From Static Formulas to Market Dynamics: Charging the time in Air Navigation

Introduction

The "time-of-use tariff" model applied in modern electricity markets is a perfect example of how a scarce infrastructure should be managed. A kilowatt of electricity consumed during peak daytime hours does not cost the same as one consumed late at night. The reason consumers consciously shift the use of high-energy household appliances or electric vehicle charging to late-night hours is the rational pricing offered by the market. The raison d'être of the time-of-use tariff is the fact that the grid's capacity is limited. Indeed, in a scenario where all demand loads onto the infrastructure simultaneously, the system congests, and it becomes mandatory to deploy highly inefficient, high-emission, and expensive backup plants to meet instantaneous needs. To overcome this bottleneck, utility companies price peak hours higher and quiet hours lower, thereby distributing consumption homogeneously across time. The most elegant way to manage a capacity-constrained infrastructure without collapse and to prevent unnecessary carbon emissions is to price "time".

Although the sky appears boundless from below, it is actually a physical network with invisible corridors, narrowing bottlenecks, and rigid capacity limits, much like metropolitan highways. However, this fundamental capacity planning solution, which works flawlessly in power grids and global logistics networks, suddenly vanishes when it comes to the sky, which is perceived as infinite.

Established fifty-six years ago by an agreement signed in Brussels in September 1970, the European air navigation charging system still operates today with the simple, static mechanism of that era. The formula at the heart of the system, Number of Service Units = $W \times D$ (Weight × Distance), was presented as a great gift to civil aviation due to the "simplicity" it offered for large-scale calculations in an era without digitalization and big data algorithms. However, in today's world where artificial intelligence can optimize flight networks instantaneously, the value of a formula should be measured not by its simplicity, but by how responsively it answers the dynamic crises of the network. This half-century-old structure takes into account the aircraft's weight and the distance it flies, but "when" it flies is not included in the formula.

For instance, suppose two aircraft of the same type plan two different flights on the same route. One takes off at eight in the morning, when the airspace is most chaotic, and the other at two in the morning, when it is quietest. Today, both flights will pay exactly the same fee for air navigation services. Yet, the service provided to them by the system and their cost to the system and nature are entirely different. The flight taking off at two in the morning is alone in the sky, completing its route completely straight and at the most optimum altitude. The flight taking off at eight in the morning, however, is forced to lengthen its route by kilometers to avoid entering a congested sector. Due to the traffic density, the controller assigns a flight level lower than the optimum cruise altitude, causing the aircraft to burn much more fuel and generate higher emissions at this altitude where friction is greater. It might even encounter a speed restriction to maintain separation within the sector, which causes an additional inefficiency that increases carbon emissions.

These are all silent costs unfelt by the passenger inside the cabin, yet they leave heavy marks on the aircraft's fuel tank and the atmosphere. You can manage the same daily traffic either as a smooth system balanced across hours or in the form of intense waves compressed into narrow windows. On paper, both yield the same total number of flights; however, every wave that overflows capacity repeats the inefficiencies experienced by that 8:00 AM flight much more frequently and intensely. Congestion does not merely cause an operational crisis; it silently writes an exponentially growing carbon bill.

This massive difference hidden behind that "same bill" paid stands precisely at the intersection of two crises that the aviation industry is trying to solve simultaneously today: the capacity crisis in the European airspace, which has become congested again with the post-pandemic demand wave returning strongly, and the climate crisis caused by the European Green Deal's "2050 Net Zero" commitment. From an economic perspective, the root of these two problems is essentially in the same place: airspace, a scarce resource, is priced without considering the time dimension at all.

Economics has a name for this problem: Peak-Load Pricing. Bolic et al. (2017) demonstrated that this could function in a revenue-neutral manner within European airspace and establish a demand-capacity balance.

To see the rationality of this mechanism, there is no need to look far; it is enough to look at the airlines' own ticketing strategies. The price of an airline ticket purchased during peak periods such as holiday seasons or summer vacations is never the same as a ticket bought on a calm winter day. Airline companies optimize their own aircraft capacity by making peak periods more expensive and directing flexible passengers to quiet periods. It is a major contradiction that this dynamic pricing logic, which has been flawlessly applied to ticket prices in the industry for years and accepted as a natural reality of the market, is completely abandoned when it comes to the "airspace infrastructure" used by that aircraft. The "Dynamic Time Pricing" proposed within the scope of this study aims to transplant this familiar market dynamic from the passenger ticket to infrastructure pricing.

However, in aviation, the matter is not as smooth as simply adding a time multiplier to the formula. Pels and Verhoef (2004) prove that in oligopolistic structures where airlines possess market power, especially in networks operating with a collect-and-distribute logic known as hub-and-spoke, carriers already internalize delay costs (for instance, by paying for it internally by holding connecting flights to wait for a delayed passenger). Imposing a blunt flat fee on peak hours in a network built on these intensive arrival-departure "waves" compressed into narrow windows risks penalizing the airline a second time for the same delay and suppressing demand to the point of causing a net welfare loss.

The empirical picture, however, tells a different story. Morrison and Winston (2007) show with US data that two-thirds of delays still remain external and that even a small economic incentive targeting time can turn into billions of dollars in net welfare gains; meanwhile, Vo et al. (2026) detect that stretching the supply side alone (such as controller shift arrangements) is insufficient, and an integrated model addressing capacity management and pricing together is still missing in the literature.

This study targets precisely this gap in the literature: a "Dynamic Time Pricing" model is proposed, which redesigns peak-load pricing not just as a capacity tool but as an active carbon mitigation mechanism. In doing so, in harmony with the outcome of Pels and Verhoef's (2004) study, the aim is not to penalize the connectivity logic but to make the true cost of peak hours visible. A well-designed tariff steering mechanism offers a smoother sky to Hub carriers by shifting flexible traffic to quiet hours; it does not ban flying, it merely reveals the price and leaves the choice to the carrier.

But the matter does not end here. When an airline chooses a quiet hour in your airspace and pays less, exactly that hour might be a peak hour in another country at the destination. In other words, the applied economic incentive can work in reverse in someone else's airspace (Waterbed Effect). This elevates "Dynamic Time Pricing" from an algorithmic issue that a single state can solve by implementing it, into a European question requiring the network-wide coordination authority that EUROCONTROL already possesses. A dynamic Bonus-Malus pricing integrated with the EUROCONTROL Network Manager can eliminate horizontal/vertical flight inefficiencies while solving congestion, naturally driving down carbon emissions as a result. This is an integrated dual-dividend scenario that transcends the classic "either capacity or environment" dichotomy. However, the realization of this scenario depends on the entire network speaking the same language, rather than a single airspace.

The Legacy of 1970

The roots of the weight-distance formula used today trace back to 1970, according to EUROCONTROL's own history. In June 1970, the Organization's Permanent Commission (CN/28, 25 June 1970) approved the principle of establishing a regional route charges system. The formula relating to the calculation of the weight-based number of service units, which forms the main backbone of the framework, was developed by Wolfgang Philipp over his Christmas holidays that year, according to an anecdote shared in the institution's history. Mr. Philipp thus gave a fine Christmas present to the civil aviation world. The result was a multilateral agreement signed in Brussels on September 8, 1970, by the plenipotentiaries of the organization's seven member states at the time; to implement this, the Central Route Charges Office (CRCO) was established in 1971. This work by EUROCONTROL subsequently guided ICAO's own principles.

The system was updated with a new multilateral agreement signed in 1981 and in force since 1986—yet the core logic of the formula has remained nearly unchanged.

At that time, the goal was for this calculation to be simple; a calculation done by hand or with primitive computers needed to be uncomplicated. But more than half a century later, in a world where algorithms optimize airspace in real time, the value of a pricing formula must no longer be measured by its computational simplicity, but by how sensitive it is to the actual needs of the network, to capacity, and to the environment.

From Holiday Tariffs to Airspace Tariffs

So how does this mechanism work in practice? To understand it, there is no need to look far; it is enough to look at your own ticket purchasing experience. If you have searched for an Ankara-Paris ticket three weeks before the Eid al-Adha holiday, you know how the price spikes compared to a calm Tuesday three months before the holiday. Airlines do this intentionally: by making high-demand periods expensive and quiet periods cheap, they spread passengers across time according to their own capacity—whether passengers want to or not. This is called "revenue management," and the industry has applied it for years, unquestioningly; no one accuses an airline of unfairness because its holiday tariff is high.

What we propose is migrating this familiar logic from the ticket price to the air navigation charge. Currently, when an airline flies during the busiest hour of a holiday period, the air navigation fee it pays is exactly the same as flying in the dead of night—it looks at weight and distance, completely ignoring the clock. "Dynamic Time Pricing" adds a small time multiplier to the existing formula: slightly higher at peak hours, slightly lower during quiet hours. To make it concrete: let's assume an average air navigation charge of 1,000 euros. While a reasonable "congestion differential" is added to this during peak hours, a similar discount can be applied during the quiet midnight hours—the system collects the same total revenue, only the direction where it channels demand changes. Airlines are forced to perform the same calculation they have been doing for years in ticket pricing, now for infrastructure costs as well.

The cost of this to the environment manifests as endless holding patterns for arriving aircraft, inexhaustible departure queues for traffic on the ground, and paths extended by the air traffic controller for en-route traffic due to separation; the passenger does not notice, but the aircraft flies more distance than necessary. Alternatively, the controller issues a speed restriction; the aircraft deviates from the optimum speed at which its engines operate most efficiently, increasing fuel consumption. Or, if the most efficient flight level is full, the aircraft continues to fly at an altitude that is now "too heavy" for it. None of these are noticed from inside the cabin; the passenger simply thinks they experienced a safe, on-time flight. However, they leave silent marks that appear on the fuel gauge and the balance sheet.

EUROCONTROL's own data shows just how widespread this is: in the summer of 2024, en-route delays rose to 5.7 minutes per flight in July, which is almost double compared to the same month of the previous year. Moreover, approximately 46% of this delay is already a reactionary (knock-on) delay, meaning a bottleneck at one point transfers to the subsequent flights it is connected to.

Economic Incentive or Administrative Rule?

So, does this time multiplier replace the tools already in EUROCONTROL's hands, or does it add something to them? Because the EUROCONTROL Network Manager can already directly issue a waiting time called CTOT when a sector is overloaded. Just like the moments when the pilot announces that we have to wait "due to density in the airspace" after accepting passengers onto the aircraft; this situation is a purely administrative and price-free intervention operating through the ATFM system in the airspace.

The answer lies in the reality that administrative restrictions are "blind," whereas price-based tariff steerings are "smart". The slot intervention implemented by the Network Manager is a reactive crisis management resorted to against capacity overloads and operational contractions emerging on the day of the flight. The administrative ruling does not look at the economic dynamic inside the aircraft; it places a massive hub carrier, which will provide connectivity to intercontinental flights and races against time, into the exact same queue as a low-cost aircraft carrying holidaymakers traveling solely based on price, and holds them indiscriminately. In contrast, "Dynamic Time Pricing" does not aim to abolish this existing administrative mechanism; it seeks to minimize the need for it.

The price mechanism provides airline companies with a proactive guidance months in advance, while they are still in the schedule and seasonal planning phase. Flexible carriers receiving this economic market indicator voluntarily shift their flights to quieter hours when no operational crisis is yet in sight. Consequently, European airspace enters the day of the flight with a much more balanced and relaxed start; the probability of the Network Manager hitting the "red button" and holding aircraft on the ground also drops dramatically.

Furthermore, it must not be forgotten that the time loss and the excess carbon released into the atmosphere caused by administrative intervention (holding the aircraft on the ground or in the air) represent a pure deadweight loss in a lose-lose format within the system; this loss does not reflect positively on anyone's balance sheet. Conversely, the price mechanism has the potential to transform this same congestion cost into funds that will strengthen the future technological infrastructure of the airspace. In economic literature, it has been known since the classic works of Weitzman (1974) that price-based incentives and quantity constraints complement each other rather than substituting for one another. In aviation too, this duo comprises two different layers of an integrated network that ensures safety and environmental efficiency together, rather than being mutually exclusive tools.

Harmonizing with Hub Carriers: Preserving Wave Operations

The harshest theoretical and practical objection facing the model will undoubtedly come from major transfer hubs. At mega-airports like Frankfurt, Amsterdam, or Istanbul, the sky repeats the same rhythm several times a day: flights arriving from dozens of cities suddenly converge within a specific morning window, rapidly make their connections, and disperse. In the afternoon, this rhythm repeats once again. This literal "inhaling and exhaling" of the sky is not an accidental bottleneck; it is the core network engineering of hub-and-spoke systems, which is the most valuable asset of hub carriers. In fact, as in the example of Istanbul Airport, the construction of a third independent runway in some mega-projects is precisely intended to serve these highly synchronized arrival and departure waves.

In the literature, Pels and Verhoef (2004) prove that in oligopolistic structures where airlines possess market power, carriers already internalize delay costs when they hold a connecting flight to wait for a delayed passenger, effectively paying for it internally. Imposing a blunt flat fee on peak hours in a network built on these intensive arrival-departure "waves" compressed into narrow windows risks penalizing the airline a second time for the same delay and suppressing demand to the point of causing a net welfare loss.

The "Dynamic Time Pricing" we propose offers an inclusive solution against this justified rebellion of hub carriers, one that does not punish them but instead promises them a "frictionless sky". The fundamental goal in the architecture of this model is not to ban the network carrier's vital "wave" operations. The true power of the system lies in shifting other traffic with high schedule flexibility (low-cost carriers, charter flights, or cargo operations) that uses the airspace into quiet hours. When dynamic pricing comes into play, flexible carriers voluntarily pull their operations into off-peak hours to capture discounts or avoid peak-hour charges. As a result, whether across Europe's continental en-route arteries or within terminal approach (TNC) areas, the peak-hour sky opens up entirely to hub carriers racing against time. The tariff steering merely puts the invisible infrastructure and carbon costs of peak hours on the table; if the value added by the connecting passenger is higher than this cost, the airline pays the difference and sustains its wave in a frictionless environment. This nominal congestion fee paid by the hub carrier actually purchases a corridor where its route is not lengthened, it is not forced into lower altitudes, and it does not enter airborne holding patterns. The decision is left entirely to market dynamics.

Cross-Border Incentives: The Waterbed Effect and System Integration

As briefly mentioned in the introduction, the model has an operational trap that could be overlooked in a fragmented geography like Europe: a carrier might take off from its departure airport, say Istanbul, during the earliest and cheapest hour of the morning to lower its air navigation bill. However, upon arriving in Frankfurt airspace at the end of this multi-hour flight, it might land precisely during the most chaotic and expensive hour of the German airspace. A traffic flow successfully managed at its origin, yielding a highly economical air navigation bill, may arrive at the destination at the most expensive hour, when it will have its highest emission intensity. In other words, the appropriate billing strategy here has worked in reverse in another airspace; the small cost savings achieved at departure have quietly evaporated both in the budget and in the amount of carbon released into the atmosphere due to flight inefficiencies.

It is clear that the solution to this phenomenon, known as the "Waterbed Effect" in economic literature, cannot be achieved by defining peak and quiet hours in isolation according to each country's own local conditions. The time bands that will form the basis of pricing must be synchronized network-wide, taking into account major traffic corridors, continental transit times, and sector capacities.

A key factor that will make the system's implementation possible is that there is no need to construct a supranational institution from scratch to establish this mechanism and ensure synchronization. EUROCONTROL's Central Route Charges Office (CRCO) already centrally handles the collection of fees for cross-border air navigation services across all of Europe. The only missing element is the integration of time bands as a "common language" into the existing calculation algorithms. This transforms the matter from an intractable international political crisis into a technical coordination issue fully suited to EUROCONTROL's raison d'être and institutional capacity.

Overlooked Gains: Human Capacity, Safety, and System Resilience

Although the main axis of debates surrounding the model generally shapes around infrastructure inadequacy and carbon emissions, one of the most vital gains the pricing mechanism will bring is frequently overlooked: operational safety and the optimization of human resources. While IATA's reports on ATC-induced delays in European airspace have for years emphasized the devastating economic impacts of bottlenecks, EUROCONTROL's most recent ACE (ATM Cost-Effectiveness) report, published on June 8, 2026, reveals the gravity of this issue with mathematical certainty.

According to system-wide data, 75.9% of total ATFM (Air Traffic Flow Management) delays in Europe originate directly en-route. Even more striking is that nearly half (49.9%) of these massive en-route delays stem precisely from "ATC capacity and staffing issues". Integrating new air traffic controllers into the system is a rigid supply-side solution that demands years of training and dramatically increases fixed system costs, thereby inflating the unit rates passed on to airlines. Conversely, "Dynamic Time Pricing" aims to optimize the existing human resource without generating new personnel costs. When the number of aircraft filling the radar screen pushes capacity limits during peak hours, the span of attention the controller can allocate to each aircraft narrows, and the cognitive workload inevitably increases. Spreading traffic through time and price parameters not only opens up this 49.9% bottleneck; it also balances the controller's workload within human limits, permanently protecting the "safety margin" that is the most fundamental building block of aviation.

This picture also features a silent gain dimension reflected on the balance sheets of airline companies. The thousands of tons of extra fuel burned due to route extensions forced upon aircraft as the price of clustering in the airspace during peak hours, radar vectors assigned for separation purposes, zigzags, and speed restrictions are the direct costs of the airline itself, not the air navigation service provider. Thanks to the pricing model, when density decreases and flights smoothen, this massive waste in the fuel item disappears as well. The "optimum flight profile" that airline companies purchase by shifting flexible traffic or paying the peak-hour premium returns to their own balance sheets as a massive net saving at the end of the day.

A network operating at the absolute limit of its capacity is completely vulnerable to operational shocks. The EUROCONTROL ACE report shows that 35.8% of en-route delays and 50.8% of airport delays stem from severe weather conditions. When poor weather conditions, a technical failure, or an unexpected emergency strike a massive ecosystem like the European airspace, the packed system has no room left to absorb this shock; the crisis turns into a domino effect spreading to the entire continent within seconds. Spreading traffic homogeneously throughout the day gives the system exactly the operational buffer and resilience capacity it needs in the face of severe weather events. This flexibility, which prevents minor disruptions from mutating into major crises, serves as the invisible shield of a sustainable aviation network.

Conclusion: Preparing the Skies for the Future

The existing navigation pricing system, whose foundations were laid in the 1970s and which has served European aviation with stability for half a century, is a highly valuable historical legacy designed to ensure "simplicity and predictability" within the limited computational possibilities of its time. Thanks to this framework, which answered the operational needs of its own era flawlessly, European airspace achieved the integrated and massive network it possesses today. However, in today's world where the aviation industry flies billion-dollar, next-generation aircraft utilizing artificial intelligence-supported algorithms, the capacity and climate crises faced by the sector make it mandatory to bridge this valuable legacy with the market dynamics of our age. The sky is no longer a void with infinite capacity; it is a living, dynamic organism with instantaneous clustering points and rigid capacity boundaries.

The "Dynamic Time Pricing" proposed in this study aims to adapt the system to the needs of modern network management without severing the aviation sector from the rules of the past. This mechanism aims to minimize the instantaneous operational need for that authority by distributing traffic system-wide months in advance through rational price steerings.

While the new Performance Reference Period (RP4) planning for European airspace is currently being discussed on the industry's table, testing this integrated pricing model through pilot applications in a few of the continent's most chaotic corridors is by no means a remote possibility. The only missing element is the integration of time bands as a common language into the centralized collection experience and established infrastructure of the CRCO. Classic economics and environmental policies generally present policymakers with a harsh dichotomy: either choose capacity (growth) or the environment. Yet, "Dynamic Time Pricing" promises a "Win-Win" scenario, drawing strength directly from the operational realities of aviation.

Returning to that eight o'clock morning aircraft outlined at the beginning of the text: in a European airspace where time is accurately priced, that aircraft will not wait for a slot at take-off, will not lengthen its route for miles to escape density upon departure, and will reach its destination gliding at the optimum altitude with the least friction under the safe management of a controller whose workload on the radar screen has been eased. It must not be forgotten that the sky is a scarce resource; and scarce resources can be shared equitably, safely, and cleanly only when they are accurately priced.