POWER ALLOCATION WITH RESPECT TO HILLS and WINDS

How hard I should pedal when going uphill or against the wind

Sometimes the most unassuming questions can be the hardest ones to answer. Suppose that someone asks you how hard you pedal when climbing a hill. "–Well…" you might answer, "…it depends on the hill gradient, doesn't it!" Indeed, it is common knowledge that the steeper the hill is, the harder it is to climb. The explanation is simple: on level ground you already spend some energy to move forward; but when going uphill, each pedal stroke requires additional energy due to elevation gain, which is proportional to the slope of the hill. Therefore, the pedaling effort must increase with increasing slope.

It makes sense…but wait a minute: the question wasn't about maintaining a *given speed* regardless of the hill slope; nor was it implied that you must stay in the same gear or maintain the same cadence. So, is it really harder to pedal on hills than on flats? Why is that the case? We know it is, but how can we explain it?

If you were putting out a "normal" effort on level ground, then, when you encounter a hill, wouldn't you shift into a lower gear and slow down just enough to maintain the same effort level? Come to think of it, that strategy makes sense, and that's why we like wide-range gearing. But wait again…does that mean I should always ride at a *constant power,* regardless of uphill slopes and head winds? We all know that this is *not* how it's done. Experience tells us to take it easier on the flats, work harder on the climbs, recover on the descents, and learn how to "listen to your body" and pace yourself accordingly.

This makes the whole situation very complicated. When you encounter a hill, on the one hand you want to work harder, but on the other hand you want to slow down so that you don't work too hard. So, every time the grade changes you must find a new sweet spot because "how hard is too hard" depends on the hill slope. But does it? Why? And if it does, doesn't it also depend on the headwind? This dilemma raises at least three questions:

1. When the hill slope increases, how do you decide how much to slow down and how much to increase your effort level?
2. Why does your preferred effort level change with hill slope in the first place?
3. If your effort level *should* change with hill slope, should it also change with headwind?

I have been asking these questions of myself a lot, lately. Of course, when I climb a hill I don't adjust my power according to a pre-determined formula. No one does. But at the same time, no one does a whole bike ride at their average power, either. Power variations are expected. Yet it is naïve to think that these variations are made randomly without any rhyme or reason. Perhaps randomness is not to be taken out of the process completely, but I believe that the process is fundamentally deterministic. Why do we vary our power output during a ride (aside from inevitable constraints like traffic, road hazards, etc)? Although there are likely to be multiple factors, hill slope must be one of them. So, it is a relevant and fundamental question to ask, how exactly the gradient of a climb (and similarly the speed of a headwind) affects a cyclist's power output, and why.

STATISTICS: We can measure how much power cyclists expend on various gradients. A single pattern would not represent the behavior of all cyclists, nor would it apply to all kinds of cycling activities. But crowd-sourced data already exists, and should be looked at. It may be possible to establish consistent relationships between hill slope, wind speed and power output by statistically analyzing the available data. In this study I look for correlations between recorded variables. I don't have access to everybody's ride data, but I want to look into how cyclists might generally respond to hills and winds, by analyzing my own ride data. I don't assume that my behavior is typical, or that it is representative of other cyclists; nor does correlation imply causality. But I do believe that valuable information can be derived from correlation studies, especially if they validate a hypothesis that links the observed behavior of experienced cyclists to practical benefits according to established laws of physics.

INSTANTANEOUS VALUES: In my previous article about bike rides and data mining, I showed that the *average* power of my bike rides increases as a function of the "hilliness" of the ride, where hilliness is defined as the elevation gain per distance. In that study, I downloaded my ride data from the Strava cloud in a format where each data point is a complete bike ride with a corresponding *hilliness*, and the analyzed quantity was the *average power* for the ride. In this study, instead of comparing average power with average slope, I want to compare instantaneous power with instantaneous slope. Since my ride data are sampled at one second intervals, and my rides are several hours in duration, each ride contains tens of thousands of instantaneous values of slope and power. And I have uploaded over one thousand rides; so, my record in the Strava cloud contains *tens of millions* of data points. At present, I don't have the tools and methods necessary to collect and organize all of this to be analyzed as a whole. Therefore, I have chosen to select a few solitary rides of long duration, and analyze them individually. My objective is to find consistent patterns in the way that I dole out power with respect to hills and winds.

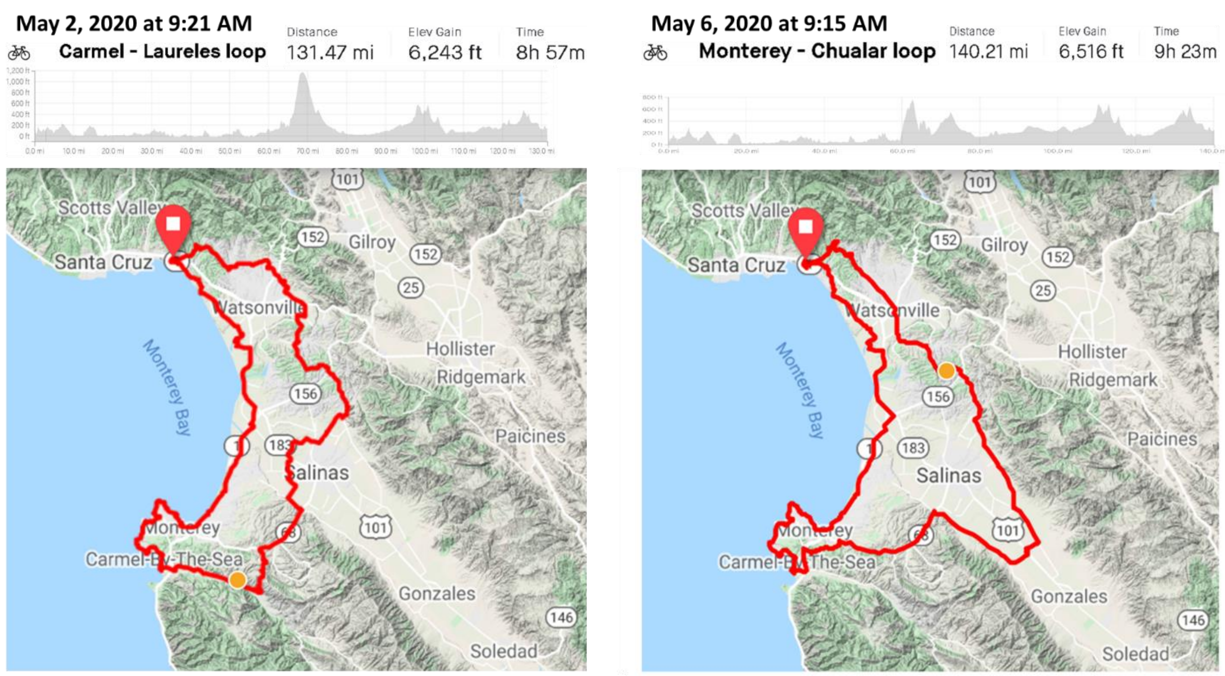

Figure 1 –Ride summaries, elevation profiles and location tracks of the analyzed activities ridden on May 2nd and May 6th.

ANALYZED RIDES: I will start by analyzing two rides that I did four days apart, on May 2nd and May 6th, 2020. Their time / distance / elevation summaries and route maps are shown in Figure 1.

As you can see, the total elevation gains and hill profiles of the routes do not differ by much. Weather conditions were also similar, and I was riding the same bicycle with the same power meter. Furthermore, many of the roads are common to both routes. The most noticeable difference between these two rides is a big hill around mile 70 on the first ride. And the second ride is a little longer. During these two rides, I made no conscious effort to change my power allocation with respect to hills or winds. I rode in a “normal” manner, and my bicycle was equipped with an instrument that measures and records wind speed, which I will explain later.

If I wanted to examine the “raw data” of a bike ride, to find out how much power I applied with respect to hills and winds, I would be looking at data points that are scattered as shown in Figure 2.

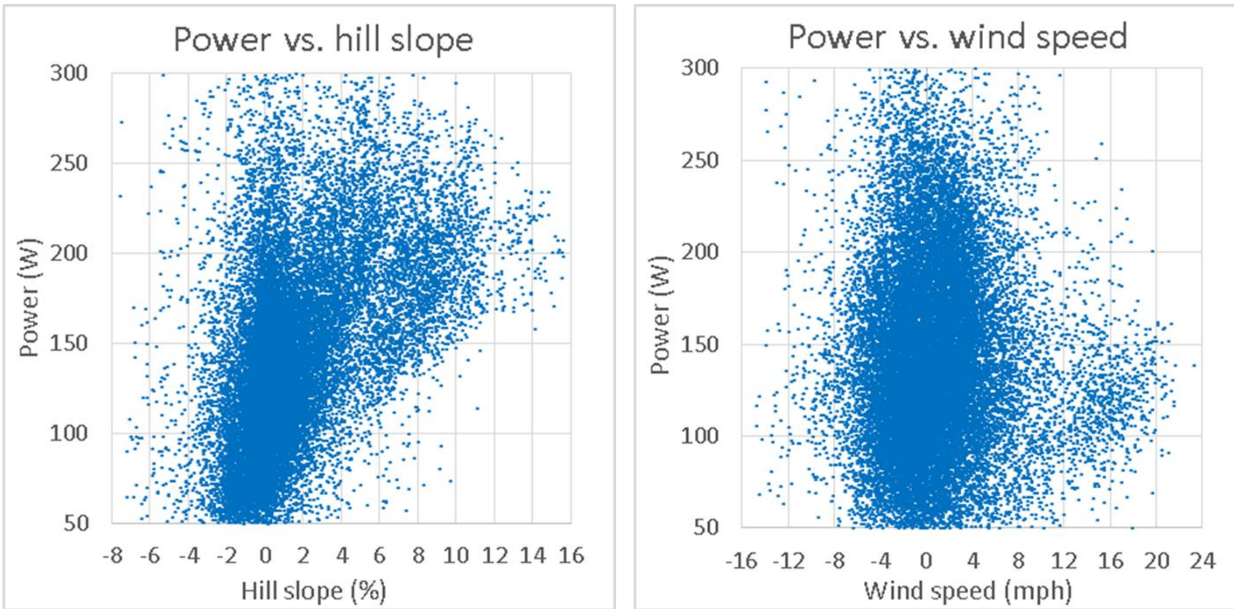

Figure 2 –Scatter plot showing instantaneous values of power, plotted with respect to hill slope (left –Fig. 2A) and wind speed (right –Fig. 2B). Negative values represent downhill and tailwind conditions.

In Figure 2, each blue dot represents a one-second snapshot (a recorded instant) taken during the ride. For each of these snapshots, power is projected onto the vertical axis, and hill slope (left) and wind speed (right) are projected onto the horizontal axes, respectively. Here, “wind speed” is a quantity that approximates the in-line component of the apparent wind velocity with respect to the ground. For example with a direct headwind, if my ground speed were 15mph and my air speed (i.e., apparent wind speed) were 18mph, then my “wind speed” would be 3mph.

These scatter graphs provide some evidence that I was generally applying greater power on steeper slopes, but it is virtually impossible to be more specific than that by visual inspection alone. Therefore, I apply data processing similar to what I did for my previous article: I divide the horizontal axes into “hill slope bins” and “wind speed bins” narrow enough for sufficient horizontal resolution, but wide enough so that each bin contains a statistically significant number of data points. This process converts the scatter graphs of Figure 2 into line graphs of Figure 3. In these graphs, hill slope bins are 0.25% wide, wind speed bins are 1mph wide, and each power value corresponding to the bin center is the *average power* of those instances where the hill slope (left - Fig. 3a) and wind speed (right -Fig. 3b) ere within the tolerance limits of the bin. Now we can clearly see the power allocation with respect to hill slope and wind speed (that is, how much power I applied, on average, on a given hill slope and with a given wind speed), during each of these two bike rides.

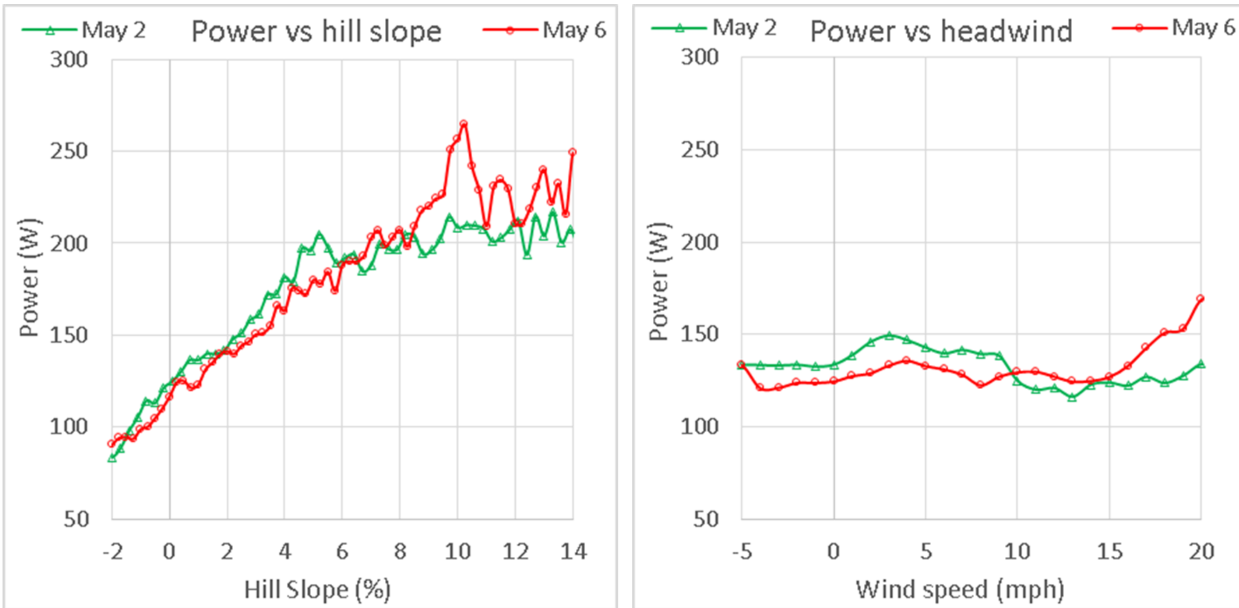

Figure 3 –Power Allocation with respect to hill slope (left –Fig. 3A) and wind speed (right –Fig.3B) on two rides, May 2 and May 6. Negative values represent downhill and tailwind conditions.

RESPONSE to HILLS: I was motivated to do this study to hopefully identify common trends (and notable differences) between rides. I cannot draw conclusions from my own behavior and extrapolate them onto other cyclists. But in my view, the information that jumps out of Figure 3 has significant relevance because, first of all, it clearly shows that my response to hills was not random, at least not on these rides. Evidently, how much power I apply to the pedals at a given instant does, indeed, depend on the slope of the hill that I am climbing at that instant. Otherwise the graphs would exhibit a less deterministic, more chaotic behavior typical of the significant influence of unrelated random factors. Also, my hill responses have in common that I increase my effort with increasing slope, quite linearly up to a characteristic slope, after which the curve more or less flattens. The meaning of such characteristics, whether they pertain to physiology, psychology, cycling experience, age, body weight, fitness level etc., or whether they reflect external factors such as terrain and wind conditions, may possibly be understood by comparing hundreds of such curves by multiple cyclists on multiple courses. This is not a small undertaking, but crowd-sourced data are already available, and this kind of scientific study can and should be done.

Differences between the two rides are also noted: on hills with gradients of up to about 6%, I evidently worked harder on May 2 than on May 6; but on hills with 8% or steeper gradients, the opposite is true. The corresponding threshold power was around 200W on May 2nd, as opposed to 220W on May 6th. In and of themselves, such observations don't amount to insightful information: on a given day, I may be tired, or simply not be in a mood to push hard. On a given ride I may be more relaxed and inclined to take it easy, but on another ride I may be in a hurry or more focused on performance. My response to the hill slopes may also be affected by certain characteristics of the terrain itself (e.g., rolling or steady hills, short or long climbs), or by prevailing winds. That there are numerous possible explanations for inconsistencies in my response to hills from ride to ride suggests that I analyze the available data from a broader perspective, focusing on the *relative* allocation of power.

POWER ALLOCATION WITH RESPECT TO HILLS AND WINDS: Just as an uphill gradient generates an opposing force, so does a headwind. Therefore, if there is any rationale for applying more power on hills, one would be compelled to think that the same rationale would also apply to headwinds. But the data presented in Figure 3 do not support this argument. To the contrary, it is striking that (during these two rides, at least), my response characteristics were fundamentally different depending on whether a given amount of pushback (i.e., the opposing force) was derived from going uphill or upwind. The reasons behind this finding are far from obvious, and this phenomenon must be further investigated. But interestingly, the finding is consistent with anecdotal evidence that cyclists find it to be "harder" to push against a headwind than against a hill slope. It is difficult to interpret such testimonials in terms of objectively measurable physical quantities, but the persistence of remarks to that effect possibly means that the ratio of Perceived Exertion Level (PEL) to actual power output might be greater when riding against a headwind than going up a hill. It is not clear why this might be the case, but I came up with an unproven theory to explain the phenomenon, to which I will return later in this article. Before digging deeper into the subject, I will start by defining the various components of the opposing force that impedes the forward motion of a bicycle.

OPPOSING FORCE COMPONENTS: In Figure 3 we can observe that on 8% gradients I was producing 200W on average, but on the flats I was producing only about 120W on average. As noteworthy as this information may be, I was producing these power levels not only against hill gradients, but also simultaneously against other opposing forces, including headwinds. On each hill slope in Figure 3A, the power is averaged over a range of wind speeds and wind yaw angles. So, I am going to analyze the data in a slightly different way so that we can look at multiple variables at once.

Cycling power is spent on inherently different types of opposing forces. These are

- hill force
- air drag
- rolling resistance
- inertia
- drivetrain losses.

There are also vibration losses, which induce an impedance that may be classified as yet another type of opposing force, or may be included in the definition (and measurement) of rolling resistance.

A positive hill slope multiplied by the combined weight of the bicycle, rider and gear, equals the opposing hill force. When going downhill, the hill force is negative, so it becomes a contributing force rather than an opposing force. Similarly, inertia can be a contributing or opposing force (help or pushback) depending on whether the ground speed is decreasing or increasing.

On different hill slopes, the relative composition of the types of opposing forces is different, and so, in Figure 3 the horizontal axis variables are not completely independent. This is so because the ground speed on steep hills is generally slower, lowering the air drag for a given headwind. Likewise, at a given wind speed, power is averaged over a range of hill slopes. But thanks to the anemometer that I carry on my bike (PowerPod by Velocomp), I always measure and record not only my power, but also my headwind and hill slope (PowerPod has multi-axis accelerometers to measure the hill slope more accurately than by estimating it from position and elevation data). Unfortunately, wind speed and hill slope variables are not included in the data files I upload to Strava. But PowerPod records them in its proprietary file format (.ibr), and allows them to be analyzed using its PC software called Isaac. This process allows power, wind and slope data to be exported in a standard file format

(.csv) so that it can be analyzed by other applications including Excel spreadsheets. That way, I can dissect my data in specific ways and look at how my power was allocated across selected *components* of the opposing force. (By the way, the PowerPod is marketed as an inexpensive power meter, but in fact it is also an anemometer and data logger that receives and records data from wireless sensors that measure speed, cadence and heart rate. It allows display of data -including wind speed- on Google Earth. When wirelessly connected to another power meter, it even measures your coefficients of air drag and rolling resistance!).

HILL FACTOR and WIND FACTOR: Cycling data recorded by PowerPod can be analyzed with respect to *basic factors* such as slopes and winds as shown in Figure 3, but analysis with respect to *derivative quantities* can provide even more useful information. One of these quantities is what I call *"Hill Factor"*. Hill Factor is defined as the ratio of the hill force to the total opposing force. Expressing cycling power as a function of hill factor (rather than hill slope as in Figure 3) is aimed at neutralizing the effects of being recovered or fatigued, fit or unfit, motivated or bored, etc. Likewise, we can replace the independent variable *wind speed* in Figure 3B with "Wind Factor", the ratio of air drag to total opposing force, as shown in Figure 4.

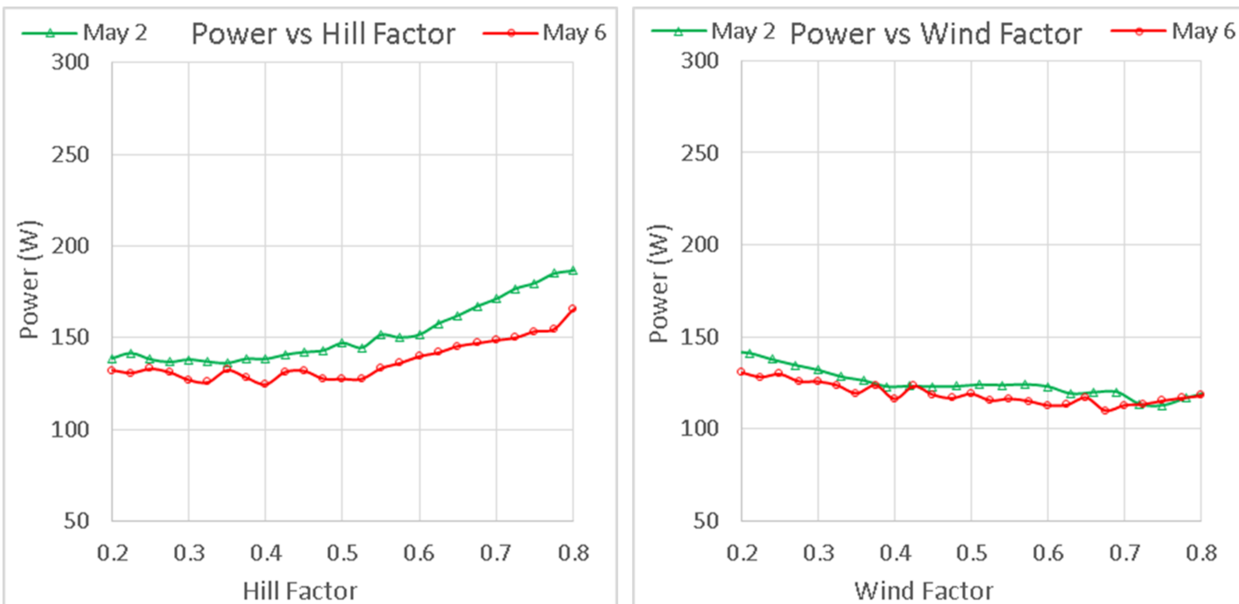

Figure 4 – Power Allocation with respect to Hill Factor (Left –Figure 4A) and Wind Factor (Right –Figure 4B) on two rides, May 2 and May 6. Hill Factor = hill-force / total-opposing-force; Wind Factor = air-drag / total-opposing-force.

Here we can see that my power output was generally higher on May 2 than on May 6, but my power response specific to hills and headwinds was consistent (i.e., the *shapes* of the curves are similar) between the two rides. It should be noted that generally, ground speed would tend to increase with wind factor and decrease with hill factor. So, Figure 4 means that on average I apply *less* power when I ride faster. If this seems counter-intuitive at first, think of it this way: when it becomes easier to go fast (low hill factor, correspondingly high wind factor), I don't work as hard.

Now, the following observation becomes even clearer: when the opposing force is increasingly due to gravity, I apply increasingly more power; and when it is increasingly due to air drag I apply increasingly *less* power. This trend cannot be observed from Figure 3 because at a given wind speed, power is averaged over a wide range of hill slopes. In other words, having a strong headwind doesn't necessarily mean that power is spent mostly against the wind. But having a large Wind Factor means exactly that. The upwards and downwards trends observed in Figures 4A and 4B, respectively, are the main subjects of this article, and the reasons for this phenomenon need to be investigated.

To better understand the phenomenon, power can be divided into *hill power* and *wind power* components, which can be normalized and expressed as dimensionless variables.

HILL-NORMALIZED and WIND-NORMALIZED POWERS: To minimize the effects of daily variations in *overall* power, to be able to compare cycling data recorded by different riders, and possibly to minimize the effects of differences in terrain, we can also normalize the vertical axis of the graphs (i.e., power). This is done by replacing for each ride, the dependent variable (power) with *hill-normalized power*, a dimensionless variable defined as the ratio of the average power at given hill factor to the average power on the flats. Similarly, *wind-normalized power* is the ratio of the average power at given wind factor to the average power in the absence of wind. Figure 5 shows how my hill-normalized power and wind-normalized power changed as functions of the hill factor and wind factor, respectively, during my rides on May 2 and May 6.

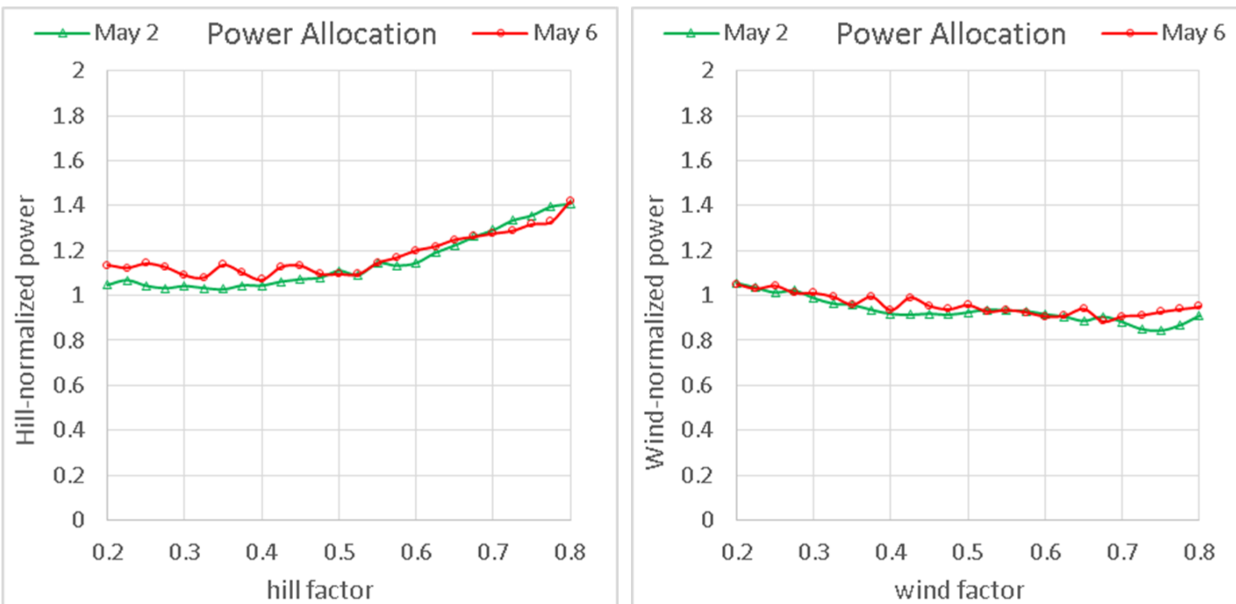

Figure 5 –Hill-normalized and wind-normalized powers as functions of Hill and Wind Factors, respectively, on two rides, May 2 and May 6.

By comparing Figure 5 with Figure 4 we can confirm that with the normalization, any differences in my overall power levels between the two rides (green and red curves) have all but disappeared, revealing the information we are seeking: how my power allocation is changing with respect to hills and winds.

Expressing a cyclist's power response to hill and wind conditions in terms of *dimensionless* variables, as is the case in Figure 5, allows comparison of the response of one cyclist to that of another. In power, hill and wind data collected from a large number of riders, some of the stronger / heavier riders would naturally produce more power than others, but their hill- and wind-normalized power values would still be directly comparable with corresponding values recorded by lighter and/or slower riders who produce less power.

The most salient observation that emerges from this analysis is that when the hill force represents increasingly greater shares of the opposing force, I apply increasingly more power to the pedals. Conversely, as the composition of the constituent components of the opposing force shifts away from hill force and increasingly towards wind force, my power output consistently declines. Is this behavior characteristic to all cyclists?

Unless I have data on other cyclists I cannot make such a generalization. But I venture to suggest that this behavior may be rooted in a subconscious effort in my part, of seeking a strategic advantage that may apply to bicycling activities in general, because it might be an effective way to ration one's energy for lasting endurance. If true, this finding (and this method of data acquisition and measurement) may possibly have some relevance for training and racing practices, and it would certainly have an application in the design and development of self-regulating e-bikes. This strategic advantage is based on a hypothesis that I will put forward without a proof.

HYPOTHESIS of NONLINEAR REWARD: In an earlier article, I had shown that in a hypothetical scenario where the goal is to complete a given course in the shortest time with a given amount of energy, (or in a given time with the least amount of energy), one must ride at constant speed (i.e., have enough power to not slow down for hills). Of course, the above scenario is highly idealized, but I believe that electric assistance may enable a cyclist to approach that ideal to a greater degree than is possible with unassisted cycling. In the present article, I am presenting actual ride data, rather than trying to validate a theory. The result from my own data is illuminating, because it suggests that cyclists may have a natural instinct or acquired skill to allocate their power according to an energy-saving strategy.

I believe that the *general* strategy by which a cyclist decides how hard to push under a given condition is to seek a tradeoff between sacrifice and reward. With more power (sacrifice) comes greater speed (reward), and the sacrifice-to-reward ratio is the power-to-speed ratio, which is, by definition, the opposing force or pushback. Averaged over the duration of a bike ride, the three biggest components of the pushback are hill force, air drag and rolling resistance. Of these three components, air drag increases with increasing speed, but the others do not. And more power always results in higher speed. Therefore, when the opposing force is predominantly air drag, more power gets more pushback; whereas when the opposing force is predominantly hill force, the pushback doesn't change with power. This means that when the wind factor is large, a given increment in sacrifice returns a small increment in reward, but when hill force is large, the same increment in sacrifice returns a greater increment in reward. Power is the time rate at which work is done, and we want to minimize the work. So, when I discover that on both of these rides my power consistently increased with the hill factor and decreased with the wind factor, I am inclined to interpret the finding as a sign of a successful application of a useful strategy of power allocation. Perhaps that's a stretch, but at least I can say that the behavior shown in Figure 5 is *consistent with* a general strategy which would be useful if the goal were to go faster and farther with natural limits on power and energy.

RELEVANCE FOR E-BIKES: While an unassisted bicycle rider decides how much power to dole out at any time during a bike ride, an e-bike rider makes separate decisions regarding human power and electric power. In some models, the e-bike decides the latter on its own, but in either case, there is a lot to be gained by researching how power, hill slope, ground speed, air speed and wind yaw affect the hybrid performance. A power management system for e-bikes based on a general strategy aimed at seeking the optimal balance between sacrifice and reward is advisable, whether the assist level is controlled manually by the rider or algorithmically by a robot. I hope that this work would inspire e-bike manufacturers to utilize crowdsourced cycling data to better understand the application and make better products.