

The start-time game in competition soaring.

John H. Cochrane*

I analyze the start time decision in competition soaring. I show how the Nash equilibrium of this game can be a large gaggle that leaves late in the day. I evaluate circumstances that can break down this equilibrium and the effect of several proposed rules changes.

*University of Chicago, 1101 E. 58th St. Chicago IL. 60637,
john.cochrane@gsb.uchicago.edu. This article appeared in *Technical Soaring* April
1998, Copyright © 1998 the Soaring Society of America.

1 Introduction

The start time decision is one of the most crucial tactical decisions in modern contest soaring. In this paper, I analyze optimal starting time. I use simple concepts from game theory to understand how the overall start time outcome depends on rules, weather and the spread of pilot/glider performance.

The results should be useful to the on-going discussion of contest format and rules. Assigned speed tasks often lead to “start gate roulette” and gaggle flying, often long after the best soaring conditions have passed. Many pilots object to this form of competition flying, either from safety concerns, or because they simply do not enjoy flying contests in which this is the outcome. In the US, this objection has led to many proposals for rules changes, including systematically longer tasks, renewed emphasis on the Pilot-Selected Task, “silent starts” monitored only by GPS, point penalties for late starts, and so forth. Whether one likes or dislikes the current system, it is clearly interesting to know whether these proposed changes will lead to the outcomes desired by their proponents.

The results may also be useful for individual pilots, to help predict when gaggles will form, what everyone else is likely to do, and hence to better optimize their own start time strategy.

2 Each pilot’s strategy

The first consideration in choosing a start time is obviously that one wishes to fly during the time of day that has the strongest lift, and will produce the greatest speed. Let us summarize this fact by a function which gives the speed the pilot could achieve if he were flying all alone, for any given start time. Denote

$t_i =$ start time of i th pilot

$V^a(t_i) =$ expected speed if flying alone.

Figure 1 graphs a typical case of the speed function. A 2:00 start time results in the best speed.




Figure 1:

I will refer to the objective as “speed,” but in reality it is “contest points.” Points are usually proportional to speed (for the individual), and I will come back to the difference below. The importance of this point for now is that a landout need not equate to zero, but rather to a low value of “speed.”

One can also fly faster by starting with or later than other pilots and using them to mark thermals. To capture this fact, let total speed be the combination of the speed one can achieve if flying alone and the speed bonus derived from flying with the others. Denote

$$\{t_j\} = \text{everyone else's start times}$$

and the extra speed

$$V^b(t_i, \{t_j\}) = \text{extra speed from following others.}$$

What does this function look like? Clearly, you can fly faster if more people start before you do. However, if everyone starts an hour before you

do, it is unlikely that you can catch them and make any use of them. The ideal arrangement from your perspective is to have everyone else start before you, at about 3 minute intervals so you can go from one well-marked thermal to the next. More people starting at once before you is a bit of a benefit, but less than linearly; 20 people marking the same thermal is not much more helpful than 10. Finally, even if you start with a gaggle or just ahead of a gaggle, that is better than flying alone.




Figure 2:

Figure 2 gives an example of what the extra speed function might look like. The figure graphs a case in which two other gliders start early and then three other gliders start a bit later. Other things equal it is better to go with the larger gaggle, but not $3/2$ times better. One can still benefit by starting a bit ahead of the other pilots, since the gaggle will go faster than any individual pilot, but it is even better to start a bit behind the other pilots and catch them the first time they slow down or have to search for a thermal. Of course if the group of three would start closer to the group of two, one could do better still by being able to jump from the late gaggle to the early gaggle. The asymmetry of the right hand part of the function is not significant, but it will help to keep this graph consistent with later ones.

The total speed is of course the sum of the speed flying alone plus the speed bonus gained from flying with others. Figure 3 shows what the total speed might look like in this example. Our pilot's objective is to maximize his speed *given* the start times of the other pilots. Thus, he picks the start time as shown. The all too familiar solution in this case is, start a bit behind the late gaggle, even though it is later than the optimum time to start a solo flight.




Figure 3:

At the optimal start time, the overall speed function is flat. (The first order condition for an optimum says to set the derivative of the speed function to zero.) This means that the slopes or time-derivatives of the two components – individual speed and speed bonus from following others – must exactly offset each other at the optimum. I have shown this fact in Figure 2 by plotting the slopes of the individual speed and speed bonus functions at the optimum.

Thus, each pilot starts at the moment at which, *if he were to wait another minute, the decline in speed due to deterioration of the day exactly matches the increase in speed due to being able to catch up with other pilots.* Each

pilot makes this calculation: “if I wait a minute longer, I will gain x miles per hour because I can catch up with the gaggle. But if I wait a minute longer I will lose y miles per hour because the day is going to die.” If $x > y$, he waits another minute. If $y > x$, he should have started already. When $x = y$, he starts.

Unless the time-of-day effect V^a is flat, the pilot does not leave at the moment which gives him the most advantage from the gaggle. This would be the top of the speed bonus function. Nor does he leave at the best time of the day, the top of the individual speed function.

3 Nash Equilibrium

If other pilots would only be so thoughtful as to spread out and start like this, one’s life would be easy. Of course, the problem is that everyone wants to start last. But once others start later, the original pilot wants to start later still. Where does it all end up? The natural definition of “where it all ends up” is an *equilibrium*. Precisely, we search for a “Nash equilibrium”. Here is a formal definition:

Nash equilibrium. The set of start times for each pilot $\{t_1, t_2, t_3, \dots, t_N\}$ is a Nash equilibrium if each pilot’s start time choice t_i is optimal *given* the start times of all the other pilots.

The above example is *not* a Nash equilibrium. Our pilot is happy to start at the indicated optimal starting time, but the *other* pilots are not at their optimal starting times. As they change their starting times, the best starting time for our pilot changes, and so on. A Nash equilibrium is a point at which all this has settled down, and nobody has an incentive to change start times.

The concept of Nash equilibrium is due to John Nash, for which he won the Nobel prize in Economics in 1994. His essay is reprinted in Nash (1996). Kreps (1990) is a standard Ph.D. level textbook that covers this material. Schelling (1980) is a delightful and readable discussion of game theory.

The most famous Nash equilibrium, which serves to illustrate the concept and the outcome for us, is the “prisoner’s dilemma.” Two prisoners are suspected of a crime. The police tells each prisoner that if he will confess

and testify against the other prisoner, he will go free. The other prisoner will get a severe 25 year sentence. If both confess, they will get 10 years. If neither confesses the most they will get is 1 year on a minor charge. The game is neatly summarized in the following diagram:

| A's choice | B's choice | |
|---------------|---------------|-----------|
| | don't confess | confess |
| don't confess | A:1 B:1 | A:25 B:0 |
| confess | A:0 B:25 | A:10 B:10 |

The outcome is inescapable: both confess and receive 10 years. No matter what the other one does, each prisoner is better off confessing. Each will cheat on the jointly desirable outcome of not confessing. This parable has been widely applied, for example in studying arms races.

3.1 The big, late gaggle equilibrium

Figure 4 shows a Nash equilibrium for the situation as graphed above. The top line gives the total speed for any pilot as a function of his start time. The smaller hump below gives the component of total speed that is the bonus for being able to fly with the gaggle. The two slanted lines give the slope of the speed functions at the optimal starting time.

All pilots start at the same time in one big gaggle. For the moment, I have assumed that all pilots and gliders are identical. Thus, if it were to any individual's advantage to start later than everybody else, then it would be to everyone else's advantage as well, and we wouldn't have a Nash equilibrium. Thus, in a Nash equilibrium everyone must end up starting (or trying to!) at exactly the same time¹.

The question is, at what time does the gaggle start? The answer is, it starts at the moment at which, with each pilot following the rule outlined above (start when the speed gained by waiting another minute and leeching equals the speed lost by delaying another minute and flying in weaker lift), all pilots want to start at the same time. Too early in the day and, with everyone else starting at the same time, the individual gains more by waiting

¹This statement assumes that the objective – total speed – is convex.




Figure 4:

and following. Too late in the day and, with everyone else starting at the same time, the individual gains more by starting earlier and using stronger lift. Graphically, we move the speed bonus function (with all the other pilots starting at the same time) to the right or left, until its slope at the point where all pilots start exactly matches the slope of the individual speed function. Then, the individual pilot chooses to start exactly at the same time as all the others.

No pilot in the end gains more than the benefit of flying together. Also, if everyone could agree to leave together at the peak time of the day, the whole gaggle could achieve a better speed. But if they did, each individual would have an incentive to cheat on the agreement and start a few minutes behind and soon the whole thing would unravel. This game is exactly an instance of the prisoner's dilemma!

This model explains why pilots hang around so long in many contests. The benefit of hanging around for that first minute after the others leave is often quite high. One can often make up the whole minute by leeching on the gaggle. The cost in total speed of starting a minute later, because the day

will die, is typically quite low. Often, the only way that the cost of waiting around an extra minute is as large as the benefit of starting a minute after a gaggle is if one has waited so late that there is a good chance of landing out. So pilots wait, and wait.

3.2 No equilibrium – mixed strategies

The last example contains an implicit assumption: each pilot can do better leaving with the late gaggle than he could if he were to leave alone at the best part of the day. The hump on the right hand side of the picture is higher than the big hump in the middle of the picture.

This may not be the case. One may be able to foresee that the “start gate roulette” crowd will fool around all day making false starts, and never make it around the course, or that they will do so in the end very slowly, catching the few dying thermals at the end of the day. Then the crafty pilots (or surprised beginners) who snuck out early will win the day. Figure 5 shows in this case that an individual pilot can achieve a better speed by starting early if everyone else leaves in the late gaggle.

In this case *there is no* (pure strategy) *Nash equilibrium*. If everyone else is going to start late in one big gaggle, any individual pilot could win by starting alone at the best solo start time (top of the graph). But the minute one pilot starts at the individually optimal time, another pilot can do even better by starting a few minutes after him. Given someone starting a few minutes after the best solo start time, our original pilot should delay until a few minutes after that, and so on until we get back to the late start. But if everyone starts late, one pilot can do better starting at the individually optimal time. And around again we go.

We model situations like this by looking for a *mixed strategy* equilibrium. The essence of this situation is that the other pilot’s start times are not predictable. For an individual pilot, it is just as if the other pilots chose their start times completely randomly. Now, given that the others will choose start times randomly, what should an individual pilot do? Perhaps there is then an optimal time to start. But if this were the case, we would not be at

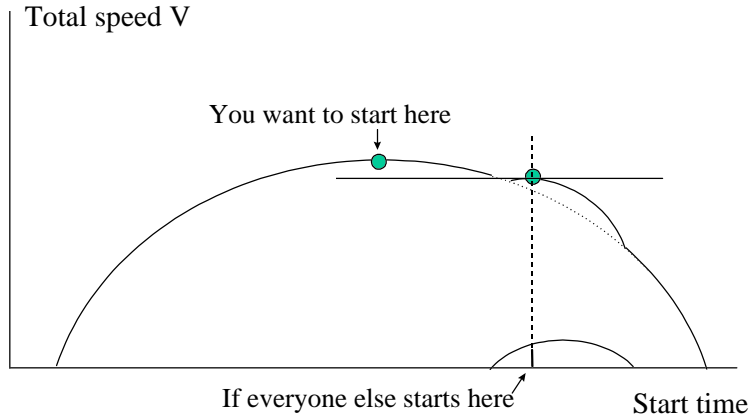


Figure 5:

a Nash equilibrium, for what is optimal for one is optimal for all, and then everyone would try to start, predictably, at this optimum time.

Hence, in the Nash equilibrium it must be the case that each pilot finds it optimal to *randomly* choose his start time, given that all the others are doing so. Each pilot now chooses a probability *distribution* over possible start times. The pure strategy equilibrium is a special case in which the probability distribution collapses to a single point. Furthermore, it must be the case that, given the probability distribution of start times chosen by the other pilots, an individual pilot is *indifferent* between start times. If he were not indifferent between start times, he could do better by starting at a single optimum time than he could by randomizing.

Figure 6 illustrates a mixed-strategy Nash equilibrium, for the same configuration I showed above in which there is no pure strategy equilibrium. The wedge-shaped curve in the lower part of the graph displays the *probability* distribution of the other pilots' unpredictable start times. We then determine the speed bonus for an individual given this probability distribution of the other pilots. A few pilots start early, then gradually more and more. At

any moment, the increasing number of markers on course balances the decay of the day. In the Nash equilibrium the probability distribution adjusts so that these two effects *exactly* counterbalance. Then the total speed for an individual pilot is flat in an interval. The individual pilot therefore is also happy to randomize across start times.

One can also see that starts now start happening at the individual best time of the day. Starts also likely to end before the late-gaggle equilibrium. The mixed strategy equilibrium emerges when the late-gaggle equilibrium is slower than the best individual time, and spreading the others around is likely to lower the benefits of following.




Figure 6:

Days like this will not settle into the steady monotony of the late-gaggle equilibrium. While gagging opponents may take heart at this outcome, start gate strategy is even more important on a day like this than on a day in which the dreaded late-start equilibrium takes hold. In the late start equilibrium, pilots can just drift around without paying much attention until the late start time approaches. In a mixed-strategy equilibrium, each pilot must be very aware of what everyone else is doing. He must be ready to boldly strike out early if it looks like enough pilots will delay. Conversely he must nervously watch his opponents and be ready to quickly follow if it looks like a group

will leave during the good part of the day. False starts, false radio messages and other attempts to bluff in order to get markers out on course will pay off in this situation. If any pilot becomes too predictable in his start times, others will hang back ready to leech.

3.3 What will happen?

To determine whether a late-pack equilibrium will form, we need to understand how weather and task affect the individual speed and the speed bonus from flying with others.

The late-gaggle equilibrium is most likely to emerge when the benefits of flying together (the hump on the right side of the graph) are large. If thermals are well-marked by cumulus clouds, for example, the benefits of following are smaller. Thus, we expect a big pack to be more likely to form on blue days, low days, or weak days with good visibility.

The late gaggle equilibrium is less likely to emerge, of course, when the benefits of sneaking out alone at the peak of the day are higher.

3.4 A spread in performance

So far, I assumed that all pilots and gliders are of the same ability and performance. Many contests however feature a spread in glider performance or pilot ability. What effect does this have on the analysis?

To analyze this case, start in the late-start Nash equilibrium, and throw into the soup one pilot/glider that is slower than everyone else. His speed bonus (benefit from gagging) is lower, and shifted forward in time. He may receive no benefit from starting after the pack, if he cannot catch up to the pack. He may also receive less benefit from starting with the pack, if he cannot keep up all the way around the course. He may receive the best benefit from the pack by starting substantially before it, and letting the pack catch him on the most difficult portion of the course.

Therefore, our slower pilot will want to start earlier. He may want to start 5 minutes ahead of the pack, let it catch him and then follow all or

partway around the course. He may even find it advantageous to ignore the pack altogether and start at the optimal individual start time. If the pack just barely makes it home, he would land out by following it.

A few such pilot/gliders will not perturb the equilibrium. One or two slow gliders out on course at 1:00 are not enough for a fast pilot/glider to use as a marker and beat the pack that starts at 3:00. Similarly, if the pack starts at 3:00, but one beginner starts at 2:50, a seasoned pilot will not make up 5 minutes on the pack by starting at 2:55 and leeching on the beginner.

But if there are many such beginners, the situation changes. If enough beginners start at 1:00, a seasoned pilot may be able to start at 1:30 using them as markers and beat a pack that starts at 3:00. The minute one can do it, all can do it so the pack start time moves up. But then it is even better to start a bit behind the pack. Is there a new equilibrium in this case, and what does it look like?

Figure 7 graphs one possible equilibrium. It is different from the late-start pack equilibrium in two respects. First, pilots leave in reverse order of glider/pilot performance. Second, the whole group leaves earlier.

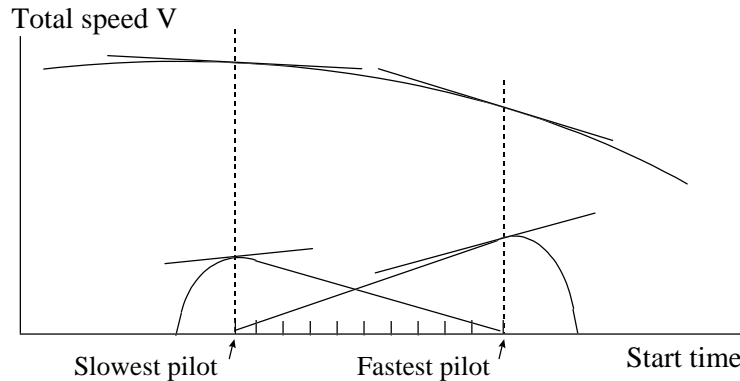


Figure 7:

The figure shows the speed boost for the slowest and fastest pilots re-

spectively, given the start times of all the remaining pilots. The slowest pilot gains the most by starting early; then he can fly with all the others for a while as they pass him. If he starts in the middle of the pack, he will still benefit from flying with those who start after him and catch him, but he will not benefit from those who start before him. The situation is reversed for the fastest pilot. He is in the delightful position that he can start last and step from marked thermal to marked thermal passing all the others. His speed boost is therefore larger than that of the slow, first pilot. Still each pilot is doing the best he can given the actions of all the others.

The speed boost functions are more drawn out than they were when all the other pilots started in a big pack. The exact timing relative to others is less important when everyone is spread out. If the fast pilot starts a minute sooner, he will still be able to use most markers, so he will finish almost a minute sooner. If there was one big gaggle and one started a minute before it, that whole minute is likely to be lost.

The fact that the speed boost function is more drawn out accounts for the fact that everyone leaves earlier. The optimal time to leave for each pilot is still dictated by the condition, leave when the increase in speed you would get by leaving a minute later and gagging more is equal to the decrease in speed due to the day dying. But with a wider speed boost function, the amount to be gained by waiting a minute longer is much less than if there is one big gaggle. Therefore, the slope of the speed boost function is lower at the Nash equilibrium start time for each pilot. It follows that the slope of the individual speed function is lower at the optimum, i.e. earlier in the day.

In sum, the analysis suggests that *contests with a wider range of glider/pilot performance should see starts that are more spread out, and earlier in the day.* This prediction seems to accord with experience. Gagging is most common at national and world contests, and less prevalent at local and regional contests. It is also more prevalent in standard and 15m classes, and less prevalent in Open and especially Sports class, which feature a wider spread of glider performance. According to this analysis, none of this comes from a more gentlemanly spirit, but rather from pilots doing their absolute self-interested best in different circumstances.

4 The effect of rules changes

In the US, much discussion of rules changes concerns whether the proposed rule will encourage or discourage pre-start gagging, and gagging and leeching on course. With the above analysis in mind, we can speculate a bit more concretely about these questions. The answers suggested by the analysis are somewhat surprising.

In general, rules changes can do one of two things. First, they can bump a given situation from the late start equilibrium to the mixed strategy equilibrium case or vice versa. This change may improve the safety situation of a huge gaggle, but it if anything enhances the importance of start-gate strategy to competitive soaring. Second, rules changes may be able to change the shape of the functions, especially the effect of other gliders on speed, in such a way that the overall speed is a much flatter function of starting time. In this case, weather knowledge, glider performance (cost!) and other factors that determine the individual speed function become more important.

I emphasize that I make no editorial recommendation. Furthermore, all these proposed rule changes have important other effects on safety and competition strategy that are not appropriate to consider here. I only consider their effect on the start game.

Fly the same course, merge the classes.

As we have seen, a wider range of glider/pilot performance should move up the start time and spreads pilots out. Pilots who want to see less gagging therefore might prefer contests with multiple classes on the same or nearby tasks. Putting the sports class on the same task is likely to have an especially strong effect. Handicapping the FAI classes is also likely to move up start times, by putting older gliders and newer pilots in those classes. The real, not handicapped performance spread is of course the relevant one for start time strategy.

Team flying and communication

The late-start equilibrium is destroyed when pilots find it advantageous to sneak off during the best time of the day. If a team of two or more pilots were to agree to start together at the best part of the day, they could make this strategy work more often, i.e. on days in which a lone pilot could not

beat the gaggle. (Soaring is like bicycle racing in this respect. Bicycle racers gain from being behind, since they can draft the rider ahead. In bicycle racing, formal or informal teams try to break out from the pack together.)

In the US, communication between pilots by radio is against the rules. Allowing such communication (or, more realistically, making legal the communication that already goes on) would strengthen the ability of such teams to form. In the context of the formal analysis, we are moving from *non-cooperative game theory* to concepts from *cooperative game theory* in which groups can coordinate their actions.

Weakening day devaluation

US contest rules strongly devalue a day in which many pilots land out. The motivation for this rule is to make contest outcomes less dependent on luck during weak days. It has an unintended consequence of strengthening the late-gaggle equilibrium and discouraging pilots from breaking away from the pack.

The vertical axis is really contest points rather than speed. Suppose the gaggle is going to delay and delay until it is quite likely that most will not make it back. You are considering whether to break off early. With day devaluation, that strategy is much less desirable. Suppose the strategy works: you make it back but a large number land out. Then you get few points for realizing what was going to happen, since the day is severely devalued or scratched altogether. Suppose that everyone does make it home at reasonable speed, however. Now it becomes a 1000 point day and you pay a heavy penalty for flying alone. In graphical terms, day devaluation lowers the “individual speed” part of the graph when everyone else starts late, making the no-equilibrium outcome less likely and the late-start equilibrium more likely.

Longer tasks

The analysis suggests that longer tasks for a given day may make the late-gaggle equilibrium *more* likely rather than less likely. If the task is short, the gaggle can sit around until 5:00 and still make it around the course. By this time, the overall speed is so low that someone who snuck out at 1:00 will beat it. Thus, the *short* task gaggle is more vulnerable to individuals breaking off and starting early. This is a surprising conclusion, since many pilots’ intuition is that longer tasks will cut down on start gate roulette and

gagging, and longer tasks have been suggested exactly with this end in mind.

The answer is, perhaps not everything else is equal. Short tasks are called of course on short and uncertain days. Thus, if more gaggles form on such days, it may reflect the increased advantage to gagging on a short and uncertain day, not the effect of calling a shorter task on an otherwise equal day. Also, the gaggle will obviously form and leave earlier in the day for a long task; if one objects to pre-start gagging rather than gagging on course a long task will certainly reduce such gagging.

Silent starts

The current US rules will change to starts that occur over a wide area monitored by GPS rather than by radio announcement and start gate. Some pilots like this change because they think it will lessen gagging.

The Nash equilibrium concept does *not* require that each pilot can see the starts of others. In fact, it was developed for simultaneous-move games in which each player does not see what the other players do, but must react to what he expects them to do. It is much harder to analyze games in which players can see moves of other players and send possibly false signals.

If starts were completely secret, perhaps pilots would venture out early on the first few days. But the pilots who started later would find markers on course and would win. Each pilot would then try to start later, trying to start after he *guesses* all the other pilots have started. By the end of the contest, each pilot would start very late, and he would discover that everyone else has started late as well.

The advantages of starting after others will still be there. Therefore, the silent start may make no difference to the late-gaggle equilibrium. In the mixed-strategy equilibrium, a silent start may make the option to leave early more desirable, if one can in fact sneak away without being followed. However, pilots may simply react by gagging *even more* tightly in the start area, so as to keep a better eye on who is leaving.

The PST

The pilot selected task is often advocated because it seems to lead to less gagging. From the point of view of the above analysis, this outcome is another puzzle. It is still advantageous to fly with thermal markers. If

we expand the strategy space to include where you go as well as when you start, pilots should all choose the same course. If everyone else is going to turnpoints 1,4 and 10, you should also go there (a little after the others) and use the markers. As with silent starts, even if you don't *know* where everyone else is going, you should try to guess, and the contest will soon settle down to a point where your guesses are on average correct.

It is possible that we do not observe this only because most contests feature unstable weather and relatively few PSTs. If a contest were to have a PST every day with stable weather and a close spread of pilot/glider performance, one might expect pilots to gravitate to a big gaggle that bashes around the same few, close-in, turnpoints near reliable house thermals.

Point penalties for late starts

An obvious idea to get pilots out on course is to give a point penalty for late start. For example, one could add one point to each pilot's score for each minute he fails to start after the gate is open. This modification would seem like a natural way to make start gate strategy less important, and break up the gaggle.

Alas, the analysis suggests that it may not work. Figure 8 presents such a case. Starting when the gate opens, the individual-speed curve is progressively lower as the day goes on. The point penalty does not change the size of the speed boost function, which is large and compact in time in this case. Therefore, all the point penalty accomplishes in this case is to slightly move up the time of the late gaggle.

If the point penalty is severe enough, it may lower the late-gaggle speed so much that we revert to the mixed-strategy case in which it benefits a pilot to leave alone early in the day. On the other hand, by moving the late-gaggle equilibrium forward in the day, it may raise its speed relative to flying alone and actually strengthen the late-gaggle equilibrium.

In any case, given that the extra speed one gains from leeching is such a strong function of the time one leaves relative to other pilots, a point penalty does not result in an overall speed function that is flatter with respect to start time, and thus will not lower the importance of start time strategy.

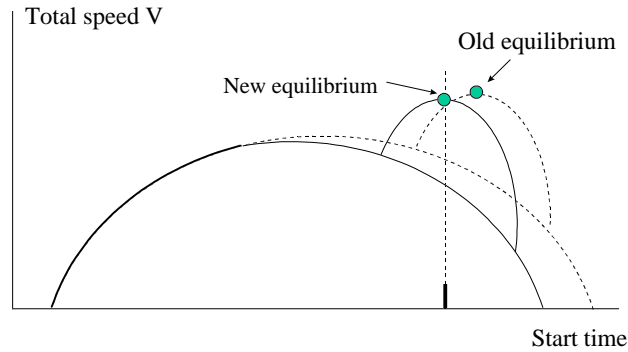


Figure 8: Effect of a point penalty for late starting.

5 Questions

It is clear that this analysis is only the beginning. First, we need a more better understanding of what the speed functions look like. We have a pretty good idea of speed flying alone vs. time of day. However, the analysis could be substantially improved with a solid quantitative understanding of the speed gained by the presence of other pilots, and how that speed changes in different weather conditions. Second, models such as this one need to be carefully and quantitatively contrasted with contest experience. Third, the analysis should be extended past the framework of one-shot, simultaneous move, non-cooperative games that I have presented here, to include the fact that some pilots can see what others are doing, and that teams may form that implicitly or explicitly coordinate their moves.

6 References

Kreps, David, 1990 *A Course in Microeconomic Theory* Princeton NJ: Princeton University Press.

Nash, John F., 1996 *Essays on game theory* Brookfield, Vt.: E. Elgar.

Schelling Thomas, 1980 *The Strategy of Conflict* Cambridge MA: Harvard University Press.