A Slot Swapping Market for European Airlines: The case of the Win-Win Economy*

by

Hugo BARRAS†

Under the Supervision of Prof. Marleen MARRA

May 31, 2023

Abstract

This paper designs an inter-carrier slot swapping market for airlines operating within the ECAC area, offering them a decentralized, high-frequency, flexible, and last-resort mechanism to either barter or trade their slots. Depending on their operational circumstances, the airline operators choose which economy they wish to participate in within the market. The paper models the Win-Win economy scenario, enabling mutually beneficial slot exchanges between companies, with a matching program. An heuristic solution algorithm, tested on simulated data, concludes that the inter-carrier market reduces the ATFM delays of buyer flights in the Win-Win economy by an average of 18 minutes per swap. This reduction represents a saving ranging from 720 to 2,646 euros per swap, while the sellers consistently secure a slot they can reliably maintain, deviating by an average of only 3 minutes from their preferred slot.

Keywords: Air Traffic Flow Management, Slot Swapping, Matching, Heuristic Algorithm.

---

*I must first express my warmest thanks to Eliane Cadoret and Nicolas Galtier, whose inspiring talks allowed me to root my subject in a specific operational reality. My deepest gratitude extends to my supervisor, Prof. Marleen MARRA, for her invaluable guidance and support throughout this journey. I would also like to express my sincere thanks to Prof. Junnan HE for his time and advice. Lastly, my heartfelt thanks go to Pascale and Eric for their unwavering love and backing.

†Master's Degree Candidate in Economics, Sciences Po, School of Research: hugo.barras@sciencespo.fr. Any remaining errors are mine.

Contents

1 Introduction

Airlines represent a challenging industry in which numerous successful investors and executives have singed their wings. High entry costs, volatile fuel prices, acute sensitivity of demand to business cycles, as well as rigid government regulations and fierce competition are among the many factors that constrain airline EBIT[1] margins to a low level, around 4% in Europe between 2009 and 2019 (IATA, 2022). Consequently, the primary aim for airlines is identified as the pursuit of maximum efficiency in operations.

Operational performance, underpinned by punctuality, forms the cornerstone of airline profitability. It is a crucial factor for carriers, as it can induce both high operational and reputational costs. In addition, significant growth is observed in passengers' demand for air travel. Despite Covid's severe negative shock, Eurocontrol's baseline forecasting scenario expects 2019's number of flights within the European airspace (11.1 million) to be recovered in 2025. After 2025, an annual growth of 1.5% per year is anticipated in the 7-years horizon forecast (Eurocontrol, 2023). This trend is perceived as sustainable in the long run, with the annual increase in flights estimated at 1.2% up to 2050, signifying an increase of 44% more flights in the European airspace in 2050 than in 2019 (Eurocontrol, 2022b). The growth in air traffic, as a result, is expected to exert additional pressure on European airspace managers and heighten the likelihood of delays.

To deal with this long-standing trend, the *Single European Sky* project was launched by the European Commission in 1999. The project aims to modernize European airspace management mainly by developing working methods for better collaborative decision-making between stakeholders. Additionally, a technological complement, labelled *SESAR*[2], was launched in 2006 to further support these initiatives. In this context, Eurocontrol became the main actor of *Air Traffic Management* (ATM) in Europe. This pan-European organization was appointed for the first time as the *Network Manager* (NM) for the Single European Sky in 2011 via decision of the European Commission (European Commission, 2011). Since then, it's role is to handle the European Air Traffic Flow Management and to enhance its performance. In addition, SESAR is shaping the future of ATM in Europe by undertaking many research based ventures to deploy new technologies in this field.

Among all these projects, the *Air Traffic Flow Management* (ATFM) *Slot Swapping* project, released in 2014, aims to address the issue of ATFM slot swapping. A decisive issue in terms of operational efficiency, ATFM slots, as opposed to airport slots (negotiated by airlines to obtain rights to land or take-off at a specific airport at a specific time), are used in case of ATFM delays (i.e., delays triggered by airspace congestion). In the event of unforeseen regulations, or airspace capacity reductions, in an airspace area or at an airport, the NM will redistribute flight slots to delayed aircraft. These slots are called ATFM slots. ATFM slot swapping is a way to reduce the impact of delays on airspace operators (AOs'), i.e., airlines. In Europe, it is reported that between 100 and 200 swaps take place each day. At present, ATFM slot exchanges are only carried out within the same airline, and only marginally with airlines belonging to an alliance. The average cost savings per swap are estimated by AOs' to be between 4,600 euros and 4,900 euros, while cost savings per flight are reported to range from less than 1,000 euros to over 10,000 euro (SESAR, 2014).

[1]Earnings Before Income and Taxes (EBIT) margins is a financial metric measuring the profitability of a firm relative to its revenue.

[2]SESAR stands for Single European Sky ATM Research Project.

This paper aims to design a decentralized, high-frequency, flexible, and last-resort marketplace for ATFM slot swapping among competitive airlines within the European airspace. This marketplace, named the *inter-carrier ATFM slot swapping market*, seeks to diminish delay costs for airlines, and is intended to operate in harmony with the solutions currently under development by Eurocontrol. The market is divided into three distinct economies: the Re-Routing economy, the mutually beneficial (Win-Win) economy, and the unilaterally beneficial (Lose-Win) economy. Each economy consists of buyers and sellers dealing with specific operational scenarios, striving to secure a particular slot. Within the Win-Win economy, the buyer is a flight burdened by an ATFM delay, while the seller is a flight impacted by both an ATFM delay and a ground delay. The simultaneous occurrence of ATFM and ground delay prevents the seller from maintaining its ATFM slot, prompting it to pursue a later departure time. As a result, this economy includes two types of flights with a mutual interest in the exchange, eliminating the need for a financial transaction. This paper thereby formalizes the inter-carrier ATFM slot swapping market and articulates a solution for the Win-Win economy via a linear programming model for optimal flight pairing. A heuristic algorithmic solution is subsequently introduced to assess the effectiveness of this model, based on simulated data. The findings show that buyer flights reduce their delay by an average of 18 minutes, while sellers consistently secure a manageable slot, typically deviating by just 3 minutes from their preferred slot. For the buyers, this reduction represents a saving ranging from 720 to 2,646 euros per swap.

The article is divided into five sections. Section 1 introduces the topic, offer a synthesis of the related literature and present the objectives of this paper in light of the most recent articles. Section 2 unpacks the operational landscape of European Air Traffic Management. Section 3 models the inter-carrier market and the solution of the Win-Win economy. Section 4 provides a heuristic solution algorithm for the model, and provides results using simulated data. Finally, Section 5 concludes and lays the groundwork for future research.

1.1 Related Literature

The literature on ATFM optimization and strategies has seen steady growth since the late 1980s. Odoni's 1987 seminal work offers a comprehensive classification of potential ATFM measures designed to alleviate air traffic congestion and thereby reduce ATFM delays. Odoni categorizes these solutions based on their implementation timelines, ranging from long-term solutions like infrastructure changes to medium-term administrative measures such as strategic airport slot assignment and congestion pricing, or even a combination of both (Odoni, 1987).

Medium-term solutions notably include auctions and congestion pricing to formulate efficient alternatives to traditional administrative resource allocation procedures. Certain authors suggest the use of auctions for long-term strategic allocation of airport slots, i.e., (Rassenti, Smith, & Bulfin, 1982) and (Ball, Donohue, & Hoffman, 2006), while others put forth a Vickerey-Clarke-Groves (VCG) sealed-bid type of auction for ATFM slots allocation in Europe, proving its efficiency in unveiling airlines' private resource values (Raffarin, 2002). However, the findings remain in the theoretical realm due to data constraints.

Furthermore, more short-term oriented solutions have emerged in operations research with the development of the Ground Holding Problem (GHP), or the problem of optimal ground delay allocation. This problem, in its simplest form (Single Airport Ground Holding Problem, or SAGHP), presumes only one airport in the system faces capacity constraints affecting only arriving flights, and it has been formulated for various scenarios. Its complex version, the Multi-Airport Ground Holding Problem (MAGHP), represents increasingly technical models for managing flight delays, involving networked airports and adjustable aircraft speeds, respectively. Seminal works by Vranas develop exact formulations and an heuristic for MAGHP – (Vranas, Bertsimas, & Odoni, 1994a) and (Vranas, Bertsimas, & Odoni, 1994b). All these models share a centralized decision-making feature, typically a Network Manager like Eurocontrol, assigning delays to minimize global costs. The Collaborative Decision Making (CDM) program by the FAA decentralizes this process, allowing airlines more control through shared information and cooperative resource allocation. The FAA's inter-airline slot exchange procedure can be seen as a mediated bartering system (Vossen & Ball, 2006).

In the past decade, Eurocontrol, under the SESAR program, has been developing the User-Driven Prioritization Process (UDPP). This concept grants airlines an influence in the sequence of flight management during times of capacity constraints or similar situations. The UDPP provides AOs' with enhanced operational flexibility and aims to boost overall air traffic management efficiency by acknowledging the preferences of each operator. UDPP is since used to investigate the optimal reallocation of ATFM baseline delay among flights, through analytical models (Ruiz, Guichard, & Pilon, 2017) and empirical studies (Cliff, 2019).

Building upon the concept of UDPP, the SESAR project *SlotMachine* leverages slot reallocation among airlines, generally over a daily timeframe. SlotMachine is devised with a genetic algorithm that reassigns slots to flights with the goal of maximizing the social utility of the participants (Schütz, Gringinger, Pilon, & Lorünser, 2021). Each flight communicates a satisfactory time interval, as well as an optimal slot within this interval, in an encrypted format. These individual preferences are then converted into weights that the genetic algorithm maximizes, producing a flight list that optimizes the social benefit. A credit mechanism is presented to compensate airlines with a low delay valuation, which would thus forfeit their slots to airlines with a higher valuation. In their 2022 paper, the same authors construct the market mechanism underlying the solution presented in their 2021 study (Schütz, Ruiz, Gringinger, Fabianek, & Lorünser, 2022). They develop a combinatorial auction based on credit transfers and expose that their market mechanism is incentive compatible, budget balanced, efficient, and strategy-proof. Finally, Hondet et al. explore the use of reinforcement learning techniques for slot swapping in case of fleet disruptions, showcasing improved performance over idle behavior through testing on historical schedules, despite not reaching optimal decision-making (Hondet, Delgado, & Gurtner, 2018).

The solution proposed in this paper relies on matching theory in economics. Indeed, in each of the economies, flights are separated into two disjoint sets based on their operational situation. Two-sided markets are particularly suitable for solving operations research or economic problems. The specific problem modeled in this paper straddles both disciplines. In the case of the Win-Win economy, the question boils down to an optimal allocation problem of slots between supply and demand to find the optimal exchange pairs. Traditional matching theory provides a solution for this situation. This configuration emerged in the literature with the study of marriage markets. The analysis of such markets using game theory was initially pioneered by Gale and Shapley, leading to the proof that a stable matching exists for every marriage market (Gale & Shapley, 1962). The authors also demonstrate that, when agents' preferences are strict, such a stable market can be both M-optimal or W-optimal, depending on whether men or women initiate the proposition. The next significant contribution came from Shapley and Shubik when they analyzed a related set of games termed assignment games, providing similar outcomes to the earlier theorems (Shapley & Shubik, 1971).

1.2 Thesis Objectives

All the European solutions for ATFM delay management that are currently in development and discussed in the aforementioned studies rely on a centralized authority, the NM, for reassigning delays to various flights. Although the SlotMachine system, relying on UDPP, allow airlines to establish preferences, these are broad and generic. Indeed, they are formulated over significant time frames and do not account for the operational subtleties of each flight. As mentionned earlier, in the United States, the rapid development of Collaborative Decision Making methods has placed airlines at the heart of managing ATFM delays. Today, CDM is a widely adopted technology in Europe (see Figure 1). It could thus be used to design a decentralized, high-frequency, flexible, and last-resort marketplace for managing ATFM slots. Decentralization would enable airlines to be in control of their decisions. They could barter, sell or purchase ATFM slots downstream of the solutions already proposed by Eurocontrol. Moreover, the solutions currently under development at Eurocontrol aim for optimization well upstream of operations (at best at the start of the day). The high-frequency dimension of the market proposed in this paper allows for exchanges up to ten minutes before a flight's departure. Additionally, the precise operational embedding of this market offers considerable flexibility to AOs, who can adopt a tailored strategy for each of their flights depending on the operational situation they face at a given time. Finally, the NM's solutions are primarily aimed at addressing extreme delays, which, as detailed in Section 2.3, constitute only a marginal fraction of total delays. In contrast, the last-resort role of the inter-carrier market offers a solution for dealing with average delays, around 15 minutes, which will not be erased by the regulator's solutions. Thus, the objective of this paper is to formalize such a market, ultimately presenting a solution and results for the case of mutually beneficial slot swapping.

2 Operational Context

In this section, the operational context in which the inter-carrier slot swapping market would operate is explained. The current European Air Traffic Management system is first described, followed by a detailed account of the slot allocation procedure, and finally, an analysis is provided on the management, causes, and costs of delays.

2.1 Air Traffic Management Environment

In the early days of commercial aviation, air traffic was sparse enough that each aircraft could manage its own navigation autonomously. However, the exponential growth of air movements quickly required a centralized coordination of operations. Thus, in the shadow of the grand narrative, a section of the Treaty of Versailles gave birth to the International Convention for Air Navigation (ICAN) and its General Rules for Air Traffic. Up until World War II, ATM developed slowly and saw few significant regulatory and technological advancements. It was in 1944 that the Chicago Convention laid the groundwork for a comprehensive regulatory model by temporarily establishing the International Civil Aviation Organization (ICAO)[3]. This United Nations specialized agency would permanently replace ICAN in 1947 to establish air traffic control procedures in all aspects. Now comprising 193 member states, ICAO produces several reference documents for the industry which outlines the norms, principles, and techniques that ensure the coordination of international air navigation.

Air Traffic Management can be divided into three categories: *Air Traffic Control* (ATC), *Airspace Management* (ASM), and *Air Traffic Flow Management.* ATC is responsible for preventing flight collisions and providing information services to aircraft, such as weather reports and facility conditions. ASM is in charge of allocating airspace to its users and can be further subdivided into three categories: control tower, approach control, and area control center. Finally, ATFM addresses airspace as a limited resource and is responsible for matching airspace demand to capacity by adjusting aircraft flows when imbalances occur due to various exogenous factors, such as ATC staffing deficits, traffic volumes, weather conditions, or special events. All these operations are based on Communication, Navigation, and Surveillance (CNS) services.

In Europe, Eurocontrol, or the European Organization for Safety of Air Navigation, is the intergovernmental institution in charge of the harmonization and coordination of Air Traffic Management. Established in 1960, and currently made of 41 members, its role is to guarantee a safe, cost efficient and environmentally friendly airspace while responding to the increase in air traffic. This concretely consists of providing services to Air Navigation Services Providers (ANSPs) and AOs' such as air traffic flow planning, technical and operational standards implementation, research expertise and trainings. ANSPs are public, private or public-private entities managing air traffic for a company, region or country. Eurocontrol centralizes the operational information delivered to ANSPs and AOs through on online platform called the NOP Portal. Its geographical perimeter of authority corresponds, approximately[4], to the European Civil Aviation Conference (ECAC) area represented on Figure 1.

[3]From 1944 to 1947, the year of ratification by the 26th country, ICAO was only a provisional organization.
[4]Approximately because the ECAC area is made of 44 countries whereas Eurocontrol is made of 41 countries.

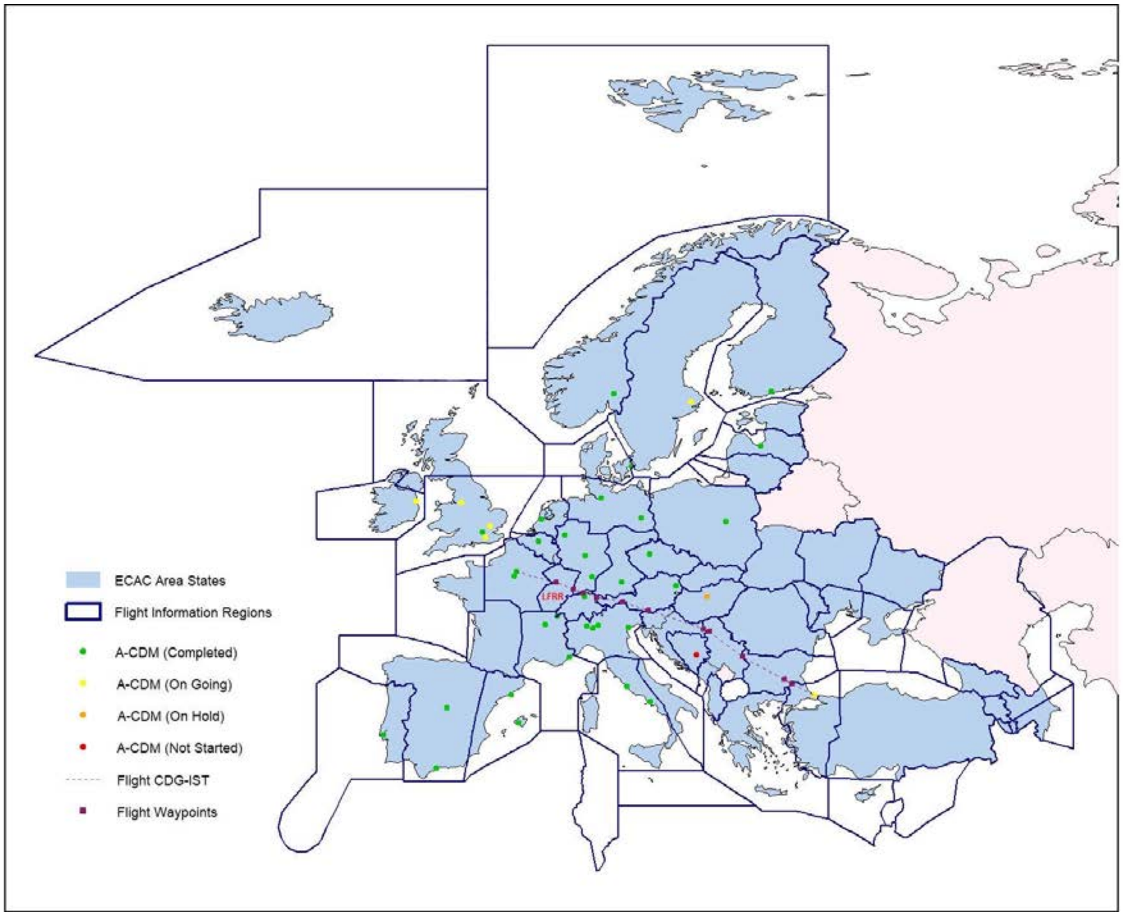

Figure 1: European Airspace Area[5]

As depicted in Figure 1, a flight connecting an origin-destination (here, CDG Paris to Istanbul) pair follows a specific route. This route is calculated by a flight dispatcher, approved by Eurocontrol, and validated by the flight captain. These routes are defined by an altitude, a speed, a separation distance between aircraft, and a heading. The emergence of Free Route Airspace (FRA) zones in some European areas, such as in western France, allows for relaxation of this structure, but FRAs remain very limited experimental measures.

If Article 1 of the Chicago Convention stipulates that states overflown by aircraft are sovereign over their airspace, this airspace is not delimited by terrestrial regions but by Flight Information Regions (FIRs). These zones, established within national borders[6], are flight volumes defined by a ceiling and a floor[7]. Figure 1 shows the European FIRs established by Eurocontrol for the year 2019. As schematically illustrated in Figure 2, FIRs are themselves divided into sectors. ATC agents are responsible for supervising these sectors to maintain the minimum horizontal and vertical separation distances between aircraft and to set the flight levels of aircraft entering and leaving the sectors. The definition of sectors is dynamic and falls under the purview of ATC agents. It relies on structural factors (e.g., technical equipment, military requirements, separation minima, etc.) and human factors (e.g., training and controller availability).

[5] The data used for Figure 1 comes from the *Aviation Data Repository for Research* from Eurocontrol.

[6] with rare exceptions, such as between Germany and Switzerland for example.

[7] Although FIR has remained the generic name, above 19,500, 24,500, or 28,500 feet (depending on the country), FIRs are called Upper Information Regions (UIR).

[8] As it is a schematic representation, no data have been used to generate this graphic.

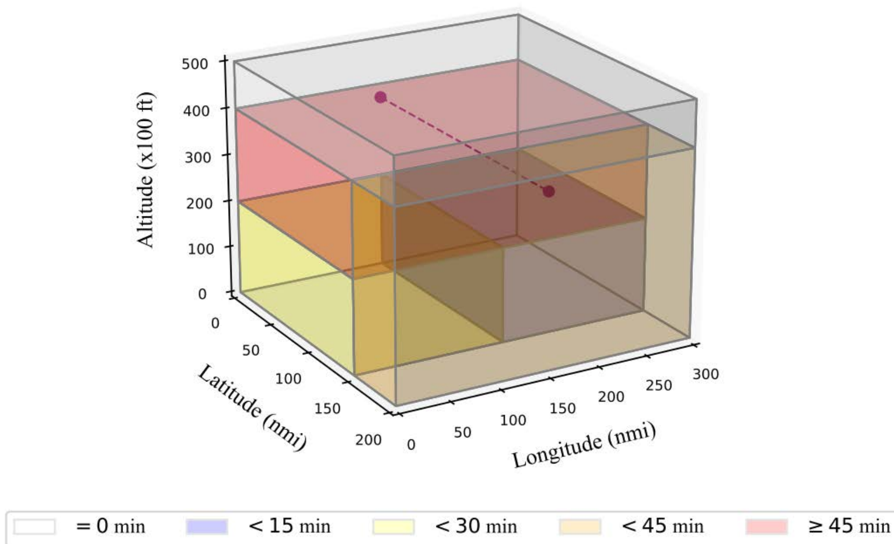

Figure 2: Zoom on Flight Information Region LFRR[8]

Figure 2 presents a 3-dimensional schematic depiction of the FIR LFRR, which surrounds the area of Reims in France. This region is also represented in Figure 1 and traversed by the exemplified flight, marked in purple in both figures. In the illustration, each sector within the FIR is subject to one of four distinct types of regulation, with one sector remaining unregulated. The exemplified flight intersects the most heavily regulated sector, denoted in red, hence will be subject to an ATFM delay (in this instance $\geq$ 45 minutes). Such circumstances are particularly common during summer days. In this situation, the flight is said to be *regulated*, and it will therefore receive an ATFM slot.

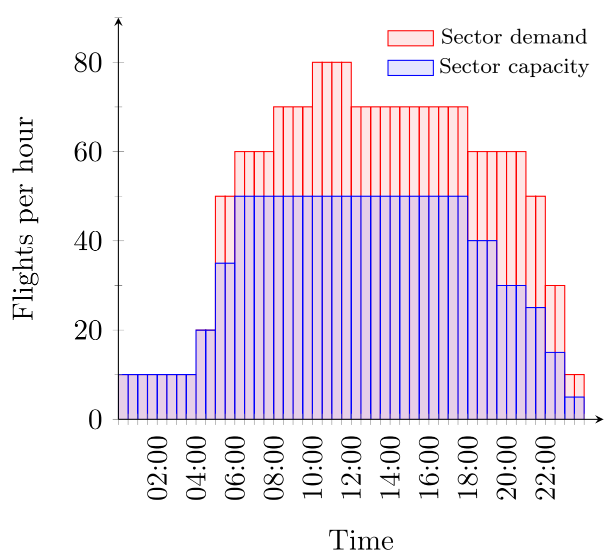

Figure 3: $\geq$ 45 min regulated sector capacity

Figure 3 represents most heavily regulated sector (in red on Figure 2) crossed by the exemplified flight at the air traffic controller level. Indeed, ATC agents, attached to Eurocontrol, manage the flows of the 44 states belonging to the ECAC area. Their role is to ensure that user demand from AOs does not overload the regulatory capacities offered by the infrastructure. In Europe, ATC sector capacity is defined as the *"maximum number of aircraft that can enter the sector in a specified period"* while still permitting an *"acceptable level of controller workload"* (A. Cook, 2007). When the volume of flights in a sector exceeds its saturation level (as illustrated in Figure 3), ATC controllers will limit the flow of flights through that area, and trigger a regulation. Regulations can be of three different types: **[1]** *departure regulations* (at origin airport), **[2]** *arrival regulations* (at destination airport), and **[3]** *en-route regulations*. Flights affected by a regulation will consequently be assigned an ATFM slot, whose allocation follows a specific procedure.

2.2 ATFM Slot Allocation Procedure

Each year, airlines negotiate departure slots for their flights with the national civil aviation authority. This negotiated slot is called an *airport slot*. In the ATFM jargon, the airport slot is called the *Scheduled Off-Block Time* (SOBT). This paper, and therefore the slot allocation procedure described bellow, is not about airport slots, but about ATFM slots. In the air transport industry, the words *slot*, *ATFM slot*, or, as detailed later, *Calculated Off-Block Time* (CTOT), are used as synonyms to name the updated slot a flight get from the NM in reaction of a regulation.

ATFM slots are the result of the failure of a long term optimisation process called *Air Traffic Flow and Capacity Managament* (ATFCM). In most context, the terms ATFM and ATFCM are perfect substitutes. The ATFCM process is divided into three phrases. **[1]** *Strategic Flow Management* occurs seven days or more before the day of operation (D-X – D-7). It aims at matching long-term demand with capacity and it results to the *Network Operations Plan*, or *Strategic Plan*. **[2]** *Pre-tactical Flow Management* occurs six days before the day of operation (D-6 – D-1). It aims at updating the Strategic Plan in the light of adjusted demand estimation based on the reception of filed flight plans and it results to the *ATFM Daily Plan*. **[3]** *Tactical Flow Management* occurs on the day of operation (D-day). It aims at providing updates to the ATFM Daily Plan as actual data informed about real time demand and capacity. The daily traffic is then managed through slot allocation and re-routings.

To maximize efficiency and accuracy in the ATFCM process, all the stakeholders (AOs', airports, military users, and traffic control centres) must collaborate in an *inclusive*, *transparent* and *trustful* manner (Eurocontrol, 2022a). The collaborative decision making techniques allow the best positioned stakeholder to ask the best dynamically optimized decision on the basis of the most comprehensive, and up-to-date accurate information. CDM relies on a fluid circulation of information, which requires specific settings and infrastructures. When an airport is able to operate in a collaborative decision making mode, it is called an airort-CDM (or A-CDM). As shown on Figure 1, most European airports are A-CDM.

The Computer Assisted Slot Allocation (CASA) is an automatic and centralized process to distribute the slot list to flights impacted by the same regulation. The algorithm functions in an aircraft operator point of view, meaning that its activation is conditional on the submission of the flight plan by the AO. Submitting a flight plan represents a request for an ATFM slot. As the algorithm follows the principle of *"First Planned – First Served"*, this develops an incentive for AOs' to submit their flight plan as soon as possible – allowing an earlier accurate view of the traffic. CASA will first compute an Estimated Take-Off Time (ETOT) based on each flight declared Estimated Off-Block Time (EOBT). The EOBT is transmitted by the AO to the NM, and must reflect the actual expected ready-to-go time of the aircraft. The ETOT is then computed by adding the taxi-time to the EOBT. This enables CASA to calculate the Estimated

Time Over (ETO), which is the time at which the flight will enter each sector on its planned route. A regulation is activated by the NM when the ETO slot list for a specific sector exceeds capacity. In this case, CASA sequences flights ETO of the regulated sector using the new flow rate and their arrival order. If the ETO is associated to a free slot, the flight receives no delay. In this case, the flight will take-off at its SOBT, i.e., its airport slot. However, if the ETO slot has already been allocated, then the flight will receive the next available slot and get a delay equal to the time difference between the two slots. It is important to note that the ETO slot list is highly volatile has new flights enter and leave constantly. The final step in this CASA slot allocation process is the calculation of the CTOT, i.e., the ATFM departure time, inferred from the flight list of ETOs', by chain reaction.

A flight can be exposed to an ATFM slot up to the very last minute before its SOBT. Slot swapping consists of exchanging the CTOT of two flights. A swap must be requested by the AO to the NM through a chat on the NOP Portal. *"Aircraft operators shall only request swaps concerning flights for which they are the responsible operator* [...] *or where there is a formal agreement between both aircraft operators for swaps to take place between their flights"* (Eurocontrol, 2022a), even though the NM will not check whether an agreement has been signed or not. For a swap to be accepted by the NM, the following conditions must *at least* be filled: **[1]** the two fights must be subject to the same most penalising regulation, **[2]** cannot have experimented more than three swaps, **[3]** must not be exposed to more than three regulations each, **[4]** must be in the same flow. If all these criterion are met, the flights will be qualified as *swappable*. This implies that it is likely that the requested swap will be accepted by the NM, but not guaranteed. Swaps can be realized from two hours before take-off. A swap is motivated by the willingness to prioritize a flight over another one. This decision is manually taken in collaboration between ATM flight dispatchers and sector managers (economic decision makers helping dispatchers to make cost-effective choices) and is driven by financial concerns, typically because one flight cost of delay is higher than another one. Hence, slot swapping is currently only an infra-carrier activity. Indeed, most of ATFM slot exchanges are carried out within the same airline, and only marginally with airlines belonging to an alliance. In Europe, between 100 and 200 swaps take place each day. At present, the average cost saved per swap is estimated by AO's between 4,600 euros and 4,900 euros, while costs saved per flight can vary between less than euros 1,000 euros to more than 10,000 euros (SESAR, 2014). As an example, from January 1 to December 1, 2022, Air France carried out 750 slot swaps, representing a saving of around 3,5 million euros (Air France, 2022).

Finally, it is important to introduce the concept of re-routing. Re-routing consists of using an alternative route, or flight level, to avoid the regulated sector. Re-routing opportunities are the main reason that can lead a flight to leave the ETO slot list of a specific sector. This is implied by the fact that each route is associated to a different ATFM slot. Therefore if a flight chooses its re-routing option to reduce its delay, it will get a new CTOT. This CTOT will be named in this article the *re-routed CTOT*. The re-routing operation is in the hands of the AO, and can be an effective way to reduce delay. However, re-routing option are far from being systematic, and have a negative environmental impact (as the initial route computed is always the shortest one). As an example, from January 1 to December 1, 2022, Air France carried out 6,600 re-routings for an estimated delay reduction of 230,000 minutes, i.e., a reduction of about 35 minutes of delay per re-routing (Air France, 2022).

2.3 Causes & Cost of Delay

Delay reduction plays a pivotal role in cost efficiency for airlines. In 2019[9], over the 30,427 IFR[10] flights within the ECAC area per day, around 50% were delayed on departure by at least 5 minutes, 25% by at least 15 minutes, 15% by at least 30 minutes, 5% by at least 60 minutes, and 1.5% by at least 120 minutes (Eurocontrol, 2022). The average delay per flight being 13.1 minutes, a statistic that is consistently around 15 minutes since 2015 (Eurocontrol, 2020). The main sources of air transport delay can be divided into 4 categories (A. Cook, 2007). On the one hand, **[1]** *Reactionary delays* are triggered by an initial and causal delay (e.g., previous flights or cascading events) whereas **[2]** *airline delays* are operational delays attributable to airlines (e.g., crew rotations, maintenance or servicing). On the other hand, **[3]** *ATFM En-Route delays* are caused by restrictions on the en-route phase of the flight (e.g., traffic congestion), and **[4]** *ATFM airport delays* are related to restrictions at airports (e.g., airport congestion).

These four categories account for 88% of the total average delay per flight of 13.1 minutes observed in 2019. Each category respectively constitutes 43%, 26%, 12%, and 7% of this share (Eurocontrol, 2019). The remaining 12% can be grouped into a fifth miscellaneous category. Again, these statistics have been consistent over the course since 2015. For the sake of the model presented in this article, the four aforementioned categories will be coupled into two groups. On one side, ATFM En-Route and ATFM airport delays will be paired together and constitute a first group call *ATFM delays*, and accounting for 19% of total delays in 2019. On the other hand, reactionary and airline delays will be paired together and constitute a second group call *ground delays*, and accounting for 69% of total delays in 2019. In this paper, flights both impacted by a ground and an ATFM delay are call *completely delayed* flights.

Measuring the cost of delay is an active field of research in air transport economics. As pivotal as it may be, it is a hard problem to quantify. Nonetheless, the average tactical delay cost[11] with network effect of a flight delayed at gate is estimated to 147 euros per minute for greater than 30 minutes delays, and 40 euros for smaller than 30 minutes delays (A. J. Cook & Tanner, 2011). This implies that, for the year 2019, the average cost of delay endorsed by ECAC area operating airlines were between 500 and 2,000 euros per flight for at gate delays only[12]. These statistics, amplified by the volume of flights, underscore the importance of delay management for airlines.

[9] Although more recent statistics exist, 2019 is kept as a benchmark throughout this article to set aside the Covid statistical disruption.

[10] Instrument Flights Rules (IFR) navigation rely on instruments, by contrast to Visual Flight Rules (VFR).

[11] Tactical delay costs incurred on the day of operations, by opposition to strategic delay costs that incurred at the planning level.

[12] Multiplying the average delay per flight (13.1 minutes) by the cost of one minute of delay at gate estimated by Cook & Tanner (2011).

3 Market Design

This section presents the model behind the inter-carrier slot swapping market. The market architecture will be first described, before formalizing the market in mathematical terms. The Win-Win econnomy model will finally be exposed.

3.1 Architecture

The inter-carrier slot swapping market is divided into three economies, where flights can either barter or trade their ATFM slots depending on their specific operational situation.

[1] The *Re-Routing* economy involves a swap between a flight that has a re-routing opportunity and a flight that do not have any possible re-routing option. To understand this case, it is important to remind that re-routing aims at improving the ATFM slot, and therefore always lead to a new and earlier slot. The AO is aware of the time associated to the re-routed CTOT. Hence, in this situation, the re-routed flight will be able to sell its original slot at auction to potential buyers before requesting the re-routed slot to the NM. By doing this, the seller flight will both reduce its delay by opting for re-routing, and also generate a financial surplus by auctioning its initial slot. Potential buyers will be the flights for which the proposed slot at the auction maximizes utility.

[2] The *Win-Win* (or mutually beneficial) economy involves a swap between an ATFM delayed flight and a completely delayed flight which cannot hold its ATFM slot because of the ground component of its delay. It is important to clarify that not all completely delayed flights are automatically unable to hold their ATFM slot. However, the Win-Win economy pertains solely to these specific flights. Thus, for the sake of simplicity, completely delayed flights will systematically refer to completely delayed flights unable to hold their ATFM slot. In such instances, the flight will be willing to exchange its ATFM slot for a later one that ideally matches, or which is at least the closest from, its Estimated Off-Block Time (EOBT) – the expected ready-to-to time. Prospective buyers for these ATFM slots would consist of flights experiencing ATFM delays that are seeking earlier slots to minimize their respective ATFM delays. The paper concentrates on this particular scenario.

[3] The *Lose-Win*[13] (or unilaterally beneficial) economy involves a swap between two ATFM delayed flights. In this situation, no external factors, such as re-routing or ground delay, would make the swap mutually beneficial in operational terms. Consequently, there will always be a winner and a loser. However, if accepting a loss at time t increases the likelihood of gaining an advantage at time $t+1$, flights with a low slot valuation might be willing to lose swaps now to boost their chances of winning a swap when they have a high slot valuation. The valuation of slots can vary greatly among flights, depending on factors such as the time of day, the day of the year, among others. This scenario would necessitate a compensation scheme for the losing party in a swap, such as credits. Moreover, the flight that loses the swap might also have the option to transfer the credit to another flight with a high slot valuation that belongs to the same airline.

Figure 4 provides a depiction of the three economies, illustrating their estimated shares and detailing the gains that arise from these economies for both the buyer and the seller sides.

[13]The semantic choice of winning or losing is defined in an operational point of view. Indeed, the first objective of airlines is to operate as efficiently as possible. Therefore the operational performance will always be more important than the financial surplus. Hence, *winning* means reducing delay, and *losing* means increasing delay.

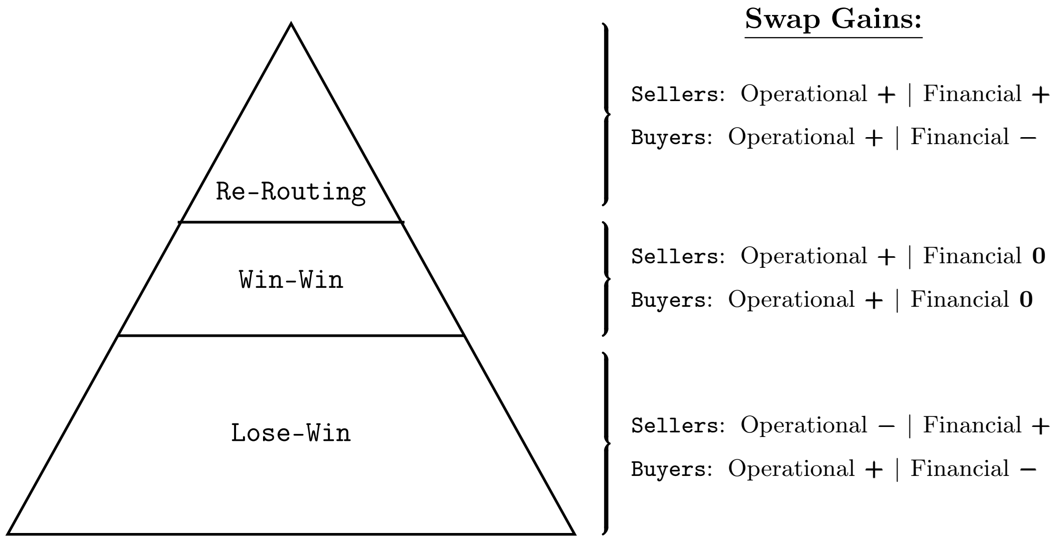

Figure 4: Architecture of the inter-carrier slot swapping market.

It should be noted that two sets of conditions must be met by a flight to enter the inter-carrier market. Firstly, it should be an IFR, commercial scheduled, medium-haul flight connecting an Origin-Destination within the ECAC area, carrying passengers (cargo flights are excluded). Scheduled commercial flights are defined by Eurocontrol as scheduled mainline, regional and low-cost flights. This market segment represented 83% of total flights in 2022, or 9, 185, 610 flights (Eurocontrol, 2022c). Secondly, exposure to an ATFM delay and thus an ATFM slot is required for all flights in the inter-carrier market, along with swappability, (i.e., they must meet the four conditions set out in Section 2.2). Lastly, the inter-carrier market is first an *ex-post* NM solutions market, and second an *ex-post* infra-carrier market. This means that flights coming on the inter-carrier must have both exhausted the solutions provided by the NM as well as their infra-carrier swap opportunities.

It is crucial to underline that the inter-carrier market is designed to be **[1]** decentralized, **[2]** high-frequency, **[3]** flexible, and **[4]** a last-resort marketplace. **[1]** The decentralization aspect is based on the fact that the decision to participate in the inter-carrier market is entirely at the discretion of the AO. **[2]** The high-frequency dimension of the market, on the other hand, lies in the fact that the algorithm is designed to operate every five minutes, thus producing swaps at a very regular interval. **[3]** Furthermore, the highly specific operational grounding of the market offers AOs significant flexibility in their strategy for choosing the economy and position that maximizes their probability of obtaining the slot they desire. **[4]** Finally, as explained in section 1.1, other solutions are being developed for reallocating slots in the case of extreme delays. With the average delay per flight consistently around 15 minutes each year (see Section 2.3), AOs are enabled to regard the inter-carrier market as a last-resort exchange market.

As referred to in the flight timeline detailed in Section 3, the activation of the inter-carrier market is anticipated approximately an hour before the Scheduled Off-Block Time. Indeed, swaps can commence two hours before the SOBT. Initially, all opportunities within their infra-airline market will be sought to be exhausted by airlines before resorting to the inter-airline market. This is based on the understanding that an infra-airline swap is always more desirable than an inter-airline swap. As a result, the inter-carrier market's activation is expected around an hour before the SOBT.

3.2 Definitions

The market horizon is set as the time span of a day. It is defined by a time variable expressed in minutes and denoted $t \in [0, T]$, where $t = 0$ and $t = T$ are respectively the opening and closing periods. This time variable is discrete, and each period is spaced by a 5-minute interval. Operationally, the time variable corresponds to the Greenwich time zone, and each period is associated to a possible departure slot (CTOT). It is assumed that the market opens at 0130z (1:30am in zoulou time[14]), and closes at 0125z 24 hours later.[15] Thus, no swap is feasible between 0125z and 0130z, as the market clears out.

Let Ω_t be a finite, discrete and static set representing flights meeting the first set of conditions exposed in Section 3.1 at a given time t[16]. Let $F_t = \{F_{1,t}, \ldots, F_{k,t}, \ldots, F_{n,t}\}$ be a subset of Ω_t containing flights meeting the second set of operational conditions defined in Section 3.1[17] Once joining the inter-carrier market, flights can enter two possible subsets of F_t, $S_t \subset F_t$, or $B_t \subset F_t$, satisfying $S_t \cap B_t = \emptyset$. Flights in S_t will systematically be dispatched into three possible subsets: $R_t^S, \subset S_t$, $G_t^S \subset S_t$, and $V_t^S \subset S_t$, satisfying $R_t^S \cap G_t^S \cap V_t^S = \emptyset$; where R^S contains the set of flights having a re-routing option $(+,+)$[18], G_t^S the set of completely delayed flights $(+, 0)$, and V_t^S the set of remaining flights with a high delay tolerance $(-,+)$. Analogously, flights in B_t will systematically be dispatched into three possible subsets: $R_t^B \subset B_t$, $G_t^B \subset B_t$, and $V_t^B \subset B_t$, satisfying $R_t^B \cap G_t^B \cap V_t^B = \emptyset$; where R_t^B contains candidates for buying a slot from a seller flight in S_t^B, G_t^B flights candidates for buying a slot from a seller in G_t^S, and V_t^B the remaining flights with a low delay tolerance. Allocation between V_t^S and V_t^B for remaining flights will be determined by the median of the distribution of the cost of a minute of delay for these flights. Figure 5 proposes a schematic representation of the agents and their circulation in the market. On this representation, economy 1 is associated to the Re-routing economy, while economies 2 and 3 respectively correspond to the Win-Win and Lose-Win economies.

For each period t, these three economies constituting the market can formally be defined as $E_t^R = (R_t^S, R_t^B)$, $E_t^G = (G_t^S, G_t^B)$ and $E_t^V = (V_t^S, V_t^B)$. Hence, the market M at time t is characterized by $M_t = (E_t^R, E_t^G, E_t^V)$. *Market* is the name given to the collection of matching equilibria in the three economies for every period throughout the day.

Flights are always defined by their CTOT, i.e., their ATFM departure time. The CTOT is a specific value of t, and is denoted t_i for flight i and t_j for flight j. In addition, flights are always associated to a *most* and *least* preferred slots, respectively denoted t^+ and t^-. The operational slots attached to t_i, t_j, t_i^+, t_j^+, t_i^-, and t_j^- depend on the economy the flight belongs to, as well as its role as either a buyer or seller within that economy. A comprehensive explanation of the operational slot associated with each variable will be provided at the start of section dedicated to the Win-Win economy.

[14]Zoulou time is the time-frame used in aeronautics.

[15]These times have been chosen so that the number of flights on the market at the time of its closure is likely to be the lowest of the day.

[16]As a reminder, a flight meets the first set of operational conditions if it is: IFR, commercial scheduled, medium-haul, connecting an Origin-Destination within the ECAC area, and transporting passengers.

[17]As a reminder, a flight meets the second set of operational conditions if it is ATFM delayed and swappable according to the defintion of swappiability from Section 2.2.

[18]This notation reminds the reader the financial and operational gains from a swap depending on the situation of the flight on the market (see Figure 4).

Moreover, flights enter in F_t at time $t^{\text{infra}} \in [0, t^+[$ to first experiment infra-carrier swaps if possible, and leave the infra-carrier market at time $t^{\text{infra}} + \Delta t$ (with $\Delta_t \in [0, t_i^+ - t_i^{\text{infra}}[$). Flights that do not observe any infra-swap opportunity after Δt periods of time can decide to join the inter-carrier market. The decision to leave F_t is left at the discretion of each flight. Flights enter the inter-carrier market at time $t_i^{\text{entry}} \in [t_i^{\text{infra}} + \Delta t, t_i^+[$ and leave it at time $t_i^{\text{exit}} \in [t_i^{\text{entry}}, t_i^-[$. A flight i or j is said to be *critical* when $t = t_i^{\text{exit}}$ or $t = t_j^{\text{exit}}$ respectively. Operationally, t_i^{entry} can realistically estimated at around 1 hour before departure time (see Section 3.1). Besides, t^{exit} will realistically be fixed at 5 minutes before flights CTOT, i.e., the last period before departure time.

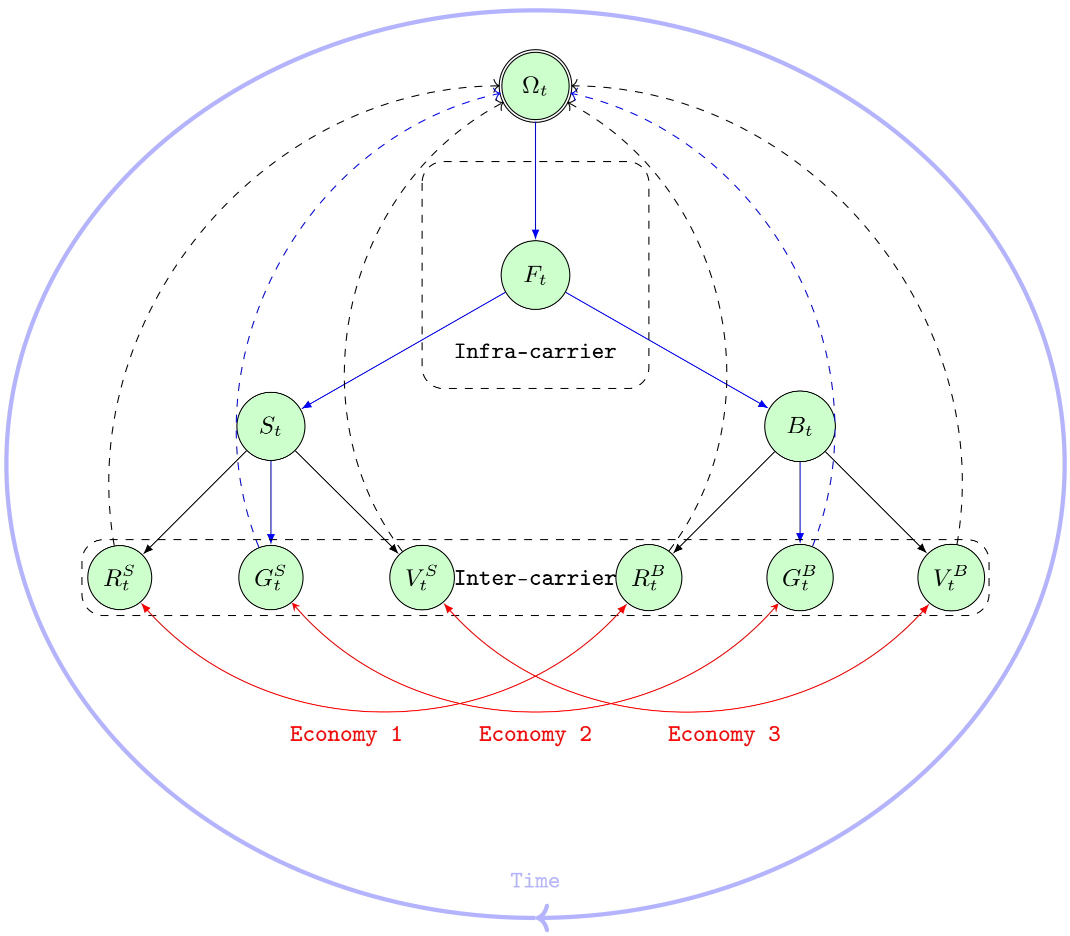

Figure 5: Tree representation of flights circulation in the market.

3.3 The Win-Win Economy

In this section, a model is proposed for matching flights for the Win-Win economy. The Win-Win economy is comprised of a static two-sided matching platform, with flights having fixed preferences for slots. This model therefore presents a matching program for one period, which can then be repeated for all periods by adjusting the parameters. This setting aligns with the classical setup of (Gale & Shapley, 1962). The distinction, in this context, is that agents' utility is not defined by monetary transfers.

3.3.1 Agents & Economy-wide Variables

This section focuses on the matching of flights belonging to the sets G_t^S and G_t^B, respectively corresponding to the buyers and sellers for the Win-Win economy. With a slight abuse of language, the semantic choice of *sellers* and *buyers* is retained even for the Win-Win economy, despite the absence of monetary transfers. Also, for the sake of clarity in notations, it is established that $G_t^S = I$ and $G_t^B = J$ for the remainder of the paper. I and J are thus observations of G_t^S and G_t^B at a particular time t. It is also stated that $|I| = P$ and $|J| = M$, with I and J referring to the *pool* of flights for each side of the Win-Win economy, while P and M refer to the *pool size* of each side. No restriction is placed on the pool sizes. Each buyer and seller is attributed a slot to swap with a flight of the opposite set, and each flight is defined by its CTOT. Finally, slots are indivisible.

As presented in Section 3.2, flights are defined by their CTOT, t_i and t_j. They also possess a most and least preferred slot, t^+ and t^-, and they enter the market at t^{enter} for an exit at t^{exit}. In the context of the Win-Win economy, a flight's i most preferred slot, $t^+i \in [0, T]$, corresponds to its SOBT, while its least preferred slot, $t^-i \in [0, T]$ (with $t^-i > t^+i$), corresponds to its CTOT. Conversely, a flight's j most preferred slot, t^+j, corresponds to its EOBT, while its least preferred slot, t^-j, corresponds to an *Upper Bound Time* (UBT). The UBT is the time after which flight j is not willing to take-off anymore. It is not a formal operational term, but a new concept that is introduced for the sake of modeling the Win-Win case. Indeed, the UBT allows the utility function of flight j not to converge asymptotically to 0, but to actually reach 0 at some point. Concretely, the UBT could either be specified by the AOs and privately communicated through a blockchain mechanism, or it could be uniformly standardized to 30 minutes after the EOBT for all flights[19].

3.3.2 Flights' Utility

The utility function for flight i is defined for every positive real number θ, representing a possible departure slot for this flight. It follows:

$$\forall t \in \mathbb{R}_+ \quad \forall \theta \in \mathbb{R}_+ \qquad u_i(t, \theta) := \begin{cases} \omega_i(t_i^{\text{exit}} - \theta) & \text{if } t_i^+ \leq t \leq t_i^{\text{exit}} \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

where ω_i is a coefficient capturing the cost of a minute of delay for flight i, and θ is a possible departure slot for flight i. θ can be interpreted as a type. It takes the value $\theta = t_j$ for every possible pair $\{(i, j) : i \in I, j \in J\}$, or $I \times J$ for short. The two cases ensure that the utility function is only defined for times t when flight i is both able and better-off to take-off.

[19]The choice of 30 minutes is motivated by the fact that delays are considered to have severe financial cost implications after 30 minutes(Eurocontrol, 2019)

If $t_i^+ \leq t \leq t_i^{\text{exit}}$, the utility function is given by the product of the flight specific weight, ω_i, with the difference between the exit time t_i^{exit} and the slot swap proposition from j, θ; and equals 0 otherwise. Hence, flights' utility function is linearly decreasing on its definition interval. It is equal to its maximum when $t = t^+$, and 0 when $t = t_i^{\text{exit}}$. All what is exposed in here reciprocally applies for flight j. A representation of each flight utility function is proposed on Figures 6 and 7[20].

In the standard (Gale & Shapley, 1962) utility is defined in monetary terms. However, in this specific set-up, utility is associated to a measure of satisfaction or preference based on the timing of the offered slot relative to the preferred one. This adaptation is reasonable as each flight's utility function creates a preference ranking over different slot times, as required by the standard model. The preference ranking can be understood as follows: for any two potential slot times, the one that results in a higher utility is preferred over the one that results in a lower utility. This implicit preference ranking is crucial to the operation of the model, as it determines the "best" slot for each flight, given the available options.

Moreover, the utility function proposed makes slots preferences *complete*, meaning that flights will always be able to make a decision as slots are comparable. Second, it guarantees that preferences are *transitive*. Weak preferences are sufficient to ensure transitivity, however, in this context, the linearly decreasing structure of flights' utility function makes that flights have *strict preferences*. This implies that flights will *never be indifferent* between any two slot offers. The two aforementioned properties make flights *rational agents*.

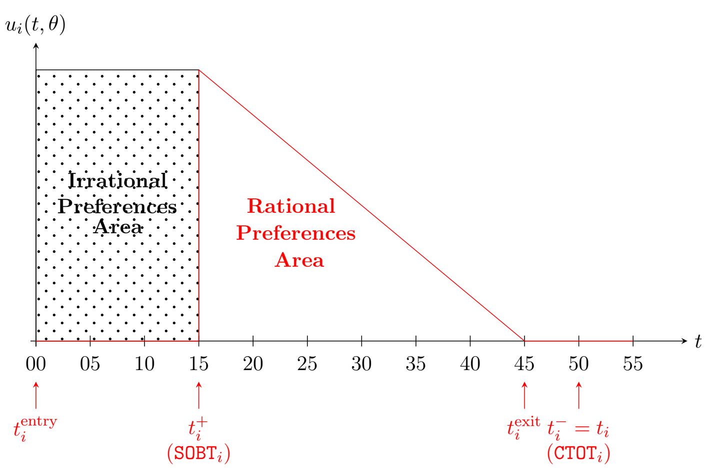

Figure 6: Buyer flight i utility

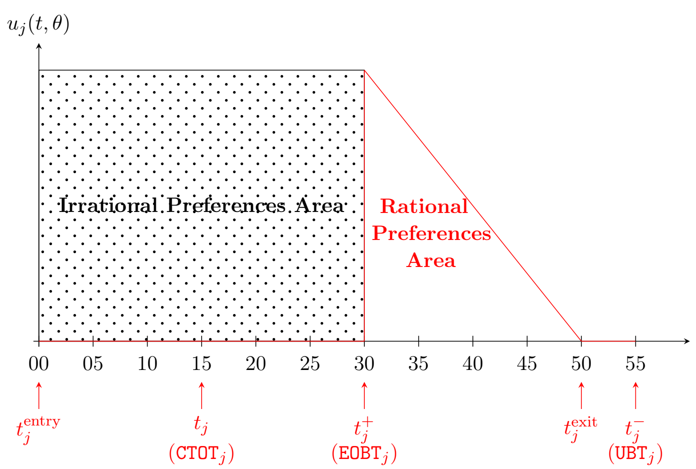

Figure 7: Seller flight j utility

[20]In Figures 6 and 7, it is assumed that flights' $t_i^{\text{entry}} = t_j^{\text{entry}} = 0$ for the sake of simplification.

3.3.3 Linear & Static Matching Program

A *swap* is defined as a match between a buyer flight i and a seller flight j. It can be represented by an $P \times M$ matrix x with $x_{i,j} = 1$ when the swap occurs, and $x_{i,j} = 0$ when it does not. The set of possible swaps (i, j) is denoted X, and thus $I \times J$: X is the correspondence of paired flights. More precisely, when a pair (i, j) belongs to X, it denotes that flight i from set I is associated with flight j from set J. For any matched pair (i, j), it is denoted as $X(i) = j$ and $X(j) = i$. Meanwhile, for any unmatched agent $a \in I \cup J$, it is written as $X(a) = a$. As described earlier, flights that have been matched receive utilities. Recall that, for flights in I, these utilities are denoted by $u_i : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$, and for flights in J, they are represented by $u_j : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$. In particular, when $(i, j) \in X$, flight i is allocated a utility of $u_i(t, \theta)$, and flight j is accorded a utility of $u_j(t, \theta)$. Any flights that remain unmatched are assigned a utility of zero, as specified by the utility function.

It therefore follows from (Shapley & Shubik, 1971) that the primal and dual linear programs ensuring *max-weight* (i.e., maximizing the social, or aggregated, utlity) and *stable* swaps between flights belonging to I and flights belonging to J can be written as:

$$\mathbf{LP}(I, J, u_i, u_j) := \max_{X \in \mathbb{R}^{|I| \times |J|}} \sum_{(i,j) \in I \times J} x_{i,j} \left[u_i(t, t_j) + u_j(t, t_i) \right] \tag{2}$$

$$\texttt{s.t.} \quad \sum_{j \in J} x_{i,j} \leq 1 \quad \forall i \in I \tag{3}$$

$$\sum_{i \in I} x_{i,j} \leq 1 \quad \forall j \in J \tag{4}$$

$$x_{i,j} \geq 0, \qquad \forall (i, j) \in I \times J \tag{5}$$

This linear program is designed to find the assignment of flights in sets I and J that maximizes the total utility, represented by the variables u_i and u_j. Constraints (3) and (4) ensure that each flight in I and J can participate in at most one swap, and that assignments are non-negative (5). It is worth noting that the program takes into account all flights in I and J. However, the optimal solutions will naturally exclude swaps bringing a null utility to one of the two flight because it seeks to maximize total utility. In addition, the heuristic algorithm solving this program will systematically exclude the solutions that bring a utility of zero to one of the two flights.

4 Simulation Analysis

This section proposes an algorithm to solve the linear and static matching problem and discusses results based on simulated data. It begins with a justification of the choices made for the data simulation, followed by an introduction of the heuristic solution algorithm. Designed to ensure both *welfare-maximizing* swaps, the algorithm is thoroughly explained. The average gains obtained from both the buyers' and the sellers' perspectives for a one period participation in the market are finally presented.

4.1 Data Simulation

The data simulation process is curcial for the future results. Its underlying assumptions must therefore rely on an accurate understanding and description of the operational situation. Each data simulation choices will thus be justified. In this section, variables are labelled by their operational names (e.g., $\texttt{SOBT}_i$ ...) rather than their modelling names (e.g., t_i^+ ...). The algorithm will need two different set of variables for flight i and j, as well as a list of parameters. Before motivating each data generation choice, it is important to note that throughout this simulation analysis it is assumed that all flights are already on the market. It is recommended to refer to Figures 6 and 7 if needed.

- **<u>Parameters:</u>**
  - `numb_slots` = 24: This parameter sets the number of slots to 24. Considering that slots are spaced by 5-minute intervals, the simulation thus focuses on a two-hour window during the day. This decision aligns with the high-frequency nature of the market. As elaborated in Section 3.1, flights are expected to enter the market one hour before their CTOT, with the primary objective of addressing last-minute adjustments and managing average delays. Therefore, concentrating our simulation within a two-hour window appropriately corresponds to the operational time horizon of the market.
  - `peak` = 0.5 × `num_slots`: This parameter is used to center the normal distributions at the midpoint of the total number of slots. Motivation for this choice will be provided when justyfing the variables $\texttt{CTOT}_i$ and $\texttt{CTOT}_j$.
  - `max_wait` = 6: This parameter determines the UBT for sellers. As explained in Section 3.3.1, $\texttt{UBT}_j$ is standardized to 30 minutes after $\texttt{CTOT}_j$.
  - `aveg_gddelay` = 3: This parameter defines the EOBT for sellers. Set at 15 minutes, it aligns with the average rounded ground delay according to (Eurocontrol, 2019).
- **<u>Variables for i:</u>**
  - $P = 100$: In the simulation, the cardinality of the set of buyers' flights I is established as 100, representing a realistically large number. This figure is chosen considering the assumption that a flight may return to the market up to three times. Furthermore, the nature of the algorithm can augment the quantity of flights in each period, as flights can remain on the market across multiple periods before acquiring a swap offer. Given the algorithm's operation every five minutes, it executes 288 times per day. With reference to the statistics in Section 2.3, a total of

15,214 flights[21] were delayed in Europe in 2019. If this distribution of delayed flights follows a uniform distribution across all possible departure slots, assuming that all delayed flights engage in the inter-carrier market, around 53 flights[22] would participate in the market at each period. However, a uniform distribution is unlikely in this scenario, so it has been chosen to double this average to emulate a period of high congestion.

- $\texttt{CTOT}_i \sim B(\texttt{num\_slots}, \texttt{peak/num\_slots}) + 2, \forall i \in I$: As discussed, the density of CTOTs' distribution over the 288 potential departure slots is unlikely to be uniform. Flights during peak hours will be more numerous than those at daybreak or late at night. Hence, a normal distribution appears to be a more fitting choice for CTOTs'. However, considering the discrete time setup, a binomial distribution centered at midday is adopted. It aligns well with the broader daily perspective. This distribution is sustained even in the simulation's two-hour window, which is consistent with the presumed distribution of departure peaks within the day. Lastly, two periods are added to every generated CTOT to ensure that each flight i has a valid market exit opportunity. This adjustment is crucial because, in the model, flight i exits the market one period prior to its actual CTOT.

- $\texttt{SOBT}_i \sim \mathcal{U}(0, \texttt{CTOT}_i - 2), \forall i \in \texttt{CTOT}_i$: The SOBT for each flight i is generated using a discrete uniform distribution. The value for $\texttt{SOBT}_i$ is chosen from all periods starting from the first market period ($t = 0$) up until two periods prior to the flight's CTOT (i.e., one period before its exit period). Given that there's no particular reason to assume a specific distribution for $\texttt{SOBT}_i$, it is reasonable to represent it as being uniformly distributed over all slots within its possible range.

- $\texttt{exit}_i = \texttt{CTOT}_i - 1$: Flight i exit period occurs one period before its CTOT (see Figure 6).

- $\omega_i \sim \mathscr{N}(40, 4^2), \forall i \in I$: The cost per minute of delay is assumed to follow a normal distribution with a mean of 40 and a standard deviation of 4. This assumption is justified by the complex nature of flight delay costs, which can be influenced by factors such as the number of passengers in transit and airport fees at the current location (A. J. Cook & Tanner, 2011). Hence, the expected heterogeneity across flights is likely to align with a normal distribution. The mean cost of 40 euros per minute is informed by Cook and Tanner's (2011) estimation of delay costs below 30 minutes. This standard is adopted given the high-frequency nature of the market model, primarly designed to tackle short to medium delays.

- **Variables for j:**

  - $M = 100$: Symmetrically, the cardinality of the sellers' flight set J is also set at 100. Introducing an asymmetry between the number of buyers and sellers reflects realistic market conditions, as there are likely to be more buyer flights (i.e., those only delayed by ATFM) than seller flights (i.e., flights delayed completely). This will be done in Section 4.3 and allows to provide a more nuanced and accurate representation of the real-world market.

[21] In 2019, there were 30,427 flights IFR flights within the ECAC area (Eurocontrol, 2019), and 50% of them were delayed by at least 5 minutes (Eurocontrol, 2022). Hence, the number of delayed flights for this year is 15,214.

[22] $15,214/288 \approx 52.8$,

– $\texttt{CTOT}_j \sim B(\texttt{num\_slots}, \texttt{peak/num\_slots}), \forall j \in J$: The distribution of CTOTs for the sellers' flights mirrors that of the buyers', with one key distinction. Unlike the buyers' side, the exit time for a seller's flight j is independent of its CTOT, as depicted in Figure 7. Consequently, there is no need to add two periods to establish the exit time for the sellers, as was necessary for the buyers.

– $\texttt{EOBT}_j \sim \texttt{CTOT}_j + \texttt{Poisson}(\texttt{avg\_delay}), \forall j \in J$: The EOBT for each flight j is modeled using a Poisson process originating from the respective CTOT of that flight. The average of this process aligns with Eurocontrol (2019) estimates of the average ground delay. The choice of a Poisson process is motivated by the fact that EOBTs are likely to cluster around this average.

– $\texttt{UBT}_j$: As exposed in Section 3.3.1, the UBT is uniformly fixed to 30 minutes after the flights' j EOBT.

– $\texttt{exit}_j = \texttt{UBT}_j - 1$: As explained in Section 3.3.1, flights' j exit time corresponds the period preceding its UBT.

– $\omega_j \sim \mathcal{N}(40, 4^2), \forall j \in J$: This variable adopts the same formulation and rationale as its counterpart on the buyers' side.

4.2 Solution Algorithm

This section presents the heuristic solution algorithm that finds *welfare maximizing* swaps between flights based on their mutual preferences. The pseudo-code of this algorithm can be written as follows:

Algorithm 1 Win-Win Swaps Algorithm

Inputs. $\texttt{SOBT}_i \in \mathbb{R}_+$: flight i's most preferred departure slot, for each buyer i in I
$\texttt{EOBT}_j \in \mathbb{R}_+$: flight j's most preferred departure slot, for each seller j in J
$\texttt{CTOT}_k \in \mathbb{R}_+$: flight k's ATFM departure slot, for each k in $I \cup J$
$\texttt{exit}_k \in \mathbb{R}_+$: slot when flight k exits the swapping market, for each k in $I \cup J$
$\omega_k \in \mathbb{R}_+$: cost of a-minute delay for flight k, for each k in $I \cup J$

Outputs. $L \in \mathfrak{P}(I \times J)$: a list of matches between buyers in I and sellers in J
$d_i \in \mathbb{R}_+$: flight i's distance in minutes between CTOT and SOBT, for each buyer in I
$\delta_j \in \mathbb{R}$: flight j's distance in minutes between CTOT and EOBT, for each seller in J

procedure ComputeMatch($(\omega_i)_{i\in I}$,$(\texttt{CTOT}_i)_{i\in I}$,$(\texttt{SOBT}_i)_{i\in I}$,$(\omega_j)_{j\in J}$, $(\texttt{CTOT}_j)_{j\in J}$,$(\texttt{EOBT}_j)_{j\in J}$)
  Initialize $L \leftarrow \varnothing$: list of Pareto-improving matches
  Initialize $d_i \leftarrow \texttt{CTOT}_i - \texttt{SOBT}_i$, $\forall i \in I$: flight i's preswap distance between CTOT and SOBT
  Initialize $\delta_j \leftarrow \texttt{CTOT}_j - \texttt{EOBT}_j$, $\forall j \in I$: flight j's preswap distance between CTOT and EOBT
  Initialize $b_{i,j} \leftarrow 0$, $\forall (i,j) \in I \times J$: utility of buyer i if it swaps with seller j
  Initialize $s_{i,j} \leftarrow 0$, $\forall (i,j) \in I \times J$: utility of seller j if it swaps with buyer i
  for $(i,j) \in I \times J$ **do**
    Compute utility of i if it swaps with j: $b_{i,j} \leftarrow \max\{\omega_i(\texttt{exit}_i - \texttt{CTOT}_j), 0\}$
    Compute utility of j if it swaps with i: $s_{i,j} \leftarrow \max\{\omega_j(\texttt{exit}_j - \texttt{CTOT}_i), 0\}$
    Compute sum of both utilities: $U_{i,j} \leftarrow b_{i,j} + s_{i,j}$
  end for
  for k in $\{1, \ldots, \min(\text{Card}(I), \text{Card}(J))\}$ **do**
    Find (i^*, j^*) that maximizes sum of utilities: $(i^*, j^*) \leftarrow \arg\max\{U_{i,j};\ (i,j) \in I \times J\}$
    Add pair to matching pairs: $L \leftarrow L \cup \{(i^*, j^*)\}$
    Add gains and delays to corresponding lists: $d_{i^*} \leftarrow \texttt{CTOT}_{j^*} - \texttt{SOBT}_{i^*}$; $\delta_{j^*} \leftarrow \texttt{CTOT}_{i^*} - \texttt{EOBT}_{j^*}$
    Remove pair (i^*, j^*) from swapping market: $I \leftarrow I \setminus \{i^*\}$; $J \leftarrow J \setminus \{j^*\}$
  end for
  return L, $(d_i)_{i\in I}$, $(\delta_j)_{j\in J}$
end procedure

First, the algorithm starts by initializing several lists to keep track of matched pairs and various measures to assess the performance of the algorithm. Next, the algorithm iterates over each possible pair of buyers and sellers. For each pair, it checks if the CTOT of the seller is greater than or equal to the SOBT of the buyer. This condition ensures that the swapped CTOTs will fall into the area rational preferences for both the buyer and the seller of each pair (cf. Figures 6 and 7). If this condition is true, it calculates the maximum possible utility of the buyer. If the condition is false, the utility is set to zero. The algorithm performs a similar process again, this time checking if the CTOT of the buyer is greater than or equal to the EOBT of the seller. If so, it calculates the utility for the seller, and if not, it sets the utility to zero. The heuristic nature of the algorithm is based on the fact that it does not consider all possible pair combinations to determine the optimal swaps. Specifically, the matching algorithm excludes all pairs that involve a flight with zero utility. The algorithm then enters a loop where it will make as many buyer-seller matches as there are buyers or sellers, whichever is smaller. In each iteration, it finds the pair that maximizes the sum of utilities, adds this pair to the list of matched pairs, records the gains and delays in the appropriate lists, and removes the matched flights from the Win-Win swapping economy. Finally, the algorithm returns all the lists, providing the matched pairs, and measures of performance that will be discussed in the next section.

4.3 Results

This section presents the results of the algorithm for both the buyers' and sellers' sides, based on the simulated data described in Section 4.1. To ensure reliable and robust results, 100 simulations of the solution algorithm are conducted for each result.

Buyers' benefits for a one period market participation

($n = 200$: $P = 100$, $M = 100$)

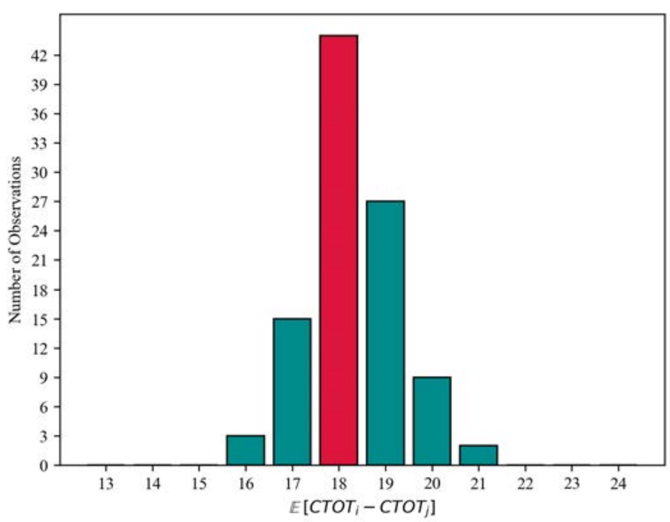

Figure 8: Average gains from slot swapping

Figure 8 presents the average time gains, in minutes, that buyers can realize by participating in the inter-carrier slot swapping market for a single period. Based on the simulated data from over 100 simulations, it is observed that matched buyers can expect an average reduction of 18 minutes in their ATFM delay.

Utilizing the estimates for the average cost per minute of delay as provided by Cook & Tanner (2011), specifically 40 euros for delays less than 30 minutes and 147 euros for delays exceeding 30 minutes, this represents average savings ranging from 720 euros to 2,646 euros for each buyer flight receiving a swap offer in the win-win economy. However, these significant gains for buyers must be interpreted with caution, as it is based on the aforementioned data distribution assumptions and results from simulations of 200 flights, evenly split between two sets of agents. Nonetheless, as detailed in Section 4.1, the assumptions used for the simulated data appear realistic, and this result is consistent across different sample sizes. Indeed, the average gain for buyers consistently hovers around 20 minutes in 100 simulations for different values of n (i.e., $n = \{10, 50, 100, 200, 500\}$). This result also holds for different levels of asymmetry between the numbers of buyers and sellers (i.e., $\frac{M}{P} = \{\frac{5}{100}, \frac{20}{100}, \frac{40}{100}, \frac{60}{100}, \frac{80}{100}, \frac{100}{100}\}$). This last element is important and likely reflective of reality, as it can be expected to have a larger number of ATFM delayed flights than completely delayed flights.

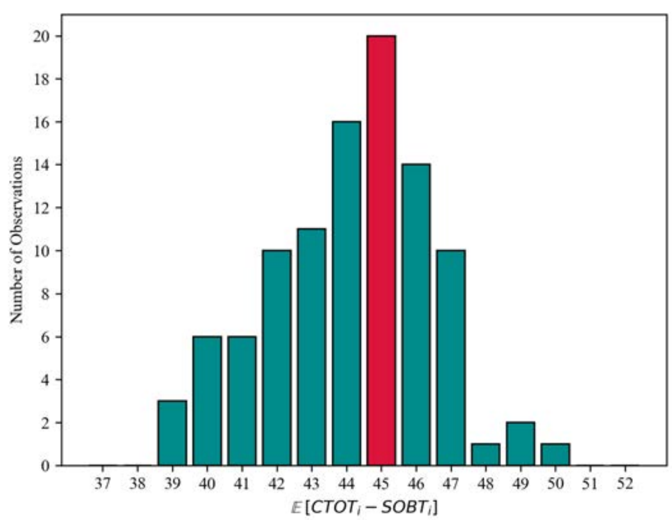

Figure 9: Average pre-swap distance to SOBT

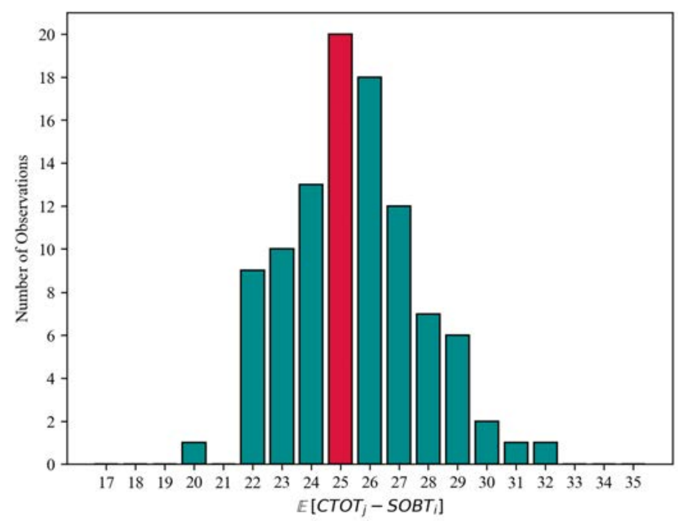

Figure 10: Average post-swap distance to SOBT

Figure 9 illustrates that the average pre-swap distance to the SOBT for buyers is approximately 45 minutes across each of the 100 simulations. This finding underscores the significance of swap gains for buyers, as elucidated by Figure 8. Indeed, a reduction of 18 minutes in delay constitutes a substantial benefit for buyers, as it would enable them to eliminate approximately 40% of their original ATFM delay. Furthermore, this indicates that within the previously produced savings range, the average savings achieved by flights based on these simulations would be far closer to 2,646 euros than to 720 euros. This is due to the fact that the average delay prior to slot swapping exceeds 30 minutes in these simulations. Conversely, Figure 10 depicts the residual distance in minutes that buyers need to traverse in order to attain their maximum satisfaction. Post-swap, the average distance to the SOBT is observed to be roughly 25 minutes. Given that flights can experiment up to three swaps, it is deduced that, on average, a flight that has not experiment any prior swaps before entering the inter-carrier market could fully achieve its goal after three swaps. Furthermore, it would be proximate to its targeted slot after merely two swaps.

Sellers' benefits for a one period market participation

$(n = 200\text{: } P = 100,\ M = 100)$

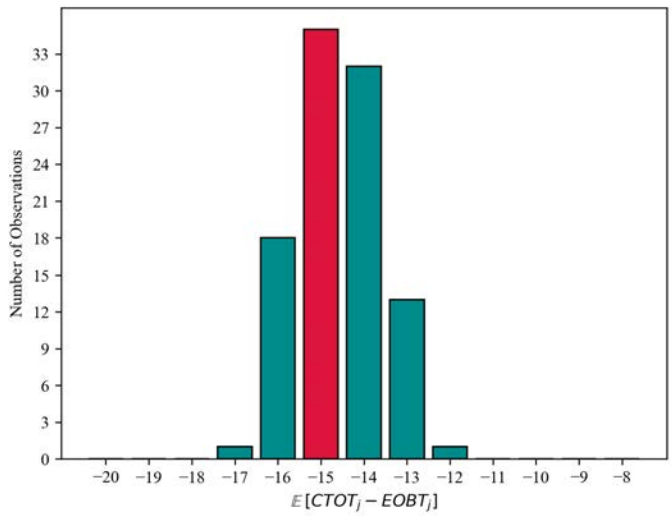

Figure 11: Average pre-swap distance to EOBT

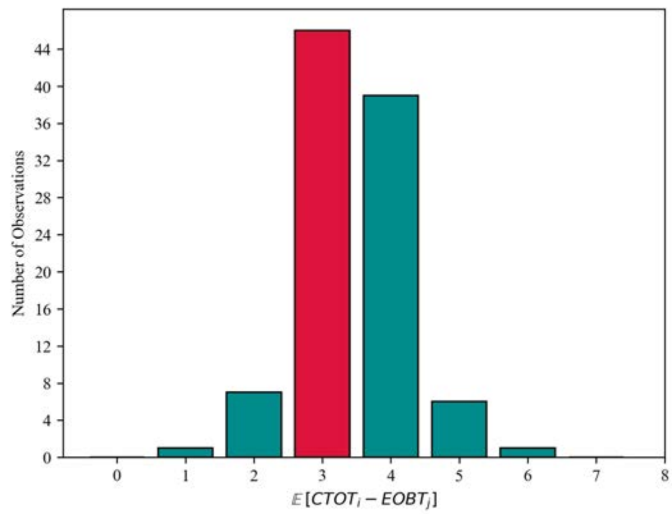

Figure 12: Average post-swap distance to EOBT

As explained in Section 3.1, unlike buyers who aim to depart as close as possible to their SOBT through slot swapping, the primary objective of sellers in the Win-Win economy is not to minimize time. Instead, they aim to maximize time, thus aligning a slot as closely as possible to their EOBT. Under these circumstances, to evaluate the performance of the algorithm in addressing sellers' needs, it is appropriate to compare the average pre-swap and post-swap distances from the sellers' EOBT. Figure 11 demonstrates that, on average, sellers are about 15 minutes away from their EOBT prior to a swap. The negative sign before the numbers on the x-axis also indicates that sellers are unable to adhere to their ATFM slot due to cumulative ground delay. Conversely, Figure 12 suggests that the algorithm also achieves excellent outcomes for sellers. Indeed, after just a single slot exchange on the inter-carrier market, the average distance of sellers from their preferred slot is only 3 minutes. Similar to buyers, these results are stable for different sample sizes and different seller-to-buyer ratios. However, a corollary of this positive outcome is that it could potentially lead sellers to exploit their swap opportunities less than buyers, thereby augmenting the inherent asymmetry that is expected between the two sets of flights. Therefore, it is crucial that the results are indeed consistent for high seller-to-buyer ratio asymmetries.

5 Conclusion

This paper posits and formalizes an innovative decentralized, high-frequency, flexible, and last-resort marketplace designed to facilitate the swapping of ATFM slots among competing airlines operating within the ECAC area. The market provides three potential economies to the AOs, depending on the operational situation a flight encounters at a given time: the Re-Routing economy, the Win-Win economy (or mutually beneficial) economy, and Lose-Win (or unilaterally beneficial) economy. Each of these economies forms a bi-sided market wherein the participants either barter or trade their slots. A heuristic solution algorithm, primarily aimed at promoting welfare maximization for flights, is subsequently presented to resolve the proposed linear programming model. The study further engages simulations based on this algorithm to gauge the performance of the inter-carrier market under the Win-Win economy framework. The results garnered are highly persuasive and robust across different tests of robustness.

Notably, slot buyers average an ATFM delay reduction of 18 minutes, representing a spare ranging from 720 euros to 2,646 euros per swap, whereas the sellers consistently secure a slot they can reliably maintain, averaging merely 3 minutes deviation from their preferred slot. These findings underscore the potential efficiency gains inherent in such an approach. This suggests that the inter-carrier slot swapping market could lead to significant efficiency improvements in air traffic management, which would be of significant benefit to airlines. This development unveils multiple avenues for future research.

Indeed, while presenting substantial insights, this paper's findings rely on simulated data only. This approach confirms the effectiveness of the algorithm in its quest to maximize the aggregate utility of the flights; however, the results must be corroborated in the future using real-world data obtained from the Network Operations Portal (NOP). The necessary data to be collected are outlined in appendix B. Furthermore, although the proposed solution is completely implementable, the linear model does not incorporate a time-dependent risk aversion factor for flights. This factor could prove beneficial in prioritizing flights that are imminently leaving the market at the expense of those with more time. This would further strengthen the high-frequency dimension of the market. To achieve this, a dynamic utility function for flights is proposed in appendix A. It could be employed by future research to construct a complex system. The most suitable technique for the development of such a system is reinforcement learning. By modeling the market in this manner, each period could be represented as a space containing information about the market's situation at instant t (e.g., number of buyers and sellers, the number of swaps already experimented for each flight). Subsequently, the reinforcement learning algorithm could learn to pair flights by maximizing the joint utility of every possible buyer-seller flight combination, in line with the solution proposed in this paper. This modeling approach could follow Jagadeesan et al. (2021) or Min et al. (2022).

Additionally, this paper opens the way for the modeling of Re-Routing and Lose-Win economies. On one hand, the rerouting economy can be modeled by employing auction theory, where the selling flight, possessing a rerouting opportunity, puts its flight slot up for auction. A dynamic component could also be added to this model. Parkes (2007) offers a comprehensive review of the dynamic mechanism design literature. In particular, the works of Budish, Cramton, and Shim (2015) on the problem of timing in high-frequency auctions, and those of Parkes and Singh (2003) on the generalization of the Vickrey-Clarke-Groves mechanism in a dynamic context, could serve as a solid foundation for further study. On the other hand, the lose-win economy could be envisaged as an extension of the win-win economy proposed in this model, this time incorporating a monetary compensation mechanism for the losers.

Furthermore, elements of game theory could be considered. It could indeed be interesting to demonstrate that the entirety of the economies guarantee participants in the exchange mechanisms to be incentive-compatible, budget balanced, efficient, and strategy-proof. Additionally, the decision-making strategy of each flight concerning which market to engage with to achieve its slot objective in the least amount of time could be explored using Mean-Field Game theory. Each flight's decision to participate in a particular economy as a buyer or a seller at a given period will influence the overall market structure (i.e., the number of players in each buyer/seller set of each economy). Consequently, this would impact the decisions of other flights. Mean-Field Game theory can provide insights into how these intricate interactions converge towards a Nash equilibrium in the market.

In sum, while the current research presents a step forward, future research building on this work could lead to further refinements and enhancements to the system, delivering even greater efficiency gains in the management of air traffic.

References

Air France. (2022). *Analyse atm air france.*

Akbarpour, M., Li, S., & Gharan, S. O. (2020). Thickness and information in dynamic matching markets. *Journal of Political Economy*, *128*(3), 783–815.

Ball, M., Donohue, G., & Hoffman, K. (2006). Auctions for the safe, efficient, and equitable allocation of airspace system resources. *Combinatorial auctions*, *1*.

Budish, E., Cramton, P., & Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, *130*(4), 1547–1621.

Burkard, R., Dell'Amico, M., & Martello, S. (2012). *Assignment problems: revised reprint.* SIAM.

Cavallo, R., Parkes, D. C., & Singh, S. (2009). Efficient mechanisms with dynamic populations and dynamic types. *Unpublished manuscript, Harvard University*.

Cliff, K. (2019). Reducing impact of delays using airspace user-driven flight prioritisation.

Condé Rocha Murça, M., & Hansman, R. J. (2019). *Data-driven modeling of air traffic flows for advanced air traffic management* (Unpublished doctoral dissertation).

Cook, A. (2007). *European air traffic management: principles, practice, and research.* Ashgate Publishing, Ltd.

Cook, A. J., & Tanner, G. (2011). European airline delay cost reference values.

Corblet, P. (n.d.). Returns to education and experience on the labor market.

Côté, J.-P. (2005). Intégration de la tarification et de l'allocation de la capacité en transport aérien: une approche bi-niveau à grande échelle.

Deschinkel, K. (2001). *Régulation du trafic aérien par optimisation dynamique des prix d'utilisation du réseau* (Unpublished doctoral dissertation). Toulouse, ENSAE.

Dupuy, A., & Galichon, A. (2014). Personality traits and the marriage market. *Journal of Political Economy*, *122*(6), 1271–1319.

Eurocontrol. (2019). *All causes of delay and concellations to air transport in europe.*

Eurocontrol. (2020). *Standard inputs for economic analyses.*

Eurocontrol. (2022). *All causes delays to air transport in europe.*

Eurocontrol. (2022a). *Atfm users manual.*

Eurocontrol. (2022b). *Eurocontrol aviation outlook 2050.*

Eurocontrol. (2022c). Eurocontrol market segment update.

Eurocontrol. (2022d). Slot swapping.

Eurocontrol. (2023). *Eurocontrol forecast update 2023-2029.*

European Commission. (2011). Commission decision of 7.7.2011.

Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, *69*(1), 9–15.

Galichon, A. (2018). *Optimal transport methods in economics.* Princeton University Press.

Hondet, G., Delgado, L., & Gurtner, G. (2018). Airline disruption management with aircraft swapping and reinforcement learning. In *Sid 2018, 8th sesar innovation days.*

IATA. (2022). *Quarterly air transport chartbook.*

Jagadeesan, M., Wei, A., Wang, Y., Jordan, M., & Steinhardt, J. (2021). Learning equilibria in matching markets from bandit feedback. *Advances in Neural Information Processing Systems*, *34*, 3323–3335.

Min, Y., Wang, T., Xu, R., Wang, Z., Jordan, M., & Yang, Z. (2022). Learn to match with no regret: Reinforcement learning in markov matching markets. *Advances in Neural Information Processing Systems*, *35*, 19956–19970.

Odoni, A. R. (1987). The flow management problem in air traffic control. In *Flow control of congested networks* (pp. 269–288).

Parkes, D. C., & Duong, Q. (2007). An ironing-based approach to adaptive online mechanism design in single-valued domains. In *Aaai* (Vol. 7, pp. 94–101).

Parkes, D. C., & Singh, S. (2003). An mdp-based approach to online mechanism design. *Advances in neural information processing systems*, *16*.

Pilon, N., Guichard, L., Bazso, Z., Murgese, G., & Carré, M. (2021). User-driven prioritisation process (udpp) from advanced experimental to pre-operational validation environment. *Journal of Air Transport Management*, *97*, 102124.

Raffarin, M. (2002). *Le contrôle aérien en france: congestion et mécanismes de prix* (Unpublished doctoral dissertation). Thèse de doctorat en sciences économiques, Université Paris 2, http://leea . . . .

Ranieri, A. (2010). Combinatorial exchange models for a user-driven air traffic flow management in europe.

Rassenti, S. J., Smith, V. L., & Bulfin, R. L. (1982). A combinatorial auction mechanism for airport time slot allocation. *The Bell Journal of Economics*, 402–417.

Roth, A. E., & Sotomayor, M. (1992). Two-sided matching. *Handbook of game theory with economic applications*, *1*, 485–541.

Ruiz, S., Guichard, L., & Pilon, N. (2017). Optimal delay allocation under high flexibility conditions during demand-capacity imbalance. In *Proc. 7th sesar innov. days* (p. 28).

Schütz, C. G., Gringinger, E., Pilon, N., & Lorünser, T. (2021). A privacy-preserving marketplace for air traffic flow management slot configuration. In *2021 ieee/aiaa 40th digital avionics systems conference (dasc)* (pp. 1–9).

Schütz, C. G., Ruiz, S., Gringinger, E., Fabianek, C., & Lorünser, T. (2022). An auction-based mechanism for a privacy-preserving marketplace for atfm slots. In *33rd congress of the international council of the aironautical sciences* (pp. 1–14).

SESAR. (2014). Air traffic flow management slot swapping.

Shapley, L. S., & Shubik, M. (1971). The assignment game i: The core. *International Journal of game theory*, *1*(1), 111–130.

Vascik, P. D., Balakrishnan, H., & Hansman, R. J. (2018). Assessment of air traffic control for urban air mobility and unmanned systems.

Vossen, T. W., & Ball, M. O. (2006). Slot trading opportunities in collaborative ground delay programs. *Transportation Science*, *40*(1), 29–43.

Vranas, P. B., Bertsimas, D., & Odoni, A. R. (1994a). Dynamic ground-holding policies for a network of airports. *Transportation Science*, *28*(4), 275–291.

Vranas, P. B., Bertsimas, D. J., & Odoni, A. R. (1994b). The multi-airport ground-holding problem in air traffic control. *Operations Research*, *42*(2), 249–261.

A Dynamic Flights' Utility

In this section, a dynamic utility function is modelled. Each flight aims now to optimize its specific, time-sensitive utility. The evolution of flight i on the market is expressed as a function of $t \in [0, T]$. The utility function of flight i, from the buyers' side, and the one of flight j, from the sellers' side, follow the same structure – keeping in mind that t_i^- and t_j^- do not correspond to the same operational concepts (see Section 3.3.1). The modeling of this utility function is made in a continuous time set-up, which will then be discretized to fit with the model proposed.
The utility function for flight i is defined by the difference between its slot valuation function and an exogenous premium function. It is expressed as follows:

$$\forall t \in \mathbb{R}_+ \quad \forall \theta \in \mathbb{R}_+ \quad \forall K \in \mathbb{N} \qquad u_i(t, \theta, K) := v_i(t, \theta) + s_i(t, K)$$

On the one hand, the valuation function $v_i : (t, \theta) \mapsto v_i(t, \theta)$ will describe how flight i prefers different departure slots associated to time θ, based on the current time t and specific time restrictions. On the other hand, the exogenous premium function $s_i : (t, K) \mapsto s_i(t, K)$ will introduce a random utility premium due to the number K of exogenous events that occurred from t_i^{entry} to current time t: these events significantly increase the cost of delay, and thus require the impacted flight to diligently avoid its ATFM slot. Instances of such events include expiration of crew working time, fine for taking-off after curfew etc.
As defined in Section 3.2, flight i enters the market at time t_i^{entry}, normalized to 0, and leave the market at time $t_i^{\text{exit}} < t_i^-$. Recall that current time is denoted t. The utility function for flight i is defined for every positive real number θ, representing a possible departure slot for flight i. It follows:

$$\forall t \in \mathbb{R}_+ \quad \forall \theta \in \mathbb{R}_+ \qquad v_i(t, \theta) := \begin{cases} \mathbb{1}_{\left[\max(t, t_i^+), t_i^-\right]}(\theta) \times f_i(t, \theta) & \text{if } 0 \leq t \leq t_i^{\text{exit}} \\ 0 & \text{otherwise} \end{cases}$$

where f_i is a sigmoid function. Indeed, $v_i(t, \theta)$ is defined in terms of the current time t, an available departure slot θ, as well as the entry and exit times of flight i. It presents two cases and an indicator function which restrict the domain of the utility in terms of both time and departure slots. The two cases ensure that the utility function is only defined for times t between the market entry time t_i^{entry} and the market exit time t_i^{exit}. If $0 \leq t \leq t_i^{\text{exit}}$, the utility function is given by the product of the indicator function and the sigmoid function, and equals 0 otherwise. The indicator function equals 1 if θ is within the interval $\left[\max(t, t_i^+), t_i^-\right]$, and 0 otherwise. It represents the feasible departure slots for flight i at the current time t by restricting the preferred departure slots for flight i to the interval between the latest possible departure time t_i^- and the maximum of the current time t and the earliest possible departure time t_i^+. This implies that the utility is nonzero only when the departure slot θ is within the feasible range and its value is determined by the sigmoid function. The sigmoid function, $f_i(t, \theta)$, is a smooth, S-shaped curve capturing the relationship between the current time and the preferred departure slot. This sigmoid function therefore modulates the utility value based on the desirability of the departure slot. The specification of the sigmoid function also allows to model a time dependent risk-aversion parameter that realistically reflects the utility withdrawn by the flights depending on their time spent on the market.

Let's determine the parameters of the sigmoid function. The central point of the sigmoid must be located in the middle of the interval $[\max(t, t_i^+), t_i^-]$. This value is equal to $m(t) := \dfrac{\max(t, t_i^+) + t_i^-}{2}$. The function $f : (t, \theta) \mapsto f(t, \theta)$ is then defined as follows:

$$\forall t \in \mathbb{R}_+ \quad \forall \theta \in \mathbb{R}_+ \quad f(t, \theta) := \omega_i \frac{\exp(a(t)(m(t) - \theta)) - \alpha(t)}{1 + \exp(a(t)(m(t) - \theta))}$$

where ω_i the coefficient capturing the cost of a minute of delay for flight i that scales its utility from slot improvement, and $a : t \mapsto a(t)$ is the function that models the change in flight preferences as the duration since the flight entered the market increases. Moreover, function $\alpha : t \mapsto \alpha(t)$ is determined such that $f(t, t_i^+)) = 0$ for every time t, i.e., $\alpha(t) = \exp(a(t)(m - t_i^-))$: this parameter enables utility function v_i to be continuous (and null) at point (t, t_i^-) for all positive number t. Its is also assumed that $\omega_i \sim \mathcal{N}(\mu, \sigma^2)$, and that $a : t \mapsto a(t)$ is a strictly positive and decreasing function. More precisely, $m(t) - \theta$ and $a(t)$ allows to change the shape of the utility over the time flight i spent on the market. $m(t) - \theta$ will linearly shift the utility density from the most preferred slot, t_i^+, to the last preferred one (i.e., the last slot before t_i^-), over time; whereas $a(t)$ will linearly flatten the sigmoid function $f(t, \theta)$. This modeling can be interpreted as a reduction in the slot improvement requirements of flight i as it approaches its ATFM slot. In other words, the further t from t_i^-, the more demanding is the flight, the smaller $a(t)$, and the more utility density is charged for slots near t_i^+. Conversely, the closer t from t_i^-, the less demanding is the flight, the higher $a(t)$, and the less utility density is charged for slots near t_i^+.

If a continuous utility modeling allows for a more accuracy in flights slot preferences determination, the model relies on a discrete representation. Therefore, one will discretize flights' utility by rounding the obtained utility for each possible value to that of the nearest available slot. It is important to recall that for medium-haul flights, the slots are spaced 5 minutes apart. As an example, the utility of the slot $\theta = 57$ minutes will be equal to that of $\theta = 55$ minutes. Ultimately, the following valuation function is:

$$\forall t \in \mathbb{R}_+ \quad \forall \theta \in \mathbb{R}_+ \quad v_i(t, \theta) = \mathbb{1}_{[0, t_i^{\text{exit}}]}(t) \times \mathbb{1}_{[\max(t, t_i^+), t_i^-]}(\theta) \times \omega_i \frac{\exp(a(t)(m - \theta^*)) - \alpha(t)}{1 + \exp(a(t)(m - \theta^*))}$$

where θ^* is the nearest 5-minute slot.

The exogenous premium function is defined as follows:

$$s_i(t, K) := \sum_{k=1}^{K} \mathbb{1}_{[t_{i,k}, t_{i,k} + \Delta t_{i,k}]}(t) \times \gamma_{i,k}$$

where $\gamma_{i,k}$ is the size of the k-th random shock for flight i, $t_{i,k}$ is the time at which this shock occurs, and $\Delta t_{i,k}$ is the duration of the k-shock. No restriction is imposed on the sign of γ, which may result in either a positive or negative premium. If the nature of the shock is negative (most plausible scenario), in the sense that it increases the cost of delay, it will generate a positive premium, subsequently shifting the utility upward. This ensures that flight i has a greater probability of securing an improved slot when a shock occurred than when it did not. Furthermore, it is assumed that the shock results in a non-zero valuation of t_i^-, which represents both the ATFM slot for flight i and the upper bound slot for flight j. This assumption is grounded in the observation that the occurrence of a shock renders flight i highly inclined to depart, to the extent that it is even willing to leave at t_i^- rather than at a later time.

Hence, for every positive real numbers t and θ and every integer K, the utility function of flight i can be re-written as:

$$u_i(t,\theta,K) = \underbrace{\mathbb{1}_{[0,t_i^{\text{exit}}]}(t) \times \mathbb{1}_{[\max(t,t_i^+),t_i^-]}(\theta) \times \omega_i \frac{\exp(a(t)(m-\theta^*)) - \alpha(t)}{1+\exp(a(t)(m-\theta^*))}}_{\text{Valuation Function } v_i(t,\theta)} + \underbrace{\sum_{k=1}^{K} \mathbb{1}_{[t_{i,k},t_{i,k}+\Delta t_{i,k}]}(t) \times \gamma_{i,k}}_{\text{Exogenous Premium Function } s_i(t,K)}$$

Simultaneously, the utility of flight j will adhere to an identical functional form, with index i replaced by j and keeping in mind that t^+j and t^-j are not associated with the same operational slots as t^+i and t^-i, despite also representing flight j's most and least preferred slots.

B Data

In order to evaluate the proposed solution algorithm with real data, it would first be necessary to identify the ATFM delayed flights and the completely delayed flights on the NOP Portal website. Subsequently, the collection of the CTOT, the SOBT, and the tracking history of the EOBT of the maximum number of flights over the longest possible period will be required. All these data points can be found on the NOP Portal website. However, web scraping this site could prove challenging, given that the connection requires a password that changes every minute. A partnership with Eurocontrol could potentially address this issue, but it is uncertain whether Eurocontrol can provide all the necessary data. Finally, the flight-specific cost per minute of delay could be estimated using variables like the cost of parking at a specific airport for a certain type of aircraft. Suggestions on how to model these variables can be found in Cook & Tanner (2011).

C Code

The Python code of the heuristic solution algorithm presented in Section 4.2 is printed below.

```

%%time
import numpy as np
import pandas as pd
from collections import Counter
from copy import copy

def matching(num_flights_I, num_flights_J, disp_omega, dist_ctot, num_slots, max_wait,
    avg_delay):

    ## Defining Parameters

    if dist_ctot == 'binomial':
        peak   = 0.5 * num_slots
        CTOT_i = np.random.binomial(num_slots, peak / num_slots, num_flights_I) + 2
        CTOT_j = np.random.binomial(num_slots, peak / num_slots, num_flights_J)

    elif dist_ctot == 'uniform': #    reprendre
        CTOT_i = np.random.randint(0, high = num_slots + 0, size = num_flights_I) + 2
```

```
        CTOT_j = np.random.randint(0, high = num_slots, size = num_flights_J)

    exit_i  = CTOT_i - 1
    SOBT_i  = [np.random.randint(0, i - 2) for i in CTOT_i]
    omega_i = [np.random.normal(50, disp_omega, size = num_flights_I)]

    EOBT_j  = [j + np.random.poisson(avg_delay) for j in CTOT_j]
    UBT_j   = [j + max_wait for j in EOBT_j]
    exit_j  = [j - 1 for j in UBT_j]
    omega_j = [np.random.normal(50, disp_omega, size = num_flights_J)]


    ## Generating Utility Functions

    buyers  = pd.DataFrame(columns = [i for i in range(num_flights_I)])
    for j in range(num_flights_J):
        line_j = []
        for i in range(num_flights_I):
            if CTOT_j[j] >= SOBT_i[i]:
                line_j.append(max(omega_i[0][i] * (exit_i[i] - CTOT_j[j]), 0))
            else:
                line_j.append(0)
        buyers.loc[len(buyers)] = line_j
    buyers.replace(0, -np.inf, inplace = True)

    sellers = pd.DataFrame(columns = [i for i in range(num_flights_I)])
    for j in range(num_flights_J):
        line_j = []
        for i in range(num_flights_I):
            if CTOT_i[i] >= EOBT_j[j]:
                line_j.append(max(omega_j[0][j] * (exit_j[j] - CTOT_i[i]), 0))
            else:
                line_j.append(0)
        sellers.loc[len(sellers)] = line_j
    sellers.replace(0, -np.inf, inplace = True)


    ## Matching Algorithm

    df                     = buyers.add(sellers, fill_value = 0)
    delta                  = 5
    matching_pairs         = []
    pre_buy_sobt_gains     = []
    buy_temp_gains         = []
    post_buy_sobt_gains    = []
    pre_sell_temp_delay    = []
    post_sell_temp_delay = []

    for i in range(min(num_flights_I, num_flights_J)):

        line_idx = df.index[0]
```

```
        col_idx   = df.columns[0]
        value     = df[col_idx][line_idx]

        for col in df.columns:
            if df[col][df[col].idxmax(axis = 0)] > value:
                value = df[col][df[col].idxmax(axis = 0)]
                line_idx = df[col].idxmax(axis = 0)
                col_idx  = col

        if value != -np.inf:
            matching_pairs.append([line_idx, col_idx])
            pre_buy_sobt_gains.append(delta  * (CTOT_i[col_idx] - SOBT_i[col_idx]))
            buy_temp_gains.append(delta  * (CTOT_i[col_idx]  - CTOT_j[line_idx]))
            post_buy_sobt_gains.append(delta  * (CTOT_j[line_idx] - SOBT_i[col_idx]))
            pre_sell_temp_delay.append(delta * (CTOT_j[line_idx] - EOBT_j[line_idx]))
            post_sell_temp_delay.append(delta * (CTOT_i[col_idx]  - EOBT_j[line_idx]))

        df = df.drop(line_idx)
        df = df.drop([col_idx], axis = 1)

    return [matching_pairs, pre_buy_sobt_gains, buy_temp_gains, post_buy_sobt_gains,
    pre_sell_temp_delay, post_sell_temp_delay]

result = matching(10, 10, 5, "binomial", 24, 6, 3)
print('Matched pairs (seller / buyer):\n', result[0], '\n')
print('Pre-Swap Distance from SOBT for buyers: \n', result[1], '\n')
print('Gain in minutes for buyers: \n', result[2], '\n')
print('Post-Swap Distance from SOBT for buyers: \n', result[3], '\n')
print('Pre-Swap Distance from EOBT for sellers:\n', result[4], '\n')
print('Post-Swap Distance from EOBT for sellers:\n', result[5], '\n')
```