UNIVERSITY OF BARCELONA
FACULTY OF ECONOMICS AND BUSINESS

Master's programme: MSc in Economics – Master Thesis

Specialization: Public Economics

Academic Year: 2025–2026

Income-Based Directional Price Discrimination Evidence from European Airline Markets

Student: Joshua Francisco Brunner

Advisor(s): Joan Calzada & Xavier Fageda

Abstract

This paper investigates income-based directional price discrimination on international airline routes. I construct a real-time dataset of directional fare differences for major European hub airports between February and July 2026 and compare mirrored roundtrips using matching strategies that either assign passengers to the same physical flight components or create closely matched counterfactual fares. The results show that passengers originating in higher-income markets systematically pay higher fares than passengers originating in lower-income markets. In the baseline national GDP specification, a 10,000 USD increase in the income difference raises the directional fare premium by approximately 3.00 to 3.43 USD. The dynamic specifications indicate that this premium tends to become larger in shorter booking windows, especially when the comparison isolates the same underlying flight components. The results also show that the income premium decreases with the number of effective competitors, suggesting that income-based directional fare premia are a signal of airline market power.

JEL Classification: D43, L13, L93, R41.
Keywords: Directional price discrimination; Airline pricing; Dynamic pricing; Income differences; Hub airports; Competition; Market power.

Barcelona, June 8, 2026

1 Introduction

Airline prices are among the most visible examples of dynamic and differentiated pricing. Passengers are used to observing that fares vary across booking dates, routes, airlines, ticket restrictions, and remaining capacity. Less visible, however, is whether airlines also charge systematically different prices depending on the departure airport where a roundtrip starts. Consider a nonstop route between origin A and origin B. A passenger may buy a roundtrip ticket A–B–A, while another passenger buys the mirrored roundtrip B–A–B. If both passengers travel on the same physical flight segments, the airline's operational costs are identical. Systematic price differences between these two otherwise comparable roundtrips can therefore not be explained by flight-specific operating costs, such as fuel, crew, airport charges, or aircraft utilization. Instead, they point to directional price discrimination.

This paper studies whether airlines use income differences between route endpoints to price discriminate on international nonstop routes. The central idea is that passengers originating in wealthier markets may have lower price elasticities of demand. Higher income can be associated with a higher opportunity cost of time, lower search intensity, lower flexibility, and a greater willingness to pay for preferred travel dates and departure times. If airlines can segment passengers by origin market, they may charge higher fares to passengers beginning their trip in the higher-income endpoint. This form of discrimination is particularly relevant in aviation because arbitrage between origin markets is practically impossible: a passenger living at origin A cannot easily purchase the B-originating roundtrip without first repositioning to B.

The paper tests three hypotheses. H1 states that airlines use income-based directional price discrimination: on a given nonstop route, passengers originating in the higher-income endpoint pay higher average fares than passengers originating in the lower-income endpoint, all else equal. H2 states that this income-based fare premium becomes larger as the booking lead time shortens. The closer the departure date, the more airlines can extract a premium from less flexible passengers at the wealthier origin. H3 states that the income premium decreases when stronger independent competitive pressure exists on the route.

The empirical setting consists of 110 international nonstop routes connected to five major European hub airports: London Heathrow, Paris Charles de Gaulle, Frankfurt am Main, Amsterdam Schiphol, and Madrid Barajas. These airports are suitable for the analysis because they combine large demand volumes, substantial income variation across destinations, and meaningful variation in route-level competition. The dataset contains real-time fare quotes retrieved from Skyscanner through a flight-search API and is restricted to economy-class nonstop roundtrips with comparable stay durations. Throughout the empirical analysis, the time between the outbound and return flight is fixed at seven days to ensure comparability across tickets and fare rules. The analysis uses national GDP per capita as the preferred income measure because it is available from a single source for all 110 routes. TL2 regional GDP per capita is used as a robustness check because it provides a more local income measure but is available only for OECD countries.

Identification requires separating directional price discrimination from fare differences driven by flight-specific costs, seat inventory, or schedule-related demand conditions. A fare difference between two directions is not informative if the compared passengers travel on different aircraft, face different seat availability, or are affected by different flight-specific cost components. For this reason, the paper constructs directional price differences using three complementary matching strategies. Strategy 1 aggregates synchronized flight chains so that passengers from both origins are compared across the same physical flight components within the chain. Strategy 2 matches mirrored roundtrips by route, airline, travel date, booking horizon, and identical flight numbers, thereby preserving daily variation for the booking-horizon analysis. The third strategy, Strategy 3, follows Lewis (2020) and constructs synthetic counterfactual prices that isolate the same intermediate flight components. Together, these strategies allow the analysis to compare passengers either on the same physical flight components or on closely matched scheduled equivalents, directly addressing the concern that directional fare differences should not be interpreted as discrimination unless flight-specific cost and capacity heterogeneity are controlled for.

The baseline model based on the national GDP specification shows that a 10,000 USD increase in the income difference raises the directional fare premium by 3.00 to 3.43 USD. This indicates that passengers starting in the higher-income endpoint systematically pay more for comparable roundtrips, consistent with the positive price gap predicted by H1. The booking-horizon results suggest that the income premium becomes larger as departure approaches. Using Strategy 3 as the main interpretation, because it most directly isolates the same underlying flight components for the booking windows it covers, the estimates show a stronger income sensitivity in shorter booking windows than in the longest booking horizon. This pattern is consistent with H2, according to which passengers from higher-income origins become less price-sensitive closer to departure. The competition results show that the income premium is highest on monopoly routes and decreases as competitive pressure becomes stronger. This supports H3 and indicates that independent competition disciplines, but does not fully eliminate, origin-based pricing power.

The empirical literature on directional price discrimination in airline markets is still limited and provides mixed evidence. Luttmann (2019) studies U.S. domestic routes using the U.S. Department of Transportation's Origin and Destination Survey (DB1B), a 10% sample of sold airline tickets, and finds that fares are higher when roundtrips originate in higher-income metropolitan areas. His results suggest that airlines may use endpoint income differences to price discriminate, but the data are based on realized ticket purchases and do not observe exact booking dates, flight times, or the same physical flight components. Lewis (2020) challenges this interpretation using more detailed fare-quote data for specific flights. He finds no systematic directional price discrimination in the U.S. domestic market once cost heterogeneity is controlled for, while showing that directional price differences are common on international routes to and from the United States. This paper contributes to the literature in three ways. First, it extends the analysis of income-based directional price discrimination from the U.S. domestic market studied by Luttmann (2019) to international airline markets in the European Union. Second, it uses more granular real-time fare data. Unlike the DB1B ticket data used in Luttmann (2019), the dataset used in this paper makes it possible to compare mirrored fares observed at the same search date and to construct outcomes based on the same physical flight components or closely matched equivalents, following the identification concern emphasized by Lewis (2020). Third, the paper extends the analysis dynamically by examining whether the income premium varies over the booking horizon and whether it is moderated by effective route-level competition.

The policy implications follow directly from the empirical results. Since the analysis documents income-based directional price discrimination, airlines appear to possess sufficient market power to segment origin markets and charge asymmetric markups across route directions. At the same time, the finding that the income premium decreases with stronger competition shows that competition is the key mechanism for disciplining this pricing power. Policy should therefore focus on strengthening effective competition in international airline markets. Competition authorities should evaluate mergers, alliances, and joint ventures based on effective competition rather than the nominal number of airline brands. This is particularly relevant for the privatization of TAP Air Portugal, where a potential acquisition by Air France-KLM (Air France-KLM, 2026) or the Lufthansa Group (Lufthansa Group, 2025) may improve network connectivity but could also reduce independent competition and reinforce origin-market segmentation. In addition, policy should address entry barriers at congested hub airports. Large hub airports are often characterized by scarce slot capacity, so high fares do not automatically induce free entry by additional competitors. An entrant can only discipline incumbent pricing if it obtains commercially viable takeoff and landing times. Capacity expansion can therefore lower structural entry barriers in the long run, although the debate around the Heathrow third runway illustrates how politically and socially difficult this channel can be (The Economist, 2026). Finally, bilateral air service agreements that increase frequencies, destination points, or reciprocal market access can improve contestability by reducing legal capacity restrictions and opening route access.

The remainder of the paper proceeds as follows. Section 2 reviews the theoretical and empirical literature on price discrimination, airline revenue management, and directional fare differences. Section 3 develops the theoretical motivation and derives the research hypotheses. Section 4 describes the data, outcome construction, matching strategies, and descriptive evidence. Section 5 presents the econometric model structure. Section 6 presents the baseline and heterogeneous results for H1. Section 7 analyzes the booking-horizon mechanism, and Section 8 analyzes the competition mechanism. Section 9 concludes and discusses policy implications.

2 Literature Review

The classical foundation of third-degree price discrimination goes back to Pigou's *The Economics of Welfare* (Pigou, 1920). A firm can increase profits by charging different prices across separable consumer groups when arbitrage is prevented and demand elasticities differ. In the present setting, passengers beginning their roundtrip at origins A and B constitute geographically segmented markets because purchasing the opposite-origin ticket requires costly repositioning. Price discrimination can remain an equilibrium pricing strategy when several firms compete in the market. Thisse and Vives (1988) show that firms strategically choose spatially discriminatory pricing rather than uniform pricing in oligopolistic markets when prices can be conditioned on geographically segmented consumer groups. This result is important for the present paper because it shows that price discrimination is not limited to monopoly settings. The competitive implications of discriminatory pricing are developed more directly by Corts (1998). His analysis shows that the effects of competitive price discrimination depend on whether firms rank consumer groups similarly. In the context of this paper, income generates a common ranking of the two origin markets: because origin A is defined as the wealthier endpoint, all airlines operating the route may identify A as the stronger market and B as the weaker market. This common ranking creates best-response symmetry, since each airline has an incentive to charge a relatively higher price in the same high-income origin market. Applied to the airline setting, this implies that competition does not necessarily eliminate origin-specific pricing. Airlines may continue to charge different fares across origins when the underlying demand conditions differ.

Market entry conditions also matter. Li and Shuai (2019) distinguish between markets with a fixed number of firms and markets with endogenous entry. When the number of firms is fixed, price discrimination increases firm profits and reduces consumer surplus relative to uniform pricing, because discriminatory markups are not competed away by additional firms. This setting is relevant for international aviation markets, where bilateral air service agreements, slot scarcity, airport capacity constraints, alliances, and joint ventures may restrict effective entry. When entry is endogenous, discriminatory pricing can induce additional entry, although this entry need not be welfare-improving and may even be excessive relative to the social optimum. Borenstein (1985) similarly shows that discriminatory prices can persist even under free entry, because competition does not necessarily eliminate differences in reservation prices, brand preferences, or demand elasticities across consumer groups. For this paper, the relevant implication is that stronger competition should reduce the income premium, but it need not fully eliminate origin-based market segmentation.

Airline pricing is shaped by high fixed costs, perishable capacity, stochastic demand, and the need to allocate a fixed number of seats across passengers with heterogeneous willingness to pay. Revenue management systems therefore play a central role in the practical implementation of price discrimination. Belobaba (1987) provides the classical foundation for seat inventory control, where airlines decide how many seats to protect for passengers expected to have a higher willingness to pay in the future. This mechanism is directly relevant for directional price discrimination. An airline does not necessarily need to post completely different fare menus in the two directions. It can also generate directional fare differences by making discounted fare classes more or less available depending on the origin market. If demand from the higher-income origin is expected to be less price elastic, fewer discounted seats may be made available for passengers starting there.

Recent evidence confirms that airline pricing should not be interpreted as the outcome of a single fully rational dynamic pricing model. Hortaçsu et al. (2023) use internal data from a large U.S. airline and show that observed prices are shaped by organizational structure, delegated decision-making, and pricing heuristics. In their setting, network planning determines routes, frequencies, and capacities; the pricing department chooses fare menus; and the revenue management department forecasts demand and allocates inventory across fare classes. This is important for the present paper because directional price differences may arise through the interaction of fare menus, demand forecasts, and seat availability rather than through an explicit rule that labels passengers by income. The observed fare premium for high-income origins can therefore be consistent with standard airline revenue-management practice, even if the airline does not directly observe individual passenger income.

This perspective also helps interpret the empirical strategy of this paper. Since airline fares depend on remaining capacity and inventory controls, comparing different flights could confound price discrimination with unobserved seat availability. This is why the outcome variables in Section 4.2 are constructed to compare the same physical flight components or closely matched equivalents. If the compared passengers are assigned to the same aircraft components, flight-specific operating costs and seat availability are held constant. The remaining price difference is therefore more informative about origin-based segmentation. This identification issue is central because observed fare differences alone are not sufficient evidence of price discrimination unless cost and capacity explanations are addressed.

Several empirical studies show that airlines use ticket restrictions and booking rules to segment demand. Stavins (2001) provides evidence that advance-purchase requirements and Saturday-night stayover restrictions separate price-sensitive leisure travelers from less price-sensitive business travelers. Lewis (2021) shows that the use of such restrictions depends on the competitive environment and on the type of competitors present on a route. Escobari and Jindapon (2014) study refundable and non-refundable tickets and show how contract design can screen passengers with different levels of demand uncertainty. Mantin and Koo (2009) further show that fare increases over the booking horizon are related to route characteristics, including income and low-cost-carrier competition. These studies are directly connected to the dynamic hypothesis of this paper: if passengers from high-income origins become less price sensitive as departure approaches, the directional income premium should increase in shorter booking windows.

A related application outside the airline industry is provided by Chang et al. (2022), who study UberX fares from major U.S. airports to nearby hotels. Their setting is useful because the origin is held constant at the airport, while hotel room rates serve as a geographic proxy for travelers' willingness to pay. They show that UberX fares increase with hotel room rates even after controlling for distance, travel duration, traffic conditions, and time-specific demand shocks. Although this is not direct evidence on airline pricing, it supports the broader mechanism that transportation providers can use observable geographic information to segment consumers.

The current empirical literature on directional price discrimination in the airline industry is limited and produces mixed results. To the best of my knowledge, only two previous papers provide comprehensive analyses of whether airlines systematically charge different fares for roundtrip itineraries depending on the point of origin. Luttmann (2019) investigates directional price discrimination by explicitly incorporating regional income as a control variable to explain fare disparities, while Lewis (2020) does not include income as a control variable in the primary analysis.

Luttmann (2019) identifies a significant relationship between origin income and pricing within the United States domestic market. The analysis demonstrates that base fares are 1.80 to 4.30 USD higher for each 10,000 USD difference in per capita income between the route endpoints. He also examines how market structure influences this effect by including interaction terms between income and competition variables in his model. The results indicate that directional price discrimination increases with the level of competition. Specifically, for every additional carrier offering nonstop service on a route, the income-driven price gap increases by approximately 0.60 USD per 10,000 USD of income disparity.

In contrast, Lewis (2020) finds an absence of systematic directional price discrimination within the U.S. domestic industry. Over 99 % of the itineraries exhibit identical prices for passengers starting at different endpoints. Unlike Luttmann (2019), however, Lewis (2020) does not use income as an explanatory control variable, because his main objective is to identify whether directional price differences remain once the same physical flight components and directional costs are controlled for. Once these directional costs are appropriately addressed, domestic price differences vanish, which suggests that the income effects reported by Luttmann (2019) may be spurious and driven by unobserved directional factors that correlate with regional income. For example, because the DB1B data contain realized ticket purchases but do not observe the booking horizon, fares may differ if passengers at the higher-income endpoint tend to book shortly before departure while passengers at the lower-income endpoint tend to book several weeks or months in advance. By contrast, Lewis (2020) finds that directional price discrimination is the norm on international flights to and from the United States, where over 95 % of fares show significant price differences.

3 Theoretical Motivation

Third-degree price discrimination requires three main conditions. First, firms must possess sufficient market power to charge prices above marginal cost. Second, consumers must be separable into identifiable groups between which arbitrage is limited. In the present setting, passengers beginning their roundtrip at origins A and B constitute geographically segmented markets because purchasing the opposite-origin ticket requires prior repositioning to the other endpoint. Third, the segmented groups must differ in their price sensitivity. The central assumption of the theoretical framework is that passengers originating in the higher-income market A have a lower price elasticity of demand than passengers originating in the lower-income market B.

To formalize this mechanism, the strategic interaction is represented as Bertrand competition with differentiated products. Following Singh and Vives (1984), firms compete strategically in prices, while product differentiation enables equilibrium prices to remain above marginal cost. Airline services on the same route are imperfect substitutes because carriers differ in departure schedules, airport presence, service quality, loyalty programs, and network connectivity. The original differentiated-duopoly framework is generalized here to N competing airlines, with $N \in \{1, \ldots, 5\}$ reflecting the maximum number of airlines observed on a route in the empirical analysis, and two origin markets, while regional income is allowed to affect consumer utility and demand. Income matters for pricing only if it affects the elasticity of demand, not merely the level of demand. If higher-income consumers are less price-sensitive, firms can charge higher markups in high-income origins; if income shifted demand without changing elasticity, the optimal markup would not necessarily change. In market $k \in \{A, B\}$, the representative consumer maximizes a quasi-linear utility function, leading to the optimization problem:

$$\max_{q_{1,k},\ldots,q_{N,k}} \mathcal{L} = V(q_{1,k}, \ldots, q_{N,k}; Y_k) + Y_k - \sum_{j=1}^{N} p_{j,k} q_{j,k} \tag{1}$$

The solution defines the direct demand function $q_{i,k}(p_{i,k}, \mathbf{p}_{-i,k}; Y_k)$, which incorporates regional income (Y_k) and the price vector of all rivals.

Airlines anticipate this demand and simultaneously set prices to maximize total profit across both segments:

$$\max_{p_{i,A}, p_{i,B}} \Pi_i = \sum_{k \in \{A,B\}} \left[(p_{i,k} - c) \cdot q_{i,k}(p_{i,k}, \mathbf{p}_{-i,k}; Y_k) \right] - F_i \tag{2}$$

Solving the first-order conditions with respect to $p_{i,k}$ and rearranging the standard inverse-elasticity rule gives the optimal price:

$$p^*_{i,k}(t) = c \left(\frac{e^F_{i,k}(t)}{e^F_{i,k}(t) - 1} \right). \tag{3}$$

Here, $e^F_{i,k}(t)$ denotes the firm-level price elasticity of demand for airline i in origin market k at booking time t, where t measures the remaining time until departure. Following Holmes (1989), firm-level elasticity can be decomposed into industry-demand elasticity and cross-price elasticity. The industry-demand component captures the extent to which passengers leave the market altogether when the fare increases, for example by not traveling, choosing another transport mode, or postponing the trip. The cross-price component captures the extent to which passengers switch to rival suppliers, such as another airline, a nearby airport, an indirect connection, or high-speed rail. A lower firm-level elasticity therefore means that passengers react less strongly to a fare increase either because they are less willing to leave the market or because they have fewer attractive substitutes. This allows the airline to charge a higher price over marginal cost. The step-by-step mathematical derivations for demand and profit maximization are provided in Appendix A.1.

The prediction underlying H1 follows directly from Equation 3. Since the optimal price is an elasticity-dependent markup over marginal cost, a lower firm-level price elasticity in the wealthier market A implies a higher optimal price than in market B. The model therefore predicts a positive directional price gap. Consequently, the average theoretical price gap over the observed booking horizon is positive:

$$\bar{p}_A^* - \bar{p}_B^* = \frac{1}{T} \int_0^T \Delta p(t)\, dt > 0. \tag{4}$$

H2 states that the income-based fare premium increases as the booking lead time shortens. The mechanism behind H2 is that firm-level price sensitivity declines as departure approaches, but that this decline is stronger for passengers from higher-income origins. Passengers originating in the wealthier market A have higher opportunity costs of time and therefore higher effective search costs. Following Stahl (1996), heterogeneous search costs are central for oligopolistic pricing because consumers with higher search costs compare fewer prices and discipline firms less through switching. In the context of this paper, this implies that wealthier passengers are expected to search less intensively for alternative fares, primarily reducing their cross-price elasticity.[1]

Another aspect not explicitly modeled here is that passengers may have already invested in expensive, non-refundable complementary goods such as hotels or other trip-related commitments. The sunk-cost mechanism documented by Arkes and Blumer (1985) implies that once money, effort, or time has been invested, individuals are more likely to continue with the chosen plan rather than cancel or substitute, primarily reducing their industry-demand elasticity.[2] Since firm-level price elasticity in Equation 3 varies with the remaining time until departure, the same inverse-elasticity rule implies that the optimal price changes over the booking horizon. As t approaches the final booking window, firm-level price sensitivity is expected to fall more strongly in market A than in market B. The theoretical price gap is therefore expected to widen as departure approaches:

$$\frac{\partial \Delta p(t)}{\partial t} < 0. \tag{5}$$

The empirical implementation tests the same mechanism in terms of income sensitivity. It asks whether a given income difference translates into a larger directional fare premium in shorter booking windows than in longer booking windows. This is the theoretical prediction tested in H2, and the complete derivation of the dynamic increase is found in Appendix A.2.

The theoretical prediction for H3 follows from the role of cross-price elasticity in Holmes (1989). Since firm-level elasticity contains a cross-price component, stronger competition makes it harder for an airline to sustain origin-specific markups. As the number of competing airlines N increases, passengers have more outside options and the aggregate cross-price elasticity $e^C_{i,k}(N)$ rises. This reduces the ability

[1] Stahl's model does not explicitly study airline booking horizons. The argument here applies the search-cost mechanism to the airline setting: as departure approaches, the remaining time to compare and switch between alternatives becomes more limited, so differences in effective search costs can become more relevant close to departure.

[2] This mechanism is not formally included in the model. It provides an additional behavioral explanation for why high-income-origin passengers may become less willing to substitute close to departure, especially if complementary trip-related commitments have already been made.

of an airline to maintain high markups in the less elastic market A. Accordingly, firm-level elasticity can be written as industry-demand elasticity plus an N-dependent cross-price elasticity. Inserting this N-dependent elasticity into the inverse-elasticity rule gives:

$$p^*_{i,k}(N) = c\left(\frac{e^I_{i,k} + e^C_{i,k}(N)}{e^I_{i,k} + e^C_{i,k}(N) - 1}\right). \tag{6}$$

The resulting static price-discrimination effect can be written as $\Delta p(N) > 0$, but it is expected to decrease as N rises:

$$\frac{\partial \Delta p(N)}{\partial N} < 0. \tag{7}$$

This result is important for the empirical analysis even though the number of effective competitors is relatively stable within a route over the sample period. The model does not require short-run variation in N within routes. Instead, it generates a cross-sectional prediction: routes with stronger effective competition should display smaller income-based fare premia than routes with weaker effective competition. Appendix A.3 provides the mathematical proof of this competition mechanism. This is the theoretical prediction tested in H3.

4 Data, Outcome Construction, and Descriptive Evidence

4.1 Data and Sample Construction

The objective of this paper is to empirically analyze whether airlines charge systematically different fares depending on the income level of the origin market, even when passengers travel on the same route and on comparable flight products. The analysis therefore focuses on international nonstop roundtrip fares and combines real-time fare data with route-, city-, airport-, airline-, and market-level characteristics. This structure allows the empirical analysis to relate directional fare differences to income differences, booking horizons, and effective route-level competition.

The sample focuses on routes connecting to the five largest European hub airports: London Heathrow (LHR), Paris Charles de Gaulle (CDG), Frankfurt am Main (FRA), Amsterdam Schiphol (AMS), and Madrid Barajas (MAD). These hubs provide large passenger volumes, substantial variation in route-level competition, and broad variation in destination income. The dataset covers 110 international nonstop routes, which are listed in Appendix B. For each hub, 20 routes were selected and grouped into five destination-income tiers using World Bank GDP per capita data (World Bank, 2026). The tiers are $\leq 10,000$ USD, $> 10,000$ to $\leq 30,000$ USD, $> 30,000$ to $\leq 50,000$ USD, $> 50,000$ to $\leq 70,000$ USD, and $> 70,000$ USD. This design ensures sufficient variation in income differences across route endpoints.

Routes were selected using frequency data from Flightsfrom.com (Flightsfrom.com, 2026). Within each income tier and hub, the four most frequently operated routes were selected. Domestic routes were excluded because income differences are limited within domestic markets and because domestic air transport markets are institutionally not comparable to international routes. The only exception is the United Kingdom, where England and Scotland are treated as distinct regional entities. The ten inter-hub routes between the five hubs were kept separately to avoid losing routes through overlap between ABA and BAB directions. The route design is intended to generate income variation without selecting routes solely on the basis of observed fares. Frequency is used as the route-selection criterion because it identifies economically relevant nonstop markets and reduces the risk that the sample is dominated by thin, irregular, or highly seasonal services.

Fare data were collected through RapidAPI using the Flights Scraper Sky tool, which retrieves real-time Skyscanner fare quotes (RapidAPI, 2026; Skyscanner, 2026). All fares were downloaded in US dollars, and the displayed lead price was used because it represents the main price signal observed by consumers on the platform. Data collection occurred in two waves. The first wave was collected on January 31, 2026, for departures between February 1 and April 30, 2026, using seven-day roundtrips. The second wave was collected on April 30, 2026, for departures between May 1 and July 31, 2026. This second wave extends the observation period from the late winter and spring schedule into the summer schedule and therefore increases variation in travel dates, seasonal demand, and booking horizons. Additional 21-day roundtrips were collected only in the second wave because the synthetic matching procedure requires a longer stay structure around the middle flight pair, and because each additional stay duration substantially increases the number of API requests. For the same reason, the data collection does not vary the stay length over many alternatives such as one-day, 14-day, or longer roundtrips. The empirical regressions use seven-day roundtrips as the common product definition to ensure comparable fare rules.

The raw API output contained 45,970,752 ticket prices. This large number reflects the Cartesian-product structure of roundtrip searches: if a query contains n outbound and m return flights, the resulting dataset contains $n \cdot m$ roundtrip prices. The observations were cleaned in Stata by harmonizing variables across waves, restricting the sample to economy-class nonstop roundtrips, requiring complete flight information, and retaining only tickets where outbound and return legs are operated by the same airline. This removes connecting flights, incomplete observations, and non-comparable cross-airline combinations. However, the Cartesian-product structure remains within each airline. After cleaning, the dataset contains 1,507,730 usable ticket observations.

The fare data were merged with route-, city-, and airline-level variables. National GDP per capita comes from the World Bank (World Bank, 2026). Regional income is measured at the TL2 level using the OECD regional database (OECD, 2026). TL2 refers to the OECD's large territorial level below the national level. For example, Barcelona is assigned to Catalonia as the relevant TL2 region. This measure is useful because the income level of a large metropolitan region can differ from the national average. However, TL2 income is used only as a robustness check because it is available only for OECD countries, which reduces the route sample and shifts it toward relatively wealthier destinations.

Additional control variables were collected to account for demand, geography, and market structure. Flight distance was taken from Great Circle Mapper (Great Circle Mapper, 2026), population from World Population Review (World Population Review, 2026), and average temperatures from WeatherSpark (WeatherSpark, 2026). March temperatures are used for the first wave covering February to April, while June temperatures are used for the second wave covering May to July. These temperature measures approximate seasonal tourism demand, which may affect fare levels independently of income differences. High-speed rail availability was checked with Omio (Omio, 2026), and alternative-airport availability was checked using Skyscanner's alternative-airports option (Skyscanner, 2026). These variables are relevant because rail and nearby airport options may provide substitutes for air travel and can therefore affect the degree of competitive pressure faced by airlines.

Competition variables were constructed from Flightsfrom.com rather than from the scraped fare data (Flightsfrom.com, 2026). This avoids measurement error for two reasons. First, an airline may operate a route but be absent from the API output if a flight is sold out or not returned by Skyscanner. Second, Skyscanner provides broad coverage as a metasearch platform, but the scraped quotes still represent fares visible through the platform rather than the full universe of tickets sold by airlines. This platform-specific limitation appears negligible in the present sample and is therefore unlikely to substantially affect the core data. To nevertheless avoid measurement error, I use Flightsfrom.com as the basis for the competition measures. The platform provides operating airlines and average flight frequencies, which are first used to measure nominal competition and the route-level HHI. These measures are then adjusted for holdings and joint ventures to construct the number of effective competitors and the adjusted HHI. Effective competition treats airlines belonging to the same corporate group or operating under a relevant joint venture as one competitor. This matters because several airline brands on a route do not necessarily imply independent pricing decisions. For example, if LHR–BCN is operated by British Airways and Vueling, it is counted as one effective competitor because both carriers belong to International Airlines Group. The definition is deliberately stricter than a simple count of airline brands because the relevant economic question is whether passengers face independent pricing constraints. This distinction is central for the interpretation of H3 and for the policy discussion because it links the measured income premium to market power rather than to the mere visibility of multiple airlines on a search platform. Hub status, route liberalization, and school holidays were added through manual research. Legacy carriers are identified through membership in oneworld, SkyTeam, or Star Alliance (Oneworld, 2026; SkyTeam, 2026; Star Alliance, 2026).

4.2 Definition of the Outcome Variable

This section defines the directional fare-difference outcomes used in the empirical analysis. The central empirical challenge is to construct this price difference without confusing origin-based fare premia with differences in operating costs, seat availability, travel dates, or fare rules.

4.2.1 Strategy 1: Aggregated Chain and Bridge Procedure

Strategy 1 is a newly developed procedure in this paper and constructs two synchronized chains of consecutive seven-day roundtrips for a given nonstop route between endpoints A and B. The general idea is to arrange mirrored roundtrips so that the return flight of one passenger simultaneously serves as the outbound flight of a passenger originating at the opposite endpoint. This creates an alternating sequence in which passengers from both origins repeatedly use the same physical flight segments. The purpose is to compare fares across directions while holding flight-specific operating costs and seat inventory as constant as possible for the shared interior segments.

To illustrate the procedure, consider the route between London Heathrow and Barcelona. London corresponds to endpoint A because the United Kingdom has the higher national GDP per capita in the route list (53,246 USD), while Barcelona corresponds to endpoint B because Spain's national GDP per capita is lower (35,326 USD; see Appendix B). The first chain starts from endpoint B in Barcelona, and the mirrored chain starts from endpoint A in London. This ordering first shows how passengers from the lower-income origin are linked to passengers from the higher-income origin through the same physical segments, before the mirrored chain reverses the assignment of origins to the alternating positions.

The Barcelona-starting chain in Table 1 illustrates the procedure. A passenger originating in Barcelona begins a seven-day roundtrip on February 3 by taking flight BA 481 from BCN to LHR. Seven days later, on February 10, this passenger returns to Barcelona on flight BA 476 from LHR to BCN. The same

February 10 flight simultaneously serves as the outbound flight of a passenger originating in London. That London-origin passenger then returns seven days later, on February 17, on flight BA 481 from BCN to LHR. The February 17 flight, in turn, simultaneously serves as the outbound flight of the next Barcelona-origin passenger. This alternating structure continues throughout the chain.

Consequently, each interior segment, meaning each segment between the first and last boundary flights of the chain, is shared by passengers originating at the two different endpoints. Since the passengers are compared using the same physical aircraft, flight-specific operating costs are held constant for the shared segments. Furthermore, because all data were downloaded on the exact same day with only seconds elapsing between the different API queries, the number of available seats did not change between requests, ensuring that the seat inventory remains constant across the compared bookings.

The only exceptions are the boundary segments at the beginning and end of the chain. In the Barcelona-starting chain, the February 3 BCN–LHR flight cannot simultaneously serve as the return flight of an earlier London-origin passenger because the required preceding roundtrip lies outside the observation window. Similarly, the May 5 LHR–BCN flight cannot serve as the outbound flight of a subsequent London-origin passenger because the corresponding return flight would occur after the observation window. These unmatched boundary segments constitute the chain's boundary problem.

Table 1: Chain Strategy Table Starting in Barcelona

Date	**Flight No.**	**Route**	**Leg Type (BCN/LHR)**	**Allocation**	**Segment**	**Price (USD)**
03. Feb	BA 481	BCN → LHR	Outbound/–	$P(B_1)$	Boundary	241.36
10. Feb	BA 476	LHR → BCN	Return/Outbound	$P(A_1)$	Interior	340.37
17. Feb	BA 481	BCN → LHR	Outbound/Return	$P(B_2)$	Interior	279.50
24. Feb	BA 476	LHR → BCN	Return/Outbound	$P(A_2)$	Interior	173.44
03. Mar	BA 481	BCN → LHR	Outbound/Return	$P(B_3)$	Interior	137.21
10. Mar	BA 476	LHR → BCN	Return/Outbound	$P(A_3)$	Interior	229.80
17. Mar	BA 481	BCN → LHR	Outbound/Return	$P(B_4)$	Interior	198.36
24. Mar	BA 476	LHR → BCN	Return/Outbound	$P(A_4)$	Interior	182.68
31. Mar	BA 481	BCN → LHR	Outbound/Return	$P(B_5)$	Interior	175.08
07. Apr	BA 476	LHR → BCN	Return/Outbound	$P(A_5)$	Interior	219.59
14. Apr	BA 481	BCN → LHR	Outbound/Return	$P(B_6)$	Interior	165.19
21. Apr	BA 476	LHR → BCN	Return/Outbound	$P(A_6)$	Interior	190.90
28. Apr	BA 481	BCN → LHR	Outbound/Return	$P(B_7)$	Interior	217.50
05. May	BA 476	LHR → BCN	Return/–	–	Boundary	–

Using only the Barcelona-starting chain would nevertheless create a systematic booking-horizon imbalance between the two origins. Within a single chain, the Barcelona-origin and London-origin roundtrips enter the sequence at alternating travel dates. Because fares generally increase as departure approaches, one origin could appear systematically more expensive merely because its tickets were observed at shorter booking horizons. For example, a fare observed three days before departure would generally be expected to exceed an otherwise comparable fare observed ten days before departure. A directional fare difference based on only one chain could therefore partly reflect the unequal assignment of booking horizons rather than origin-based pricing.

To remove this imbalance, Strategy 1 also constructs the mirrored London-starting chain shown in Table 2. The logic is identical, but the assignment of origins to the alternating positions is reversed. A London-origin passenger departs on February 3 on flight BA 476 from LHR to BCN and returns on February 10 on flight BA 481 from BCN to LHR. The February 10 return flight simultaneously serves as the outbound flight of a Barcelona-origin passenger, who returns to Barcelona on February 17 on flight BA 476. Thus, the London-starting chain reverses the sequence used in the Barcelona-starting chain.

Table 2: Chain Strategy Table Starting in London

Date	**Flight No.**	**Route**	**Leg Type (LHR/BCN)**	**Allocation**	**Segment**	**Price (USD)**
03. Feb	BA 476	LHR → BCN	Outbound/–	$P(A_1^*)$	Boundary	573.93
10. Feb	BA 481	BCN → LHR	Return/Outbound	$P(B_1^*)$	Interior	352.08
17. Feb	BA 476	LHR → BCN	Outbound/Return	$P(A_2^*)$	Interior	402.28
24. Feb	BA 481	BCN → LHR	Return/Outbound	$P(B_2^*)$	Interior	243.93
03. Mar	BA 476	LHR → BCN	Outbound/Return	$P(A_3^*)$	Interior	150.16
10. Mar	BA 481	BCN → LHR	Return/Outbound	$P(B_3^*)$	Interior	236.14
17. Mar	BA 476	LHR → BCN	Outbound/Return	$P(A_4^*)$	Interior	240.09
24. Mar	BA 481	BCN → LHR	Return/Outbound	$P(B_4^*)$	Interior	239.74
31. Mar	BA 476	LHR → BCN	Outbound/Return	$P(A_5^*)$	Interior	368.00
07. Apr	BA 481	BCN → LHR	Return/Outbound	$P(B_5^*)$	Interior	347.70
14. Apr	BA 476	LHR → BCN	Outbound/Return	$P(A_6^*)$	Interior	248.31
21. Apr	BA 481	BCN → LHR	Return/Outbound	$P(B_6^*)$	Interior	146.33
28. Apr	BA 476	LHR → BCN	Outbound/Return	$P(A_7^*)$	Interior	214.40
05. May	BA 481	BCN → LHR	Return/–	–	Boundary	–

Combining the Barcelona-starting and London-starting chains balances the booking-horizon distribution across the two origins. Each origin occupies both positions in the alternating sequence and is therefore observed under comparable booking-horizon conditions. Apart from the unavoidable boundary segments, passengers originating in Barcelona and London are compared while traveling on the same physical flight segments. The resulting directional fare difference is therefore not systematically driven by one origin being observed closer to departure than the other. The detailed matching rules are reported in Appendix E.1.

The Strategy 1 outcome is the average price of the high-income-origin roundtrips across both balanced chains minus the corresponding average price of the low-income-origin roundtrips:

$$\Delta P_{\text{chain}} = \left(\frac{1}{n} \sum_{i=1}^{n} P_{A,i} \right) - \left(\frac{1}{n} \sum_{i=1}^{n} P_{B,i} \right). \tag{8}$$

Here, n is the number of matched weekly roundtrips assigned to each origin after combining the two chains. A positive value indicates that passengers starting in the higher-income endpoint pay more on average.

4.2.2 Strategy 2: Daily Flight-Pair Matching

Strategy 2 matches a roundtrip originating in A with the mirrored roundtrip originating in B on the same route, airline, travel dates, booking horizon, and flight numbers. The matched passengers do not occupy the same physical aircraft. However, identical flight numbers ensure that the matched flights correspond to the same scheduled service pattern, including time of day and day of week. This matters because airport charges, slot scarcity, and peak-hour demand can differ systematically across the day. Matching on flight numbers therefore controls more tightly for schedule-related cost and demand conditions than a comparison based only on route and date.

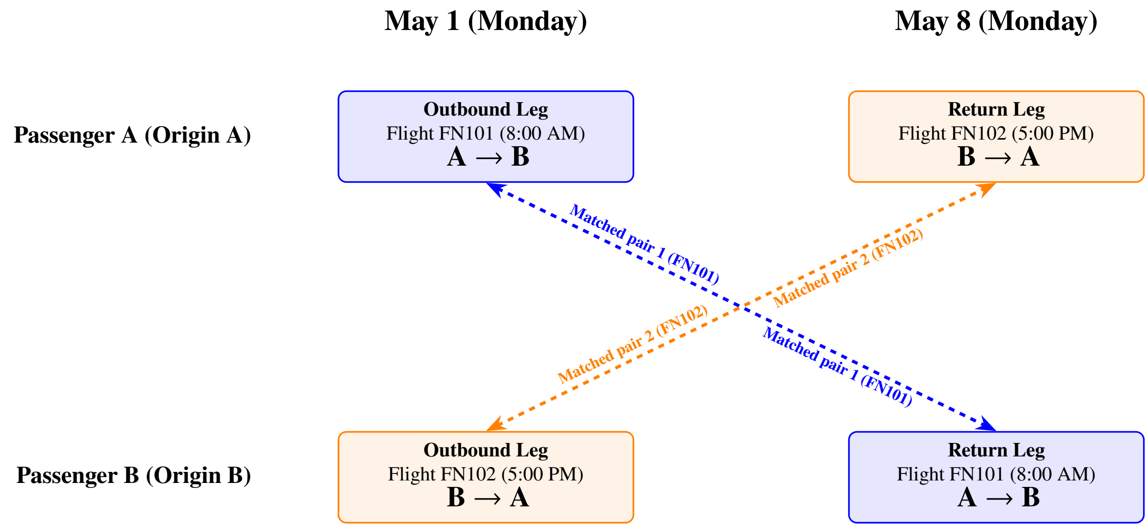

Figure 1: Daily flight-pair matching by mirrored flight numbers.

The Strategy 2 outcome is

$$\Delta P_{\text{pair}} = P_A - P_B. \tag{9}$$

Here, P_A is the seven-day roundtrip fare from the higher-income origin and P_B is the matched mirrored fare from the lower-income origin. A positive value indicates a directional fare premium for the higher-income origin. The operational ranking and tie-breaking rules used for daily matching are reported in Appendix E.2.

4.2.3 Strategy 3: Synthetic Matching

Strategy 3 follows the synthetic matching logic developed by Lewis (2020). It constructs a counterfactual price for the same middle flight components by combining two seven-day roundtrips with one 21-day roundtrip. Because this matching procedure requires a complete 21-day sequence after the middle flight pair, the observation window is shorter than in the seven-day strategies and runs from May 1 to July 10. The fare data cover seven-day roundtrips only up to July 31; middle flight pairs after July 10 would therefore require return flights in August, for which no data are available. This construction restores a flight-component comparison similar in spirit to Strategy 1, but it requires an additional fare-rule assumption.

Figure 2 shows the logic of the synthetic construction. Trip 1 covers the first seven-day roundtrip from t_1 to t_2, and Trip 2 covers the second seven-day roundtrip from t_3 to t_4. Trip 3 is a 21-day roundtrip from

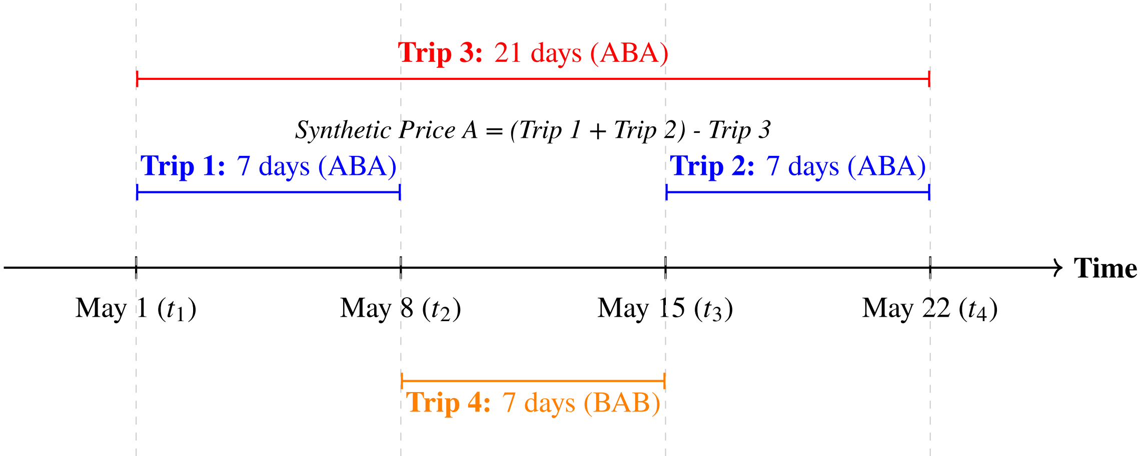

Figure 2: Synthetic isolation of the middle flight pair.

t_1 to t_4 and contains the two outer flight segments. Subtracting Trip 3 from the sum of Trips 1 and 2 removes the outer segments and leaves the implied price of the two middle segments between t_2 and t_3. The synthetic price is

$$P_{\text{Synthetic } A} = P_A(t_1, t_2) + P_A(t_3, t_4) - P_A(t_1, t_4). \tag{10}$$

The detailed priority rules used to select the underlying itineraries are reported in Appendix E.3. The construction rotates the observed and synthetic directions across routes to prevent either direction from being systematically assigned to the observed or synthetic side. For even route IDs, the A-origin roundtrip (ABA) is synthetic and the B-origin roundtrip (BAB) is observed. The Strategy 3 outcome is therefore calculated as the synthetic higher-income-origin fare minus the observed lower-income-origin fare. For odd route IDs, the assignment is reversed: the A-origin roundtrip (ABA) is observed and the B-origin roundtrip (BAB) is synthetic, so the outcome is calculated as the observed higher-income-origin fare minus the synthetic lower-income-origin fare. This rotation ensures that the observed and synthetic sides are balanced across routes, while the resulting outcome is always oriented as the higher-income-origin fare minus the lower-income-origin fare. Positive values therefore have the same interpretation as in Strategy 1 and Strategy 2.

The identifying assumption is that roundtrip fares are approximately additive in latent directional segment prices. With two outbound flights X and X' and two return flights Y and Y', the observable implication is

$$P_{X,Y} - P_{X',Y} = P_{X,Y'} - P_{X',Y'}. \tag{11}$$

This condition matters because Strategy 3 combines seven-day and 21-day tickets. If the two stay durations imply different fare restrictions or ticket rules, the synthetic price could partly reflect tariff structure rather than the latent price of the middle flight pair. Table 3 illustrates the additivity condition using British Airways fares for the Barcelona (BCN)–London Heathrow (LHR) corridor. The example was collected from the British Airways website on May 15, 2026 and is used only to demonstrate the observable test (British Airways, 2026).

In this example, both sides of Equation 11 yield the same difference: $(309 - 276 = 419 - 386 = 33)$ EUR. In the empirical implementation, I allow a 2% tolerance in this additivity test because the scraped Skyscanner lead price can differ slightly from airline-website prices when online travel agencies apply

Table 3: Example of Prices that Support the Additivity Assumption

British Airways: Barcelona (BCN) to London Heathrow (LHR) Departing May 18, 2026, returning May 25, 2026; accessed May 15, 2026				
Itinerary	**Outbound Departure time**	**Inbound Departure time**	**Airfare**	**Airfare difference**
1	(X): BA475, 12:40 PM	(Y): BA482, 09:10 AM	309 EUR	
2	(X'): BA477, 4:35 PM	(Y): BA482, 09:10 AM	276 EUR	33 EUR
3	(X): BA475, 12:40 PM	(Y'): BA480, 3:05 PM	419 EUR	
4	(X'): BA477, 4:35 PM	(Y'): BA480, 3:05 PM	386 EUR	33 EUR

discounts, service fees, or platform-specific markups. This tolerance is documented in Appendix D. The direct additivity test requires at least two flights per day on the same directional route. For airline-route observations with only one daily flight, the test cannot be implemented directly; in these cases, fare-basis codes were checked manually in ITA Matrix for a subsample of seven-day and 21-day tickets (ITA Matrix, 2026). Appendix F explains this check.

4.2.4 Empirical Use of the Three Outcome Strategies

The three outcome strategies are complementary and differ in the degree to which they control for flight-specific costs, seat availability, fare rules, and time-varying route conditions.

Strategy 1 provides the cleanest static comparison with respect to flight-specific cost and inventory heterogeneity. The chain construction links mirrored roundtrips through the same physical flight segments, except for the boundary links at the beginning and end of the chain. As a result, flight-specific operating costs, such as fuel, crew, airport charges, aircraft utilization, and unobserved seat availability at the time of the data request are held constant to a large extent. Since all compared tickets are seven-day roundtrips, the stay duration and the main fare-rule structure are also kept constant. The limitation of Strategy 1 is that prices are aggregated over the full chain. Consequently, days-before-departure effects and other time-varying controls, such as school holidays, cannot be identified at the route-day level.

Strategy 2 provides greater empirical flexibility. It does not compare passengers who travel on the same physical aircraft, so unobserved seat inventory and direction-specific demand shocks may still differ across the two matched roundtrips. However, it matches route, airline, travel dates, booking horizon, and flight numbers. Identical flight numbers are important because they capture the same scheduled service pattern and therefore control for schedule-related characteristics, including time of day, day of week, slot conditions, and peak-hour airport-cost components more tightly than an unmatched directional fare comparison. As in Strategy 1, all tickets are seven-day roundtrips, so the main stay-duration and fare-rule structure is kept constant. Because Strategy 2 is defined at the matched route-day level, it can be used for booking-horizon groups and for time-varying controls such as holidays. Its limitation is that remaining seat inventory may differ across the two physical aircraft. In addition, direction-specific demand shocks can still affect one side of the matched pair. Such shocks are relevant in airline pricing because event-related demand can raise fares on the affected directional route without implying income-based price discrimination.[3]

Strategy 3 provides a flight-component comparison by using the synthetic additivity logic introduced

[3]For example, if Atletico Madrid plays Arsenal in London on May 5, demand for flights from Madrid to London may increase in one direction because supporters travel to the match. This type of one-sided demand shock is conceptually related to Piga et al. (2025), who study UEFA Euro 2016 and show that flights suitable for transporting supporters to match locations experienced substantial fare adjustments after the tournament draw revealed the relevant fixtures.

above. Because the observed and synthetic tickets isolate the same middle flight components, operating costs are held constant and the unobserved seat-inventory concern is substantially reduced. The main limitation is that this design requires the additivity assumption across seven-day and 21-day roundtrips. Moreover, it does not cover the final one to seven days before departure because the 21-day construction must be observed around the middle flight pair.

For the subsequent empirical analysis, the three strategies are used according to these identification properties. For the static specifications, Strategy 1 and Strategy 3 are preferred because both control most directly for flight-specific operating costs and seat inventory. Strategy 1 is not suitable for the dynamic booking-horizon analysis because its outcome is aggregated over the full chain. The dynamic specifications therefore rely on Strategy 2 and Strategy 3. Among these, Strategy 3 is preferred for the booking windows it covers because its synthetic construction isolates the same physical flight components, whereas Strategy 2 matches comparable but physically distinct flights. Strategy 2 nevertheless remains important because, unlike Strategy 3, it also covers the final booking window of one to seven days before departure.

4.3 Descriptive Statistics

The raw dataset obtained from the API contained 45,970,752 ticket observations. Since the API constructed ticket prices as combinations of available outbound and return flights, direct descriptive statistics on the raw Cartesian product would be misleading. The data also included marketed flights operated by partner airlines. Therefore, the descriptive statistics are reported only after applying the filtering and matching logic described in Section 4.2. Except for Table 4, the descriptive statistics are based on Strategy 2 because it preserves daily variation and provides the largest matched sample.

Table 4 summarizes the distribution of the strategy-specific outcome variables after applying the outlier restriction. Each outcome is defined as the ticket price for the higher-income origin (A) minus the ticket price for the lower-income origin (B), calculated according to the matching logic of the respective strategy. After excluding the lower and upper five percent of the price-difference distribution, the mean price difference is positive in all strategies: 6.51 USD in Strategy 1, 7.33 USD in Strategy 2, and 6.25 USD in Strategy 3. The positive means indicate that, in the restricted sample, tickets originating in the higher-income endpoint are on average more expensive across all three outcome constructions. All regression estimates in the main analysis are based on this restricted sample. This restriction is necessary because a small number of matched tickets in the underlying raw data display extreme price differences of several thousand US dollars. Such observations may reflect unusual fare-class availability, platform-specific price displays, rare demand shocks, or data artifacts, and they would otherwise exert disproportionate influence on the estimated income premium.

Table 4: Matched Sample and Price-Difference Distribution (Restricted Sample)

Strategy	Sample	Observations	Mean	Std. dev.	Minimum	Maximum
Strategy 1	Restricted sample	4,460	6.51	38.35	-107.09	105.03
Strategy 2	Restricted sample	115,540	7.33	101.15	-263.20	293.00
Strategy 3	Restricted sample	36,113	6.25	47.15	-142.26	143.00

Table 5 shows the distribution of effective competition in the matched sample using Strategy 2. The largest group consists of routes with two effective competitors, accounting for 51.5 percent of the matched observations. Monopoly routes account for 11.9 percent, while routes with three or more effective competitors account for 36.6 percent. Detailed route-level information is reported in Section B.

Table 5: Effective Competition in the Matched Sample

Competition group	Observations	Share
1 effective airline	15,272	11.9%
2 effective airlines	66,089	51.5%
3 or more effective airlines	47,004	36.6%
Total	128,365	100.0%

Table 6 reports the role of the five main hub carriers in the dataset. British Airways, Lufthansa, Air France, KLM, and Iberia together account for approximately 53.29 percent of all observations. The remaining observations are distributed across other European, international, and low-cost carriers. Legacy carriers account for roughly 81 percent of the matched sample, while non-legacy carriers account for roughly 19 percent. This composition is relevant because pricing behavior differs systematically between legacy and low-cost carriers, which motivates the inclusion of carrier-type controls in the competition specifications.

Table 6: Main Hub Carriers in the Sample

Carrier	Observations	Share
British Airways	17,267	13.45%
Air France	14,681	11.44%
Lufthansa	14,484	11.28%
KLM	11,575	9.02%
Iberia	10,400	8.10%
Five carriers combined	68,407	53.29%
Other carriers	59,958	46.71%
Total	128,365	100.00%

Figure 3 compares average matched roundtrip prices in the restricted Strategy 2 sample across booking-horizon groups. The higher-income-origin roundtrip is more expensive in most booking windows, namely 50–89, 29–49, 15–28, and 8–14 days before departure. Moreover, the descriptive fare gap increases up to the 8–14 day window, which is consistent with the idea that the income premium becomes stronger as departure approaches. However, this pattern breaks down in the final 1–7 day window, where the lower-income-origin roundtrip is more expensive. The figure therefore provides descriptive motivation for H1 and H2, but it is not sufficient to establish whether the income premium increases systematically across the booking horizon. The formal tests are provided by the regression models below.

5 Empirical Strategy

This section presents the econometric framework used to test whether airlines engage in income-based price discrimination across the two directions of the same international nonstop route. The dependent variable is denoted by Y_i^s, where $s \in \{1, 2, 3\}$ indicates the empirical strategy used to construct the directional price difference. The construction and empirical use of Strategy 1, Strategy 2, and Strategy 3 are described in Section 4.2. In all cases, Y_i^s measures the fare difference between the higher-income origin A and the lower-income origin B. All baseline and interaction regressions are estimated on the restricted

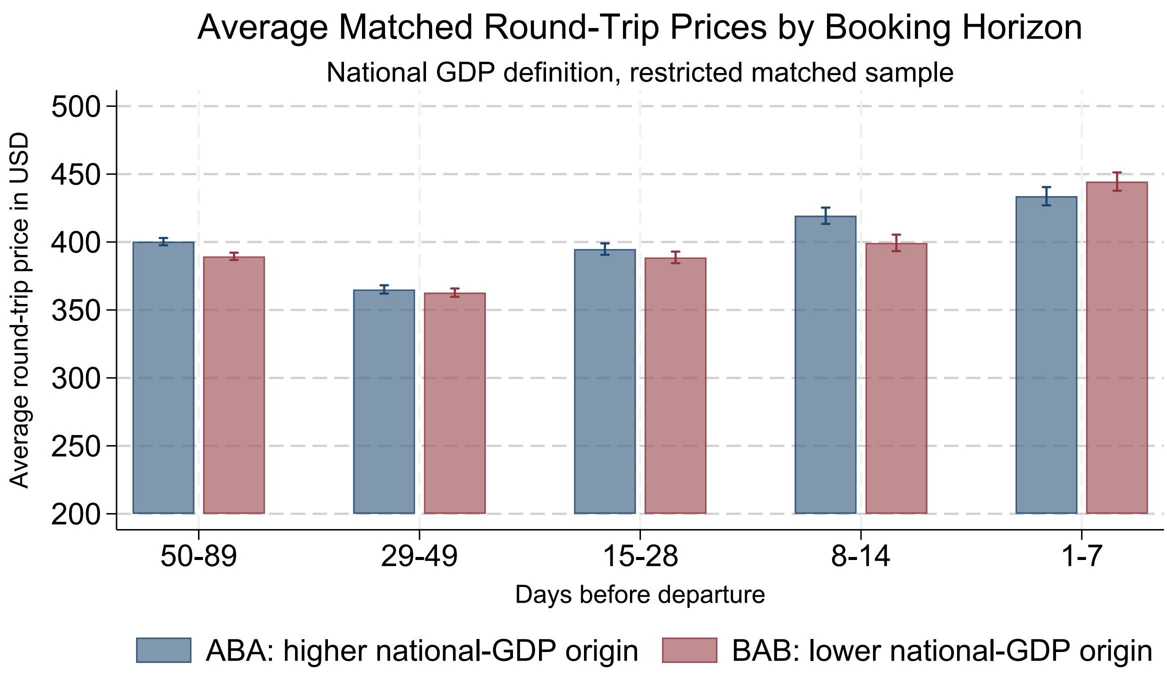

Figure 3: Average Matched Roundtrip Prices by Booking Horizon (Strategy 2)

sample, which excludes the lower and upper five percent of the strategy-specific outcome distribution.

The baseline model estimates the relationship between income differences and directional fare differences. The main explanatory variable is ΔInc_i, defined as income in A minus income in B. In the main specification, income is measured by national GDP per capita. Regional TL2 GDP per capita is used as an alternative income definition in the robustness analysis. The same directional logic is applied to the continuous control variables, which are also defined as differences between A and B. The baseline specification is

$$Y_i^s = \alpha + \beta_1 \Delta Inc_i + \lambda' X_i^s + \varepsilon_i^s, \qquad s \in \{1, 2, 3\}. \tag{12}$$

Here, X_i^s contains population differences, carrier-specific hub differences, temperature differences, and holiday differences between the two route endpoints. These variables control for alternative explanations of directional fare differences that may be correlated with income. The population, hub, and temperature controls follow the empirical logic of Luttmann (2019), while the holiday variable is added to capture temporary differences in travel flexibility and peak-period demand conditions.

$DPOP\_mio_i$ is defined as the population of city A minus the population of city B, measured in millions. It controls for differences in market size and service availability. While flight frequency on the specific route between A and B is mechanically the same in both directions, passengers departing from the larger endpoint may benefit from a broader airport network, more overall flight options, and better schedule connectivity. This can increase the probability of finding a preferred departure time or convenient connection, making demand less price elastic at the more populous endpoint. Airlines may therefore charge higher average fares to passengers originating from the larger city.

$DHUB_i$ measures the difference in carrier-specific hub status. It equals one if city A is a hub for the operating airline, minus one if city B is a hub for the operating airline, and zero otherwise. This control is important because carriers may charge a premium for roundtrips originating at their own hub, where they have stronger airport presence, better network connectivity, and a more captive customer base. Hub dominance and airport presence have been shown to increase fares, making it necessary to separate hub-related market power from the income effect (Borenstein, 1989; Lijesen et al., 2001).

$DTEMP_i$ is the temperature difference between city A and city B. It controls for differences in travel purpose and substitution possibilities. Brueckner et al. (2013) use a similar temperature-difference variable to identify leisure-oriented markets. The intuition is that passengers traveling from colder to warmer destinations may be more leisure-oriented and have several substitute destinations, making their demand more price elastic. In the opposite direction, passengers may be more likely to travel for specific reasons such as business, family visits, or fixed appointments, which reduces flexibility.

$DHOLIDAYS_i$ measures holiday differences between the two endpoints. It equals one if city A is in a holiday period and city B is not, minus one if city B is in a holiday period and city A is not, and zero if neither or both cities are in a holiday period. This variable controls for temporary differences in travel flexibility. If only one endpoint is in a holiday period, passengers departing from that endpoint may face a narrower feasible travel window and therefore lower price elasticity. Evidence on holiday-related peak-load pricing supports the inclusion of this control (Gaggero and Luttmann, 2025).

6 Main Results

The baseline test of H1 estimates whether income differences are systematically associated with directional fare differences across the three outcome definitions. The analysis then checks whether the result is robust to a regional income measure and to route-distance and carrier-type subsamples.

6.1 Baseline Results

Table 7 reports the baseline model results across the three identification strategies. The coefficient on Δ Income is positive and statistically significant at the 1% level in all specifications, supporting H1 and the positive price-gap prediction in Equation 4. Since income is measured in units of 10,000 USD, the estimates imply that a 10,000 USD increase in the income difference raises the directional fare premium by 3.00 USD in Strategy 1, 3.43 USD in Strategy 2, and 3.37 USD in Strategy 3. This magnitude is highly consistent with Luttmann (2019), who finds a fare increase of 1.8 to 4.3 USD per 10,000 USD income difference. Economically, this implies that passengers originating in richer cities systematically pay higher prices for the same or a closely matched roundtrip route than passengers starting in the poorer endpoint. This interpretation is particularly robust for Strategy 1 and Strategy 3, which control most directly for flight-specific operating costs and seat inventory.

The auxiliary controls do not overturn this interpretation. The hub coefficient is positive and statistically significant in all three strategies, ranging from 5.24 USD in Strategy 2 to 6.69 USD in Strategy 3. This indicates that passengers starting their roundtrip at the operating airline's hub pay higher fares than passengers starting at the non-hub endpoint. This can also be interpreted as a form of directional price discrimination, but one that is based on carrier-specific airport dominance rather than regional income. A hub origin gives the airline stronger local market power because passengers at that endpoint are more likely to value the carrier's dense network, schedule convenience, loyalty program, and connecting options. These passengers may therefore have fewer attractive substitutes from the perspective of the operating airline, allowing the carrier to charge a hub-origin premium. Population differences are slightly negative,

Table 7: Baseline Estimates Across Identification Strategies

	S1	S2	S3
Identification strategy	Aggregated Chains	Daily Pair Matching	Synthetic Matching
Δ Income, \$10k	2.998***	3.428***	3.373***
	(0.344)	(0.181)	(0.156)
Δ Population, million	-0.077	-0.102**	-0.084**
	(0.093)	(0.049)	(0.043)
Δ Hub	5.832***	5.240***	6.693***
	(0.698)	(0.356)	(0.306)
Δ Temperature	-0.737***	-0.151**	0.005
	(0.135)	(0.069)	(0.054)
Δ Holidays	–	15.337***	-2.069***
	–	(0.622)	(0.512)
Constant	-4.300***	-1.574***	-2.654***
	(0.985)	(0.520)	(0.413)
Observations	4,460	115,540	36,113
R^2	0.0526	0.0107	0.0274

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Restricted samples exclude the bottom and top 5% of the outcome variable. Holiday differences are not included in S1 because the Strategy 1 outcome is not defined at the route-day level.

and temperature differences are negative in two specifications but insignificant in Strategy 3. Holiday effects are mixed across the two strategies where they can be included.

6.2 Heterogeneous Analysis

This section examines whether the positive income coefficient from the baseline model remains stable across alternative definitions and subsamples. I consider three robustness and heterogeneity checks: replacing national GDP per capita with regional TL2 GDP per capita, splitting the sample by route distance, and estimating the baseline model separately for legacy and non-legacy carriers. These exercises do not introduce separate hypotheses. They assess whether the baseline evidence for H1 depends on a particular income measure, distance segment, or carrier business model. The analysis focuses on Strategy 1 and Strategy 3 because they provide the strongest static identification: passengers are assigned either physically or synthetically to the same aircraft components.

6.2.1 Regional Measure of Income

This subsection replaces national GDP per capita with the OECD TL2 regional income definition. The motivation is that national GDP per capita may hide substantial within-country income differences. This is relevant for international routes because passengers do not originate from an abstract national average, but from a specific regional market around the departure airport. TL2 income therefore provides a more geographically precise measure of local purchasing power and local willingness to pay. In this specification, all directional variables are redefined according to the richer TL2 region: the outcome variable, the income difference, and the control variables are consistently measured as richer TL2 region minus poorer TL2 region.

Table 8 reports the results using the regional TL2 measure. Since these data are only available for a smaller set of countries, the number of usable routes decreases from 110 to 84. The income coefficient remains positive and statistically significant in both strategies: a 10,000 USD increase in the TL2 income difference raises the directional fare premium by 2.12 USD in Strategy 1 and 2.34 USD in Strategy 3.

Table 8: Robustness Check: Regional TL2 Income Definition

	S1	S3
Identification strategy	Aggregated Chains	Synthetic Matching
Δ Income, 10k, TL2	2.123*	2.341**
	(0.383)	(0.179)
Δ Population, million, TL2	−0.583*	−0.639**
	(0.109)	(0.043)
Δ Hub, TL2	4.592***	7.061***
	(0.757)	(0.316)
Δ Temperature, TL2	0.039	0.461***
	(0.174)	(0.079)
Δ Holidays, TL2	–	−3.936***
	–	(0.516)
Constant	−0.688	−2.467***
	(1.141)	(0.486)
Observations	3,833	32,530
R^2	0.0310	0.0302

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Restricted samples exclude the bottom and top 5% of the outcome variable. Holiday differences are not included in S1 because the Strategy 1 outcome is not defined at the route-day level.

Economically, this means that passengers starting in the richer regional endpoint pay more for the same or closely matched roundtrip product, even when income is measured closer to the airport market rather than at the national level. The estimates are smaller than in the national GDP specification, but the sign and statistical significance remain unchanged. This strengthens the interpretation of the baseline result: the estimated income premium is not only driven by broad cross-country income differences, but is also visible when the richer and poorer endpoints are identified using regional income. National GDP per capita remains preferable for the full sample because it avoids losing non-OECD routes, but the TL2 specification supports the conclusion that origin-market income is related to directional pricing.

6.2.2 Difference Between Long-Haul and Short-Haul Routes

This subsection splits the sample by route distance. Long-haul routes are defined as routes with a distance of at least 1,200 km, while short-haul routes are below 1,200 km. This distinction is economically relevant because the set of outside options differs strongly by distance. On short-haul routes, passengers may substitute toward trains, buses, cars, alternative airports, or connecting flights. On long-haul routes, these alternatives are less attractive or not available. The split therefore tests whether the income premium is stronger in markets where passengers have fewer substitutes and airlines may have greater pricing power.

Table 9 shows that the income coefficient is positive and statistically significant in all four specifications. For long-haul routes, a 10,000 USD increase in the income difference raises the directional fare premium by approximately 4.38 USD in Strategy 1 and 5.99 USD in Strategy 3. For short-haul routes, the corresponding effects are smaller, at approximately 2.91 USD in Strategy 1 and 3.20 USD in Strategy 3. Income-based price discrimination therefore appears in both distance groups, but it is clearly stronger on long-haul routes.

The stronger long-haul effect is consistent with the role of substitution possibilities in price discrimination. If passengers have fewer attractive alternatives, their demand is less elastic and airlines have more scope to charge different prices across segmented origin markets. This is likely to be the case on long-haul routes, where passengers cannot easily replace the nonstop flight by rail or road transport and where

Table 9: Robustness Check: Long-Haul and Short-Haul Routes

	Long-haul routes		Short-haul routes	
	S1: Chains	S3: Synthetic	S1: Chains	S3: Synthetic
Δ Income, \$10k	4.379***	5.994***	2.907***	3.203***
	(0.626)	(0.299)	(0.445)	(0.192)
Δ Population, million	-0.071	-0.022	0.086	-0.175***
	(0.116)	(0.061)	(0.141)	(0.056)
Δ Hub	2.575**	3.242***	8.601***	9.126***
	(1.003)	(0.487)	(0.963)	(0.388)
Δ Temperature	-1.044***	-0.212***	-0.119	0.727***
	(0.178)	(0.065)	(0.222)	(0.071)
Δ Holidays	–	-1.326	–	-2.799***
	–	(0.865)	–	(0.566)
Constant	-12.547***	-13.849***	-1.168	0.548
	(1.995)	(0.962)	(1.159)	(0.437)
Observations	2,145	15,256	2,315	20,857
R^2	0.0743	0.0330	0.0495	0.0416

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Long-haul routes are at least 1,200 km; short-haul routes are below 1,200 km. Holiday differences are not included in S1 because the Strategy 1 outcome is not defined at the route-day level.

indirect connections involve substantial time costs. A passenger originating in a high-income long-haul market may therefore be less willing to search extensively or accept inconvenient alternatives. Airlines can exploit this by charging a higher fare premium at the richer origin.

The fact that the income coefficient remains positive on short-haul routes is also informative. It shows that income-based price discrimination is not restricted to markets with very limited outside options. Even where substitutes are more available, passengers starting in richer markets still pay higher fares. However, the smaller coefficient suggests that short-haul competition and intermodal alternatives reduce the scope for discriminatory pricing. In economic terms, outside options do not eliminate the income premium, but they appear to discipline it.

6.2.3 Difference Between Legacy and Non-Legacy Carriers

This subsection splits the sample by carrier type. The purpose is to examine whether income-based price discrimination is mainly driven by traditional network carriers or whether it also appears among non-legacy carriers, including low-cost and charter airlines. This distinction matters because legacy and non-legacy carriers differ in network structure, fare construction, and revenue-management practices. Legacy carriers more often operate hub-and-spoke networks and use complex fare-class systems, while non-legacy carriers more often rely on point-to-point networks and simpler fare structures.

Table 10 presents the results separately for legacy and non-legacy carriers. The income coefficient remains positive and statistically significant for both carrier groups. For legacy carriers, a 10,000 USD increase in the income difference raises the directional fare premium by approximately 3.49 USD in Strategy 1 and 3.98 USD in Strategy 3. For non-legacy carriers, the corresponding effects are smaller, at approximately 2.50 USD in Strategy 1 and 1.68 USD in Strategy 3. Thus, income-based price discrimination is not limited to legacy carriers, but it is more pronounced among them.

This pattern is economically meaningful. Legacy carriers typically operate more complex networks and use more differentiated fare products. Their pricing systems can segment passengers by origin, booking horizon, fare class, roundtrip structure, and network position. This gives them more instruments to translate differences in willingness to pay into directional fare differences. In contrast, non-legacy carriers often rely more heavily on simpler point-to-point pricing and unbundled ancillary revenues. If

Table 10: Robustness Check: Legacy and Non-Legacy Carriers

	Legacy carriers		Non-legacy carriers	
	S1: Chains	S3: Synthetic	S1: Chains	S3: Synthetic
Δ Income, \$10k	3.486***	3.983***	2.503***	1.682***
	(0.448)	(0.192)	(0.379)	(0.199)
Δ Population, million	-0.119	-0.300***	0.783***	0.931***
	(0.107)	(0.049)	(0.138)	(0.070)
Δ Hub	6.148***	7.472***	-0.255	-7.633***
	(0.725)	(0.314)	(2.128)	(1.714)
Δ Temperature	-0.889***	-0.331***	0.650***	0.779***
	(0.150)	(0.062)	(0.197)	(0.088)
Δ Holidays	–	-2.354***	–	-1.344*
	–	(0.604)	–	(0.701)
Constant	-5.906***	-4.720***	0.484	1.637***
	(1.288)	(0.508)	(0.965)	(0.482)
Observations	3,440	28,859	1,020	7,254
R^2	0.0626	0.0360	0.0708	0.0644

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Non-legacy carriers include low-cost and charter airlines. Holiday differences are not included in S1 because the Strategy 1 outcome is not defined at the route-day level.

roundtrip fares are close to, or exactly equal to, the sum of two one-way tickets, there is less or no room to price the same route differently depending on the starting point. The smaller income coefficient for non-legacy carriers is therefore consistent with a pricing model that allows less directional segmentation.

At the same time, the positive and significant coefficient for non-legacy carriers is important. It shows that income-based price discrimination is not exclusively a legacy-carrier phenomenon. Even carriers with simpler fare structures may charge higher fares to passengers starting in richer markets. This can happen through route-level demand forecasts, different fare availability, market-specific competitive conditions, or the timing of fare changes. The result therefore suggests that the income premium is a general feature of airline pricing, although traditional network carriers appear to implement it more strongly.

7 Price Discrimination and Booking Horizon

This section tests H2, which predicts that the income premium becomes larger as the departure date approaches. The economic logic is that passengers differ not only in income, but also in their flexibility over the booking horizon. Far from departure, passengers can still compare alternative airlines, nearby airports, indirect connections, or alternative travel dates. Closer to departure, these options become more limited. If passengers from higher-income origins have a higher opportunity cost of time and lower willingness to adjust their travel plans, their effective price elasticity should fall more strongly as departure approaches. Airlines may therefore be able to extract a larger fare premium from the richer origin in shorter booking windows. To test this mechanism, I extend the baseline specification by interacting the income difference with booking-horizon groups. Let DBD_{ki} denote a dummy variable for booking-horizon group k. The dynamic specification is

$$Y_i^s = \alpha + \beta_1 \Delta Inc_i + \sum_k \gamma_k DBD_{ki} + \sum_k \delta_k \left(\Delta Inc_i \cdot DBD_{ki}\right) + \lambda' X_i^s + \varepsilon_i^s. \tag{13}$$

This model allows the income coefficient to differ across booking-horizon groups. The omitted category is the longest booking horizon, 50–91 days before departure. The coefficients on the interaction terms therefore show whether the income premium in shorter booking windows differs from the income premium far from departure. The main interpretation is based on the marginal effects of ΔInc_i, which report the income effect within each booking-horizon group. If H2 is correct, these marginal effects should be larger in shorter booking windows than in the longest booking-horizon group.

Table 11: Interaction Effect: Income Difference and Booking Horizon

	S2: Daily Pair	S3: Synthetic
Regression coefficients		
Δ Income, $10k	4.093***	2.457***
	(0.248)	(0.243)
1–7 days before departure	-15.954***	–
	(2.439)	–
8–14 days before departure	6.633***	-10.765***
	(2.334)	(1.654)
15–28 days before departure	-1.195	-10.347***
	(1.596)	(1.144)
29–49 days before departure	-2.590**	-5.644***
	(1.139)	(0.976)
Δ Income × 1–7 days	-2.605***	–
	(0.819)	–
Δ Income × 8–14 days	1.264	1.691***
	(0.792)	(0.596)
Δ Income × 15–28 days	-1.828***	1.959***
	(0.530)	(0.413)
Δ Income × 29–49 days	-2.044***	1.106***
	(0.384)	(0.354)
Δ Population, million	-0.154***	-0.065
	(0.050)	(0.044)
Δ Hub	5.080***	6.667***
	(0.359)	(0.316)
Δ Temperature	-0.071	-0.023
	(0.070)	(0.056)
Δ Holidays	15.525***	-1.455***
	(0.630)	(0.562)
Constant	-0.025	2.637***
	(0.699)	(0.663)
Marginal effect of Δ Income by booking horizon		
1–7 days before departure	1.487*	–
	(0.787)	–
8–14 days before departure	5.356***	4.149***
	(0.758)	(0.549)
15–28 days before departure	2.264***	4.417***
	(0.478)	(0.343)
29–49 days before departure	2.048***	3.563***
	(0.311)	(0.268)
50–91 days before departure	4.093***	2.457***
	(0.248)	(0.243)
Observations	115,540	36,113
R^2	0.0150	0.0311

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. S2 = daily pair matching; S3 = synthetic matching. The reference category is the longest booking horizon, 50–91 days before departure. The marginal effects report the effect of a $10,000 income difference by booking-horizon category. Restricted samples exclude the bottom and top 5% of the outcome variable.

Table 11 reports the interaction between income differences and booking-horizon groups. The interpretation differs from a simple comparison of average fares across booking windows. The empirical question is whether the sensitivity of the directional fare premium to income differences becomes stronger when departure is closer. This corresponds to the theoretical prediction in Equation 5: the price gap should widen as firm-level price sensitivity falls more strongly in the high-income origin.

In Strategy 3, the preferred specification for the dynamic test, the marginal effects are larger in all shorter booking windows than in the longest booking window. A 10,000 USD increase in the income difference raises the directional fare premium by 2.46 USD at 50–91 days before departure, 3.56 USD at 29–49 days, 4.42 USD at 15–28 days, and 4.15 USD at 8–14 days. The pattern is not perfectly monotonic, but it supports H2 by showing that the income premium becomes substantially larger once departure is closer and passenger flexibility declines. The strongest increase appears between the longest booking horizon and the middle booking windows, which is consistent with the idea that airlines can extract more surplus once passengers have fewer attractive opportunities to adjust their itinerary.

Strategy 2 adds the 1–7 day booking window, which is not available in Strategy 3. The evidence is less monotonic in this specification: the marginal effect is largest in the 8–14 day window and smaller in the final week. This does not invalidate the dynamic interpretation, but it makes it more nuanced. In the last week before departure, remaining inventory, fare-class closures, and route-specific demand shocks may dominate the income-gradient mechanism. The final-week coefficient should therefore be interpreted cautiously, especially because Strategy 2 does not compare the same physical aircraft. By contrast, the 8–14 day window still shows a strong income premium while preserving a more regular matching structure. Overall, the dynamic evidence supports H2 most clearly in the synthetic specification, while the daily matching specification provides a broader but less clean view of the final booking window.

8 Price Discrimination and Competition

This section tests H3, which predicts that effective competition reduces the income premium. The economic logic is that additional competitors increase passengers' outside options and raise the aggregate cross-price elasticity on a route. If passengers can switch more easily to another airline, the operating carrier should have less scope to charge a higher fare to passengers originating in the richer endpoint. Competition should therefore weaken income-based price discrimination, although it may not eliminate it completely if products remain differentiated, airports are slot constrained, or airlines coordinate through alliances and joint ventures.

To test this mechanism, I extend the baseline specification by interacting the income difference with the number of effective competing airlines on the route. The competition specification is

$$Y_i^s = \alpha + \beta_1 \Delta Inc_i + \beta_2 NComp_i + \beta_3 \left(\Delta Inc_i \cdot NComp_i\right) + \lambda' X_i^s + \delta' C_i + \varepsilon_i^s. \tag{14}$$

Here, $NComp_i$ denotes the number of effective competitors on the route as defined in the data section. The control structure consists of the baseline controls X_i^s and the additional competition-related controls C_i, which capture substitution possibilities and institutional features that may affect pricing power beyond the number of effective competitors. $Legacy_i$ equals one if the operating airline is a member of one of the three global airline alliances, oneworld, SkyTeam, or Star Alliance, and zero otherwise.

$AllianceRoute_i$ equals one if the route connects two hub airports served by airlines belonging to the same alliance, with one alliance carrier having a hub at city A and another at city B. This variable controls for the possibility that nominally separate airlines may not compete fully on such routes. Alliance cooperation, antitrust immunity, or joint ventures can reduce effective competition because partner airlines may coordinate pricing and scheduling decisions. Brueckner and Singer (2019) show that, on gateway-to-gateway routes, alliance cooperation can be equivalent to removing a competitor and can therefore increase fares.

HST_i equals one if a direct high-speed train connection operates between city A and city B. It directly captures intermodal competition. High-speed rail can be a close substitute for air travel on short- and medium-haul routes and may therefore limit airlines' ability to charge higher fares. Evidence from the Madrid–Barcelona corridor and from German domestic routes shows that high-speed rail can reduce air demand (Román et al., 2007; Oesingmann and Ennen, 2025).

$AltAirport_i$ equals one if at least one alternative nonstop route between the two cities is available through another airport in the same metropolitan area. This captures additional competitive pressure because passengers may switch to nearby airports and airlines may shift capacity across airports when market conditions change. Thelle and Sonne (2018) emphasize that airline and passenger switching possibilities are central competitive constraints in European airport markets.

Finally, $LiberalRoute_i$ equals one if airlines from the relevant countries can offer unlimited frequencies between city A and city B, and zero otherwise. This variable controls for the legal openness of the route. If price discrimination is observed on a restricted route, the resulting profits may persist because entry by additional airlines is legally constrained. In liberalized or Open-Skies-type markets, entry and capacity restrictions are reduced, increasing the threat of entry and limiting the persistence of excessive markups. Button (2009) describes how Open Skies arrangements remove capacity and pricing restrictions and thereby shift international aviation toward a more competitive market environment.

Table 12: Interaction Effect: Income Difference and Market Competition

	S1: Chains	S3: Synthetic
Regression coefficients		
Δ Income, \$10k	4.364***	6.531***
	(0.691)	(0.435)
Number of competing airlines	2.165***	4.593***
	(0.541)	(0.363)
Δ Income × Number of competing airlines	-0.287*	-1.154***
	(0.150)	(0.184)
Δ Population, million	0.082	-0.128***
	(0.100)	(0.044)
Δ Hub	5.078***	6.977***
	(0.705)	(0.307)
Δ Temperature	-0.798***	0.090
	(0.144)	(0.057)
Δ Holidays	–	-2.144***
	–	(0.509)
High-speed rail	4.771***	6.760***
	(1.590)	(0.685)
Alternative airport	-2.463*	9.340***
	(1.276)	(0.595)
Legacy carrier	-1.823	2.888***
	(1.116)	(0.467)
Alliance route	3.130*	8.101***
	(1.777)	(0.786)
Liberalized route	-6.623***	0.991
	(2.205)	(1.120)
Constant	-6.357*	-26.159***
	(3.258)	(1.757)
Marginal effect of Δ Income by market competition		
1 Airline	3.849***	5.377***
	(0.479)	(0.276)
2 Airlines	3.338***	4.223***
	(0.356)	(0.176)
3+ Airlines	2.509***	2.513***
	(0.515)	(0.297)
Observations	4,460	36,113
R^2	0.0638	0.0416

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. S1 = aggregated chain strategy; S3 = synthetic matching. The marginal effects report the effect of a \$10,000 income difference by number of effective competing airlines. Holiday differences are not included in S1 because the Strategy 1 outcome is not defined at the route-day level. Restricted samples exclude the bottom and top 5% of the outcome variable.

Table 12 reports the results of the interaction model between income differences and market competition. The interaction coefficient is negative in both specifications: an additional effective competitor reduces the directional income premium by 0.29 USD in Strategy 1 and by 1.15 USD in Strategy 3. This mitigation effect is statistically significant at the 10% level in Strategy 1 and at the 1% level in Strategy 3, supporting

H3.

The marginal effects provide the clearest interpretation and map directly into the theoretical condition in Equation 7. In Strategy 1, a 10,000 USD increase in the income difference raises the directional fare premium by 3.85 USD on monopoly routes, 3.34 USD on duopoly routes, and 2.51 USD on routes with three or more effective competitors. Strategy 3 shows the same pattern, but with a stronger decline: the corresponding marginal effects are 5.38 USD, 4.22 USD, and 2.51 USD. Competition therefore disciplines income-based price discrimination, but does not eliminate it. Even on routes with at least three effective competitors, the income premium remains positive and statistically significant.

Economically, the decline in the income premium means that an additional independent competitor reduces the airline's ability to separate the two origin markets by willingness to pay. A passenger from the high-income origin becomes less captive when another independent carrier offers a credible substitute. The fact that the premium remains positive even with three or more effective competitors shows that competition in airline markets is not equivalent to homogeneous-product Bertrand competition. Carriers remain differentiated by departure times, airport convenience, loyalty programs, alliance networks, corporate contracts, and service attributes. Competition therefore compresses the directional income premium, but persistent product differentiation and entry barriers prevent it from eliminating the premium entirely.

The competition-related controls are also informative for the interpretation of market structure. First, the alliance-route coefficient is positive in both specifications and statistically significant, although more strongly in Strategy 3 than in Strategy 1. This suggests that alliance structures are associated with higher directional fare premia, consistent with the idea that cooperation between alliance partners can weaken independent competitive pressure even when several airline brands are visible to passengers. Second, the liberalized-route coefficient is negative in Strategy 1 and statistically insignificant in Strategy 3. The negative sign in Strategy 1 is consistent with the interpretation that liberalized routes reduce pricing power because additional airlines can enter the market, expand frequencies, and thereby compete away part of the profits generated by directional price discrimination. The weaker result in Strategy 3 indicates that liberalization alone is not sufficient to eliminate directional fare premia once route-specific product differentiation, airport capacity constraints, and alliance structures are taken into account. Among the remaining competition controls, high-speed rail is positively associated with fare premia in both strategies, while alternative airports and legacy-carrier status are positive mainly in Strategy 3. These coefficients should not be interpreted as separate causal effects, but they show that additional substitution channels and carrier characteristics are correlated with the structure of directional pricing.

The adjusted-HHI robustness check in Appendix G supports the same competition mechanism. Unlike the number of effective competitors, the adjusted HHI also accounts for unequal flight shares and economically integrated carriers. The results show that the income premium is larger on more concentrated routes, consistent with the main finding that the premium decreases as effective competition increases. Appendix H further shows that this disciplining effect is strongest in booking windows where passengers still have enough time to switch across competing airlines. Together, these robustness checks indicate that competition limits income-based directional pricing, but only where passengers face sufficiently strong and usable alternatives.

These results show that competition constrains the income premium, but they also raise the question why profitable directional fare differences persist instead of being fully competed away through entry or substitution. A first explanation is hub dominance combined with slot scarcity. The positive hub coefficients across the main specifications suggest that dominant carriers at airports such as LHR, FRA, CDG, AMS, and MAD enjoy route-specific market power. This is consistent with the slot-allocation problem described by Starkie (1998): when airport capacity is scarce and slots are allocated administratively rather than through market-clearing prices, incumbents retain valuable access rights and can earn rents from scarce peak-hour slots. For a potential entrant, high fares alone are not sufficient to make entry attractive; the entrant also needs access to commercially viable slots. Without peak-time access, it cannot offer a schedule that is competitive for business travelers or other time-sensitive passengers.

A second explanation is that airline products are not perfect substitutes. Network carriers combine nonstop service with connecting options, corporate contracts, alliance benefits, and frequent-flyer programs. Lederman (2008) shows that frequent-flyer programs can explain a substantial part of the hub premium because they create lock-in effects that reduce passengers' willingness to switch carriers. Thus, a low-cost entrant may be price competitive but still unable to discipline the incumbent if passengers value loyalty benefits, lounge access, status, and alliance-wide redemption opportunities.

A third explanation is regulatory entry barriers. This mechanism is consistent with the liberalization coefficient discussed above, but it becomes most visible in bilateral third-country markets where entry may be legally capped. The FRA–DXB route illustrates this mechanism: access between Germany and the United Arab Emirates is shaped by bilateral traffic-rights restrictions, and UAE carriers have historically been limited to four German destination points (Emirates, 2026). In such markets, directional price discrimination can persist not only because entry is costly, but because additional entry may simply not be legally permitted.

9 Conclusion and Policy Implications

This paper investigates whether airlines engage in income-based directional price discrimination on international nonstop routes. Using real-time fare data from five major European hub airports, the analysis constructs directional price differences through three empirical strategies. Strategy 1 links passengers across synchronized flight chains, Strategy 2 matches mirrored daily flight pairs with identical flight numbers and travel dates, and Strategy 3 constructs synthetic counterfactual prices from the same underlying flight components. These strategies make it possible to interpret observed price differences as directional fare differences conditional on the matching design and the controls, rather than as simple unmatched price gaps.

The results provide strong evidence for income-based directional price discrimination. In the national GDP specifications, the income coefficient is positive across all three empirical strategies: a 10,000 USD increase in the income difference raises the directional fare premium by approximately 3.00 to 3.43 USD. The effect also remains positive when TL2 regional income is used as an alternative income measure. The heterogeneous analysis confirms this pattern: the income premium appears in both long-haul and short-haul markets, but is larger on long-haul routes. It is also present for both legacy and non-legacy carriers, although it is stronger among legacy carriers. These results support H1, which implies that passengers originating in higher-income markets pay systematically higher fares than passengers originating in lower-income markets.

The paper also presents several extensions that show the mechanisms behind this result. First, the booking-horizon interaction shows that the income premium tends to be larger in shorter booking windows, especially in the synthetic specification. This supports H2 and suggests that income-based price discrimination has a dynamic component: as departure approaches, passengers become less flexible, their search options shrink, and airlines can charge higher directional income premia to passengers from richer origins. Second, the competition interaction shows that income premia decrease as the number of effective competitors increases, supporting H3. In Strategy 3, for example, the marginal effect of a 10,000 USD income difference falls from 5.38 USD on monopoly routes to 2.51 USD on routes with three or more effective competitors. Competition therefore moderates income-based price discrimination, but does not eliminate it.

The analysis has limitations that can be addressed in future research. First, the data cover a six-month period. Although this window includes both winter and summer schedules, a longer observation period would be desirable to test whether the estimated income premium is stable across seasons, years, and broader demand conditions. Second, the booking-horizon analysis is cross-sectional and does not follow the same aircraft, itinerary, and fare product repeatedly over time; a true panel dataset would therefore be valuable for the dynamic analysis. Third, the dataset uses the minimum fare displayed on Skyscanner, which may differ from airline-website prices if intermediaries offer discounts, service fees, or platform-specific markups. Future research could use multi-year fare panels, direct airline-website prices, and richer information on fare classes and remaining seat inventory. Such data would make it possible to separate more precisely whether directional fare premia arise through posted fare menus, fare-class availability, or dynamic seat protection rules. It would also allow researchers to test whether the same route changes its income premium after entry, exit, alliance changes, or slot reallocations.

The policy conclusion is straightforward. Because the income premium declines with competition, the main policy priority is to keep route-level competition effective rather than merely nominal. Merger control, alliance review, slot allocation, and bilateral traffic rights should therefore be evaluated according to whether they preserve independent entry and credible substitution possibilities. This does not imply that dynamic pricing should be prohibited. It implies that institutional barriers should not allow carriers to protect origin-market power from competitive pressure. Bilateral air service agreements that increase frequencies, destination points, or reciprocal market access can improve contestability because liberalization reduces capacity restrictions and opens route access. Airport capacity expansion can also lower entry barriers at congested hubs, although the debate around the Heathrow third runway illustrates how politically and socially difficult this channel can be (The Economist, 2026).

A Detailed Mathematical Derivations and Proofs

This appendix provides the explicit mathematical proofs, demand-function derivations, optimal pricing rules, and theoretical foundations of the research hypotheses discussed in the main text. The derivations build on the representative-consumer framework of Singh and Vives (1984) and on the decomposition of firm-level elasticity into industry-demand and cross-price components in Holmes (1989).

A.1 Utility Maximization and Direct Demand Derivation

A representative consumer in market $k \in \{A, B\}$ maximizes a quasi-linear utility function. The consumer allocates regional income Y_k across air travel from N possible airlines and a composite numeraire good m_k. The budget constraint is $m_k = Y_k - \sum_{j=1}^{N} p_{j,k} q_{j,k}$. Substituting this into the utility function yields the unconstrained Lagrangian formulation:

$$\max_{q_{1,k},\ldots,q_{N,k}} \mathcal{L} = V(q_{1,k}, \ldots, q_{N,k}; Y_k) + Y_k - \sum_{j=1}^{N} p_{j,k} q_{j,k} \tag{15}$$

To find the optimal consumption of flights from airline i, the first derivative of the Lagrangian with respect to the quantity $q_{i,k}$ is set to zero:

$$\frac{\partial \mathcal{L}}{\partial q_{i,k}} = \frac{\partial V}{\partial q_{i,k}} - p_{i,k} = 0 \qquad \Longrightarrow \qquad p_{i,k} = \frac{\partial V}{\partial q_{i,k}} \tag{16}$$

This generates a system of N inverse demand functions where the price is a function of the quantities. Solving this system simultaneously yields the direct demand function for airline i:

$$q_{i,k} = q_{i,k}(p_{i,k}, \mathbf{p}_{-i,k}; Y_k) \tag{17}$$

Airline i maximizes profit across both market segments. Let c be the constant marginal cost and F_i the fixed cost. The maximization problem is:

$$\max_{p_{i,A}, p_{i,B}} \Pi_i = \sum_{k \in \{A,B\}} \left[(p_{i,k} - c) \cdot q_{i,k}(p_{i,k}, \mathbf{p}_{-i,k}; Y_k) \right] - F_i \tag{18}$$

The first-order condition with respect to the own price $p_{i,k}$ requires the application of the product rule:

$$\frac{\partial \Pi_i}{\partial p_{i,k}} = q_{i,k} + (p_{i,k} - c) \frac{\partial q_{i,k}}{\partial p_{i,k}} = 0 \tag{19}$$

Dividing both sides by $p_{i,k} \frac{\partial q_{i,k}}{\partial p_{i,k}}$ and rearranging establishes the standard inverse elasticity rule:

$$\frac{p_{i,k} - c}{p_{i,k}} = \frac{1}{e^F_{i,k}} \tag{20}$$

Solving the equation for $p_{i,k}$ gives the optimal price:

$$p^*_{i,k} = c \left(\frac{e^F_{i,k}}{e^F_{i,k} - 1} \right) \tag{21}$$

Following Holmes (1989), at a symmetric oligopoly equilibrium, the firm-level elasticity can be decomposed. By adding and subtracting the cross-price effects of the $N-1$ rivals, the total firm elasticity expands to:

$$e_{i,k}^{F} = -\left(\frac{\partial q_{i,k}}{\partial p_{i,k}} + \sum_{j\neq i}\frac{\partial q_{i,k}}{\partial p_{j,k}}\right)\frac{p_k}{q_{i,k}} + \sum_{j\neq i}\frac{\partial q_{i,k}}{\partial p_{j,k}}\frac{p_k}{q_{i,k}} \tag{22}$$

This simplifies exactly to $e_{i,k}^{F}(t) = e_{i,k}^{I}(t) + e_{i,k}^{C}(t)$.

A.2 Dynamic Increase in the Price Difference (Derivation of H2)

H2 states that the directional price penalty $\Delta p(t) = p_{i,A}^{*}(t) - p_{i,B}^{*}(t)$ increases as departure approaches. To show this, the time derivative of the optimal price path is calculated using the quotient rule on Equation (21):

$$\frac{\partial p_{i,k}^{*}}{\partial t} = c\cdot\frac{\frac{\partial e_{i,k}^{F}}{\partial t}(e_{i,k}^{F}-1) - e_{i,k}^{F}\frac{\partial e_{i,k}^{F}}{\partial t}}{(e_{i,k}^{F}-1)^2} = c\cdot\frac{-\frac{\partial e_{i,k}^{F}}{\partial t}}{(e_{i,k}^{F}-1)^2} \tag{23}$$

Because t represents the days remaining until departure, price sensitivity falls as time runs out, meaning $\frac{\partial e_{i,k}^{F}}{\partial t} > 0$. Therefore, $\frac{\partial p_{i,k}^{*}}{\partial t} < 0$, confirming that fares rise as departure nears.

In the high-income Market A, opportunity and sunk costs are larger, causing the elasticity to collapse much faster: $\frac{\partial e_{i,A}^{F}}{\partial t} > \frac{\partial e_{i,B}^{F}}{\partial t}$. Furthermore, the baseline elasticity in Market A is lower ($e_{i,A}^{F} < e_{i,B}^{F}$), which makes the denominator $(e_{i,A}^{F}-1)^2$ strictly smaller. This dictates that the absolute rate of price increase is strictly greater in the wealthy market:

$$\frac{\partial \Delta p(t)}{\partial t} = \frac{\partial p_{i,A}^{*}}{\partial t} - \frac{\partial p_{i,B}^{*}}{\partial t} < 0 \tag{24}$$

A.3 Competition as a Moderating Mechanism

Let the optimal price in market $k \in \{A, B\}$ be given by the generalized Lerner rule

$$p_{i,k}^{*}(N,t) = c\left(1 + \frac{1}{e_{i,k}^{F}(N,t)-1}\right), \qquad e_{i,k}^{F}(N,t) > 1. \tag{25}$$

Market A is the higher-income origin and is assumed to be less price-elastic than Market B:

$$1 < e_{i,A}^{F}(N,t) < e_{i,B}^{F}(N,t). \tag{26}$$

Since the markup term $1/(e_{i,k}^{F}-1)$ is strictly decreasing in $e_{i,k}^{F}$, the lower elasticity in Market A implies a higher markup than in Market B. Therefore,

$$\Delta p(N,t) = p_{i,A}^{*}(N,t) - p_{i,B}^{*}(N,t) = c\left[\frac{1}{e_{i,A}^{F}(N,t)-1} - \frac{1}{e_{i,B}^{F}(N,t)-1}\right] > 0. \tag{27}$$

For the static competition effect, suppress the time index. The directional price gap is

$$\Delta p(N) = c\left[\frac{1}{e^F_{i,A}(N) - 1} - \frac{1}{e^F_{i,B}(N) - 1}\right] > 0. \tag{28}$$

This expression is positive for every N as long as Market A remains less elastic than Market B, that is,

$$1 < e^F_{i,A}(N) < e^F_{i,B}(N). \tag{29}$$

Thus, competition does not eliminate directional price discrimination mechanically; it reduces its magnitude. To see this, assume that additional competitors increase firm-level elasticity through the aggregate cross-price elasticity:

$$\frac{\partial e^F_{i,k}(N)}{\partial N} = \frac{\partial e^C_{i,k}(N)}{\partial N} > 0. \tag{30}$$

Differentiating $\Delta p(N)$ with respect to N gives

$$\frac{\partial \Delta p(N)}{\partial N} = -c\left[\frac{\partial e^F_{i,A}/\partial N}{\left(e^F_{i,A} - 1\right)^2} - \frac{\partial e^F_{i,B}/\partial N}{\left(e^F_{i,B} - 1\right)^2}\right]. \tag{31}$$

Because Market A is less elastic, its markup is more sensitive to an increase in elasticity. If competitive pressure raises firm-level elasticity at least as strongly in Market A as in Market B, then

$$\frac{\partial e^F_{i,A}/\partial N}{\left(e^F_{i,A} - 1\right)^2} > \frac{\partial e^F_{i,B}/\partial N}{\left(e^F_{i,B} - 1\right)^2}, \tag{32}$$

which implies

$$\frac{\partial \Delta p(N)}{\partial N} < 0. \tag{33}$$

Together, these inequalities show that

$$\Delta p(N) > 0 \qquad \text{and} \qquad \frac{\partial \Delta p(N)}{\partial N} < 0. \tag{34}$$

Hence, the static directional price-discrimination effect remains positive, but decreases as the number of competitors rises.

The same logic applies dynamically. Let t denote the number of days remaining until departure. Since passengers are more price-sensitive when departure is farther away, elasticities increase with t. The dynamic mechanism assumes that this recovery in price sensitivity is stronger in Market A:

$$\frac{\partial e^F_{i,A}(N,t)}{\partial t} > \frac{\partial e^F_{i,B}(N,t)}{\partial t} \geq 0. \tag{35}$$

Differentiating the dynamic price gap with respect to t gives

$$\frac{\partial \Delta p(N,t)}{\partial t} = -c\left[\frac{\partial e^F_{i,A}/\partial t}{\left(e^F_{i,A} - 1\right)^2} - \frac{\partial e^F_{i,B}/\partial t}{\left(e^F_{i,B} - 1\right)^2}\right] < 0. \tag{36}$$

Since t measures days remaining until departure, this negative derivative means that the directional price gap increases as departure approaches.

Finally, competition weakens this dynamic increase. Assuming that the direct competition effect on elasticity and the time effect on elasticity are separable, so that $\partial^2 e^F_{i,k}/(\partial N \partial t) = 0$, differentiating the dynamic effect with respect to N yields

$$\frac{\partial^2 \Delta p(N,t)}{\partial N \partial t} = 2c \left[\frac{\left(\partial e^F_{i,A}/\partial N\right)\left(\partial e^F_{i,A}/\partial t\right)}{\left(e^F_{i,A} - 1\right)^3} - \frac{\left(\partial e^F_{i,B}/\partial N\right)\left(\partial e^F_{i,B}/\partial t\right)}{\left(e^F_{i,B} - 1\right)^3} \right]. \tag{37}$$

Given $e^F_{i,A} < e^F_{i,B}$, the denominator is smaller in Market A. If competitive pressure and the recovery of price sensitivity are at least as strong in Market A as in Market B, the first term dominates the second, implying

$$\frac{\partial^2 \Delta p(N,t)}{\partial N \partial t} > 0. \tag{38}$$

The positive cross-partial means that the negative time derivative becomes less negative as N increases. Therefore, competition reduces the static directional price gap and weakens the late-booking increase of directional price discrimination.

B Complete List of Flight Routes, Market Concentration, and Connectivity

The following table provides a comprehensive overview of the 110 flight routes analyzed in this study. The columns are defined as follows: **ID** is the sequential numbering of the routes. **Tier** represents the income category of the destination based on GDP per capita (World Bank, 2026): Tier 1 (>$70,000), Tier 2 ($50,000–$70,000), Tier 3 ($30,000–$50,000), Tier 4 ($10,000–$30,000), and Tier 5 (<$10,000). **Destination** specifies the target city and country. **GDP** refers to the destination's GDP per capita. **Diff.** indicates the economic disparity, calculated as the GDP of the origin minus the GDP of the destination. **N** denotes the effective number of airlines operating on the route. This means that if a joint venture or a holding company operates on the route, it is counted as one airline. For example, British Airways and Vueling both fly on the BCN-LHR route, but both airlines are part of the IAG group, so the effective number of competitors is one. **HHI** represents the Herfindahl-Hirschman Index for market concentration, calculated based on the effective competition structure of the route.

Table 13: Detailed Overview of Sampled Flight Routes, Competition (N), and Market Concentration (HHI)

ID	Tier	Route	Destination (Country)	GDP (Dest.)	Diff.	N	HHI
HUB: LHR - London Heathrow (UK) | Ref. GDP: $53,246							
1	Tier 1	LHR-DUB	Dublin (IE)	$112,894	+$59,648	1	1.0000
2	Tier 1	LHR-GVA	Geneva (CH)	$103,998	+$50,752	2	0.5102
3	Tier 1	LHR-ZRH	Zurich (CH)	$103,998	+$50,752	2	0.5372
4	Tier 1	LHR-JFK	New York (USA)	$84,534	+$31,288	3	0.4351
5	Tier 2	LHR-MUC	Munich (DE)	$56,103	+$2,857	2	0.5868
6	Tier 2	LHR-GLA	Glasgow (UK)	$53,246	$0	1	1.0000
7	Tier 2	LHR-EDI	Edinburgh (UK)	$53,246	$0	1	1.0000
8	Tier 2	LHR-DXB	Dubai (UAE)	$50,273	-$2,973	3	0.5938
9	Tier 3	LHR-LIN	Milan (IT)	$40,385	-$12,861	2	0.7551
Continued on next page							

Table 13 – continued from previous page

ID	Tier	Route	Destination (Country)	GDP (Dest.)	Diff.	N	HHI
10	Tier 3	LHR-FCO	Rome (IT)	$40,385	-$12,861	2	0.6916
11	Tier 3	LHR-BCN	Barcelona (ES)	$35,326	-$17,920	1	1.0000
12	Tier 3	LHR-RUH	Riyadh (SA)	$35,121	-$18,125	2	0.5918
13	Tier 4	LHR-LIS	Lisbon (PT)	$29,292	-$23,954	2	0.5165
14	Tier 4	LHR-WAW	Warsaw (PL)	$25,103	-$28,143	2	0.5041
15	Tier 4	LHR-ATH	Athens (GR)	$24,626	-$28,620	2	0.5247
16	Tier 4	LHR-IST	Istanbul (TR)	$15,892	-$37,354	2	0.5703
17	Tier 5	LHR-CPT	Cape Town (ZA)	$6,267	-$46,979	1	1.0000
18	Tier 5	LHR-CAI	Cairo (EG)	$3,338	-$49,908	2	0.5200
19	Tier 5	LHR-BOM	Mumbai (IN)	$2,694	-$50,552	4	0.2813
20	Tier 5	LHR-DEL	Delhi (IN)	$2,694	-$50,552	4	0.2909
HUB: AMS - Amsterdam Schiphol (NL) | Ref. GDP: $67,520							
21	Tier 1	AMS-DUB	Dublin (IE)	$112,894	+$45,374	3	0.3368
22	Tier 1	AMS-ZRH	Zurich (CH)	$103,998	+$36,478	2	0.5102
23	Tier 1	AMS-OSL	Oslo (NO)	$86,785	+$19,265	3	0.4810
24	Tier 1	AMS-CPH	Copenhagen (DK)	$71,026	+$3,506	4	0.3521
25	Tier 2	AMS-VIE	Vienna (AT)	$58,268	-$9,252	2	0.5017
26	Tier 2	AMS-ARN	Stockholm (SE)	$57,117	-$10,403	3	0.4879
27	Tier 2	AMS-MUC	Munich (DE)	$56,103	-$11,417	2	0.5011
28	Tier 2	AMS-MAN	Manchester (UK)	$53,246	-$14,274	2	0.5050
29	Tier 3	AMS-LIN	Milan (IT)	$40,385	-$27,135	3	0.3684
30	Tier 3	AMS-FCO	Rome (IT)	$40,385	-$27,135	2	0.5918
31	Tier 3	AMS-BCN	Barcelona (ES)	$35,326	-$32,194	2	0.5171
32	Tier 3	AMS-AGP	Malaga (ES)	$35,326	-$32,194	4	0.3765
33	Tier 4	AMS-LIS	Lisbon (PT)	$29,292	-$38,228	2	0.5703
34	Tier 4	AMS-WAW	Warsaw (PL)	$25,103	-$42,417	2	0.5000
35	Tier 4	AMS-OTP	Bucharest (RO)	$20,080	-$47,440	2	0.5200
36	Tier 4	AMS-IST	Istanbul (TR)	$15,892	-$51,628	2	0.5918
37	Tier 5	AMS-BKK	Bangkok (TH)	$7,346	-$60,174	2	0.6536
38	Tier 5	AMS-BOM	Mumbai (IN)	$2,694	-$64,826	2	0.5000
39	Tier 5	AMS-DEL	Delhi (IN)	$2,694	-$64,826	2	0.5000
40	Tier 5	AMS-NBO	Nairobi (KE)	$2,132	-$65,388	1	1.0000
HUB: MAD - Madrid Barajas (ES) | Ref. GDP: $35,326							
41	Tier 1	MAD-DUB	Dublin (IE)	$112,894	+$77,568	2	0.5062
42	Tier 1	MAD-ZRH	Zurich (CH)	$103,998	+$68,672	3	0.3516
43	Tier 1	MAD-GVA	Geneva (CH)	$103,998	+$68,672	2	0.7397
44	Tier 1	MAD-JFK	New York (USA)	$84,534	+$49,208	3	0.4400
45	Tier 2	MAD-VIE	Vienna (AT)	$58,268	+$22,942	2	0.6543
46	Tier 2	MAD-BRU	Brussels (BE)	$56,614	+$21,288	4	0.2909
47	Tier 2	MAD-MUC	Munich (DE)	$56,103	+$20,777	3	0.3580
48	Tier 2	MAD-TLV	Tel Aviv (IL)	$54,176	+$18,850	1	1.0000
49	Tier 3	MAD-ORY	Paris (FR)	$46,103	+$10,777	3	0.3865
50	Tier 3	MAD-FCO	Rome (IT)	$40,385	+$5,059	5	0.2489
51	Tier 3	MAD-MXP	Milan (IT)	$40,385	+$5,059	4	0.2664
52	Tier 3	MAD-VCE	Venice (IT)	$40,385	+$5,059	4	0.3245
53	Tier 4	MAD-LIS	Lisbon (PT)	$29,292	-$6,034	5	0.2579
54	Tier 4	MAD-OPO	Porto (PT)	$29,292	-$6,034	4	0.2968
55	Tier 4	MAD-MEX	Mexico City (MX)	$14,185	-$21,141	2	0.5041
56	Tier 4	MAD-EZE	Buenos Aires (AR)	$13,969	-$21,357	3	0.4050

Continued on next page

Table 13 – continued from previous page

ID	Tier	Route	Destination (Country)	GDP (Dest.)	Diff.	*N*	HHI
57	Tier 5	MAD-LIM	Lima (PE)	$8,452	-$26,874	4	0.2800
58	Tier 5	MAD-BOG	Bogota (CO)	$7,919	-$27,407	4	0.3845
59	Tier 5	MAD-CMN	Casablanca (MA)	$4,153	-$31,173	3	0.4167
60	Tier 5	MAD-RAK	Marrakech (MA)	$4,153	-$31,173	3	0.4235
HUB: CDG - Paris Charles de Gaulle (FR) | Ref. GDP: $46,103							
61	Tier 1	CDG-DUB	Dublin (IE)	$112,894	+$66,791	2	0.5000
62	Tier 1	CDG-GVA	Geneva (CH)	$103,998	+$57,895	1	1.0000
63	Tier 1	CDG-ZRH	Zurich (CH)	$103,998	+$57,895	2	0.5022
64	Tier 1	CDG-JFK	New York (USA)	$84,534	+$38,431	3	0.6600
65	Tier 2	CDG-MUC	Munich (DE)	$56,103	+$10,000	3	0.4195
66	Tier 2	CDG-MAN	Manchester (UK)	$53,246	+$7,143	2	0.5102
67	Tier 2	CDG-HEL	Helsinki (FI)	$53,149	+$7,046	2	0.6800
68	Tier 2	CDG-DXB	Dubai (UAE)	$50,273	+$4,170	2	0.6250
69	Tier 3	CDG-LIN	Milan (IT)	$40,385	-$5,718	3	0.4136
70	Tier 3	CDG-FCO	Rome (IT)	$40,385	-$5,718	2	0.5283
71	Tier 3	CDG-BCN	Barcelona (ES)	$35,326	-$10,777	2	0.5679
72	Tier 3	CDG-PRG	Prague (CZ)	$31,823	-$14,280	2	0.5800
73	Tier 4	CDG-WAW	Warsaw (PL)	$25,103	-$20,999	2	0.5000
74	Tier 4	CDG-ATH	Athens (GR)	$24,626	-$21,477	3	0.3846
75	Tier 4	CDG-OTP	Bucharest (RO)	$20,080	-$26,023	5	0.2833
76	Tier 4	CDG-IST	Istanbul (TR)	$15,892	-$30,211	2	0.7689
77	Tier 5	CDG-ALG	Algiers (DZ)	$5,752	-$40,351	3	0.5612
78	Tier 5	CDG-TUN	Tunis (TN)	$4,181	-$41,922	2	0.5139
79	Tier 5	CDG-CMN	Casablanca (MA)	$4,153	-$41,950	2	0.5740
80	Tier 5	CDG-RBA	Rabat (MA)	$4,153	-$41,950	2	0.6250
HUB: FRA - Frankfurt am Main (DE) | Ref. GDP: $56,103							
81	Tier 1	FRA-DUB	Dublin (IE)	$112,894	+$56,791	2	0.6033
82	Tier 1	FRA-ZRH	Zurich (CH)	$103,998	+$47,895	2	0.6250
83	Tier 1	FRA-GVA	Geneva (CH)	$103,998	+$47,895	1	1.0000
84	Tier 1	FRA-CPH	Copenhagen (DK)	$71,026	+$14,923	2	0.6450
85	Tier 2	FRA-VIE	Vienna (AT)	$58,268	+$2,165	2	0.6543
86	Tier 2	FRA-ARN	Stockholm (SE)	$57,117	+$1,014	1	1.0000
87	Tier 2	FRA-BRU	Brussels (BE)	$56,614	+$511	1	1.0000
88	Tier 2	FRA-DXB	Dubai (UAE)	$50,273	-$5,830	1	1.0000
89	Tier 3	FRA-LIN	Milan (IT)	$40,385	-$15,718	2	0.6184
90	Tier 3	FRA-FCO	Rome (IT)	$40,385	-$15,718	3	0.4206
91	Tier 3	FRA-BCN	Barcelona (ES)	$35,326	-$20,777	2	0.6033
92	Tier 3	FRA-PRG	Prague (CZ)	$31,823	-$24,280	2	0.5200
93	Tier 4	FRA-LIS	Lisbon (PT)	$29,292	-$26,811	2	0.5102
94	Tier 4	FRA-WAW	Warsaw (PL)	$25,103	-$30,999	2	0.5139
95	Tier 4	FRA-ATH	Athens (GR)	$24,626	-$31,477	3	0.4306
96	Tier 4	FRA-IST	Istanbul (TR)	$15,892	-$40,211	2	0.5400
97	Tier 5	FRA-BKK	Bangkok (TH)	$7,346	-$48,757	2	0.5556
98	Tier 5	FRA-CPT	Cape Town (ZA)	$6,267	-$49,836	2	0.5794
99	Tier 5	FRA-CAI	Cairo (EG)	$3,338	-$52,765	2	0.5556
100	Tier 5	FRA-DEL	Delhi (IN)	$2,694	-$53,409	2	0.5200
INTER-HUB ROUTES							
101	Hubs	LHR-AMS	Amsterdam (NL)	$67,520	+$14,274	2	0.5007
102	Hubs	LHR-FRA	Frankfurt (DE)	$56,103	+$2,857	2	0.6450

Continued on next page

Table 13 – continued from previous page

ID	Tier	Route	Destination (Country)	GDP (Dest.)	Diff.	N	HHI
103	Hubs	AMS-FRA	Frankfurt (DE)	\$56,103	-\$11,417	2	0.5000
104	Hubs	LHR-CDG	Paris (FR)	\$46,103	-\$7,143	2	0.5009
105	Hubs	AMS-CDG	Paris (FR)	\$46,103	-\$21,417	1	1.0000
106	Hubs	FRA-CDG	Paris (FR)	\$46,103	-\$10,000	3	0.3690
107	Hubs	LHR-MAD	Madrid (ES)	\$35,326	-\$17,920	1	1.0000
108	Hubs	AMS-MAD	Madrid (ES)	\$35,326	-\$32,194	3	0.3719
109	Hubs	FRA-MAD	Madrid (ES)	\$35,326	-\$20,777	3	0.3580
110	Hubs	CDG-MAD	Madrid (ES)	\$35,326	-\$10,777	2	0.5433

C Data Extraction Procedure and Data Processing

The fare data were collected through RapidAPI using the *Flights Scraper Sky* endpoint. The Python script generated route-date-specific roundtrip queries for each airport pair, direction, departure date, return date, and duration type. Each request specified the origin airport, destination airport, departure date, return date, nonstop restriction, and USD as the common currency. The API workflow consisted of two steps. First, the script submitted a `search-roundtrip` request. Second, if the returned search status was incomplete, the corresponding session ID was passed to the `search-incomplete` endpoint until the search was completed or the maximum number of polling attempts was reached. To avoid exceeding the API limit, requests were controlled by an asynchronous rate limiter and a semaphore. The JSON response was then flattened recursively so that nested itinerary, leg, segment, carrier, timing, and price information could be stored in rectangular tabular form. Each flattened itinerary was saved together with route ID, search direction, queried departure date, queried return date, duration type, and API status.

Data collection was conducted in two waves. The first wave was downloaded on January 31, 2026 and covered departures from February 1 to April 30, 2026. This wave contains only seven-day roundtrips. The second wave was downloaded on April 30, 2026 and covered departures from May 1 to July 31, 2026. In this second wave, seven-day roundtrips were collected for the full period, while additional 21-day roundtrips were collected from May 1 to July 10, 2026 because these observations are required for the synthetic matching strategy. Although each wave was collected on a single calendar day, the individual API calls were not perfectly simultaneous. The extraction process took several hours because the task list contained thousands of route-date-direction requests and incomplete searches required repeated polling. Fully parallel execution was not feasible. The relevant identifying variation is nevertheless preserved because all observations within a wave share the same collection date and the later matching procedures compare tickets within the same wave, route, carrier, and booking-horizon structure.

The raw API output contained 45,970,752 quoted roundtrip prices. This number follows directly from the structure of Skyscanner roundtrip results. For a given query, if n outbound flights and m return flights are available, the API returns the Cartesian product of $n \cdot m$ roundtrip combinations. The raw files were first imported into Stata and appended into one master file. Because the API output used slightly different variable names across extraction waves, the Stata script harmonized the naming convention before appending. Price variables, stop indicators, airport codes, carrier names, carrier codes, flight numbers, departure and arrival timestamps, query dates, and route identifiers were standardized. Price and stop-count variables were converted to numeric format, and the resulting file was saved as the raw merged pricing dataset.

The cleaned ticket-level sample was then constructed by applying the restrictions required for the empirical design. Only economy-class roundtrips were retained. Observations were dropped if relevant flight numbers or timestamps were not observed, or if outbound and return legs were not operated by the same airline. This filter removes cross-airline combinations and incomplete itineraries, but it does not collapse the Cartesian-product structure within an airline. For example, if one airline offers three feasible outbound flights and two feasible return flights for a route-date query, the cleaned dataset still contains the resulting six same-airline roundtrip combinations. After these restrictions, the dataset was reduced from 45,970,752 raw prices to 1,507,730 usable ticket observations.

The Stata preparation script then merged the cleaned fare data with route-, city-, and airline-level variables. Airport codes were matched to city-level information on GDP per capita, population, March and June temperatures, and hub status. Distances were merged by unordered airport-pair identifiers. The script also generated airline-level and route-level controls, including the legacy-carrier dummy, carrier-specific hub indicators, school-holiday indicators by city and date, high-speed-train availability, alternative-airport availability, alliance-route status, and the liberalized-route dummy. Competition variables were not inferred from the scraped fare data. Instead, the effective number of competitors and the HHI were constructed from Flightsfrom.com so that airlines serving a route were counted even if a specific flight did not appear in the fare API because it was sold out or not returned by Skyscanner. Holdings and joint ventures were treated as one effective competitor.

After the merge, the script generated the directional variables used in the regressions. For each route, the higher-income endpoint was defined as city A and the lower-income endpoint as city B. Income, population, temperature, hub, and holiday differences were then constructed as $A - B$. National GDP per capita is used as the main income definition, while TL2 regional income is generated analogously for the regional robustness specifications. The booking horizon was calculated relative to the extraction date of the corresponding wave: January 31, 2026 for wave 1 and April 30, 2026 for wave 2. These values were then converted into booking-horizon dummies used in the dynamic specifications.

The final estimation datasets are smaller than the cleaned ticket-level file because the empirical strategies construct price differences from multiple ticket observations rather than using individual tickets directly. Thus, the reduction in observations is a direct consequence of the matching design: the final restricted-sample datasets used in the national-GDP regressions contain 115,540 matched observations for Strategy 2, 36,113 observations for Strategy 3, and 4,460 observations for Strategy 1.

D Skyscanner Price Display and Booking Options

Figure 4 illustrates how the fare data were observed on Skyscanner. For the example route London Heathrow (LHR) to Barcelona (BCN), departing on June 15, 2026 and returning on June 22, 2026, Skyscanner reports a roundtrip price of 126 EUR in the main result card. This displayed fare corresponds to the lowest available booking option for the selected itinerary and flight numbers. After clicking on the price card, Figure 5 shows that several booking channels are available for the same itinerary. In this example, the minimum price is offered by Kiwi.com at 126 EUR, while the airline website, Vueling, lists the same itinerary at 130 EUR. This illustrates why the Skyscanner price should be interpreted as the lowest observed market price rather than necessarily the airline-posted fare. Online travel agencies or booking platforms may offer small discounts, service-fee differences, or platform-specific markups relative to the airline website. This matters for the empirical strategy in two ways. First, the scraped price may not perfectly reproduce the airline's own pricing strategy if the lowest available fare is displayed through an intermediary. Second, this motivates the 2% error tolerance in the Lewis-style additivity test discussed in Section 4.2. Small deviations between official airline prices and Skyscanner fares can arise from OTA discounts, service fees, or platform-specific pricing, and should therefore not mechanically invalidate the identifying assumption.

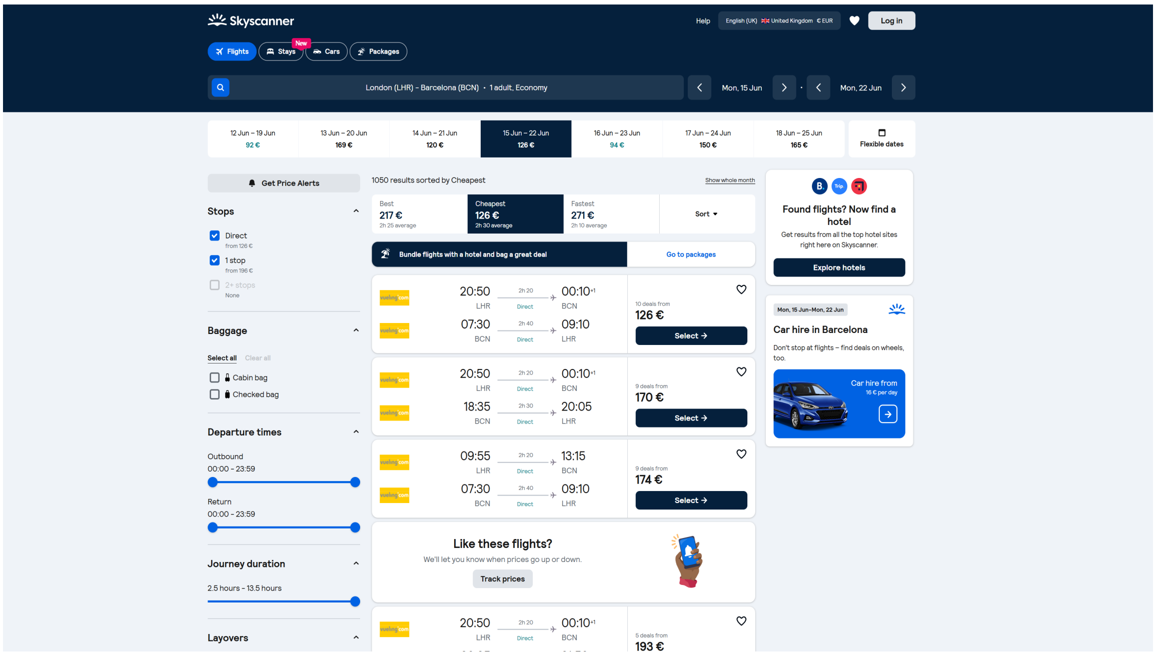

Figure 4: Skyscanner Buy-Box Price for an LHR–BCN Roundtrip

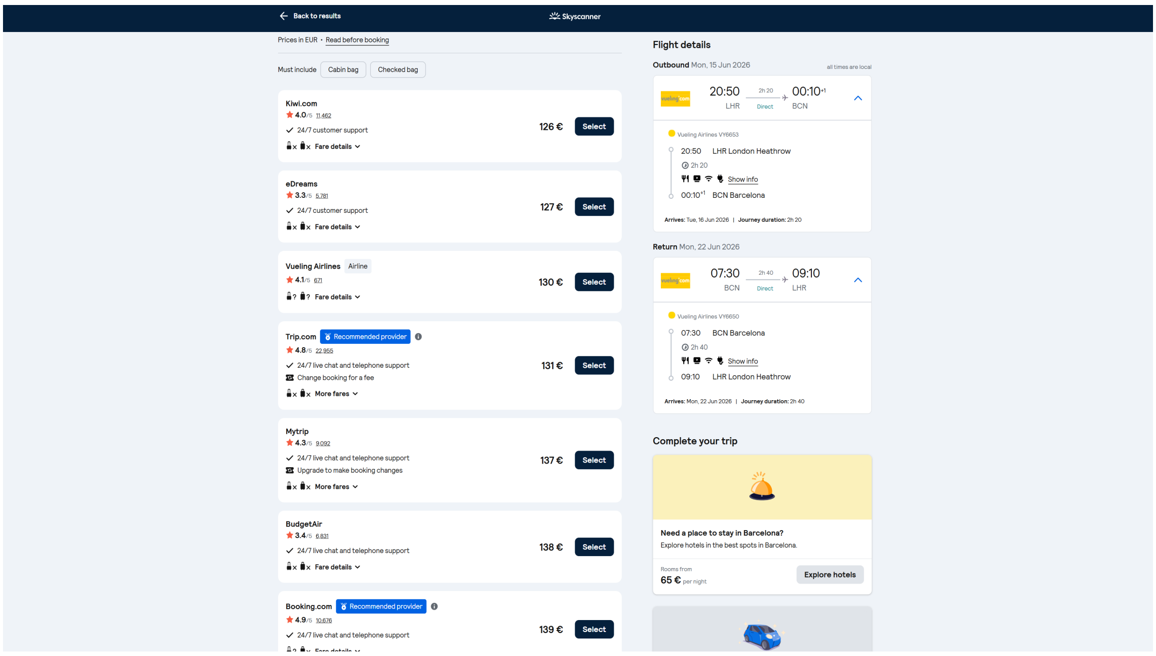

Figure 5: Booking Options After Selecting the Skyscanner Buy-Box Fare

E Appendix: Detailed Matching Rules and Calculations

E.1 Strategy 1: Matching Rules

Constructing a flight chain requires matching specific flights across all weeks of the observation period. A chain is only retained if it can be completed over the full three-month period; incomplete chains are discarded because otherwise the average price difference would be based on different booking horizons or different numbers of flight segments. The matching is implemented separately for the outbound direction $A \rightarrow B$ and the return direction $B \rightarrow A$.

The first-best rule requires that the exact same flight number is available in every week of the chain. In the example below, FN103 is available in all twelve weeks for the outbound direction $A \rightarrow B$. Therefore, FN103 forms a first-best outbound chain. Analogously, FN104 is available in every week for the return direction $B \rightarrow A$, so FN104 forms a first-best return chain. These perfectly consistent chains are marked in **blue**. The first-best match is therefore the preferred case because every weekly roundtrip is constructed from the same pair of flight numbers throughout the full observation period.

If no complete first-best chain can be formed, a second-best chain is constructed. The second-best rule is applied independently in each week: among all available flight numbers that are not already used by the first-best chain, the algorithm selects the lowest available flight number in that week. Therefore, the second-best path does not try to keep the same flight number from the previous week. Even if the previous week's flight number is available again, it is replaced whenever a lower feasible flight number exists in the current week. In the outbound example, the second-best path starts with FN105 in week 1 and week 2. In week 3, FN101 becomes available and is selected because it is the lowest available non-first-best flight number, even though FN105 is also available. In week 4, FN101 is unavailable, so the algorithm selects FN107. In week 5, FN109 is selected because it is the lowest available non-first-best option in that week. The same rule is applied week by week until the full chain is completed. This dynamically constructed second-best path is marked in **orange**. Red entries mark weeks in which a given flight number is unavailable and therefore cannot be selected in that week.

The same procedure is applied to the return direction. In the return matrix, FN104 forms the first-best chain because it is available in every week. The second-best return chain is then constructed from the remaining available flights. The rule is again applied independently in each week: among all available return-flight numbers that are not part of the first-best chain, the algorithm selects the lowest available flight number. Therefore, the second-best return path starts with FN102 in week 1. In week 2, FN102 is unavailable, so the algorithm selects FN108. In week 3, FN106 is selected because it is the lowest available non-first-best return flight in that week. In week 4, FN102 becomes available again and is therefore selected because it is lower than FN106, FN108, and FN110. The path can therefore switch back to a lower flight number whenever it becomes available. The key requirement is that the resulting path contains exactly one feasible non-first-best flight in every week of the three-month period. Only such complete chains enter Strategy 1.

Tables 1 and 2 provide an illustrative example of how the first-best price chains are constructed. The tables report the two synchronized chains for British Airways on the Barcelona (BCN)–London Heathrow (LHR) route. Since the same flight numbers can be followed consistently over the full observation period, the example represents a first-best match.

Table 14: Availability Matrix: Direction A → B (Outbound)

Flight No.	W1	W2	W3	W4	W5	W6	W7	W8	W9	W10	W11	W12
FN101	-	-	X	-	-	X	-	-	X	-	-	X
FN103 (1st-Best)	X	X	X	X	X	X	X	X	X	X	X	X
FN105	X	X	X	-	-	X	X	-	X	X	X	-
FN107	X	X	X	X	-	X	X	X	X	X	X	X
FN109	X	X	X	X	X	X	X	X	X	X	X	X

Table 15: Availability Matrix: Direction B → A (Return)

Flight No.	W1	W2	W3	W4	W5	W6	W7	W8	W9	W10	W11	W12
FN102	X	-	-	X	-	-	-	-	-	X	-	-
FN104 (1st-Best)	X	X	X	X	X	X	X	X	X	X	X	X
FN106	X	-	X	X	-	X	-	-	X	X	X	-
FN108	X	X	X	X	X	X	X	-	X	X	X	X
FN110	X	X	X	X	X	X	X	X	X	X	X	X

To avoid booking-horizon bias, both chain directions in Tables 1 and 2 are combined. First, all roundtrips originating in London (LHR) from both chains are added together:

$$\sum P_{LHR} = 573.93 + 340.37 + 402.28 + 173.44 + 150.16 + 229.80 + 240.09$$
$$+ 182.68 + 368.00 + 219.59 + 248.31 + 190.90 + 214.40 = 3533.95 \text{ USD}$$

Next, all roundtrips originating in Barcelona (BCN) are summed up:

$$\sum P_{BCN} = 241.36 + 352.08 + 279.50 + 243.93 + 137.21 + 236.14 + 198.36$$
$$+ 239.74 + 175.08 + 347.70 + 165.19 + 146.33 + 217.50 = 2980.12 \text{ USD}$$

Since the chain contains 13 observed roundtrip prices for each origin, the average roundtrip price is obtained by dividing each total by 13:

$$\bar{P}_{LHR} = \frac{3533.95}{13} = 271.84 \text{ USD}$$

$$\bar{P}_{BCN} = \frac{2980.12}{13} = 229.24 \text{ USD}$$

Subtracting the average Barcelona-origin price from the average London-origin price gives the directional price difference:

$$\Delta P = \bar{P}_{LHR} - \bar{P}_{BCN} = 271.84 - 229.24 = 42.60 \text{ USD}$$

The values in the equations therefore come directly from the first-best chain. In this example, passengers originating in London pay on average 42.60 USD more than passengers originating in Barcelona for the matched roundtrip construction.

E.2 Strategy 2: Matching Rules

Strategy 2 applies a strict daily flight-pair matching rule. The objective is to compare a roundtrip originating in city A with the mirrored roundtrip originating in city B for the same travel dates. Within each search direction, all available outbound and return flights are first sorted by flight number in ascending order. The lowest available outbound flight is paired with the lowest available return flight, the second-lowest outbound flight with the second-lowest return flight, and so on. This ranking rule is applied separately to both search origins. A valid observation is retained only if the resulting pair is perfectly mirrored across the two searches. This means that the outbound flight number in the A-originating roundtrip must appear as the return flight number in the B-originating roundtrip, and the return flight number in the A-originating roundtrip must appear as the outbound flight number in the B-originating roundtrip.

Table 16: Flight Availability and Mirrored Matching (May 1 to May 8)

Search 1: Origin A		**Search 2: Origin B**	
A → B (May 1)	**B → A (May 8)**	**B → A (May 1)**	**A → B (May 8)**
FN120	FN121	FN121	FN120
FN122	FN123	FN123	FN122
FN124	FN125	FN127	FN124
FN126	FN127	-	FN126
FN128	-	-	-

Table 16 illustrates this rule. The first valid match is marked in **blue**. A passenger starting in A flies on FN120 from A to B on May 1 and returns on FN121 from B to A on May 8. The mirrored passenger starting in B flies on FN121 from B to A on May 1 and returns on FN120 from A to B on May 8. The two roundtrips therefore use the same two flight numbers in opposite order and form a valid mirrored pair. The second valid match, marked in **green**, follows the same logic with FN122 and FN123.

All combinations that do not satisfy this mirrored-flight-number condition are excluded. For example, the third ranked A-originating combination pairs FN124 with FN125. However, the third ranked B-originating combination pairs FN127 with FN124. This is not a valid mirror because FN125 is not used as the outbound flight in the B-originating search. Similarly, the combination involving FN126 and FN127 is discarded because the corresponding mirrored pair is incomplete.

E.3 Strategy 3: Priority Matching Rules

The reliability of the synthetic price depends on a priority-based matching system. Strategy 3 requires that each synthetic observation consists of four internally consistent roundtrip prices: Trip 1, Trip 2, Trip 3, and Trip 4. These four components are reported in Tables 17–20. A match is retained only if the relevant flight numbers can be linked across all four trips. The highest priority is given to First-Best matches, marked in **blue**. In this case, the same flight numbers are observed in all required positions. In the example, LH902 and LH905 form such a First-Best match: Table 17 shows LH902 on t_1 and LH905 on t_2, Table 18 again contains LH902 and LH905 on t_3 and t_4, Table 19 links LH902 from t_1 to LH905 on t_4, and Table 20 contains the mirrored reverse-direction combination LH905 and LH902. This creates a complete synthetic matching structure.

If no First-Best match is available for a remaining flight combination, the Second-Best rule is applied. Second-Best matches are constructed from the remaining flights by selecting the lowest available flight-number combination that can still form a complete four-trip structure. These matches are marked in **green** and **plum** in Tables 17–20. For example, the green match combines LH900/LH901 in Trip 1, LH906/LH907 in Trip 2, LH900/LH907 in Trip 3, and the mirrored LH901/LH906 combination in Trip 4. Although the same flight numbers are not identical across all periods, the four required components are jointly available and therefore allow the synthetic price to be constructed. The plum match follows the same logic using LH908/LH911, LH910/LH915, LH908/LH915, and LH911/LH910.

Flights are dropped if they cannot be linked into a complete four-trip structure. For example, the observations involving LH913 are marked as Dropped in Tables 17–20 because at least one required counterpart is missing in the corresponding trip. A one-sided or incomplete combination cannot be used, since the synthetic price $P_A(t_1, t_2) + P_A(t_3, t_4) - P_A(t_1, t_4)$ would no longer isolate the same flight components that are compared to Trip 4. Therefore, only complete First-Best or Second-Best structures enter the Strategy 3 sample.

Table 17: Trip 1: Origin A (Out: t_1, Return: t_2) – 7 Days

Outbound (t1)	Return (t2)	Price	Status
LH900	LH901	\$410	**Matched**
LH902	LH905	\$425	**Matched**
LH908	LH911	\$440	**Matched**
-	LH913	\$480	Dropped

Table 18: Trip 2: Origin A (Out: t_3, Return: t_4) – 7 Days

Outbound (t3)	Return (t4)	Price	Status
LH902	LH905	\$405	**Matched**
LH906	LH907	\$390	**Matched**
LH910	LH915	\$415	**Matched**
LH912	LH913	\$450	Dropped

Table 19: Trip 3: Origin A (Out: t_1, Return: t_4) – 21 Days

Outbound (t1)	Return (t4)	Price	Status
LH900	LH907	\$550	**Matched**
LH902	LH905	\$565	**Matched**
LH908	LH915	\$580	**Matched**
-	LH913	\$610	Dropped

F Manual Fare-Basis Check Using ITA Matrix

For airline-route observations with only one daily flight, the within-day additivity test in Equation 11 cannot be implemented because two alternative outbound and return flights are required. In these cases, I manually checked the fare-basis information using ITA Matrix (ITA Matrix, 2026). The purpose of

Table 20: Trip 4: Origin B (Out: t_2, Return: t_3) – Reverse Direction

Outbound (t2)	Return (t3)	Price	Status
LH901	LH906	$310	**Matched**
LH905	LH902	$320	**Matched**
LH911	LH910	$305	**Matched**
LH913	LH912	$610	Dropped

this check is to verify whether the common flight segment enters the 7-day and 21-day roundtrip tickets under the same fare rule. If the same physical flight segment has the same fare-basis code in both tickets, the segment is treated as following the same tariff rule and can be retained for the Lewis-style synthetic matching procedure.

Figures 6 and 7 illustrate this check for a Vueling itinerary between London Heathrow (LHR) and Barcelona (BCN). Figure 6 shows a 7-day roundtrip from August 1 to August 8, while Figure 7 shows a 21-day roundtrip from August 1 to August 22. The relevant common segment is the outbound flight from LHR to BCN on August 1. In both tickets, ITA Matrix reports the same fare basis for this segment: `OPLRYV`. The corresponding fare component is also identical at 76.70 EUR. This indicates that the outbound segment is priced under the same fare rule in both the 7-day and the 21-day ticket. The total roundtrip prices differ, 227.15 EUR for the 7-day ticket and 238.15 EUR for the 21-day ticket, because the return segments are on different dates and use different fare-basis codes. This difference does not invalidate the check, since the relevant requirement is that the overlapping segment used in the synthetic construction is governed by the same tariff rule. The example therefore supports the use of the observation in the synthetic matching strategy.

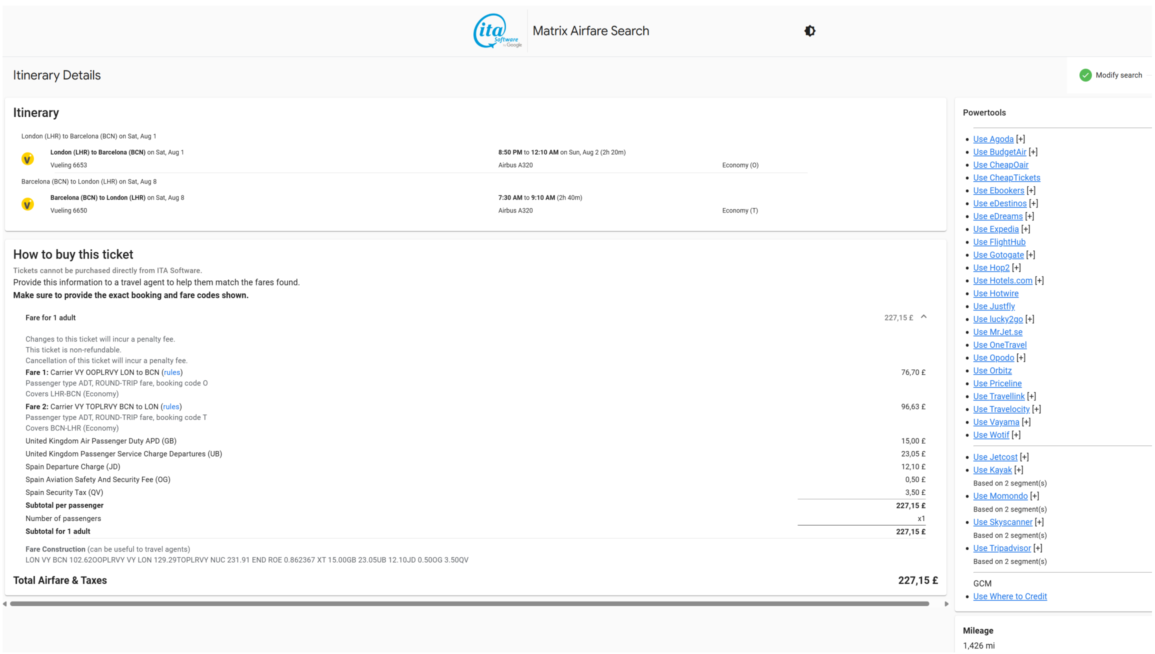

Figure 6: ITA Matrix Fare-Basis Check for a 7-Day Roundtrip

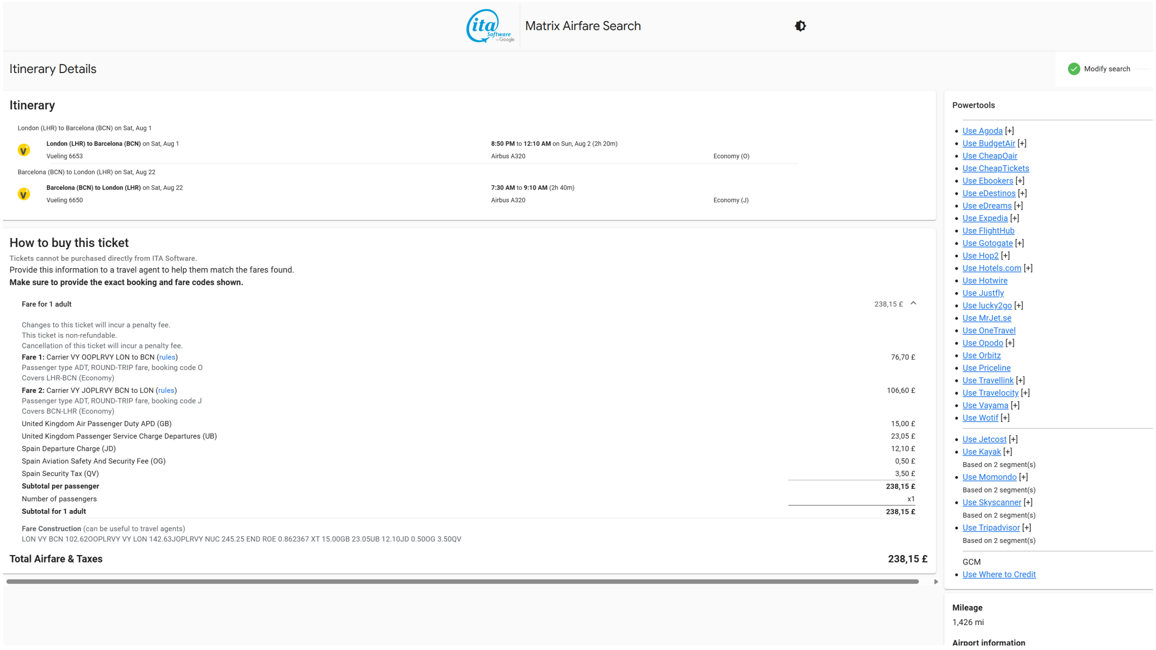

Figure 7: ITA Matrix Fare-Basis Check for a 21-Day Roundtrip

G Robustness Check: Market Concentration Measured by Adjusted HHI

The main competition specification uses the number of competing airlines because this measure has a direct and intuitive interpretation: if one additional airline operates between A and B, the market becomes more competitive. However, this measure does not account for asymmetries in market shares. For example, a route served by two airlines may still resemble a near-monopoly if one airline operates most flights while the second airline offers only very limited service. Therefore, I use the adjusted HHI as an additional robustness check. This measure captures market concentration more directly because it accounts for the distribution of flight activity across airlines. As in the construction of the number of competitors, airlines that belong to the same holding company, such as British Airways and Iberia, or airlines that operate under a joint venture on a given route are treated as one economic competitor. A higher adjusted HHI therefore indicates a more concentrated market after accounting for both unequal flight shares and economically integrated carriers.

The results in Table 21 support the competition mechanism from the main analysis. The interaction between Δ Income and adjusted HHI is positive and statistically significant in both strategies. This means that the income premium becomes larger when the market is more concentrated. In Strategy 1, a 10,000 USD increase in the income difference raises the directional fare premium by 2.19 USD in markets with an adjusted HHI below or equal to 0.4, by 3.10 USD in markets with an HHI between 0.4 and 0.7, and by 5.02 USD in markets with an HHI above 0.7. Strategy 3 shows the same pattern: the corresponding effects are 2.85 USD, 3.72 USD, and 5.61 USD. Thus, passengers from the richer origin pay a higher fare for the same matched roundtrip product when the route is more concentrated. This is consistent with the main competition specification, where the income premium decreases as the number of competing airlines increases.

The HHI robustness check therefore confirms the interpretation that competition limits income-based directional price discrimination. The fact that both measures lead to the same qualitative conclusion

Table 21: Robustness Check: Income Difference and Market Concentration

	S1: Chains	S3: Synthetic
Regression coefficients		
Δ Income, \$10k	0.718	1.405***
	(0.922)	(0.431)
Adjusted HHI	-13.779***	-23.675***
	(4.971)	(1.957)
Δ Income × Adjusted HHI	4.500***	4.354***
	(1.592)	(0.717)
Δ Population, million	0.100	-0.093**
	(0.097)	(0.044)
Δ Hub	5.557***	6.874***
	(0.713)	(0.307)
Δ Temperature	-0.801***	0.123**
	(0.137)	(0.057)
Δ Holidays	–	-2.061***
	–	(0.510)
High-speed rail	7.176***	7.063***
	(1.706)	(0.688)
Alternative airport	-1.568	8.585***
	(1.261)	(0.586)
Legacy carrier	1.358	2.855***
	(0.943)	(0.465)
Alliance route	2.889	8.586***
	(1.767)	(0.791)
Liberalized route	-5.980***	0.836
	(2.196)	(1.120)
Constant	5.936*	-1.547
	(3.544)	(1.701)
Marginal effect of Δ Income by adjusted HHI		
HHI ≤ 0.4	2.186***	2.846***
	(0.490)	(0.231)
0.4 < HHI ≤ 0.7	3.103***	3.723***
	(0.366)	(0.167)
HHI > 0.7	5.017***	5.606***
	(0.768)	(0.338)
Observations	4,460	36,113
R^2	0.0605	0.0422

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. S1 = aggregated chain strategy; S3 = synthetic matching. The adjusted HHI measures market concentration. The marginal effects report the effect of a \$10,000 income difference by HHI category. Holiday differences are not included in S1 because the Strategy 1 outcome is not defined at the route-day level. Restricted samples exclude the bottom and top 5% of the outcome variable.

strengthens the result: airlines have more scope to segment passengers by origin income when market concentration is high, while stronger competition reduces the directional income premium.

H Triple Interaction: Income, Booking Horizon, and Competition

This section extends the dynamic analysis by estimating whether the income-based directional price gap varies jointly with the booking horizon and the number of competing airlines, as implied by the moderating competition mechanism derived in Appendix A.3. The specification is estimated for Strategy 3, which is the preferred strategy for this exercise because the synthetic matching construction isolates the relevant physical flight components within the roundtrip. The specification is estimated for Strategy 3, which is the preferred strategy for this exercise because the synthetic matching construction isolates the relevant physical flight components within the roundtrip. The comparison therefore remains tied to the same matched roundtrip product while allowing the income premium to vary both over time and across competitive environments. This directly relates to the theoretical mechanism in Section 3: the price gap is expected to widen as departure approaches, but this dynamic widening should be weaker when the number of competitors increases.

The estimated specification is

$$
\begin{aligned}
Y_i^{S3} = \alpha &+ \beta_1 \Delta Inc_i + \sum_k \gamma_k DBD_{ki} + \beta_2 NComp_i \\
&+ \sum_k \delta_k \left(\Delta Inc_i \cdot DBD_{ki}\right) + \beta_3 \left(\Delta Inc_i \cdot NComp_i\right) \\
&+ \sum_k \theta_k \left(DBD_{ki} \cdot NComp_i\right) + \sum_k \phi_k \left(\Delta Inc_i \cdot DBD_{ki} \cdot NComp_i\right) \\
&+ \lambda' X_i^{S3} + \eta' C_i + \varepsilon_i^{S3}.
\end{aligned}
\tag{39}
$$

The coefficient ϕ_k is the central term in this appendix specification. It indicates whether the dynamic change in the income premium over the booking horizon differs with the number of competing airlines.

Table 22: S3: Marginal Income Effects by Competition and Booking Horizon

	S3: Synthetic Matching
Control variables	
Δ Population, million	-0.105**
	(0.046)
Δ Hub	6.968***
	(0.316)
Δ Temperature	0.051
	(0.059)
Δ Holidays	-0.905
	(0.558)
HST	7.771***
	(0.703)
Alternative airport	9.457***
	(0.611)
Legacy carrier	2.924***
	(0.486)
Alliance route	8.601***
	(0.810)
Liberalized route	0.761
	(1.151)
Constant	-12.566***
	(2.160)
Marginal effect of Δ Income, \$10k	
1 Airline, 50–91 days	4.201***
	(0.438)
1 Airline, 29–49 days	6.085***
	(0.463)
1 Airline, 15–28 days	6.774***
	(0.581)
1 Airline, 8–14 days	5.199***
	(0.946)
2 Airlines, 50–91 days	3.156***
	(0.259)
2 Airlines, 29–49 days	4.487***
	(0.286)
2 Airlines, 15–28 days	5.459***
	(0.364)
2 Airlines, 8–14 days	5.385***
	(0.583)
3+ Airlines, 50–91 days	1.627***
	(0.518)
3+ Airlines, 29–49 days	2.095***
	(0.509)
3+ Airlines, 15–28 days	3.514***
	(0.626)
3+ Airlines, 8–14 days	5.660***
	(0.963)
Observations	36,113
R^2	0.0501

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The table reports S3 synthetic matching estimates. Marginal effects show the effect of a \$10,000 income difference by booking horizon and number of competing airlines. Restricted samples exclude the bottom and top 5% of the outcome variable.

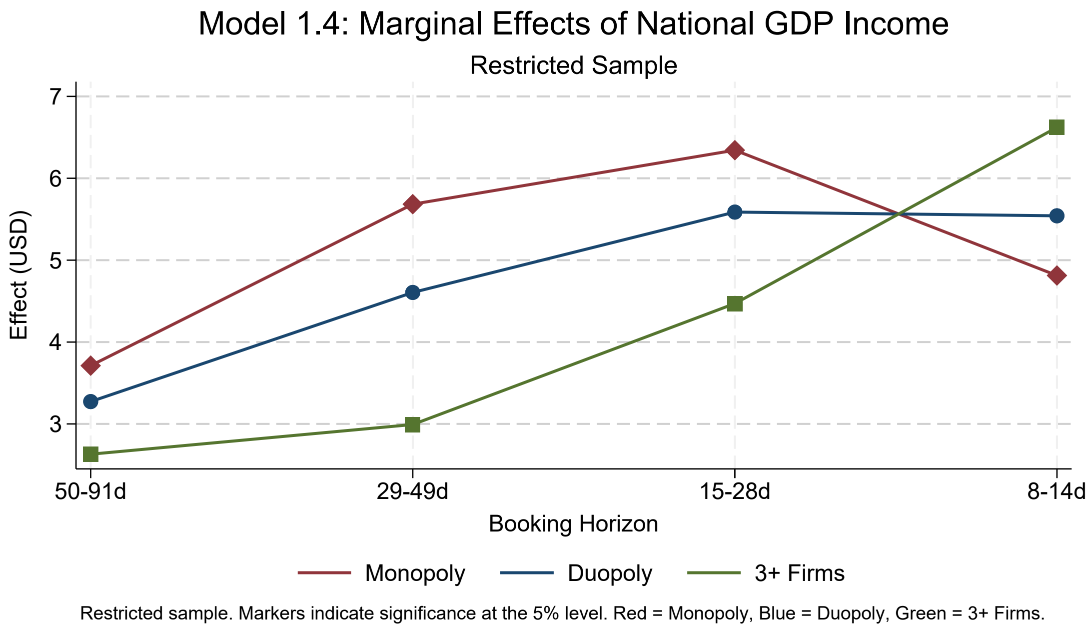

Figure 8: S3 Marginal Income Effects by Booking Horizon and Competition

Table 22 and Figure 8 show how the income premium changes jointly with booking horizon and competition. The marginal effects are positive and statistically significant in every cell. Thus, even in the most flexible specification, passengers from the richer origin pay more for the same synthetically matched roundtrip product. The size of this premium, however, depends strongly on both the number of competing airlines and the time remaining until departure.

For monopoly routes, a 10,000 USD increase in the income difference raises the directional fare premium by 4.20 USD at 50–91 days before departure, 6.09 USD at 29–49 days, 6.77 USD at 15–28 days, and 5.20 USD at 8–14 days. For routes with two airlines, the premium rises from 3.16 USD at 50–91 days to 4.49 USD at 29–49 days and 5.46 USD at 15–28 days, before remaining high at 5.39 USD in the 8–14 day window. For routes with three or more airlines, the income premium starts much lower at 1.63 USD in the longest booking window, but then increases steadily to 2.10 USD, 3.51 USD, and finally 5.66 USD in the 8–14 day window.

This pattern is closely linked to the theoretical model. The theory predicts that the directional price gap widens as departure approaches because passengers become less flexible, search less intensively, and face higher switching costs. This corresponds to the dynamic condition that the price gap increases when the booking horizon becomes shorter. The empirical results are consistent with this mechanism: in all competition groups, the income premium is generally higher in the shorter booking windows than in the longest booking window. This indicates that airlines increase income-based directional price discrimination when passengers have less time to adjust their travel plans.

At the same time, the model predicts that competition should weaken this dynamic increase because additional airlines increase outside options and raise cross-price elasticity. The results support this mechanism mainly in the earlier booking windows. At 50–91 days before departure, the marginal income effect is 4.20 USD on monopoly routes, 3.16 USD on duopoly routes, and only 1.63 USD on routes with three or more airlines. A similar ordering is visible at 29–49 days and 15–28 days. Thus, when passengers still have enough time to search and compare alternatives, competition clearly limits the income premium.

Closer to departure, however, this disciplining effect becomes weaker. In the 8–14 day window, the marginal effects are high across all competition groups: 5.20 USD for monopoly routes, 5.39 USD for duopoly routes, and 5.66 USD for routes with three or more airlines. This suggests that the time dimension can dominate the competition channel close to departure. Even when several airlines operate on a route, passengers from the richer origin may become sufficiently time-constrained and less price elastic that airlines can still extract a high directional premium.

The figure illustrates this interaction visually. The line for monopoly routes starts at a high level and peaks in the 15–28 day window. The duopoly line increases steadily and remains high close to departure. The line for routes with three or more firms starts much lower, which reflects the disciplining role of competition early in the booking horizon, but it rises sharply as departure approaches. Hence, the figure supports both parts of the theoretical prediction: competition reduces the income premium when passengers still have time to substitute, but this competitive constraint weakens in the final booking windows.

Declaration of AI Use

Generative artificial intelligence was used as a supporting tool during the preparation of this thesis for brainstorming and refining ideas, as well as for correcting syntax errors in Stata and Python. Additionally, the entire manuscript was proofread using AI tools to correct spelling and grammatical errors. All final content, analysis, and conclusions remain my own responsibility.

References

Air France-KLM Group. (2026). *Air France-KLM submits a non-binding offer for the acquisition of a minority stake in TAP Air Portugal*. Air France-KLM Corporate Newsroom, April 2, 2026. Available at: `https://www.airfranceklm.com/en/newsroom/air-france-klm-submits-non-binding-offer-acquisition-minority-stake-tap-air-portugal` (Accessed: June 4, 2026).

Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes*, 35(1), 124–140.

Belobaba, P. P. (1987). Survey paper—Airline yield management: An overview of seat inventory control. *Transportation Science*, 21(2), 63–73.

Borenstein, S. (1985). Price discrimination in free-entry markets. *The RAND Journal of Economics*, 16(3), 380–397.

Borenstein, S. (1989). Hubs and high fares: Dominance and market power in the U.S. airline industry. *The RAND Journal of Economics*, 20(3), 344–365.

British Airways. (2026). *Flight availability and fare quotes for the Barcelona (BCN) to London Heathrow (LHR) corridor*. Available at: `https://www.britishairways.com` (Accessed: May 15, 2026).

Brueckner, J. K., & Singer, E. (2019). Pricing by international airline alliances: A retrospective study. *Economics of Transportation*, 20, 100139.

Brueckner, J. K., Lee, D., & Singer, E. S. (2013). Airline competition and domestic U.S. airfares: A comprehensive reappraisal. *Economics of Transportation*, 2(1), 1–17.

Button, K. (2009). The impact of US–EU "Open Skies" agreement on airline market structures and airline networks. *Journal of Air Transport Management*, 15(2), 59–71.

Chang, Y., Winston, C., & Yan, J. (2022). Does Uber benefit travelers by price discrimination? *Journal of Law and Economics*, 65(S2), S433–S452.

Corts, K. S. (1998). Third-degree price discrimination in oligopoly: All-out competition and strategic commitment. *The RAND Journal of Economics*, 29(2), 306–323.

The Economist. (2026). Heathrow's third runway is turning into another infrastructure fiasco. *The Economist*, 26 February 2026.

Emirates. (2026). *Bilateral regulatory frameworks and network constraints: Additional points of call in the German aviation market*. Emirates Public Affairs Briefing. Available at: `https://www.emirates.com/de/english/about-us/our-communities/our-global-presence/emirates-and-germany/` (Accessed: May 25, 2026).

Escobari, D., & Jindapon, P. (2014). Price discrimination through refund contracts in airlines. *International Journal of Industrial Organization*, 34, 1–8.

Flightsfrom.com. (2026). *Route frequencies, operating airlines, and flight schedules*. Available at: `https://www.flightsfrom.com/` (Accessed: May 15, 2026).

Gaggero, A. A., & Luttmann, A. (2025). Systematic peak-load pricing during holiday periods: Evidence from the U.S. airline industry. *Economics of Transportation*, 41, 100395.

Great Circle Mapper. (2026). *Airport distance calculator*. Available at: `http://www.gcmap.com/` (Accessed: May 15, 2026).

Holmes, T. J. (1989). The effects of third-degree price discrimination in oligopoly. *American Economic Review*, 79(1), 244–250.

Hortaçsu, A., Natan, O. R., Parsley, H., Schwieg, T., & Williams, K. R. (2023). Organizational structure and pricing: Evidence from a large U.S. airline. *NBER Working Paper* No. 29508.

ITA Matrix. (2026). *ITA Matrix Airfare Search* (Version 3.0) [Software]. Available at: `https://matrix.itasoftware.com` (Accessed: May 12, 2026).

Lederman, M. (2008). Are frequent-flyer programs a cause of the "Hub Premium"? *Journal of Economics & Management Strategy*, 17(1), 35–66.

Lewis, M. S. (2020). On the absence of directional price discrimination in the U.S. airline industry. *The Journal of Industrial Economics*, 68(3), 556–581.

Lewis, M. S. (2021). Identifying airline price discrimination and the effect of competition. *International Journal of Industrial Organization*, 78, 102761.

Li, Y., & Shuai, J. (2019). Monopolistic competition, price discrimination and welfare. *Economics Letters*, 174, 114–117.

Lijesen, M. G., Rietveld, P., & Nijkamp, P. (2001). Hub premiums in European civil aviation. *Transport Policy*, 8(3), 193–199.

Lufthansa Group. (2025). *Lufthansa Group intends to play an active role in the privatization of TAP Air Portugal*. Lufthansa Group Newsroom, November 20, 2025. Available at: `https://newsroom.lufthansagroup.com/en/lufthansa-group-intends-to-play-an-active-role-in-the-privatization-of-tap-air-portugal/` (Accessed: June 4, 2026).

Luttmann, A. (2019). Evidence of directional price discrimination in the U.S. airline industry. *International Journal of Industrial Organization*, 62, 291–329.

Mantin, B., & Koo, B. (2009). Dynamic price dispersion in airline markets. *Transportation Research Part E: Logistics and Transportation Review*, 45(6), 1020–1029.

OECD. (2026). *Regional income and GDP statistics: TL2 regions*. OECD Regional Statistics. Available at: `https://www.oecd.org/regional/regional-statistics/` (Accessed: May 15, 2026).

Oesingmann, K., & Ennen, D. (2025). High-speed rail and domestic air transport: Evidence from Germany. *Transport Policy*, forthcoming.

Omio. (2026). *Rail connections and high-speed train availability*. Available at: `https://www.omio.com/` (Accessed: May 15, 2026).

oneworld. (2026). *Member airlines*. Available at: https://www.oneworld.com/members (Accessed: May 15, 2026).

Piga, C. A., Gaggero, A. A., & Alderighi, M. (2025). Demand-shock characteristics and pricing behavior: A natural experiment from UEFA Euro 2016. *Economics of Transportation*, 41, 100393.

Pigou, A. C. (1920). *The Economics of Welfare*. London: Macmillan.

RapidAPI. (2026). *Flights Scraper Sky API*. Available at: https://rapidapi.com/ (Accessed: May 15, 2026).

Román, C., Espino, R., & Martín, J. C. (2007). Competition of high-speed train with air transport: The case of Madrid–Barcelona. *Journal of Air Transport Management*, 13(5), 277–284.

Singh, N., & Vives, X. (1984). Price and quantity competition in a differentiated duopoly. *The RAND Journal of Economics*, 15(4), 546–554.

Skyscanner. (2026). *Flight search, fare quotes, and alternative airport information*. Available at: https://www.skyscanner.com/ (Accessed: May 15, 2026).

SkyTeam. (2026). *Member airlines*. Available at: https://www.skyteam.com/en/about/our-members (Accessed: May 15, 2026).

Stahl, D. O. (1996). Oligopolistic pricing with heterogeneous consumer search. *International Journal of Industrial Organization*, 14(2), 243–268.

Star Alliance. (2026). *Member airlines*. Available at: https://www.staralliance.com/en/members (Accessed: May 15, 2026).

Starkie, D. (1998). Allocating airport slots: A role for the market? *Journal of Air Transport Management*, 4(2), 111–116.

Stavins, J. (2001). Price discrimination in the airline market: The effect of market concentration. *The Review of Economics and Statistics*, 83(1), 200–202.

Thelle, M. H., & Sonne, M. L. (2018). Airport competition in Europe. *Journal of Air Transport Management*, 67, 232–240.

Thisse, J.-F., & Vives, X. (1988). On the strategic choice of spatial price policy. *American Economic Review*, 78(1), 122–137.

WeatherSpark. (2026). *Average monthly temperature data*. Available at: https://weatherspark.com/ (Accessed: May 15, 2026).

World Bank. (2026). *GDP per capita, current US dollars*. World Development Indicators. Available at: https://data.worldbank.org/indicator/NY.GDP.PCAP.CD (Accessed: May 15, 2026).

World Population Review. (2026). *City population data*. Available at: https://worldpopulationreview.com/ (Accessed: May 15, 2026).