $5000p$  and variance  $5000p(1-p)$  exceeds 3150, which is negligible.

The standard error of the value 0.63 can be estimated using the central limit theorem to be approximately 0.0035. So we can be 97.5% confident that the hybrid policy will win at least 62.3% of the races (i.e. 2 standard errors less than 0.63).

In order to quantify the tactical improvement on the policy we compare the results obtained by boat A and boat B with a third boat C that has perfect knowledge of the future behaviour of the wind. In this case we simulated 1000 races. Obviously the boat with perfect knowledge of the wind scenario always wins and the increases in arrival time of A and B are always positive. The sample average difference in time of arrival is 133 s for boat A while for boat B the sample average difference is 149 s. The difference is not significant because of high variance and low sample size. However this experiment confirms a theoretical result: the expected time difference for boat A relative to C is never more than the expected time difference for boat B relative to C (see Appendix for proof).

#### 4. Conclusions

In this paper we have presented a method for approximating a solution of a stochastic shortest path problem with applications to yacht racing. We showed that with an adequate subdivision of the problem it is possible to find a solution that minimises the expected time needed to reach an upwind mark during a race.

Moreover, we introduce for the first time a model of the risk attitude of the sailor. We showed that if a skipper of a trailing boat has a risk-seeking attitude it enhances the chance to win the race. An important result of the simulations run to simulate races was that aiming at minimising the expected time to finish is not always the best approach: being on average slower might allow a bigger probability of winning against an opponent following a fixed policy.

The results presented in this paper underline that when trying to optimise a policy in order to win a competition, looking at average values is rarely the best approach, and accounting for differing risk attitudes might give policies that perform significantly better. Further work is being carried out in order to validate the model with data registered during America's Cup races, and we are developing methodologies for learning risk parameters that yield maximum win probabilities.

#### Acknowledgements

This research has been performed within the SAILING FLUIDS project (PIRSES-GA-2012-318924), which is funded by the European Commission under the 7th Framework Programme through the Marie Curie Actions, People, International Research Staff Exchange Scheme. The authors would like to thank Newcastle University for providing wind data.

#### Appendix

**Proposition 1.** Minimising the expected arrival time over all strategies will give a policy that is slower than a perfect skipper by the least amount on average.

**Proof.** Suppose a perfect skipper sails races in wind that she predicts perfectly. Each race is a random sample of wind and so her time to finish is an independent identically distributed random variable  $T$ .

Suppose she now sails a strategy  $s$  that is not clairvoyant in each of these same wind conditions. The time to finish under this strategy is an independent identically distributed random variable  $S(s)$ .

Now the delay in finishing under strategy  $s$  versus the perfect strategy is also an independent identically distributed random variable  $D(s) = S(s) - T$ . The expected delay from sailing  $s$  is then

$$\mathbb{E}[D(s)] = \mathbb{E}[S(s)] - \mathbb{E}[T].$$

To minimise this we should minimise  $\mathbb{E}[S(s)]$  as  $\mathbb{E}[T]$  is a constant. So the strategy that minimises expected delay after a clairvoyant skipper is the one that minimises expected arrival time.  $\square$

## References

- Anderson, E., 2013. *Business Risk Management: Models and Analysis*. John Wiley & Sons Inc., United States, ISBN: 978-1-118-34946-5.
- Anderson, E.J., 2012. Ranking games and gambling: when to quit when you're ahead. *Oper. Res.* 60 (5), 1229–1244.
- Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., 1999. Coherent measures of risk. *Math. Financ.* 9 (3), 203–228.
- Bellman, R.E., 1957. *Dynamic Programming*, 1st edition Princeton University Press, Princeton, NJ, USA.
- Bertsekas, D.P., 1995. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA, United States.
- Bertsekas, D.P., Tsitsiklis, J.N., 1991. An analysis of stochastic shortest path problems. *Math. Oper. Res.* 16 (3), 580–595.
- Bitner-Gregersen, Elzbieta M., Bhattacharya, Subrata K., Chatjigeorgiou, Ioannis K., Eames, Ian, Ellermann, Katrin, Ewans, Kevin, Hermanski, Greg, Johnson, Michael C., Ma, Ning, Maisondieu, Christophe, Nilva, Alexander, Rychlik, Igor, Waseda, Takuji, 2014. Recent developments of ocean environmental description with focus on uncertainties. *Ocean Engineering* 86, 26–46. <http://dx.doi.org/10.1016/j.oceaneng.2014.03.002>, ISSN 0029-8018.
- Costa, A., Crespo, A., Navarro, J., Lizcano, G., Madsen, H., Feitosa, E., 2008. A review on the young history of the wind power short-term prediction. *Renew. Sustain. Energy Rev.* 12 (6), 1725–1744.
- Fleischmann, B., Gnutzmann, S., Sandvoß, E., 2004. Dynamic vehicle routing based on online traffic information. *Transp. Sci.* 38 (4), 420–433.
- Letcher, J.S., Marshall, J.K., Oliver, J.C., Salvesen, N., 1987. Stars and Stripes. *Sci. Am.* 257, 34–40.
- Philpott, A.B., 2005. Stochastic optimization and yacht racing. In: *MOS-SIAM Series on Optimization*, SIAM, Philadelphia, PA, United States, pp. 315–336.
- Philpott, A.B., Henderson, S.G., Teirney, D., 2004. A simulation model for predicting yacht match race outcomes. *Oper. Res.* 52 (1), 1–16.
- Philpott, A.B., Mason, A.J., 2001. Optimising yacht routes under uncertainty. In: *The 15th Chesapeake Sailing Yacht Symposium*, Annapolis, USA, pp. 89–98.
- Resch, C.L., Piatko, C., Pineda, F.J., Pistole, J., Wang, I.-J., 24 March 2003. Path planning for mine countermeasures. in: *AeroSense*, Orlando, Florida. International Society for Optics and Photonics, pp. 1279–1286.
- Sahin, A.D., Sen, Z., 2001. First-order Markov chain approach to wind speed modelling. *J. Wind Eng. Ind. Aerodyn.* 89 (34), 263–269.
- Shamshad, A., Bawadi, M.A., Hussin, W.M.A.W., Majid, T.A., Sanusi, S.A.M., 2005. First and second order Markov chain models for synthetic generation of wind speed time series. *Energy* 30 (5), 693–708.
- Shuxia, L., 2012. Study on routing optimization problem of the logistics center. In: *World Automation Congress (WAC)*, IEEE, Puerto Vallarta, Mexico, 24–28 June 2012, pp. 1–4.
- Yamada, T., 1996. A network flow approach to a city emergency evacuation planning. *Int. J. Syst. Sci.* 27 (10), 931–936.