


# THE PRICE OF ORBIT: HEDONIC ESTIMATES FROM THE SATELLITE LAUNCH MARKET

Sébastien Rouillon and Marc-Alexandre Sénégas

University of Bordeaux



SPACE ECONOMICS

**Working Paper n. 10/2026**

# The Price of Orbit: Hedonic Estimates from the Satellite Launch Market

September 4, 2026

Sébastien Rouillon<sup>1</sup> & Marc-Alexandre Sénégas<sup>2</sup>

**Keywords:** economics of space; launch market; hedonic regression.

**JEL Codes:** C23; L11; L93.

**Abstract:** We analyse the determinants of commercial satellite launch prices and the structure of the global launch market over 1996–2025 using a new database combining launch records, expert cost estimates, and regional industry revenue data. This approach provides more granular launch price estimates than previous studies. Econometric results show that payload mass and target altitude are the main determinants of launch prices, while real prices declined by 3.3% annually, reflecting technological progress. Market segmentation identifies three distinct service categories with differentiated price tiers, suggesting that competition is driven by provider specialisation across market segments rather than within a single global market.

## 1. Introduction

The space industry is experiencing a period of rapid growth, unprecedented since the end of the Cold War, marked by the relative decline of traditional institutional players and the entry of new private firms (Weinzierl, 2018). The pace of rocket launches has accelerated sharply (an average of 200 launches per year since 2020, compared with 75 between 1996 and 2019), as has the number of satellites placed in orbit (an average of 2,623 per year since 2020, compared with 173 between 1996 and 2019).

These changes have drawn growing academic attention, particularly efforts to understand the competitive dynamics at play (Dos Santos Paulino and Le Hir, 2016; Weinzierl, 2018; Guyot, Rao & Rouillon, 2023; Rouillon, 2026) and to warn against the resulting environmental risks (Adilov, Alexander & Cunningham, 2015; Adilov, Alexander & Cunningham, 2018). Yet the lack of public information on the costs and profitability of firms in the space industry greatly complicates such work. For example, the figures used in the recent economic literature to estimate the unit cost of a satellite (including design, manufacturing, launch and operation) range from less than \$1 million to over \$400 million (Bongers & Torres, 2023; Dos Santos & Le Hir, 2016; Guyot & Rouillon, 2023; Guyot, Rao & Rouillon, 2023; Nozawa et al., 2023; Rouillon, 2020; Rouillon, 2025, Rouillon, 2026). This dispersion reflects, in part, the heterogeneity of missions, launch vehicles, and satellites, but also the fact that firms in the sector do not disclose their prices for strategic reasons.

This article focuses on the factors influencing launch prices and their evolution since 1996.<sup>3</sup> It pursues two objectives: first, to identify their main determinants; second, to characterise whether and, if so, how those determinants shape the structure of the market for launch services.

---

<sup>1</sup> Bordeaux School of Economics (UMR CNRS 6060), University of Bordeaux, sebastien.rouillon@u-bordeaux.fr (corresponding author).

<sup>2</sup> Bordeaux School of Economics (UMR CNRS 6060), University of Bordeaux, marc-alexandre.senegas@u-bordeaux.fr.

<sup>3</sup> Adilov et al. (2023) note that launch costs are the second-largest item of expenditure for satellite operators, after satellite manufacturing costs.

To our knowledge, three other articles offer empirical analyses in a similar vein (Adilov et al., 2023; Kim, 2025; Su et al., 2026).

Adilov et al. (2023) retrieve data on the mass, altitude, and launch vehicle of 2,787 satellites placed in low Earth orbit between 2000 and 2020 from the database maintained by the Union of Concerned Scientists (UCS).<sup>4</sup> They cross-reference these data with expert estimates of the launch costs per kilogram of the main launch vehicles (Roberts, 2022; GAO, 2017).<sup>5</sup> Finally, they apply this framework to assign launch costs to the 2,787 satellites in their database. Their econometric analysis reveals a sustained decline in the cost of placing satellites into orbit over the period (around 6% per year per kilogram and 10% per year per satellite), which they attribute to both technical progress and the miniaturisation of satellites. By contrast, target altitude does not emerge as a determinant of launch costs. These results should be read with caution, however: the variables constructed by Adilov et al. (2023) are fragile, both because of the nature of the underlying launch-cost information (Kim, 2025) and because these costs are mapped to satellites without distinguishing the date or specific characteristics of each mission.

Kim (2025) draws 88 launches conducted by NASA between 1996 and 2024 from the confidential One NASA Cost Engineering (ONCE) database. For each mission, he records two quantitative variables (launch service cost and launch mass) and three qualitative ones (delivery location, mission complexity, and launch vehicle). His econometric analysis shows that the cost per kilogram of NASA missions rises markedly over the period (approximately 3% per year) and falls with launch mass. The other variables are not significant.

Su et al. (2026) differ from the other two references in that their primary objective is to develop an applied economic model to analyse changes in the U.S. launch services market since the early 2000s. In this regard, their findings confirm the driving role of U.S. government contracts in an industry characterized by a strong learning-by-doing dynamic (Weinzierl, 2018). This article, however, particularly caught our attention for the hedonic regression they conducted beforehand to calibrate their theoretical model. Their technical data, drawn from the General Catalog of Artificial Space Objects (GCAT), includes various characteristics of launch vehicles (size, mass, category, reusability, year of market introduction) used in 95% of worldwide launches during the period from 1985 to 2024. Their economic data comes from several sources (Federal Aviation Administration reports, the U.S. Department of Defense, news reports), from which the authors reconstruct launch prices by launcher family. As the authors themselves acknowledge, the granularity and reliability of the launch prices obtained in this way remain limited due to missing data, data available only in batches, and/or data available only for a limited set of launchers.<sup>6</sup> Using this database, Su et al. (2026) conducted a hedonic regression analysis of the launch prices per kilogram common to the different variants prevailing within each launch vehicle family, which then allowed them to assign, by extrapolation, a price to each launch recorded during the period covered by their study. The various econometric specifications they tested identify the launch vehicle's payload capacity to low Earth orbit, its length, and its reusability as key factors of the cost of launching a satellite.

---

<sup>4</sup> This database has not been updated since 2023.

<sup>5</sup> Adilov et al. (2023) compiled average cost estimates per kilogram for a total of 49 launch vehicles (Cf. Table A1).

<sup>6</sup> FAA reports do not disclose the prices of launch contracts, but rather provide order-of-magnitude figures by launch vehicles (sometimes in the form of ranges), based on expert estimates and/or past data. Furthermore, they contain numerous missing values starting in 2009. The reports published by the U.S. Department of Defense cover only a handful of launch vehicles (Atlas 5, Delta 4, and Falcon 9) and typically disclose only the contract values for batches of multiple launches.

In this article, we build on the empirical analyses of Adilov et al. (2023), Kim (2025) and Su et al. (2026), drawing on alternative and complementary data sources. First, we use the General Catalog of Artificial Space Objects (GCAT), continuously updated by Jonathan McDowell, which records the number and characteristics of all satellite launches conducted each year since the first launch, that of Sputnik in 1957. Second, following Adilov et al. (2023), we draw on expert estimates of the average launch costs of the main launch vehicles (Boone and Miller, 2016; Roberts, 2022; Kulu, 2025; and a handful of online sources), but extend these to a broader set of 99 launch vehicles (instead of the 49 covered by Adilov et al. (2023)). Third, and more importantly, we combine these expert estimates with economic data published annually on the aggregate revenues of satellite launchers by region. These data come from two main sources: the State of the Satellite Industry Reports (SSIR), published by the Satellite Industry Association (SIA), give the aggregate revenues of satellite launchers for the U.S. and the rest of the world, while the State of the European Space Industry Reports (SESIR), published by ASD-Eurospace and supplemented by financial statements and various online sources, give the corresponding figures for Europe.

Combining the two sources (expert and revenue) yields greater variability in individual launch prices than studies relying on expert data alone. This gain is worth stressing because, in the absence of publicly available information on individual launch contracts, there is no alternative but to rely on estimates rather than actual transaction prices, which, in general, limit the granularity of the price data. In our case, we assign a single price to all launches operated with the same launcher, in the same year, and in the same region. Even so, the cross-reference we implement performs relatively well, yielding 678 distinct launch price entries for 2,471 launches (19,171 satellites), or 3.64 launches per price on average.<sup>7</sup> By comparison, Adilov et al. (2023), who assign the same expert-estimated price to all satellites carried by a given launch vehicle, across all years and regions,<sup>8</sup> match 49 distinct launch price entries to (by our own calculations<sup>9</sup>) a total of 854 launches (3,880 satellites), or 17.43 launches per price on average.<sup>10</sup> This is a significant difference and the greater variability in launch prices, that our database exhibits, allows in turn for a potentially clearer picture of market trends and provides firmer ground for relating the former prices to launch-specific explanatory variables.

Based on the comprehensive database we have built, our analysis proceeds in two steps.

First, we run an econometric analysis of the determinants of the prices of all commercial satellite launches conducted in Europe, the United States, and the rest of the world between 1996 and 2025, focusing on the role played by the physical, technical and institutional attributes of the launches. We base inference on cluster-robust standard errors to account for the within-launcher cluster correlation of the measurement error that is potentially implied by the method used to build the individual launch prices.

---

<sup>7</sup> It should be noted that this average masks a significant disparity among launch vehicles. For example, Falcon 9 has a total of 543 launches but only 14 price data points, resulting in an average of 38.78 launches per launch price.

<sup>8</sup> Adilov et al. (2023) acknowledge this directly: their cost data give a single average per-kilogram figure per launch vehicle and do not distinguish among the satellites it carries, so the same figure applies to all of them. The use of expert data only in Adilov et al. (2023) has also been questioned by Kim (2025).

<sup>9</sup> Considering missions to low Earth orbit between 2000 and 2020, as in Adilov et al. (2023), we identify 854 launches (3,880 satellites) in our database. Adilov et al. (2023) do not provide this information, as their unit of analysis is the satellite. For the sake of comparison, their database includes 2,787 satellites.

<sup>10</sup> Note also that Adilov et al. (2023) limit their analysis to low Earth orbit, whereas we consider all destination orbits.

We then turn to the analysis of the structure of the market for launch services. Based on a k-means clustering of the physical attributes of each launch, we partition the sample into a small number of segments. The price equation is then recast in terms of “specialisation” cells combining orbital regime and payload mass, so as to test whether positioning along these dimensions is associated with distinct price tiers.

Several findings stand out.

Payload mass and target altitude (apogee) are the two robust attribute determinants of price, with positive elasticities of about 0.44 and 0.15 that survive our most demanding inference. Launch prices, computed in real terms, also fell steadily, at roughly 3.3% per year over 1996–2025, consistent with ongoing, cost-reducing technical progress. Regional differentials are large as well: relative to the United States, European launches command a premium of about 0.55 log points and Rest-of-the-World launches a discount of about 0.60, mirroring the segmentation of the industry.

As for the market structure, the k-means classification partitions the launches into three groups—light institutional missions to low orbit, heavy missions to geostationary orbit, and recent low-orbit constellations—each tied to distinct launch service providers and eras. These groups align in turn with a coherent ladder of price tiers defined jointly by orbital regime and payload mass. Strategies of specialisation and positioning thus appear to be associated with the price differentials we observe—the second of the two questions raised at the outset.

The latter findings complement the picture drawn by existing work on the organisation of the launch industry. Institutional accounts place launch within an upstream segment supplied by a small number of providers (OECD, 2022), and the empirical literature treats this segment as an oligopoly in which inelastic institutional demand limits the competitive effects of entry (Kim, 2025) and in which procurement and learning-by-doing have repeatedly reshaped the field around a few players (Su et al. 2026; Weinzierl et al., 2021). Our results are consistent with this account but read it at a finer grain. Competition does not take place in a single market: the relevant arena is the satellite service segment, delineated by the physical characteristics of the launch, so that providers specialise across segments rather than competing uniformly, and a given provider may face one set of rivals on one segment while being absent from another. The resulting positioning resolves what these accounts treat as a single oligopoly into segment-specific contests—Arianespace and Russian operators on heavy geostationary launches, Chinese and Russian actors on institutional low-orbit launches, and SpaceX on the recent commercial multi-satellite segment.

The remainder of the paper is organised as follows. Section 2 presents the database and sets out the procedure by which we reconstruct individual launch prices from expert costs and aggregate regional revenues, together with a first descriptive overview of launch activity. Section 3 specifies the hedonic price equation and reports the baseline estimates of the price determinants. Section 4 turns to market structure, combining an unsupervised classification of launches with a parsimonious specialisation model to relate orbital and mass positioning to distinct price tiers. Section 5 concludes and outlines directions for further work.

## 2. Data: database and summary statistics

Our database combines technical and economic information from multiple sources. In the description below, we isolate prices from the other data available as the former are subject to a specific reconstruction process that needs to be set out (see Section 2.2).

### 2.1 Technical, physical and institutional data

Those data come from the General Catalog of Artificial Space Objects (GCAT), continuously updated by Jonathan McDowell, which lists and describes all satellite launches conducted each year. We describe which database fields have been considered in building our own database.

- **LAUNCHES**: The set of launches included in our database

As we focus on launches dedicated to *commercial* satellite missions, we have to remove from the GCAT database launches of objects other than satellites, as well as failed and non-commercial missions.<sup>11</sup>

To eliminate non-satellite missions (flights with humans aboard, resupply missions to the International Space Station, tests, etc.), we used two GCAT fields (*Type* and *Category*), which specify the type and category of each payload. This data allows us to identify the satellites among all the objects carried. In the case of mixed missions (carrying satellites and other objects), we include only launches whose payload consists of at least 50% satellites.

To exclude failed missions, we used a GCAT field (*Launch\_Code*), which indicates the type and success of the mission. These fields allow us to identify and exclude missions with an orbital insertion rate of less than 50%.<sup>12</sup>

To filter out non-commercial missions, we used four GCAT fields (*Launch\_Code*, *Owner*, *Agency*, and *Group*), which provide the mission type, the identities of the satellite and launch vehicle owners, and the nature of the launch service. These fields allow us to identify and exclude non-orbital missions and launches conducted by government operators carrying satellites they own.

- **APOGEE**: the maximum insertion altitude of the payload (km)

Here we use three GCAT fields (all labelled *Apogee*, drawn from three different databases respectively dedicated to objects, launches, and launch vehicles), which provide the apogee of each payload, the flight apogee, and the launch vehicle's theoretical maximum apogee. When the mission carries multiple objects, we retain the highest apogee so that the constructed variable measures the maximum power required for the mission. In the event of missing or

---

<sup>11</sup> « Not included in SSIR: Launches of non-satellite missions (crew and cargo to ISS, other space vehicles); Launches not procured commercially, i.e. a government agency providing a launch of a satellite owned by the same agency. » (quote of SIA reports, 1996 to 2025).

<sup>12</sup> Contracts between the customer (satellite owner) and the launch service provider sometimes include a condition precedent in the event of a failure to place the satellite into orbit. Without information on the type of contract, we have chosen to count missions with a success rate exceeding 50% as successful. It should be noted that this assumption has no significant effect on our results, as missions with a 100% success rate account for 99.26% of the total sample.

inconsistent values for the apogee of the object, we correct the result by substituting the launch apogee and, as a last resort, the theoretical maximum apogee of the launch vehicle.<sup>13</sup>

- **MASS**: the total mass of the onboard satellites (kg)

Here we use a GCAT field (*Mass*), which specifies the mass of each carried object. When the mission carries multiple objects, we sum the masses of the various components.<sup>14</sup>

- **NBS**: the number of satellites on board

For each launch, we count the number of satellites carried.

- **LATITUDE**: the latitude of the launch site

Here we use a GCAT field (*#Site, Code and Latitude*), which gives the launch site, the launch pad within the launch site and its latitude.

- **INSTIT**: a binary variable indicating institutional missions

Here we use a GCAT field (*Class*), which specifies the class of each payload object. Civil and defence objects are counted as institutional.<sup>15</sup> A mission is classified as institutional if at least 50% of the payload objects fall into these two categories.

- **REGION**: the 3 regions (EU, ROW, US)

Launches are assigned based on a GCAT field (*LVState*), which associates a country with each agency. Launches operated by Arianespace using the Russian Soyuz launch vehicle are handled separately. They are assigned to Europe if they launch from the Guiana Space Centre and to Russia if they launch from Baikonur.

## 2.2 Rebuilding prices

Ideally, relating launch characteristics to prices would call for the market price of each individual launch — the price agreed in each transaction. In the absence of publicly disclosed information on those prices, we need to resort to an estimate that is as individualised as the available data may allow.

Expert estimates are the usual source for launch costs in the literature. However, they are essentially constant over time and carry poor market information (in terms of structure and evolution).

---

<sup>13</sup> We do not use the launch apogee directly, due to the uncertain quality of this data: « In many cases for non-orbital launches the value is a guess based on known capabilities of the launch vehicle. » (quote of GCAT database). It should be noted that outliers account for only 0.28% of the sample.

<sup>14</sup> In the case of Space Shuttle flights, several components of the launch vehicle (external tank, boosters, etc.), which are included in the cargo by GCAT, were excluded because they were irrelevant to our analysis. Furthermore, we do not directly use the total cargo mass (the *OrbPay* field in GCAT) due to the uncertain quality of this data: « Current values are autogenerated from the payload tables and there are likely some special cases not treated properly yet. » (quote of GCAT database).

<sup>15</sup> « Civil (government but not military); Defense (govt. military, intelligence) » (quote of GCAT database).

On the other hand, we have access to annual revenue figures for the launch services sector by region. Since this aggregate revenue can be seen as the product of quantity and an average price, the revenue per launch in a region-year can serve as an observed average price reflecting demand and competition, not cost alone.<sup>16</sup>

The two sources (expert and revenue data) can be considered as complementary as each carries information the other lacks. Expert estimates discriminate among launch vehicles but, attached to the vehicle, are constant across both regions and years; regional revenues discriminate across regions and years but, being averages, assign a common price to every launch within a given region-year. Combining them to produce a reconstructed price makes it vary along all three margins at once — vehicle, region, and year — where each source alone would hold at least one of them fixed.

As for the first source, we collected publicly available estimates of launch costs for most of the launch vehicles used between 1996 and 2025 (Boone and Miller, 2016; Roberts, 2022; Kulu, 2025; and a handful of online sources).<sup>17</sup> As for the second source, the data for annual revenues come from SIA reports for the U.S. and the rest of the world, and from ASD-Eurospace reports (supplemented by financial statements) for Europe. The compiled data covers the period from 1996 to 2025. We have harmonised these figures by converting them to dollars<sup>18</sup> and then deflating them to constant 2025 dollars.<sup>19</sup> We calculate the average revenue per launch for each region and each year, by dividing annual revenues by the number of launches operated for commercial satellite missions. Since the latter figures are not provided in the SIA and ASD-Eurospace reports (with some exceptions, or only in aggregate form), we reconstruct them using the GCAT database, following the methodology explained in Section 2.1.<sup>20</sup>

The combination operation consists in cross-referencing the experts' estimates of the average launch cost per launch vehicle with the average annual revenue per region in order to reconstruct the prices for individual launches. In this operation, we let the revenues set the price level in each region-year pair, while expert costs determine only the relative positioning of the vehicles actually flown.

More specifically, for each region and each year, we list, for each launch, the launch vehicle used and its associated estimated expert cost. This allows us to calculate an average launch cost for the launches operated, according to expert estimates, by dividing the sum of the estimated costs by the number of launches in that region and year. By calculating the ratio of each launcher's estimated cost to the average (regional and yearly) estimated cost, we obtain an index that measures the launcher's "comparative advantage" with respect to other launchers in the same region and year. Finally, we calculate the price of each launch by multiplying the known

---

<sup>16</sup> Note that the regional resolution is a property of our revenue sources, not a claim that the launch market is segmented by region: what the revenues add, beyond their regional and annual granularity, is information on the prices actually realised, which reflects market conditions beyond expert estimates alone.

<sup>17</sup> We have consulted the following websites: <http://www.astronautix.com/c/costpriceanholeदारnthing.html>; <https://en.wikipedia.org/wiki/Gravity-1>; [https://en.wikipedia.org/wiki/Space\\_Shuttle](https://en.wikipedia.org/wiki/Space_Shuttle); <https://nextspaceflight.com/rockets/105/>; <https://nextspaceflight.com/rockets/126/>.

<sup>18</sup> ASD-Eurospace publishes annual revenue figures for the launch services sector in euros. The exchange rate used for conversion to dollars is available on the INSEE website at: <https://www.insee.fr/en/statistiques/serie/010002053#Telechargement>.

<sup>19</sup> The deflator is available on the Federal Reserve Bank of St. Louis website at: <https://fred.stlouisfed.org/series/GDPDEF>.

<sup>20</sup> We used the launch data provided by the SIA and ASD-Eurospace, when available. For the years for which launch data are available, we calculated a coefficient of determination of 0.99 and 0.93, respectively, between the series published by the SIA and ASD-Eurospace and those that we derived from GCAT.

average price for the region and year (derived from reports published by the SIA and ASD-Eurospace) by the former launcher's index.

For clarity, we illustrate the various steps of the procedure using a hypothetical example. It is described in Table 1, listing region A's launches in years  $t$  and  $t+1$ . The number of launches is 3 (indexed #1 to #3) in year  $t$  and 2 in year  $t+1$  (indexed #1 to #2). The launch vehicles used are X, Y, and Z, associated with costs estimated by experts at 60, 100, and 50, respectively. In year  $t$ , launch vehicles X, Y, and Z are each used once. In year  $t+1$ , launch vehicles X and Y are each used once, and launch vehicle Z is not used. From this, we derive an average cost estimated by experts to be 70 ( $60 + 100 + 50 = 210$ , divided by 3 launches) in year  $t$  and 80 ( $60 + 100 = 160$ , divided by 2 launches) in year  $t+1$ . The cost ratios for each launcher, relative to other launchers from the same year, are calculated by dividing the cost estimated by experts for each launcher by the estimated average cost for that year. They provide the comparative advantage index for launchers. The calculation is explained in the sixth column of the table. The known revenue from launches in region A is  $RT_t = 270$  in year  $t$  and  $RT_{t+1} = 240$  in year  $t+1$ . From this, we derive a revenue per launch of  $RM_t = 90$  ( $270$ , divided by 3 launches) in year  $t$  and  $RM_{t+1} = 120$  in year  $t+1$  ( $240$ , divided by 2 launches). Finally, the reconstructed price for each launch is obtained by multiplying each launcher's index for that year by the average revenue for the year. This final calculation is shown in the last column of the table.

**Table 1.** *Price construction (hypothetical example)*

| Launch data |        |        | Experts' cost data |        |          | Economic data |        | Result              |
|-------------|--------|--------|--------------------|--------|----------|---------------|--------|---------------------|
| Year        | Region | Launch | Launcher           | Expert | Index    | $RT_t$        | $RM_t$ | Rec. Price          |
| $t$         | A      | #1     | X                  | 60     | $60/70$  | 270           | 90     | $60/70 \times 90$   |
| $t$         | A      | #2     | Y                  | 100    | $100/70$ | 270           | 90     | $100/70 \times 90$  |
| $t$         | A      | #3     | Z                  | 50     | $50/70$  | 270           | 90     | $50/70 \times 90$   |
| $t+1$       | A      | #1     | X                  | 60     | $60/80$  | 240           | 120    | $60/80 \times 120$  |
| $t+1$       | A      | #2     | Y                  | 100    | $100/80$ | 240           | 120    | $100/80 \times 120$ |

Source: authors' calculations.

Two properties deserve emphasis. First, the relative-cost index averages to one within each region-year, so the sum of the reconstructed prices corresponds to the observed sector revenue: the expert data redistribute the observed total without altering it, and the level stays anchored in observed market magnitudes. Second, although each vehicle's expert cost is a fixed constant, the same vehicle is priced differently across region-years through two channels — the evolving revenue per launch, and the cost index, whose denominator shifts with the mix of vehicles flown.

Finally, the single margin our reconstruction leaves flat — launches sharing a vehicle within the same region and year — generates a correlated measurement error that we address through cluster-robust inference (see Section 3.2 for details).

## 2.3 Launch activity, 1996–2025: a regional overview

We provide here a first, descriptive overview of launch activity over the period covered by our sample (1996–2025). For each region — Europe (EU), the United States (US), and the rest of the world (ROW) — we report in Tables 2 to 4 how launches are distributed across three characteristics taken one at a time: the altitude reached (denoted as apogee), the payload mass (denoted as mass) and the number of payload objects deployed (denoted as the number of

satellites).<sup>21</sup> Each table uses the same class boundaries throughout while every cell gives the share of a region's own launches that fall in the class. Because the shares are read over the whole period, they describe an average regional profile rather than the pattern of any single year.

**Table 2.** *Distribution of launches by apogee (km), as a share of each region's launches (1996–2025).*

| Region (n) \ Apogee | < 400             | [400, 600)        | [600, 800) | [800, 2000) | [2000, 35000) | ≥ 35000           |
|---------------------|-------------------|-------------------|------------|-------------|---------------|-------------------|
| EU (224)            | 0%                | 3%                | 12%        | 4%          | 8%            | <b><u>73%</u></b> |
| US (998)            | <b><u>36%</u></b> | 12%               | 10%        | 7%          | 6%            | <b><u>29%</u></b> |
| ROW (1249)          | 4%                | <b><u>29%</u></b> | 17%        | 15%         | 6%            | <b><u>29%</u></b> |

**Table 3.** *Distribution of launches by mass (kg), as a share of each region's launches (1996–2025).*

| Region (n) \ Mass | < 500 | [500, 1000) | [1000, 2000) | [2000, 4000)      | [4000, 10000)     | ≥ 10000           |
|-------------------|-------|-------------|--------------|-------------------|-------------------|-------------------|
| EU (224)          | 1%    | 4%          | 10%          | 26%               | <b><u>57%</u></b> | 1%                |
| US (998)          | 8%    | 4%          | 11%          | 17%               | 22%               | <b><u>39%</u></b> |
| ROW (1249)        | 17%   | 15%         | 20%          | <b><u>24%</u></b> | <b><u>23%</u></b> | 1%                |

**Table 4.** *Distribution of launches by number of satellites, as a share of each region's launches (1996–2025).*

| Region (n) \ Nb of satellites | 1                 | 2                 | 3  | [4, 6) | [6, 10) | ≥ 10              |
|-------------------------------|-------------------|-------------------|----|--------|---------|-------------------|
| EU (224)                      | <b><u>36%</u></b> | <b><u>50%</u></b> | 3% | 6%     | 3%      | 2%                |
| US (998)                      | <b><u>42%</u></b> | 6%                | 3% | 5%     | 3%      | <b><u>42%</u></b> |
| ROW (1249)                    | <b><u>57%</u></b> | 13%               | 8% | 10%    | 6%      | 6%                |

**Reading:** classes are left-closed, right-open (e.g. [400, 600) includes 400 but not 600). Shares are computed over each region's own launches; rows may not sum exactly to 100% because of rounding.

For each region, the following picture emerges.

**Europe.** Tables 2–4 point to a fairly concentrated profile. In apogee terms, 73% of launches reach beyond 35,000 km, that is, geostationary-transfer orbit (Table 2). Turning to mass, most European launches carry medium-to-heavy satellites — 57% in the 4–10 t class and a further 26% in the 2–4 t class — while the lighter classes are almost empty (Table 3). The number of satellites comes in small values: half of launches carry two satellites and a further 36% a single one (Table 4). Read together, the three tables are consistent with a small number of medium-to-heavy satellites delivered to high orbit — a profile in line with the commercial and institutional telecommunications missions that dominate European activity and with the dual-passenger geostationary model operated by Arianespace, which conducts 96% of European launches.

**United States.** The US tables are more dispersed and, on two of the three dimensions, markedly bimodal. Apogee splits between low orbit (36% below 400 km) and geostationary transfer (29% above 35,000 km), with little in between (Table 2). Mass divides between very heavy launches (39% above 10 t) and heavy ones (22% in the 4–10 t class), with a non-negligible light tail (8% below 500 kg) (Table 3). The number of satellites is almost evenly divided between single-satellite launches (42%) and large batches of ten or more (42%) (Table 4). These contrasts are consistent with the coexistence of two phases of US activity: legacy single-satellite missions to

<sup>21</sup> In Tables 1 through 3, values greater than 20% are highlighted and in bold.

high orbit (Delta 7000-series and Atlas 5) and, from 2019 onward, very heavy multi-satellite batches to low orbit — Falcon 9, which alone accounts for 54% of US launches over the period, deploying Starlink, which accounts for 33%.

**Rest of the world.** The ROW tables are flatter. Apogee leans toward low-to-medium orbits — 29% in the 400–600 km band and 17% in the 600–800 km band — while retaining a sizeable geostationary share (29%) (Table 2). Mass is spread fairly evenly across classes, from below 500 kg (17%) to the 4–10 t class (23%), with no single class dominating (Table 3). The number of satellites is dominated by single-satellite launches (57%), complemented by small multi-satellite rideshares (Table 4). This more diffuse pattern is consistent with the heterogeneity of the providers active in the region — Long March (CASC), Proton and Soyuz, Electron, and H-2 among others — and points to a generalist rather than a single dominant type of launch.

Table 5 completes this picture by reporting, for each region, the dominant launch vehicles and the leading launch service providers behind these profiles.

**Table 5.** Launchers and launch service providers: regional features.<sup>22</sup>

| Region (n)        | Dominant launchers [LSP] (share of region)                                                                                                                                                                                                 | Leading LSPs (share of region)                                        | Main features                                                                                                                                                                                                                                      |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>EU (224)</b>   | Ariane 5G — 48% [ <i>Arianespace</i> ]<br>Ariane 4 family — 28% [ <i>Arianespace</i> ]<br>Soyuz, from Kourou — 12% [ <i>Arianespace</i> ]<br>Vega / Vega C — 10% [ <i>Arianespace</i> ]<br>Ariane 6 — 2% [ <i>Arianespace</i> ]            | Arianespace — 96%<br>European Space Agency — 4%                       | One or two medium-to-heavy satellites delivered to high orbit (geostationary transfer), chiefly commercial telecommunications and European institutional satellites.                                                                               |
| <b>US (998)</b>   | Falcon 9 — 54% [ <i>SpaceX</i> ]<br>Delta 7000-series — 11% [ <i>Boeing, later ULA</i> ]<br>Atlas 5 — 10% [ <i>United Launch Alliance</i> ]<br>(plus Pegasus XL, Zenit 3SL, Delta 4) [ <i>Orbital, Sea Launch, ULA</i> ]                   | SpaceX — 56%<br>United Launch Alliance — 15%<br>Boeing — 7%           | Two superimposed patterns: in the earlier years (Delta, Atlas), single heavy satellites to geostationary orbit; from 2019 onward (Falcon 9), very heavy batches of many satellites to low Earth orbit — the Starlink constellation.                |
| <b>ROW (1249)</b> | Long March family — 41% [ <i>CASC: CALT &amp; SAST</i> ]<br>Proton — 13% [ <i>Khrunichev / ILS</i> ]<br>Soyuz — 12% [ <i>Roscosmos / military</i> ]<br>Electron — 5% [ <i>Rocket Lab</i> ]<br>(plus H-II and others) [ <i>Mitsubishi</i> ] | CASC (CALT + SAST) — 41%<br>Khrunichev / ILS — 14%<br>Rocket Lab — 5% | Predominantly single satellites across the full range of orbits — to geostationary orbit and to low / Sun-synchronous orbits — with a growing recent share of multi-satellite launches to low orbit (Chinese constellations, Electron rideshares). |

**Reading:** the entity in brackets is the launch service provider (LSP). Orbit terms: low Earth orbit (below ~600 km); Sun-synchronous orbit (a low/polar orbit used for Earth observation); geostationary transfer / geostationary orbit (~36,000 km, telecommunications). Providers: CASC = China Aerospace Science and Technology Corporation (parent of CALT and SAST); ULA = United Launch Alliance (Boeing–Lockheed Martin joint venture); ILS = International Launch Services (markets Proton). Shares are computed over each region’s own launches.

Turning to prices, we plot in Figure 1, for each region, the reconstructed launch price (vertical axis) against the expert launch cost (horizontal axis), one point per launcher and year, each

<sup>22</sup> In Table 5, providers are consolidated at the corporate-group level for readability (CASC pools CALT and SAST; ULA, Khrunichev/ILS and the European Space Agency likewise pool their respective agency codes). The econometric analysis clusters either on (individual) launch service providers or launchers.

marker's area proportional to the number of launches it represents<sup>23</sup>; the dashed line is the 45° line and the red markers locate each vehicle's expert cost along it. Because the expert cost is constant for a given vehicle, each launcher forms a vertical column whose spread shows how its reconstructed price moves from year to year, while points above the line indicate a price above expert cost and points below the reverse.

The three panels differ in a telling way. In the United States the points hug the 45° line across the whole range — sitting on or just below it, with wider scatter only at the few highest-cost, low-volume vehicles — so that the reconstructed price stays close to expert cost, which is reassuring for the consistency of the two independent sources there. In Europe, they lie predominantly above the line, most visibly in the tall Ariane 5 column (near  $L\_COST \approx 283$  M\$), which carries the region's largest markers and where the reconstructed price varies widely from year to year around a fixed expert cost and sits, more often than not, above it. In the rest of the world the configuration is reversed and systematic: the reconstructed price lies below expert cost for the great majority of launches, and by a wide margin, the largest markers — the highest-volume vehicles — lying furthest below the line.

Taking the expert estimates as reliable, where the two series diverge, the gap points to market conditions at the time of launch departing from those underlying the expert estimates, chiefly the balance of supply and demand. On the other hand, the persistent, one-sided shortfall in the rest of the world could, however, be interpreted more as an indicator of data quality than as a discrepancy between prices and expert estimates. This asymmetry is itself another rationale for relying on neither source alone — the revenue data anchor the level of the price, the expert costs the relative positioning of the vehicles, and the comparison of the two serves as a built-in check on data quality.

### 3. Empirical Strategy and Results

#### 3.1 Econometric specification

##### Econometric approach

The data take the form of a pooling of yearly (unbalanced) cross-sections of individual launch observations over the 1996–2025 window.<sup>24</sup>

We use them to examine the relationship between the (log) launch price and a parsimonious set of service characteristics recorded for each launch.

We consider the following, baseline specification:


$$\ln(P_{it}) = \beta_0 + \beta_1 \cdot \text{TIMELAPSE}_t + \beta_2 \cdot \ln(MASS_{it}) + \beta_3 \cdot \ln(APOGEE_{it}) + \beta_4 \cdot \ln(NBS_{it}) + \beta_5 \cdot \ln(1 + LAT_{it}) + \beta_6 \cdot \text{INSTIT}_{it} + \beta_7 \cdot \text{STARLINK}_{it} + \varepsilon_{it} \quad (1)$$

<sup>23</sup> Plotting one point per launch rather than per launcher and year would leave the figure qualitatively unchanged: since the reconstructed price is identical for all launches sharing a launcher, region and year, the additional observations fall exactly on existing points; rather than accumulate as overplotting, their number is reflected in the area of each marker.

<sup>24</sup> We therefore do not have a panel in the econometric sense of the term. Also note that the data have a natural grouped structure: launches nest within launch vehicles (LV\_LAUNCHERS) and within operating companies (LV\_COMPANIES). The operating company is the launch service provider (LSP), and we use the two terms interchangeably below. This grouped structure is central to the inference strategy that follows.

where  $i$  indexes the launch operated and  $t$  the year of the launch;  $TIMELAPSE = YEAR - 1996$  is a linear time trend;  $MASS$  is the total mass injected (kg);  $APOGEE$  is the target altitude (km);  $NBS$  is the number of satellites embarked and deployed;  $LAT$  is the launch-site latitude with the related regressor  $\ln(1 + LAT)^{25}$ ;  $INSTIT$  is an indicator for institutional missions; and  $STARLINK$  is a dummy equal to one for SpaceX Starlink batch launches.<sup>26</sup> Continuous regressors enter in logarithms, so their coefficients read as elasticities.

**Figure 1.** Reconstructed launch price (vertical axis) against the expert launch cost. Marker area is proportional to the number of launches represented.



<sup>25</sup> It aims to highlight the presumed difference in launch costs associated with proximity to the equator.

<sup>26</sup> Falcon 9 missions to deploy the Starlink constellation are subject to internal pricing, which may differ from that of the rest of the commercial market.

As such, this specification places the analysis within the hedonic tradition of Court (1939), Griliches (1961) and Rosen (1974), in which the price of a differentiated good is expressed as a function of its characteristics and the estimated coefficients recover the implicit prices that the market attaches to them. A launch is precisely such a good: buyers do not purchase "a launch" in the abstract but a bundle of attributes — mass to orbit, apogee, number of satellites — each of which is valued in the market.<sup>27</sup> Following Rosen (1974), the hedonic equilibrium price function is the envelope of buyers' bid functions and sellers' offer functions, and therefore relates the launch price both to the buyers' valuation of attributes and to the sellers' cost of providing them. In our setting the cost reading is especially transparent, since the technical attributes are physically grounded cost drivers: the positive coefficients on  $\ln$  APOGEE and  $\ln$  MASS thus admit a direct reading as the marginal cost gradient of the launch service.<sup>28</sup> By the same envelope logic, these coefficients also admit a demand-side reading: launch demand is derived from the demand for orbital services, so that the buyer's valuation of a given apogee, injected mass and number of satellites reflects the downstream service that each such bundle enables.

This framing has a direct consequence for whether or not we include fixed effects. Indeed, identification of a hedonic function rests on variation in attributes *across* products. The standard practice in the hedonic literature is consequently not to absorb seller or producer heterogeneity through entity fixed effects, since doing so would purge precisely the cross-product variation that the hedonic regression exploits. Thus, our decision not to include launcher or company fixed effects conforms the canonical specification for a hedonic model.

This positive justification is reinforced, in our particular case, by a mechanical one rooted in the way prices are constructed. As established earlier (cf. Section 2.2), the price is, by construction, the same for a given launch vehicle in a given year and in a given region, for all launches that share those characteristics. It follows that launcher dummies—launcher fixed effects—would absorb essentially all between-launcher variation in the dependent variable. Company fixed effects are less extreme but would sweep out the between-company dispersion in price levels that constitutes one of the objects of interest of the study.

Regional effects, by contrast, are admissible and informative. The hedonic literature routinely includes time or market fixed effects to capture common conditions that shift the price surface without being attributes of the good itself (housing hedonics with location and period effects are the textbook case; see, e.g., Kuminoff, Parmeter and Pope, 2010). In our case, the variable Region—US, EU, ROW—could play exactly this role, as it captures an exogenous market factor (industrial policy, captive institutional demand, exchange-rate and cost environments) without absorbing the attribute-based differentiation we seek to measure. We therefore report a regionally augmented specification alongside the baseline. As shown below, the regional

---

<sup>27</sup> On the demand side, the differentiation among launches is essentially horizontal—the type of orbit and the payload configuration serve distinct objectives and do not fall along a single quality scale. On the supply side, however, the cost of launches should vary monotonically depending on the characteristics of the launch service—higher target orbits require a greater velocity increment; and a larger mass requires a larger rocket. This reading fits perfectly within the hedonic framework, whose feature space (Lancaster, 1966) and multi-attribute matching formulations (Chernozhukov et al., 2021; Bajari and Benkard, 2005) constitute models of product differentiation, not quality ranking.

<sup>28</sup> Our study anchors methodologically in this hedonic tradition and its modern econometric development: the identification results of Ekeland, Heckman and Nesheim (2004) and Heckman, Matzkin and Nesheim (2010), their multi-attribute extension by Chernozhukov, Galichon, Henry and Pass (2021), and the hedonic demand framework of Bajari and Benkard (2005).

differentials are large and economically meaningful, which both vindicates their inclusion as market controls and leaves the attribute coefficients qualitatively intact.

### **Inference: cluster-robust standard errors**

The coefficients of the baseline equation (1) are estimated by the standard pooled OLS estimator. Inference, however, must be handled with particular care, owing notably to the particular nature of the dependent variable.

Because the price variable is reconstructed from expert sources and aggregated launch service providers' revenues rather than quoted transaction by transaction, the dependent variable carries a measurement component. Moreover, the measurement errors, and in turn the regression disturbances, are likely to be correlated across observations sharing a given launch vehicle, region and year, due to the features of the price reconstruction process. Ignoring this dependence would mechanically understate standard errors and inflate  $t$ -statistics.

We accordingly base inference on cluster-robust variance estimators, of the sandwich form:

$$V_{CR}(\beta) = (X'X)^{-1} \left[ \sum_{g=1}^G X_g' A_g \hat{u}_g \hat{u}_g' A_g' X_g \right] (X'X)^{-1} \quad (2)$$

where  $g = 1, \dots, G$  indexes the clusters,  $X_g$  is the matrix of regressors for the observations belonging to cluster  $g$ , and  $\hat{u}_g$  the corresponding vector of OLS residuals; the adjustment matrices  $A_g = (I - H_{gg})^{-1/2}$ , with  $H_{gg} = X_g(X'X)^{-1}X_g'$ , implement the CR2 bias reduction.

We use the CR2 estimator: the clustered analogue of HC2, which implements the bias reduction of Bell and McCaffrey (2002) and removes the downward bias of the unadjusted (CR0/CR1) cluster-robust estimator when clusters are few or unbalanced. Following Cameron and Miller (2015) and Pustejovsky and Tipton (2018), CR2 is paired with a Satterthwaite degrees-of-freedom approximation to guard against over-rejection in finite samples. Cameron and Miller's (2015) practitioner's guide is explicit that, with a moderate number of clusters, the appropriate response is to combine a residual-adjusted estimator with a small-sample reference distribution—exactly the CR2/Satterthwaite combination used here. Our setting has a moderate cluster count (83 launch vehicles, 76 companies), squarely the regime for which CR2 is recommended.

Clustering at the launcher level appears to be relevant given the price-construction mechanism while strategic pricing would invite to consider the company layer. Thus, neither clustering strategy dominates a priori; if anything, the usual advice is to cluster at the most aggregate level at which dependence may plausibly operate (see Cameron and Miller (2015)). This tilts marginally towards the company (LSP) level, which encompasses the launcher level.

Taking these elements into account, we report, for every coefficient, four sets of standard errors: (i) classical OLS; (ii) heteroskedasticity-robust (White, HC3); (iii) CR2 clustered by launch vehicle (the most conservative); and (iv) CR2 clustered by company, the level of the strategic decision-maker. Reading inference across the four columns makes transparent how the assessed precision of each coefficient evolves as the dependence structure is progressively acknowledged.

### 3.2 Baseline estimates

Table 6 reports the estimation of equation (1) on the full sample of 2,471 launches, with the four standard-error specifications in adjacent columns. Panel A is the baseline; Panel B adds regional market controls (reference region = US).

**Table 6.** *Hedonic regression of (log) launch price*

| ln(PRICE)                                                                               | Coef.   | SE (OLS)    | SE (White)  | SE (CR2, launch.) | SE (CR2, LSPs) |
|-----------------------------------------------------------------------------------------|---------|-------------|-------------|-------------------|----------------|
| <b>Panel A — (no regional controls)</b>                                                 |         |             |             |                   |                |
| Intercept                                                                               | 0.7460  | (0.1534)*** | (0.1827)*** | (0.7687)          | (0.7158)       |
| <i>TIMELAPSE</i>                                                                        | −0.0327 | (0.0018)*** | (0.0019)*** | (0.0070)***       | (0.0070)***    |
| <i>ln MASS</i>                                                                          | 0.4390  | (0.0136)*** | (0.0181)*** | (0.0557)***       | (0.0531)***    |
| <i>ln APOGEE</i>                                                                        | 0.1480  | (0.0091)*** | (0.0093)*** | (0.0280)***       | (0.0331)***    |
| <i>ln NBS</i>                                                                           | 0.0553  | (0.0194)*** | (0.0203)*** | (0.0465)          | (0.0351)       |
| <i>ln (1+LAT)</i>                                                                       | −0.2719 | (0.0226)*** | (0.0229)*** | (0.1365)*         | (0.1477)       |
| <i>INSTIT</i>                                                                           | 0.2023  | (0.0342)*** | (0.0366)*** | (0.0904)**        | (0.1021)*      |
| <i>STARLINK</i>                                                                         | −0.4095 | (0.0677)*** | (0.0608)*** | (0.2126)          | (0.2452)       |
| Adj. R <sup>2</sup> = 0.588 N = 2,471 F = 504.6 G <sub>l</sub> = 83 G <sub>p</sub> = 76 |         |             |             |                   |                |
| <b>Panel B — (regional controls; reference = US)</b>                                    |         |             |             |                   |                |
| Intercept                                                                               | 0.2120  | (0.1590)    | (0.1658)    | (0.6763)          | (0.6201)       |
| <i>TIMELAPSE</i>                                                                        | −0.0269 | (0.0017)*** | (0.0018)*** | (0.0096)**        | (0.0099)**     |
| <i>ln MASS</i>                                                                          | 0.4099  | (0.0128)*** | (0.0174)*** | (0.0644)***       | (0.0620)***    |
| <i>ln APOGEE</i>                                                                        | 0.1351  | (0.0085)*** | (0.0086)*** | (0.0228)***       | (0.0227)***    |
| <i>ln NBS</i>                                                                           | 0.0064  | (0.0182)    | (0.0198)    | (0.0451)          | (0.0345)       |
| <i>ln (1+LAT)</i>                                                                       | 0.0555  | (0.0298)*   | (0.0219)**  | (0.0914)          | (0.0755)       |
| <i>INSTIT</i>                                                                           | 0.1754  | (0.0319)*** | (0.0349)*** | (0.1067)          | (0.1054)       |
| <i>STARLINK</i>                                                                         | −0.6442 | (0.0655)*** | (0.0593)*** | (0.4059)          | (0.4179)       |
| EU (vs US)                                                                              | 0.5513  | (0.0671)*** | (0.0463)*** | (0.2685)          | (0.1682)**     |
| ROW (vs US)                                                                             | −0.5953 | (0.0348)*** | (0.0350)*** | (0.2372)**        | (0.2506)**     |
| Adj. R <sup>2</sup> = 0.643 N = 2,471 F = 495.6 G <sub>l</sub> = 83 G <sub>p</sub> = 76 |         |             |             |                   |                |

Notes: Each panel is a single OLS estimation; the four right-hand columns report alternative standard errors for the same point estimates, with significance stars attached to each (\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ ). CR2 is the bias-reduced cluster-robust estimator (Bell and McCaffrey, 2002) with Satterthwaite degrees of freedom.  $G_l$  and  $G_p$  are the numbers of launcher and provider (LSP) clusters. Effective (Satterthwaite) degrees of freedom vary across coefficients, from 2.1 to 21.9. They are 13 or above for *TIMELAPSE*, *ln MASS*, *ln APOGEE* and *INSTIT*, but fall below six for *ln LAT* (2.1–5.7), *STARLINK* (4.1–5.9) and the EU dummy (3.5–4.5), whose identification rests on very few clusters. Source: authors' calculations.

Several features stand out, and each is sharpened by reading across the standard-error columns. First, the magnitude of the clustering correction remains striking. The CR2 standard errors are three to five times the OLS values — for *ln MASS*, from 0.014 (OLS) to 0.056 (CR2, launchers); for *TIMELAPSE*, from 0.002 to 0.007 — a direct quantitative illustration of within-launcher dependence.

Second, the hedonic core of the model is robust. The time trend is negative and significant under every standard-error specification — at the 1% level in the baseline, and at the 5% level

under cluster-robust inference once regional controls are added: the real price of access to orbit fell at roughly 3.3% per year over the period in the baseline, declining to 2.7% once regional controls are added. This downward trajectory survives the most demanding (launcher-clustered) inference and is consistent with ongoing, cost-reducing technological progress in launch provision, holding service attributes fixed.

Mass and apogee remain the two structural attributes:  $\ln MASS$  (elasticity  $\approx 0.44$  in Panel A,  $\approx 0.41$  in Panel B) and  $\ln APOGEE$  ( $\approx 0.15$  and  $\approx 0.14$  respectively) are both positive and significant at the 1% level even under CR2. Reaching higher orbits and injecting greater mass command higher implicit prices, as a hedonic reading predicts.

Third, the coefficient of the number of satellites is small and statistically indistinguishable from zero under CR2 in both panels. Multi-satellite launches are largely heavy launches, so the number of satellites carries little independent pricing information once mass and apogee are held fixed.

Fourth, launch-site latitude displays a negative and significant coefficient in the baseline (Panel A:  $-0.272$ , significant at the 10% level under launcher clustering and insignificant under company clustering), but loses all significance once regional fixed effects are included (Panel B). This pattern reflects a confounding between latitude and region rather than a latitude effect proper. Across regions the latitude ordering runs exactly opposite to the price ordering (for this, see below): all European launches depart from a single equatorial site (the Guiana Space Centre, at Kourou  $5.2^\circ N$ ), the lowest latitude in the sample, and Europe commands the largest positive premium, whereas the Rest-of-the-World region combines the highest latitudes with the lowest prices.<sup>29</sup> The negative baseline coefficient thus captures this between-region gradient. Once the regional dummies absorb it, latitude is identified from within-region variation alone — and carries essentially no pricing information there (it is insignificant within both the US and the Rest-of-the-World, and, European latitude being constant, Europe contributes nothing to its identification). The latitude coefficient in Panel A should therefore be read with caution, as its identification is confounded with regional composition.

Fifth, the institutional indicator is positive and significant under OLS and White inference in both panels. Under cluster-robust inference, it remains significant at the 5% level in Panel A under launcher clustering and at the 10% level under company clustering, but loses significance under both schemes once region is controlled for (Panel B). The institutional premium — approximately 0.20 log points in the baseline — is thus fragile: once the regional fixed effects absorb the between-region variance, it is no longer separately identified from regional composition under cluster-robust inference.

Sixth, the *STARLINK* dummy warrants specific discussion. Its OLS coefficient is negative and large in both panels ( $-0.41$  in Panel A,  $-0.64$  in Panel B), but it is never significant under CR2, in either panel and under either clustering scheme. This pattern is not primarily an economic finding; it reflects a structural feature of the cluster design. Of the 330 Starlink launches in the sample, all are operated by a single launcher (Falcon 9) from a single company (SpaceX), which together account for 543 and 556 observations respectively. Consequently, virtually all the within-sample variation in *STARLINK* is concentrated within a single cluster. The CR2 estimator, recognising that identification rests on a near-single cluster, inflates the standard error accordingly and reduces the effective Satterthwaite degrees of freedom to a very small

---

<sup>29</sup> The correlation between  $\ln(1+LAT)$  and the European Region indicator is  $-0.71$ , and  $+0.43$  with the Rest-of-the-World indicator.

number (below six in both panels). The loss of significance under cluster-robust inference should therefore be interpreted as a limitation of the identification strategy for this particular regressor, rather than as evidence that Starlink missions are not priced differently from comparable commercial launches.

Seventh, the regional differentials are large and make economic sense (Panel B). Relative to the United States, European launches carry a price premium of approximately 0.55 log points, while Rest-of-the-World launches are priced about 0.60 log points lower. The Rest-of-the-World discount is significant at the 5% level under both clustering schemes. The European premium is significant at the 5% level under company clustering but not under launcher clustering: Europe operates a single vehicle family through essentially one provider, so the effective number of clusters identifying the EU dummy is very small — the Satterthwaite degrees of freedom fall to about four — and the CR2 estimator reflects that limitation. The ordering — EU above the US, and the Rest of the World (driven by Chinese and Russian operators in the institutional-LEO segment) well below — mirrors the regional segmentation documented descriptively in Section 2.3. We go no further at this stage, since whether that segmentation reflects structural differences in market conduct is a question taken up in the remainder of the analysis. Including region lifts the adjusted  $R^2$  from 0.588 to 0.643. The attribute coefficients remain positive and significant, with magnitudes that decline modestly, indicating that part of the apparent attribute coefficients in the baseline reflected the regional composition of the sample; the pattern is precisely the role a market fixed effect is expected to play in a hedonic specification.

Note that the baseline treats physical and technical attributes as exerting separate, (log-)linear marginal effects on price; by construction it captures neither the effects of their combinations nor any interaction between attributes.

The remainder of the analysis asks a complementary and, for an industrial-organization reading, more informative question: whether the market organizes itself into discrete *segments*—combinations of attributes associated with distinct price tiers and identifiable operators. We approach this first through unsupervised classification, then through a set of hedonic interaction (“specialization”) variables motivated by that classification.

## **4. Mapping strategic positioning in launch services**

### **4.1 Unsupervised classification of launch services**

The objective of this section is to examine whether clustering launches on their physical characteristics yields conclusive groupings which can be in turn associated with distinct pricing strategies.

We classify launches using k-means clustering on three variables describing the physical service delivered:  $\ln MASS$ ,  $\ln APOGEE$ , and  $\ln NBS$ . The variables are standardised prior to clustering, since k-means is sensitive to scale.

The number of clusters is selected by combining the elbow (within-cluster sum of squares) and average silhouette criteria. The within-cluster dispersion falls most sharply between  $K = 1$  and  $K = 3$  and flattens markedly thereafter; the silhouette is maximised at  $K = 3$  (average silhouette  $\approx 0.56$ ), and both criteria therefore converge on the same solution. We retain  $K = 3$ . The choice

is defended on three grounds.  $K = 2$  is statistically clean but too coarse: it merely isolates the heaviest launches from the rest and yields two clusters whose mean log prices are close (3.98 versus 3.79), so it fails to structure prices informatively.  $K = 4$  splits the first cluster into a very-light and a medium subgroup but introduces a high-variance cell without sharpening the economic reading.  $K = 3$  is the most informative partition: it produces three clusters with clearly distinct price levels and an immediately interpretable sectoral meaning.

Table 7 profiles the three clusters we obtained through the latter method. The top block lists the variables on which the k-means partition was built (standardised  $\ln$  *MASS*,  $\ln$  *APOGEE*,  $\ln$  *NBS*); the bottom block lists descriptors computed ex post. Note that the cluster labels are interpretive.

**Table 7.** *Profile of the three k-means clusters ( $K = 3$ ): attributes*

|                                                                             | <b>Cluster 1</b><br><i>Light-LEO</i><br>(n = 988) | <b>Cluster 2</b><br><i>Heavy-GEO</i><br>(n = 967) | <b>Cluster 3</b><br><i>LEO-Const</i><br>(n = 516) |
|-----------------------------------------------------------------------------|---------------------------------------------------|---------------------------------------------------|---------------------------------------------------|
| <b>Clustering variables (standardised <math>\ln</math> used in k-means)</b> |                                                   |                                                   |                                                   |
| Mean mass (kg)                                                              | 1,632                                             | 4,186                                             | 13,136                                            |
| Mean apogee (km)                                                            | 713                                               | 68,074                                            | 422                                               |
| Mean NBS                                                                    | 2.9                                               | 1.3                                               | 29.1                                              |
| <b>Post-hoc descriptors</b>                                                 |                                                   |                                                   |                                                   |
| Mean $\ln$ price                                                            | 3.23                                              | 4.72                                              | 3.84                                              |
| Price s.d. ( $\ln$ )                                                        | 1.17                                              | 0.73                                              | 0.49                                              |
| Median year                                                                 | 2018                                              | 2011                                              | 2024                                              |
| Years (min–max)                                                             | 1996–2025                                         | 1996–2025                                         | 1997–2025                                         |
| Institutional (%)                                                           | 67.0                                              | 60.9                                              | 8.5                                               |

*Notes: Clusters in columns, attributes in rows. Mean  $\ln$  price and its s.d. are in logs. Source: authors' calculations).*

Overall, three main results emerge. First, clusters defined purely by the physical characteristics of the service correspond to distinct price levels (3.23 versus 4.72 versus 3.84 in logs). Second, the clusters are carried by distinct sets of launch service providers (LSPs)—Chinese and Russian LSPs in the institutional-LEO cluster, Arianespace, CASC and Khrunichev/ILS (Proton) in the geostationary cluster and SpaceX in the commercial constellation cluster. The correspondence is close but not exclusive: Arianespace, Khrunichev/ILS and Rocket Lab each concentrate more than 80% of their launches in a single cluster, whereas CASC and SpaceX operate across two. Third, they trace three eras of the market: an intermediate institutional-LEO cluster spanning the full period, a legacy heavy-GEO market centred on the 2000s, and a recent commercial-constellation market concentrated in the 2020s.

In the following, we describe in more detail how the three clusters can be characterized.

**Cluster 1 — light, low-Earth-orbit, institutional.** This cluster gathers 988 launches. They carry a moderate mean mass (about 1,600 kg) to a low mean apogee (about 700 km) and embark only a handful of satellites (mean 2.9). The cluster is two-thirds institutional and is centred on the late 2010s (median year 2018), dominated by Chinese operators (SAST and CALT, together 32% of the cluster), with Russian state operators (14%) and Rocket Lab (RLABN, 6%) also present. Its mean log price is the lowest of the three (3.23) and also the most dispersed (s.d. 1.17), reflecting the heterogeneity of pricing across the Chinese and Russian segments.

**Cluster 2 — heavy, geostationary, legacy market.** This cluster gathers 967 launches. They carry a high mass (about 4,200 kg) to a high mean apogee (about 68,000 km) and typically embark a single satellite (mean 1.3). Note that the median apogee, 35,799 km, sits precisely at the geostationary belt.<sup>30</sup> It is the historical core of the commercial launch industry—Ariane 5, Proton, Long March 3B and Falcon 9—roughly 61% institutional and temporally centred on the 2000s (median year 2011). It is the most expensive cluster, with a mean log price of 4.72 and a moderate dispersion (s.d. 0.73).

**Cluster 3 — low-Earth-orbit constellations, recent market.** This cluster gathers 516 launches. They carry a very high aggregate mass (about 13,000 kg) to a low mean apogee (about 420 km) and embark a large number of satellites (mean 29). The cluster is overwhelmingly commercial (8.5% institutional), dominated by SpaceX/Falcon 9, and concentrated in the most recent years (median year 2024). Its mean log price (3.84) sits between the other two, and its price dispersion is by far the lowest (s.d. 0.49), pointing to strong tariff standardisation.

To sum up, the clustering reveals clear groups of launches associated with distinct mass and apogee profiles. Direct evidence of price differentiation associated with these groups suggests in turn specialization strategies with distinct effects on prices. This might be interpreted as one empirical signature of strategic groups in the industrial-organisation sense (Caves and Porter, 1977): bundles of service characteristics that map onto homogeneous price levels and onto identifiable sets of operators.

## 4.2 Orbital specialization

Accordingly, as a further step, we establish a parsimonious parametrization of the clustering in terms of specialization types—combining mass and apogee— so as to take it to the regression. We will then investigate therein whether such specialization bundles have significant effects on pricing. If they do, the clusters, by “signing” specifically in terms of these specialization bundles, would have a genuine impact on prices.

To build this parametrization, we examine how each cluster projects onto the apogee×mass plane, reading each cluster as a distribution over the cells of an apogee×mass cross-tabulation.<sup>31</sup> For the sake of simplicity, we chose a coarse grid. On the apogee dimension we retain a single boundary, the geostationary frontier—the economically salient orbital distinction, and also the one that carries discriminating power: within the non-geostationary regime Clusters 1 and 3 are told apart by mass, so a further low/medium apogee split would add cells without separating any pair of clusters. On the mass dimension we keep the three standard classes used throughout the paper (light, medium and heavy). A cluster may still occupy several cells; what matters is that each is well distinguished from the others by the cells it concentrates in.

The classes are defined as follows. For apogee: non-GEO ( $\leq 35,000$  km) and GEO ( $> 35,000$  km), the latter being the geostationary regime. For mass: light (*DLM*,  $< 1,000$  kg), medium

---

<sup>30</sup> The gap between the mean and the median reflects a right tail of very high-apogee launches, on highly elliptical trajectories. These fall well inside the cluster's interior rather than near any boundary, so they do not affect the cluster frontiers. Re-estimating the partition with median-based methods that are insensitive to such values—k-medoids (PAM) and k-medians—reassigns fewer than 1% of launches and leaves the three clusters, their membership and their interpretation intact.

<sup>31</sup> Cells in % of the cluster's launches, with row and column totals.

(*DMM*, 1,000–4,000 kg) and heavy (*DHM*,  $\geq 4,000$  kg). Crossing the two gives the six-cell apogee $\times$ mass scheme.

Cluster 2 is captured by the geostationary cells (GEO  $\times$  heavy 45%, GEO  $\times$  medium 35%); Cluster 3 by the single non-GEO  $\times$  heavy cell (92%); Cluster 1 by the non-GEO light and medium cells (48% and 43%) (Tables 8a–8c).

The correspondence between clusters and cells is thus sharp and the overlap between clusters minimal: each cluster signs on cells that the others scarcely occupy, so the partition gives a faithful image of the sample. Cluster 1, the least homogeneous, is the only one that spreads over more than one cell.

**Table 8a. Cluster 1 (Light-LEO) (n = 988) — parsimonious grid (6 cells)**

*Distribution by apogee  $\times$  mass, collapsed classes (% of cluster):*

| MASS \ APOGEE         | Non-GEO ( $\leq 35000$ ) | GEO ( $> 35000$ ) | TOTAL  |
|-----------------------|--------------------------|-------------------|--------|
| Heavy ( $\geq 4000$ ) | 9.3%                     | 0.0%              | 9.3%   |
| Medium [1000, 4000)   | <u>42.9%</u>             | 0.0%              | 42.9%  |
| Light ( $< 1000$ )    | <u>47.8%</u>             | 0.0%              | 47.8%  |
| TOTAL                 | 100.0%                   | 0.0%              | 100.0% |

**Table 8b. Cluster 2 (Heavy-GEO) (n = 967) — parsimonious grid (6 cells)**

*Distribution by apogee  $\times$  mass, collapsed classes (% of cluster):*

| MASS \ APOGEE         | Non-GEO ( $\leq 35000$ ) | GEO ( $> 35000$ ) | TOTAL  |
|-----------------------|--------------------------|-------------------|--------|
| Heavy ( $\geq 4000$ ) | 4.4%                     | <u>45.0%</u>      | 49.4%  |
| Medium [1000, 4000)   | 10.2%                    | <u>35.3%</u>      | 45.5%  |
| Light ( $< 1000$ )    | 1.1%                     | 3.9%              | 5.1%   |
| TOTAL                 | 15.8%                    | 84.2%             | 100.0% |

**Table 8c. Cluster 3 (LEO-Const.) (n = 516) — parsimonious grid (6 cells)**

*Distribution by apogee  $\times$  mass, collapsed classes (% of cluster):*

| MASS \ APOGEE         | Non-GEO ( $\leq 35000$ ) | GEO ( $> 35000$ ) | TOTAL  |
|-----------------------|--------------------------|-------------------|--------|
| Heavy ( $\geq 4000$ ) | <u>91.5%</u>             | 0.0%              | 91.5%  |
| Medium [1000, 4000)   | 7.6%                     | 0.0%              | 7.6%   |
| Light ( $< 1000$ )    | 1.0%                     | 0.0%              | 1.0%   |
| TOTAL                 | 100.0%                   | 0.0%              | 100.0% |

Notes: Apogee in two classes (non-GEO  $\leq 35,000$  km; GEO  $> 35,000$  km), mass in three (light  $< 1,000$  kg; medium 1,000–4,000 kg; heavy  $\geq 4,000$  kg); cells are % of the cluster. Source: authors' calculations on *df\_original*.

In the following, we parametrize the specialization on this six-category apogee $\times$ mass scheme and carry it to the price regression.

### 4.3 Specialisation and price differentials

The regression we consider is a variant of the price equation (equation (1)) with the specialisation dummies ( $DSPEC_{k,i}$ )<sup>32</sup> replacing the continuous technical regressors<sup>33</sup>. It is specified as:

$$\ln(P_{it}) = \gamma_0 + \gamma_1 \cdot TIMELAPSE_t + \gamma_2 \cdot INSTIT_{it} + \sum_{k=2}^6 \delta_k \cdot DSPEC_{k,i} + \varepsilon_{it} \quad (3)$$

where the omitted (reference) category is the non-geostationary  $\times$  light cell, and inference relies on CR2 standard errors clustered alternately by launcher and by company (LSP).

Table 9 reports the specification without regional controls (Panel A) and with them (Panel B).

**Table 9.** Price levels by specialisation cell

| Cell / variable                                                                     | Coef.   | SE (OLS)    | SE (White)  | SE (CR2, launch.) | SE (CR2, LSPs) |
|-------------------------------------------------------------------------------------|---------|-------------|-------------|-------------------|----------------|
| <b>Panel A — specialisation only (reference: non-GEO <math>\times</math> light)</b> |         |             |             |                   |                |
| Intercept                                                                           | 3.3513  | (0.0573)*** | (0.0619)*** | (0.2081)***       | (0.2026)***    |
| Non-GEO $\times$ medium                                                             | 1.0064  | (0.0494)*** | (0.0609)*** | (0.2547)***       | (0.2958)***    |
| Non-GEO $\times$ heavy                                                              | 1.4631  | (0.0498)*** | (0.0602)*** | (0.2791)***       | (0.2297)***    |
| GEO $\times$ light                                                                  | 1.0928  | (0.1349)*** | (0.1407)*** | (0.2313)***       | (0.2037)***    |
| GEO $\times$ medium                                                                 | 1.6126  | (0.0583)*** | (0.0577)*** | (0.2217)***       | (0.1879)***    |
| GEO $\times$ heavy                                                                  | 2.0611  | (0.0526)*** | (0.0598)*** | (0.3686)***       | (0.3159)***    |
| <i>TIMELAPSE</i>                                                                    | -0.0369 | (0.0020)*** | (0.0020)*** | (0.0077)***       | (0.0069)***    |
| <i>INSTIT</i>                                                                       | 0.1770  | (0.0342)*** | (0.0357)*** | (0.0949)*         | (0.1067)       |
| <i>Adj. R<sup>2</sup> = 0.497    N = 2,471</i>                                      |         |             |             |                   |                |
| <b>Panel B — adding region (reference region = US)</b>                              |         |             |             |                   |                |
| Intercept                                                                           | 3.6133  | (0.0565)*** | (0.0622)*** | (0.1937)***       | (0.2219)***    |
| Non-GEO $\times$ medium                                                             | 0.8827  | (0.0448)*** | (0.0565)*** | (0.2240)***       | (0.2536)***    |
| Non-GEO $\times$ heavy                                                              | 1.1390  | (0.0487)*** | (0.0616)*** | (0.2637)***       | (0.2468)***    |
| GEO $\times$ light                                                                  | 0.8534  | (0.1226)*** | (0.1306)*** | (0.2095)***       | (0.1877)***    |
| GEO $\times$ medium                                                                 | 1.4676  | (0.0528)*** | (0.0558)*** | (0.2342)***       | (0.1967)***    |
| GEO $\times$ heavy                                                                  | 1.6957  | (0.0498)*** | (0.0546)*** | (0.2514)***       | (0.2075)***    |
| EU (vs US)                                                                          | 0.5879  | (0.0569)*** | (0.0466)*** | (0.3043)          | (0.1726)*      |
| ROW (vs US)                                                                         | -0.5778 | (0.0341)*** | (0.0353)*** | (0.2013)**        | (0.2139)**     |
| <i>TIMELAPSE</i>                                                                    | -0.0303 | (0.0018)*** | (0.0019)*** | (0.0085)***       | (0.0083)***    |
| <i>INSTIT</i>                                                                       | 0.2709  | (0.0311)*** | (0.0330)*** | (0.0901)***       | (0.0957)**     |
| <i>Adj. R<sup>2</sup> = 0.593    N = 2,471</i>                                      |         |             |             |                   |                |

Notes: Coefficients are differences in mean log price relative to the non-GEO  $\times$  light reference cell. SE (White) is the HC3 heteroskedasticity-robust standard error; CR2 standard errors are clustered by launcher and by company (LSP), with Satterthwaite degrees of freedom. Significance: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ . Effective (Satterthwaite) degrees of freedom range from 2.0 to 21.5; they fall below nine for the GEO  $\times$  light cell (5.3–8.1) and for the EU dummy (2.0–5.8), and exceed nine for every other coefficient. Source: authors' calculations.

<sup>32</sup> The six apogee $\times$ mass cells defined above are treated as mutually exclusive and exhaustive categories  $CAT_k$ ,  $k = 1, \dots, 6$ , ordered by apogee then by mass:  $CAT_1$  = non-GEO $\times$ light,  $CAT_2$  = non-GEO $\times$ medium,  $CAT_3$  = non-GEO $\times$ heavy,  $CAT_4$  = GEO $\times$ light,  $CAT_5$  = GEO $\times$ medium,  $CAT_6$  = GEO $\times$ heavy. For each launch  $i$  we set  $DSPEC_{k,i} = \mathbf{1}\{i \in CAT_k\}$ , i.e. 1 if launch  $i$  belongs to  $CAT_k$  and 0 otherwise. Each launch falls in exactly one category, so one dummy is dropped as the reference — here  $CAT_1$  (non-GEO $\times$ light) — and the coefficients  $\delta_k$ ,  $k = 2, \dots, 6$ , in equation (3) measure mean log-price differences relative to it.

<sup>33</sup> We do not retain STARTLINK and LAT as additional regressors, given the results of the baseline equation (1) estimation (see section 3.2).

The estimates display an economically coherent ladder of price levels (Panel A). Within the non-geostationary regime, moving from light to medium mass raises the log price by about 1.01, and to heavy mass by about 1.46 — the latter being the tier in which the constellation Cluster 3 concentrates. The geostationary regime commands a premium at every mass class: 1.09 (light), 1.61 (medium) and 2.06 (heavy). These geostationary cells are precisely those in which the legacy Cluster 2 sits. A telling internal comparison is preserved from the baseline: a medium-mass non-GEO launch (+1.01) and a light-mass GEO launch (+1.09) carry near-identical premia, so reaching geostationary orbit with a light mass embarked costs roughly as much as lifting a medium mass to a lower orbit. The specialisation ladder thus lines up, cell by cell, with the three clusters obtained from the unsupervised classification, and every premium survives the most demanding cluster-robust inference at the 1% level.

The time trend in Panel A ( $-0.037$  per year) is somewhat attenuated relative to the continuous-attribute specification, consistent with the specialisation dummies absorbing part of the secular compositional shift in the sample. The institutional indicator is positive (0.18) but remains fragile under clustering — significant at the 10% level under launcher CR2 and insignificant under company CR2 — mirroring the pattern observed in the baseline without regional controls.

Adding region (Panel B) leaves the specialisation ladder qualitatively intact but compresses it: every cell premium shrinks, the geostationary-heavy cell falling from 2.06 to 1.70. The regional differentials are themselves large — European launches carry a premium of about 0.59 log points over the United States, while Rest-of-the-World launches are priced about 0.58 lower. The Rest-of-the-World discount remains significant at the 5% level under both clustering schemes, whereas the European premium survives only under company clustering, and only at the 10% level, for the small-cluster reason set out in Section 3.2. Two further changes are notable. First, the institutional premium in Panel B (0.27) is now significant at the 1% level under launcher clustering and at the 5% level under company clustering, a meaningful improvement in robustness relative to the earlier estimates. Second, the time trend moderates further to  $-0.030$ , suggesting that part of the Panel A trend captured the changing regional composition of launches over time.

Because region and specialisation are strongly aligned — the geostationary-heavy cell is largely European, the non-GEO-heavy cell largely American — the two sets of controls share much of the same variation, and Panel B should be read as a robustness check on the specialisation ladder rather than as a clean decomposition of regional from specialisation effects.

In sum, the unsupervised classification and the structured interaction model tell a consistent story. The launch market is organised into a small number of price tiers defined jointly by orbital regime and mass; these tiers are populated by distinct sets of operators and correspond to distinct eras of the industry. The pricing information lives in these discrete tiers across clusters, which the DSPEC interactions are designed to capture.

## 5. Conclusion

This article has examined the determinants of commercial satellite launch prices worldwide and their evolution since 1996, and what these patterns reveal about the structuring of the market and the specialisation strategies at work within it. Its empirical basis is an original reconstructed price series, built on the combination of two sources — the aggregate regional revenues of launch service providers and expert launch-cost data — and covering the near-totality of launches recorded in GCAT.

On this basis, and drawing on a hedonic regression relating log price to service characteristics, two robust determinants emerge: payload mass and target altitude, with positive elasticities that survive our most demanding cluster-robust inference. Real launch prices fell steadily, by roughly 3.3% per year over 1996–2025, consistent with ongoing, cost-reducing technical progress, while regional differentials — a European premium and a Rest-of-the-World discount relative to the United States — proved large, the latter robust to our most demanding inference.

Turning to the second question — the organisation of the market and the patterns of specialisation within it — the unsupervised classification we implement and the specialisation model partly derived from it converge on a common reading: launch services are organised not as a single arena but into a small number of price tiers defined jointly by orbital regime and payload mass, each populated by distinct providers and tied to a distinct era of the industry. Strategies of specialisation and positioning thus appear associated with the price differentials we observe.

This reading invites a closer look at how competition operates within, rather than across, these segments: how each is contested, and how positioning along them has shifted as the industry has been reshaped. We take up these questions in subsequent work, notably by bringing measures of concentration to bear on the segment structure identified here.

## 6. References

- Adilov, N., P. J. Alexander and B. M. Cunningham (2015). An economic analysis of earth orbit pollution. *Environmental and Resource Economics*, 60(1), 81-98.
- Adilov N., P. J. Alexander and B. M. Cunningham (2018), “An economic “Kessler Syndrome”: A dynamic model of earth orbit debris”, *Economics Letters*, 166: 79–82.
- Adilov, N., P. J. Alexander, B.M. Cunningham and N. Albertson (2023), « An Analysis of Launch Cost Reductions for Low Earth Orbit Satellites », *Economics Bulletin*, 42(3): 1561–1574.
- Bajari, P., and C. L. Benkard (2005), "Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach", *Journal of Political Economy*, 113(6): 1239–1276.
- Bell, R. M., and D. F. McCaffrey. 2002. “Bias Reduction in Standard Errors for Linear Regression with Multi-Stage Samples.” *Survey Methodology* 28 (2): 169–181.
- Bongers, A. and J. L. Torres (2023). Orbital debris and the market of satellites. *Ecological Economics*, 209, 107831.
- Boone, T. R., and D. P. Miller (2016), « Capability and Cost-Effectiveness of Launch Vehicles », *New Space*, 4(3): 168-189.
- Cameron, A. C., and D. L. Miller. 2015. “A Practitioner’s Guide to Cluster-Robust Inference.” *Journal of Human Resources* 50 (2): 317–372.

Caves, R. E., and M. E. Porter. 1977. "From Entry Barriers to Mobility Barriers." *Quarterly Journal of Economics* 91 (2): 241–261.

Chernozhukov, V., A. Galichon, M. Henry and B. Pass (2021), "Identification of Hedonic Equilibrium and Nonseparable Simultaneous Equations", *Journal of Political Economy*, 129(3): 842–870.

Court, A. T. (1939), "Hedonic Price Indexes with Automotive Examples", in *The Dynamics of Automobile Demand*, New York: General Motors Corporation, 99–117.

Dos Santos, V. P. and G. Le Hir (2016), « Industry structure and disruptive innovation: the satellite industry », *Journal of Innovation Economics and Management*, 20(2): 37–60.

Ekeland, I., J. J. Heckman and L. Nesheim (2004), "Identification and Estimation of Hedonic Models", *Journal of Political Economy*, 112(S1): S60–S109.

Government Accountability Office (2017) "Surplus Missile Motors: Sale Price Drives Potential Effects on DOD and Commercial Launch Providers," Report GAO-17-609, August 2017, available at <https://www.gao.gov/assets/gao-17-609.pdf>.

Griliches, Z. (1961), "Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change", in *The Price Statistics of the Federal Government*, NBER General Series No. 73, New York: NBER, 173–196.

Guyot, J., A. Rao and S. Rouillon (2023), "Oligopoly competition between satellite constellations will reduce economic welfare from orbit use", *Proceedings of the National Academy of Sciences*, 120(43): e2221343120.

Guyot, J. and S. Rouillon (2023), « Sustainable management of space activity in low Earth orbit », *Journal of Environmental Economics and Policy*, 1–25.

Heckman, J. J., R. L. Matzkin and L. Nesheim (2010), "Nonparametric Identification and Estimation of Nonadditive Hedonic Models", *Econometrica*, 78(5): 1569–1591.

Kim, M. J. (2025). "Counting Stars and Costs: An Empirical Examination of Space Launch Cost Trend at NASA." *\*Acta Astronautica\**, vol. 232 (July), pp. 633–639. <https://doi.org/10.1016/j.actaastro.2025.04.011>

Kulu, E. (2025), « Small Launchers - 2025 Survey and Competitive Landscape », 76th International Astronautical Congress, IAC-25-D2.1.7.

Kuminoff, N. V., C. F. Parmeter and J. C. Pope (2010), "Which Hedonic Models Can We Trust to Recover the Marginal Willingness to Pay for Environmental Amenities?", *Journal of Environmental Economics and Management*, 60(3): 145–160.

Lancaster, K. J. (1966), "A New Approach to Consumer Theory", *Journal of Political Economy*, 74(2): 132–157.

MacKinnon, J. G., and H. White. 1985. "Some Heteroskedasticity-Consistent Covariance Matrix Estimators with Improved Finite Sample Properties." *Journal of Econometrics* 29 (3): 305–325.

Nozawa, W., K. Kurita, T. Tamakic and S. Managi (2023), « To What Extent Will Space Debris Impact the Economy? », *Space Policy*, 66. <https://doi.org/10.1016/j.spacepol.2023.101580>

OECD (2022), *OECD Handbook on Measuring the Space Economy*, 2nd Edition. Paris: OECD Publishing. <https://doi.org/10.1787/8bfef437-en>

Pustejovsky, J. E., and E. Tipton. 2018. "Small-Sample Methods for Cluster-Robust Variance Estimation and Hypothesis Testing in Fixed Effects Models." *Journal of Business & Economic Statistics* 36 (4): 672–683.

Rao, A., M. G. Burgess and D. Kaffine (2020), "Orbital-use fees could more than quadruple the value of the space industry", *Proceedings of the National Academy of Sciences* 117(23): 12756-12762.

Roberts, T. G. (2022) "Space Launch to Low Earth Orbit: How Much Does It Cost?" Aerospace Security website, available at <https://aerospace.csis.org/data/space-launch-to-low-earth-orbit-how-much-does-it-cost/>.

Rosen, S. (1974), "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition", *Journal of Political Economy*, 82(1): 34–55.

Rouillon, S. (2020), "A physico-economic model of space debris management", *Environmental and Resource Economics*, 77: 695-723.

Rouillon, S. (2025), « Monopolistic Competition in a Limited Orbital Space », *Environmental and Resource Economics*, 88(4): 1009-1038.

Rouillon, S. (2026), « Strategic entry deterrence in a finite orbital environment », *Revue d'Economie Politique*, forthcoming.

Su, R., Yang, C., & Sweeting, A. (2026). "Competition, Procurement and Learning-by-Doing in the Space Launch Industry." \*NBER Working Paper\* No. 34766.

Weinzierl, M. (2018), « Space, the final economic frontier », *Journal of Economic Perspectives*, 32(2), 173-192.

Weinzierl, Matthew C., Kylie Lucas, and Mehak Sarang (2021), "SpaceX, Economies of Scale, and a Revolution in Space Access." *Harvard Business School Case 720-027*, April 2020. (Revised October 2021.)

### **Alphabetical list of abbreviations:**

ASD-Eurospace : AeroSpace and Defence Industries Association of Europe (Eurospace)

CALT : Chinese Academy of Launch Vehicle Technology

CASC : China Aerospace Science and Technology Corporation (China)

CR2 : Bias-reduced cluster-robust variance estimator

EU : Europe/European Union  
FAA : Federal Aviation Administration (US)  
GAO : Government Accountability Office (US)  
GCAT : General Catalog of Artificial Space Objects  
GEO : Geostationary Earth Orbit  
HC3 : Heteroskedasticity-consistent variance estimator (type 3)  
ILS : International Launch Services  
INSEE : Institut National de la Statistique et des Études Économiques (France)  
KVR : Kosmicheskiye Voyska Rossii (Russian Space Forces)  
LEO : Low Earth Orbit  
LSP : Launch Service Provider  
NASA : National Aeronautics and Space Administration (US)  
OECD : Organisation for Economic Co-operation and Development  
OLS : Ordinary least squares  
ONCE : One NASA Cost Engineering  
RLABN : Rocket Lab Ltd.  
ROW : Rest Of World  
SAST : Shanghai Academy of Spaceflight Technology (China)  
SESIR : State of the European Space Industry Reports  
SIA : Satellite Industry Association (US)  
SSIR : State of the Satellite Industry Reports  
UCS : Union of Concerned Scientists (US)  
ULA : United Launch Alliance (US)  
US : United States