

Bayes' beacon: avalanche prediction, competence, and evidence for competence

Modelling the effect of competent and incompetent predictions of highly improbable events

Philip A. Ebert*¹ and Theoni Photopoulou²

¹Law and Philosophy, School of Arts and Humanities, University of Stirling

²Department of Statistical Sciences, University of Cape Town, South Africa

ABSTRACT: In this paper, we will discuss how competence can affect a person's ability to avoid avalanches and present a way of modelling such competence. Given that the prior probability of getting caught in avalanches is fairly low for any skier (competent or not), we draw some consequences from the model using Bayes' theorem for "everyday" situations.

KEYWORDS: Decision-making in avalanche-terrain, Modelling competence, Bayes' Theorem

1. INTRODUCTION

In this paper, we will discuss how competence can affect a person's ability to avoid avalanches and present a way of modelling such competence. Given that the prior probability of getting caught in avalanches is fairly low for any skier (competent or not), we draw some consequences from the model using Bayes' theorem for "everyday" situations. The paper is structured as follows: in section two, we highlight some basic assumptions that motivate our model. In section three, we outline the specifics behind our model and, in section four, present the results of our model. In section five, we draw on Bayes' theorem and apply it to "everyday" scenarios. In the last section, we summarise the general lessons to draw from the model for avalanche education and decision-making in avalanche terrain.

2. AVALANCHE PREDICTION: TESTING THE TERRAIN

Available data suggests that it is fairly improbable to get caught in an avalanche. In Switzerland alone, for example, there are about 200 people who get caught in an avalanche per year¹ and the death toll in Switzerland, Italy, Austria and France is between 20 to 30 people per year per country.² Overall, given the great number of skiers—ski tourers and off-piste skiers—the probability of get-

ting caught in an avalanche seems not that high.³ Moreover, there is no denying that being a competent avalanche decision-maker—however far from being a perfect predictor—will involve skills that will help reduce this probability further. There is much we can learn in avalanche courses and such knowledge lowers the overall probability of getting caught in an avalanche. In particular, guided groups—involving professional mountain guides with extensive training in avalanche safety—are less likely to get avalanched.⁴ These observations give rise to the following assumptions we will use in our model:

1. Competent and incompetent decision-makers are both subject to false positive judgements: scenarios when they judge a slope to be safe, yet they will get caught in an avalanche if they cross it.
2. Being competent helps to reduce the number of false positive judgements.
3. Irrespective of competence, the probability of getting caught in an avalanche is fairly low for any (non-reckless) skier.

Let us say a little more about how we understand the competences involved in predicting avalanches. When we speak of an "incompetent" decision-maker, we do *not* mean to imply that he/she is a "reckless" skier who skis in any type of condition. Rather, what we have in mind here with "incompetent decision-maker" is a broadly competent skier—let's call him *Ingo*—who at least considers the basic avalanche warning issued by the local avalanche forecasters. Nearly all mountain

*Corresponding author: Philip Ebert, Law and Philosophy, University of Stirling, Stirling, FK9 4LA, Scotland; email: philipebert@me.com

¹Compare (Harvey and Zweifel, 2008b).

²Compare (Valt et al., 2009). This suggests a fatality rate of roughly 1 in 10 people caught in an avalanche, compare (Schweizer and Lutschg, 2001). See also (Brugger et al., 2007) who present a fatality rate of 1 in 5 without so-called avalanche airbags and 1 in 33 with such airbags. In contrast (Jamieson et al., 2009) suggest a fatality rate of 1 in 100 in Canada. These numbers are notoriously difficult to quantify, see again (Jamieson et al., 2009).

³According to (Zweifel et al., 2006)—as presented in (Jamieson et al., 2009)—the fatality rate is roughly 1:70,000 in an area around Davos. So using a rate of 1 in 10 fatalities, there is a probability of getting caught in an avalanche of around 1:7,000 per day per skier in the backcountry.

⁴See for example, (Harvey and Zweifel, 2008a) where this claim is made.

areas have an avalanche warning system which rates the avalanche risk on a level from 1 (low) to 5 (extreme). So, let's assume that Ingo and any so-called incompetent decision-maker has no proper training for predicting avalanches, yet he heeds these warnings and avoids backcountry skiing on days when the warning is 4 (high risk) or 5 (extreme risk). Hence, the skier might reasonably take himself to be careful and not reckless.⁵

In contrast let us introduce *Connie*, she is also a competent skier who over a period of five years takes part in extensive avalanche safety training. She completes a number of courses, involving at least a week of training. The skills she acquires are multifarious: for example, she learns how to prepare for her trips by reading detailed avalanche forecasts; she learns how the wind and weather can affect the stability of a slope depending on aspect and exposure, and so on. Given these skills, we can reasonably assume that based on her 5 years of training and experience in how to apply these new skills, her ability to predict unsafe avalanche terrain has improved in contrast to Ingo's. After 5 years of training she decides to stop improving her skills (we will assume that the learning curve is linear and we assume that Connie does not lose any competence after the first 5 years).

The question we want to answer is the following:

Given that it is fairly improbable to get caught in an avalanche in the first place, and given that the skills that Connie learns will not guarantee that she will not get avalanched, and given that it takes time and experience to learn such skills, how do these competences actually affect the probability of getting avalanched over a longer period of time?

Before outlining the details of the model, we want to highlight certain simplifications and draw attention to our chosen methodology: First, we focus only on the probability of getting caught in an avalanche and ignore fatality rates. Second, we ignore naturally triggered avalanches and focus only on specific cases where a skier triggers an avalanche in which she is caught. Third, it is worth emphasising an important aspect of our methodology: we are *not* relying on *confirmed empirical data* to estimate the values for competence. Rather, our interest is, shall we say, a *philosophical* one: we engage in a theoretical exercise where we assume that a competent decision-maker has a certain probability of getting caught in an avalanche over a ten year

⁵In fact ~81% of avalanche-fatalities occur when the level is between 1-3. Compare (Zweifel, 2002). One way to interpret this statistic is that many skiers heed these warnings and don't ski when the avalanche warning is high-risk or extreme.

period (1 in 28) which has empirical backing.⁶ We will then work back to what the probability for an incompetent decision-maker will be, *given* our assumptions—taken as rough estimates—about how much avalanche safety training can affect the reliability of a competent decision-maker. Hence, in effect, our model is itself a prediction that is subject to confirmation or disconfirmation and subject to further improvement given more data.

3. THE SEARCH: MODELLING COMPETENCE IN AVALANCHE DECISION-MAKING

Our model uses the following four components to calculate the probability of a competent and incompetent decision-maker getting avalanched:⁷

Ski days per year We assume that a skier has 30 ski days in the backcountry a year. And so we assume a total of 300 days for a ten year model.

Avalanche warning We assume that skiers ski at either avalanche warning 1, 2, or 3. We assign different probabilities to the event that an avalanche occurs depending on these avalanche warnings.

Avalanche slope crossings We assume that a skier has, on average, five avalanche slopes that she/he crosses per day.

Competence A skier has a choice with respect to the crossing of the avalanche slope which can exhibit her or his competence. She can decide that it is safe to cross, or she can decide that it is unsafe to cross. The *incompetent* skier is modelled on the basis of no specific competence—the analogy being that of a coin-tosser crossing on, say heads, and not crossing otherwise using a fair coin. The competent skier's competence rises over a period of 5 years to what is, using the analogy, in effect a heavily biased coin.

⁶If there is roughly one avalanche fatality in 70,000 ski touring days—compare (Zweifel et al., 2006)—and there is a fatality rate of 1 in 10, we have a probability of getting caught in an avalanche of about 1 in 7000. If we consider, as we do, 30 skiing days per year over 10 years, this suggest a probability of 1 in 23 of getting caught over that period. However, this statistic does not individuate competent, incompetent or reckless skiing. Yet, assuming that there are very few completely reckless skiers, and mostly competent and incompetent backcountry skiers, the assumption we make about competent decision-maker seems appropriate. Compare here also (Kristensen et al., 2012) where the probabilities used are overall lower—importantly, the main results in section 5 will be even more striking if the probabilities are lower than we assume! It is also worth noting that there is a *strong* variability of the relevant risk given other factors which we are ignoring—such as geography, age group, etc.

⁷Some initial inspiration was drawn from (Tremper, 2008) but there are crucial differences.

In our model, we first assign probabilities to the risk of avalanche depending on “objective” factors that are captured in the avalanche warning issued by local experts. So on a low avalanche warning, we assign the probability 1 in 100,000 to the event *that a slope is disposed to avalanche*. The probability progressively increases as the avalanche warning gets raised; we used factor 10 (see (Jamieson et al., 2009)). That is, on avalanche warning 2 (moderate) we have 1 in 10,000 and, if it is considerable, we use 1 in 1,000. We here assume that if a slope is *disposed to avalanche*, a skier who crosses the slope, will trigger an avalanche.

We modelled *competence* by having a skier make a decision with respect to the slope. She has a simple choice: she can either cross the slope or decide not to cross it. The incompetent decision-maker, who in our model is sufficiently competent to actually identify a potential avalanche slope—clearly a non-trivial assumption—is modelled simply as tossing a coin whether to cross the slope or not cross the slope. In our model the competence function outputs a 0 (*not crossing the slope*) or 1 (*cross a possibly avalanche disposed slope*), which is then multiplied by the probability of the slope being disposed, i.e. the “objective factors”, to give the overall probability of getting caught in an avalanche in that crossing. A more competent decision-maker will more often make the objectively right decision and thus be less likely to cross a slope that is avalanche disposed. Hence, in the competent case the function outputs more 0's and fewer 1's. In effect, the 0 output can be interpreted not merely to stand for *not cross* but also, in the case of the competent decision-maker as standing for *cross and the slope is correctly judged not avalanche disposed*—hence there is zero probability to get avalanched. Importantly, we here assume that both subjects have overall the same level of risk-aversion (or risk-seeking) so that the difference between crossing or not crossing cannot simply be explained by different risk attitudes; rather, the thought here is that it is *in virtue of* the competences involved that the competent decision-maker can reduce the number of unsafe crossings.⁸

⁸In (Kristensen et al., 2012) different levels of risk aversion are modelled. Using Munter's risk reduction model (RM), see (Munter, 2003), different fatality rates are assigned to skiers with different levels of risk acceptance—however differences between competence and incompetence are not modelled there. In effect, one can view our model as relative to a specific RM level. So one could vary the probability for an incompetent skier to cross—set in our model simply as a 50-50 chance—as (partly) dependent on the RM level the skier pertains to. The point of our model would then be that given a specific value for the incompetent decision-maker, the competent decision-maker will have a more favourable value and will overall have more safe crossings. We will leave this complication aside and

As mentioned before, the values we assign to competence are not directly confirmed by empirical data. To the best of our knowledge there is no such data available.⁹ The so-called incompetent skier's “competence” is, however, straightforwardly assigned. Remember that we don't simply assume our incompetent skier, Ingo, to be silly or reckless (he only skis on avalanche warning level 1-3) but he has a choice to either ski or not ski a given slope. Given that Ingo has no relevant skill, he is given a 50-50 chance of crossing or not crossing.

On the other hand, the competent skier is assigned much better odds of getting it right. Our competent skier, Connie, is modelled so that in only 1 out of 5 choices, she is exposed to a non-zero probability of getting caught in an avalanche. So, the way we perceive competence is that it manifests itself in making some crossings completely safe. That is, by exercising her competence, Connie manages to reduce the number of unsafe choices from 5 out of 10 in Ingo's case, to 2 out of 10 slopes. Later, we will also consider the case of an expert—we call her *Ezra*—and in her case competences further reduce her exposure to a potential slide to only 1 out of 20 choices.

Given the type of training a competent person has received we think these assumptions are not unfounded; but, as mentioned before, empirical data is needed to be more confident in our competence values. Also, we have not yet taken into account that competence could manifest itself in other situations—not merely in judging an avalanche slope as safe. A competent skier might ski with more foresight and take various measures that can help his survival chances, e.g. she may “cut the slope”, i.e. intentionally trigger an avalanche while being in a safe spot. Also we ignore that incompetence can manifest itself by not being able to identify a potentially dangerous slope and thus have no choice whether to cross a given slope. Sceptics of the competence values used are invited to change the relevant values in our model and see how more competences can have an effect on the probability of getting avalanched and, most importantly, use the relevant values in our calculations in section 5.

4. PROBING: RESULTS OF THE MODEL

We modelled the cumulative risk of getting caught in an avalanche for a competent and incompetent skier over a period of ten years (300 days).

deal with it in future research. We also ignore the observation that self-efficacy can in turn lead to a higher level of risk acceptance. See e.g. in rock climbing (Llewellyn et al., 2008).

⁹It would be great to use the model of *citizen science* to collect such data, e.g. in the form of apps that record backcountry skiers' exposure days, formal training and how often they get caught in an avalanche.




Figure 1: The cumulative risk of getting caught in an avalanche over 10 years for an incompetent and competent decision-maker (left) and the associated risk ratios (right).

Figure 1 demonstrates the effects of increasing competence.

The odds of an incompetent skier getting caught is 1 in 15, or roughly 7 out of 100 incompetent skiers (6.7%). The ratio of the respective cumulative risk is 1.84 which means that the incompetent decision-maker is nearly twice as likely to get caught in an avalanche in that period compared to the competent one. On our model while (roughly) 7 out of 100 incompetent skiers get avalanched over 10 years, roughly 4 out of 100 competent skiers (3.6%) get caught as well. The effects of competence are pronounced but maybe not as pronounced as one may have expected.

This, of course, raises an important question about competence. Maybe we have underestimated the effects of competence. So let's consider our third skier Ezra mentioned above. Similar to Connie, Ezra is keen to learn more about avalanches, and, like Connie, she learns a fair bit over the first five years, but in contrast to Connie, she continues to take advanced courses so that her competences keep on developing further. So, while her development over the first five years is like Connie's, over the next 5 years, she becomes what we call an *expert*:¹⁰ in her case competences further reduce her exposure to a potential slide to only 1 out of 20 choices, in contrast to 10 out of 20 in Ingo's case and 4 out of 20 in Connie's case. Figure 2 has the results for Ezra (high comp), Connie (med comp), and Ingo (low comp).

1 in 32 experts (3.1%) will get caught in an avalanche compared to 1 in 28 for the competent one over a period of 10 years. It seems more marked in contrast to the 1 in 15 for the incom-

¹⁰(Tremper, 2008) suggests that acquiring the relevant competences will take up to 10 years of training in the case of an expert.

petent decision-maker. So, while 4 out of 100 competent decision maker will get caught in an avalanche, we have 3 out of 100 expert decision-maker who will get caught (compared to 7 out of 100 incompetent skiers). Of course, we have to remember that over the first five years the competent and expert decision-maker have an equal probability. Also note it is only after ten years, having reached a very high level of competence, that Ezra might reasonably be called an expert.¹¹

It is worth remembering the practical limitation of this model: until we have solid data to confirm our competence values, the applicability of our model is limited. Still, having this model in place, it is time now to draw some more theoretical consequences from it.

5. BAYES BEACON: BAYES' THEOREM

In this section, we will investigate how our model can affect or, rather, should affect a subject's reasoning about competent or incompetent avalanche decision-making. Using Bayes' theorem, we hope to shed light on how strongly certain pieces of evidence (e.g. "having so far avoided avalanches") should affect a subject's judgement. Naturally, when making such normative claims—about how one *ought* to reason—we make certain idealisations and assume that the notion of evidential sup-

¹¹Thus, one may object that this graph doesn't give the answer to the question how probable it is that an expert decision-maker will get avalanched; after all, we take the whole training into account to become an expert. Fair enough: without taking training into account the cumulative risk of getting caught in an avalanche over 10 years *being* competent and *being* expert, assuming no improvement or loss in their skills is: 3 in 100 for the competent (or: 1 in 37) and 7 in 1000 for the expert (1 in 150). Given that experts do get avalanched (while admittedly also spending a lot of time in the snow) the assigned value for competence is very likely an upper limit.




Figure 2: The cumulative risk of getting caught in an avalanche over 10 years for an incompetent, competent and expert decision-maker (left) and the associated ratios of these risks (right).

port is best modelled in a standard probabilistic framework. Nonetheless, we think the following “real-life” scenarios have relevance and should be taken into account in everyday decision-making about avalanche safety.

Scenario 1

Imagine a skier, call him Alexander, who has skied for 10 years. He has studied some books on avalanche safety and has taken some courses, so he has a fairly high degree of confidence that he is competent (on a scale from 0 to 1, with 0 being certain that he is not competent, to 1 being certain that he is competent, let’s say that Alexander confidence is 0.8). Yet, after 10 years, he makes the wrong call and gets caught in an avalanche that he himself triggers.¹² He is lucky and he survives the slide. Aspiring to always be an ideal rational agent, Alexander now considers what consequences to draw from the event: Should Alexander change his mind and consider himself as an incompetent decision-maker given that he got caught in an avalanche? Or, to put it in more academic slang: Given Alexander’s prior confidence about his own competence in avalanche decision-making, how should the fact *that he got caught in an avalanche* affect Alexander’s confidence about his competence?

Given our results, we can calculate how confident Alexander should be in his competence taking account of the new piece of evidence using

¹²This is important, we thereby assume here that there are no other overriding reasons for Alexander to ignore the evidence. For example, we ignore the case where he got caught in an avalanche triggered by other skiers. We assume instead that he did judge the slope as safe but it did avalanche; in that sense he made the wrong call.

Bayes’ theorem. With *C* for *that one is a competent decision-maker*, and *L* for *that one got avalanched after 10 years of skiing*, and $P(C|L)$ for *the probability that one is a competent decision-maker given that one got caught in an avalanche after 10 years of skiing* we get:

$$P(C|L) = \frac{P(L|C)P(C)}{P(L)} \tag{1}$$

which is equivalent to:

$$P(C|L) = \frac{P(L|C)P(C)}{P(L|C)P(C) + P(L|\neg C)P(\neg C)} \tag{2}$$

Assuming a fairly high degree of prior confidence on Alexander’s part that he is competent (0.8) on a scale from 0 to 1, we can calculate the updated confidence that he is competent given he has been caught in an avalanche using our model results, i.e. $P(L|C) = 0.036$ and $P(L|\neg C) = 0.067$.¹³

$$P(C|L) = \frac{(0.036) \cdot (0.8)}{(0.036) \cdot (0.8) + (0.067) \cdot (0.2)} = 0.69 \tag{3}$$

That is, the event of getting caught in an avalanche should have some effect on the confidence that one is competent but not as much as one might anticipate. Alexander should still be reasonably confident that he is competent (after all the degree of confidence is still considerably greater than

¹³We are interpreting $\neg C$ as “incompetent” as defined above and read off the relevant data $P(L|\neg C)$ from the model for the incompetent decision-maker. Strictly speaking, this is incorrect since there are reckless skiers as well as experts that fall under $\neg C$. We thus take the value as an approximation on the assumption that nearly all backcountry skiers who are not competent will be incompetent, as we defined it.

0.5).¹⁴ The following situation highlights a different aspect:

Scenario 2

You arrive at a new backcountry ski area and you look for a ski partner. You don't know how many competent or incompetent decision-makers are in the area. So, you are unprejudiced whether a person is competent (or incompetent) and your confidence is 0.5 that a skier is competent (or incompetent).¹⁵ Given you want to ski fast, you choose a ski partner who has at least 10 years of experience.¹⁶ Having found a partner, you ask her whether she has ever been caught in an avalanche. Your partner tells you that she skied for ten years and has never been caught in an avalanche. How should this fact affect your confidence that your potential partner is competent.

It's time to draw on Bayes' theorem again and see how this information should change your confidence that she is competent.

$$P(C|\neg L) = \frac{P(\neg L|C)P(C)}{P(\neg L|C)P(C) + P(\neg L|\neg C)P(\neg C)} \quad (4)$$

and with a prior probability of being perfectly unprejudiced (the technical term is: *indifferent*), i.e. 0.5, we get using again our model results ($P(\neg L|C) = 0.96$ and $P(\neg L|\neg C) = 0.93$):

$$P(C|\neg L) = \frac{(0.96) \cdot (0.5)}{(0.96) \cdot (0.5) + (0.93) \cdot (0.5)} = 0.508 \quad (5)$$

As a result, that a person did not get caught in an avalanche should *not* affect your judgement very much at all. That is, if you are unprejudiced/indifferent about whether a person is competent (0.5), receiving the information that she has never been caught hardly provides any additional evidence that she is competent (0.508) and you should remain (very close to) indifferent upon receiving this new information. Hence, the fact that a skier has never been caught in an avalanche in a ten year period, provides nearly no evidence to think that she is a competent decision-maker. Figure 3 provides a matrix for scenario 1 (left) and scenario 2 (right) and the effects of the relevant evidence given different prior probabilities.

¹⁴One way to explain the change in attitude is to use bets: prior to the incident, Alexander would have been justified to take a bet with the odds of 4–1 against him being competent, his updated credences (still) justify him to take a bet of 11–5 against him being competent.

¹⁵This assignment of priors follows from an application of the *principle of indifference* and the assumption that all skiers are either competent or incompetent.

¹⁶Importantly you don't treat this as evidence for thinking that the person is competent; just that they are fast skiers.

6. WIDENING THE (RE)SEARCH

We think the second scenario contains an important lesson for avalanche education and it is worth considering in more detail. According to the model, even though an incompetent decision-maker is nearly twice as likely to get caught in an avalanche, there is still a fairly low probability that an incompetent person will get caught at all. As a result, the fact that a skier didn't get avalanched provides hardly any evidence for whether she is competent. Yet, all too often, one hears of people speaking of *experienced* backcountry skiers, and it is thereby suggested that the person is a competent ski-tourer with the ability to safely negotiate avalanche terrain. We regard the results of scenario 2, to offer a warning not to treat experience, or, more specifically, *having successfully avoided avalanches over 10 years*, as significant evidence for judging that person to be competent—based on our model, it simply isn't!

More worryingly, we also seem to have all the required ingredients to generate an *illusion of competence*: as is well-known and discussed widely in popular psychology books¹⁷ humans often aim to offer reasons and explanations why events happened, even if they are best regarded as random events. So, if a skier manages to avoid avalanches over a ten year period, there has to be, one might be tempted to think, a good reason for it—it can't simply be by luck. And the illusion might become even more extensive when judging one's own history: I might be easily led to think that it must be for good reason that I have avoided avalanches over such a long period of time. That is, I must have developed a competence. What our model shows is that one should really not be that surprised—from a statistical point of view—that incompetent (yet non-reckless) decision-makers manage to avoid avalanches. We maintain that competent decision-making is grounded in *knowing* how to minimise risk, while simply having avoided bad outcomes (even over a long time) is no guarantee of being competent.

Hence, we want to draw attention to the fact that although experience is necessary to develop certain competences, that is, it is an *enabling* condition to gain competences and know how to apply them, having experience, i.e. having been exposed to the relevant risks and "got a way with it" is itself *no good evidence* to having acquired these competences. More generally, we want to caution against the inference that in most mountain sports, experience is by itself a reliable indicator of competence. Our model and the use of Bayes'

¹⁷See in particular (Kahneman, 2011) and (Chabris and Simons, 2009); both books provided much inspiration for this work.




Figure 3: Prior confidence in competence (black) and updated confidence (red). Scenario 1 (left) and scenario 2 (right).

Theorem is a way to substantiate this warning.¹⁸

In this context, it is worth drawing attention to the following unwarranted inference from our model: given that the probability for an incompetent skier to get caught in an avalanche is so low, there is no point in taking avalanche courses. However, it is, firstly, worth remembering that we assumed certain values for competence and more empirical data is required to confirm our model. Hence, it might well be that the difference between competence and incompetence is more pronounced than we assume. Secondly, in our model, incompetent decision-makers are nearly twice as likely to get caught in an avalanche within the first 10 years. This is surely relevant! And over a longer period of time, this ratio will continue to go up. Competence, therefore, manifests itself not within the first few years; rather it manifests itself only over a longer period of time and its effects will continuously become more and more pronounced. So, the lesson is instead that if you plan to ski into "old age", and you plan to reduce the overall risk of getting caught in an avalanche, the sooner you start informing yourself about avalanche safety, the better the odds will be in the long run when compared to an incompetent decision-maker.¹⁹ Figure 4 shows

¹⁸See also the excellent article (McCammon, 2004) who outlines a number of *heuristic traps* for avalanche decision-making. Our paper can be seen as providing an explanation of a potential further trap: that one is easily misled by an allusion to experience to overestimate one's own or other people's competences—it relates to what he calls *expert halo* trap.

¹⁹Of course the best way to avoid avalanches is not to go backcountry-skiing at all; we here assume that this is not really an option. For an article which aims to justify competent risk-taking in mountain sports, see (Ebert and Robertson, 2013).

the cumulative risk of our three skiers over a 20 year period. It seems to us that the premium to pay to gain competences is very low to effect such significant change and thereby lower your *future risk* which, we maintain, should be of relevance to you *now*; remember, about 1 in 10 people who get caught in an avalanche do not survive.²⁰

References

- Brugger, H., Etter, H., Zweifel, B., Mair, P., M., H., Ellerton, J., Elsensohn, F., Boyd, J., Sumann, G., and Falk, M. (2007). The impact of avalanche rescue devices on survival. *Resuscitation*, 75(3):476–83.
- Chabris, C. F. and Simons, D. J. (2009). *The invisible gorilla*. Crown Publishing Group.
- Ebert, P. A., Photopoulou, T., Hanson, N., and Hancock, P. J. B. (2013). Avalanche avoidance: Competence and luck. Manuscript—available on request.
- Ebert, P. A. and Robertson, S. (2013). A plea for risk. In O'Hear, A., editor, *Philosophy and Sport*, London Lecture Series. Cambridge University Press.
- Harvey, S. and Zweifel, B. (2008a). Neue trends in der lawinenunfallstatistik. *Die Alpen*.

²⁰We would like to thank Charlie Boscoe, Nora Hanson, Will Hoppitt, Peter Milne, and Martin Smith for comments and/or stimulating discussion of the core ideas of this paper. This paper is based upon a longer manuscript (Ebert et al., 2013).




Figure 4: Cumulative risk of competent and incompetent decision-makers over 20 years (left) and the associated risk ratios (right).

Harvey, S. and Zweifel, B. (2008b). New trends of avalanche accidents in backcountry terrain in Switzerland. In *Proceedings Whistler 2008 International Snow Science Workshop September 21-27, 2008*.

Jamieson, B., Schweizer, J., and Shea, C. (2009). Simple calculations of avalanche risk for backcountry skiing. In *International Snow Science Workshop, Davos 2009, Proceedings*.

Kahneman, D. (2011). *Thinking, fast and slow*. Allen Lane Publishers, London.

Kristensen, K., Genswein, M., and Munter, W. (2012). Perception of risk in avalanche terrain. In *Proceedings, International Snow Science Workshop, Anchorage, Alaska*.

Llewellyn, D. J., Sanchez, X., Asghar, A., and Jones, G. (2008). Self-efficacy, risk taking and performance in rock climbing. *Personality and Individual Differences*, 45(1):75–81.

McCammon, I. (2004). Heuristic traps in recreational avalanche accidents: Evidence and implications. *Avalanche News*, 68:1–10.

Munter, W. (2003). *3x3 lawinen. Risikomanagement im wintersport*. Pohl & Stellhammer, 3rd edition.

Schweizer, J. and Lutschg, M. (2001). Characteristics of human-triggered avalanches. *Cold Reg. Sci. Technol.*, 33:147–62.

Tremper, B. (2008). *Staying alive in avalanche terrain*. Mountaineers Book, second edition.

Valt, M., Chiambretti, I., and Zasso, R. (2009). 1985–2009 twenty-five years of avalanche accidents in Italy. In *International Snow Science Workshop, Davos 2009, Proceedings*.

Zweifel, B., Raez, A., and Stucki, T. (2006). Avalanche risk for recreationists in backcountry and in off-piste area: Surveying methods and pilot study at Davos, Switzerland. In *Proc. International Snow Science Workshop 2006, Telluride CO.*, pages 733–741.

Zweifel, S. (2002). Skifahrerlawinen und lawinenbulletin in der Schweiz: Ein vergleich. *Bergundsteigen*, 4.

Corrigendum to Ebert, P.A. & Photopoulou, T. (2013). Bayes' Beacon: Avalanche Prediction, Competence, and Evidence for Competence.

Philip A. Ebert, Ian Durbach, and Theoni Photopoulou

August 13, 2018

1 Error in cumulative risk calculation

We found an error in the R-code underlying the cumulative risk calculation of triggering an avalanche presenting in our article:

“Bayesian Beacon: Avalanche Prediction, Competence, and Evidence for Competence, Modelling the Effect of Competent and Incompetent Predictions of Highly Improbable Events.” *Proceedings of the International Snow Science Workshop, 2013*, pp.363-370. http://arc.lib.montana.edu/snow-science/objects/ISSW13_paper_O5-04.pdf

Given numerous guesstimates and assumption the risk of triggering an avalanche was calculated to be around 6.7% over 10 year for a non-competent decision and 3.6% for a competent decision over the same period. In one line of the code, we made a mistake when calculating the probability of getting avalanched on a given day for five different slope crossings. This resulted in the stated risk to be significantly less than it should be. Using the same assumptions as in our paper, the cumulative risk of triggering an avalanche for a non-competent decision maker are closer to 28% while a competent decision maker will have a cumulative risk of 12% of triggering an avalanche over that period.

While this might seem more in line with intuitive expectation, it is important to note that the main lessons from the paper, in particular that a track record of avalanche avoidance alone is not a good indicator of competence, remains correct. After all, 72 out of 100 non-competence decision makers will avoid triggering an avalanche, while 88 of competent decision-makers avoid avalanches. According to Bayes' Theorem, learning of a person (who you assign a prior of 0.5 to be competent) that she has never been in an avalanche over the relevant period increases your confidence that the person is competent from 0.5 to 0.55, and so, as before, this evidence is not as strong as one may expect, when judging competence of a decision-maker.

Rather than correcting the original model, we offer an alternative way to calculate the cumulative risk of triggering an avalanche over a given period, assuming a certain number of slope-crossings per day and avalanche risk per slope. In this model, these assumptions, as well as the effects of competence in reducing the avalanche risk, are easily adjusted and readers can calculate their own cumulative risk of triggering an avalanche given their own guesstimates. We make that R-code available for download on researchgate and personal websites.

2 An alternative way to calculate the the cumulative risk of triggering an avalanche

The probability of experiencing at least one avalanche in n slope crossings is given by $1 - P(X = 0) = 1 - (1 - p)^n$, where $X \sim \text{Bin}(n, p)$ is a binomially-distributed random variable representing the number of avalanches experienced in n crossings, and p is the probability that a single crossing results in an avalanche.

The probability p depends on the competence of the decision maker and the riskiness of the environment. For an avalanche to be triggered, two events need to happen – the slope must be predisposed to avalanche, in the sense that if it is crossed, it will avalanche; and the skier must choose to cross the slope. Defining these events as G and A respectively, the joint probability that a single crossing results in an avalanche $p = P(G \cap A) = P(G|A)P(A) = \zeta\alpha$, where $\zeta = P(G|A)$ is a natural (inverse) measure of competence and $\alpha = P(A)$ measures the riskiness of the crossing. Thus, a choice of values for ζ , α , and n leads directly to an estimate for the probability of being avalanched in n crossings.

Realistically, neither the background risk parameter α nor the decision-making parameter ζ will be constant over time. Skiers cross slopes of varying riskiness and also learn or unlearn over time. Our probability model can easily be adapted for this by defining blocks of crossings in which these parameters are constant. Then the probability of being avalanched in n crossings is $1 - \prod_{b=1}^B \prod_{i=1}^{n_b} (1 - \alpha_b \zeta_b)$ where n_b is the number of crossings made in block b , and α_b and ζ_b are the risk and decision parameters in that block.

To calculate how confident a skier should be in his/her competence given the triggering of an avalanche, we make use of Bayes' theorem. Let C denote that one is a competent decision-maker, and L_n denote an experience of at least one avalanche after n slope crossings. Then Bayes' theorem stipulates that the probability that one is a competent decision-maker given an avalanche event is $P(C|L_n) = P(L_n|C)P(C)/P(L_n)$ where $P(L_n) = P(L_n|C)P(C) + P(L_n|\bar{C})P(\bar{C})$. The probabilities $P(L_n|C)$ and $P(L_n|\bar{C})$ can be obtained from our model by choosing appropriate values for ζ . Values for $P(C|\bar{L}_n)$ can be computed similarly. Our estimates of cumulative risk are based on $n = 1500$ crossings, with $\alpha = 1/10000$ and $\zeta = 0.25$ and 0.46 for competent and untrained skiers respectively.

3 New R-code

Please consult the additional material on researchgate to download the R-code. An image of the code can be found below:

```
1  ### Risk calculations for avalanche study
2
3  # user specifies inputs for PRE-TRAINING PHASE
4  ut.p.nca = .8      # probability of not crossing an avalanching slope
5  ut.N = 1500       # total number of slopes crossed
6  ut.prop = c(5/20,14/20,1/20) # proportion of slopes of type 1,2,3
7  ut.avrisk = c(1/100000,1/10000,1/1000) #probability slope type 1,2,3 will avalanche
8
9  # user specifies inputs for POST-TRAINING PHASE (so can have different parameters if want)
10 t.p.nca = .5      # probability of not crossing an avalanching slope
11 t.N = 0          # total number of slopes crossed
12 t.prop = ut.prop # proportion of slopes of type 1,2,3
13 t.avrisk = c(1/100000,1/10000,1/1000) #probability slope type 1,2,3 will avalanche
14
15 # concatenate pre- and post-training phases
16 p.nca <- c(rep(ut.p.nca,3),rep(t.p.nca,3))
17 avrisk = c(ut.avrisk,t.avrisk)
18 N = c(ut.N*ut.prop, t.N*t.prop)
19
20 # Prob of getting into avalanche = Prob slope avalanches * Prob do wrong thing
21 p.av = (1 - p.nca) * avrisk
22
23 # cumulative risk of not being in avalanche (for slopes of each type crossed pre/post training)
24 # 1) Let X be number of avalanches experienced in n trials, then  $X \sim \text{Binom}(n,(1-p.dca)*avrisk)$ 
25 # 2) Here n trials relates to N slope crossings, and we work out  $P[X = 0]$ 
26 # 3) I work this out for each set of slopes with different prob of getting into avalanche
27 # 4) Prob of getting into avalanche is affected by pre/post training, and slope type (so 6 sets)
28 # 5) Note the order that slopes are crossed is unimportant, risks multiply up
29 rr = c()
30 for(i in 1:6){rr = c(rr,dbinom(0,rep(1,N[i]),p.av[i]))}
31
32 # final cumulative risk of being in an avalanche
33 tail(1-cumprod(rr),1)
34 |
```