Competition, Procurement and Learning-by-Doing in the Space Launch Industry

Ruibing Su
University of Maryland

Chenyu Yang*
University of Maryland

Andrew Sweeting
University of Maryland, CEPR and NBER

Wednesday 21st January, 2026

Abstract

We estimate a dynamic model of the U.S. space launch industry. The model allows past launches to improve rocket reliability and lower launch costs. It also allows the government to make forward-looking procurement choices. We use the model to analyze policy-relevant issues in the recent history of the industry: the 2006 United Launch Alliance "merger-to-monopoly" and the effects of efficiencies in the form of learning synergies; innovations, such as SpaceX's Falcon 9 and ULA's recent introduction of Vulcan Centaur; the costs and benefits of forward-looking procurements; and, the trade-offs between the advantages of centralized control and possible inefficiencies.

Keywords: dynamic competition, learning-by-doing, procurement, forward-looking buyers, mergers, antitrust, space launch.

JEL codes: C73, D21, D43, L13, L41.

*Contact: rbsu@umd.edu, cyang111@umd.edu or atsweet@umd.edu. We are very grateful to seminar participants at National University of Singapore, Singapore Management University, Temple University, the University of Maryland and attendees at the 2025 Texas IO Day. We thank Brian Pierce and Thomas Hedberg of UMD's ARLIS Center and Christopher Reynolds of UMD's Joint Space Science Institute for useful discussions about the industry, and Shanglyu Deng, Uli Doraszelski, Brad Larsen, Mario Leccese and Huseyin Yilidirim for insightful comments on earlier drafts. The usual disclaimer applies.

1 Introduction

> "*It is the policy of the United States to enhance American greatness in space by enabling a competitive launch marketplace and substantially increasing commercial space launch cadence and novel space activities by 2030.*" [Presidential Executive Order "Enabling Competition in the Space Launch Industry", August 2025]
>
> "*A space revolution is coming. The 21st century will be the century of space. Europe needs to be at the forefront.*" [EU Commissioner for Defense and Space, June 2025]

Industries of strategic national importance, such as aerospace, defense, nuclear power, and aircraft manufacturing, often share common features. On the supply-side, one or two sellers, who can benefit from large economies of scale or learning-by-doing (LBD) economies, have to make large investments to develop better products. On the demand-side, a government buyer knows that it will be in the market repeatedly and cares about not only procurement costs, but also the reliability of the products being purchased. Having a domestic industry that matches or outperforms the industries of rival countries can also be vital to the country's national interests.

The sensitivity and concentration of these industries make it difficult to answer important economic questions, such as how to design procurement and evaluate whether a merger is beneficial, using descriptive empirical methods. For example, even if we can observe prices after a consummated merger, a credible control industry, also with available data, is likely to be almost impossible to find.

This paper therefore takes a model-based approach to answering policy-relevant questions, using data from the space launch industry. This industry is critical to intensifying international competition in space, and, in some readings, the future of humanity (Rees (2003), Vance (2015), Guthrie (2017), Troutman (2020)).

We assemble a dataset of unmanned medium-to-heavy-rocket launches which place objects into stable orbits (160 km or more above the earth), and another new dataset on the prices paid by the U.S. government for launches. We then take a parsimonious model of dynamic competition, which has the following basic structure, to these data. Each rocket family has a state variable that defines its current level of know-how or experience. Know-how can be accumulated through launches, but it can also depreciate. There are twenty possible government launch procurements each year, and, in each procurement, the launch suppliers, whose costs are functions of their know-how and idiosyncratic cost shocks, submit bids. The government's flow utility from choosing a particular launch rocket depends on the rocket's bid and a logit choice shock, but the government may also choose an outside option of no launch.

As we do not directly observe the costs of different launches, or the costs of introducing a new rocket, we choose the parameter that determines the rate of cost-side LBD, and the distribution of innovation costs, based on industry sources. We estimate the demand parameters and the remaining cost parameters using the simulated method of moments, where we match predicted price moments and choice probability moments, to moments from our data. When we simulate our model, we adjust know-how to account for the fact that the rockets are also used for commercial launches, which should also affect the evolution of know-how.

While our model is simple, it allows us to capture at least four significant features of the industry since 2002. First, the model allows for three different market structures, covering the period before 2006 when Boeing and Lockheed Martin formed the United Launch Alliance joint venture (ULA JV); after the JV, when the government faced a multi-product monopolist seller; and, after the certification of SpaceX's Falcon 9 rockets. Second, our model allows for accumulated seller know-how to not only lower sellers' production costs (i.e., the type of cost-reducing LBD and know-how depreciation considered in the previous literature, such as Benkard (2000) and Benkard (2004)), but also to increase the buyer's expected valuation of a launch. Consistent with the engineering literature on launches, government policy and patterns that we find in our own data, we interpret the effect on buyer valuations as reflecting the value of expected launch reliability. Third, our model assumes a forward-looking government buyer, so that dynamically optimizing agents on both sides of the market interact. The difference to assuming that each period's buyer is atomistic, in the sense of simply maximizing its current flow payoff, which is more common in the literature, is that our buyer accounts for how its current procurement choices affect the evolution of sellers' know-how states, and therefore the buyer's future surplus. This change in the buyer's incentives can lead to changes in prices, as well as choice probabilities, in equilibrium. Fourth, in the latter period of our data, we extend our model to allow for the possibility of innovation by ULA, in the form of a new, lower cost rocket family (ULA's Vulcan family was certified in 2025). This allows us to evaluate the differences between the private and social returns to innovation, when compared to lowering the costs of existing technologies by accumulating know-how.

Our creation of new datasets, which lets us develop a new empirical application to an important industry, and our extensions to existing models of dynamic competition are contributions of our paper. However, we view our primary contributions as coming from substantive counterfactual analyses of four welfare-relevant and policy-relevant questions. First, we compare the performance of alternative market structures, paying particular attention to the 2006 ULA JV, which was, at the time, a "merger-to-monopoly" in the launch market. The Federal Trade Commission (FTC) allowed the JV to proceed based on its understanding that it would enhance the reliability of the parties' Atlas V and Delta IV rocket families (Kovacic (2019)), a particular form of merger synergy. Our analysis can be viewed as a model-based retrospective (Miller and Weinberg (2017), An and Zhao (2019)) on a high-profile and, at the time, controversial enforcement decision. Our model predicts that, in the presence of the synergy, which, in our model, takes the form of LBD benefiting both of ULA's rockets, the decision to allow the JV was the correct one, in the sense that it is expected to raise both buyer (i.e., government) and total surplus.[1]

Second, we compare market outcomes with those under a social planner. This not only allows us to quantify the inefficiency of dynamic market outcomes (Besanko et al. (2019)), but it also allows us to address the policy question of whether the U.S.'s market-focused approach should be expected to perform better than the more centrally (i.e., buyer)-controlled space industries of other countries and the U.S. historically. This question is especially pertinent given that the current U.S. administration is taking controlling "golden share" interests in some strategically important firms (Mishra (2025)), and international competition in space, especially with Russia and China, whose industries are tightly controlled,[2] is increasing.

[1]Kovacic (2019) argues that the experience of the 15 years after the approval has been "astonishingly positive." Our approach allows us to come to a similar conclusion based on counterfactual predictions about what would have happened if the JV had been blocked.

[2]Since 2014, China has also opened the nation's launch and small satellite sectors to the private firms (Jones (2019)).

Weinzierl (2018) provides an overview of the growing role of commercial companies in providing launch services in the U.S., and argues that "[t]he vulnerabilities of centralized control will be familiar to any economist: weak incentives for the efficient allocation of resources, poor aggregation of dispersed information, and resistance to innovation due to reduced competition."[3] Motivated by this argument, we assess the trade-offs by quantifying how much production cost inefficiency under planner control is needed to equalize welfare under the planner and market outcomes. Absent innovation, we find that a 38% production cost inefficiency is required.

Third, we compare the performance of the government's forward-looking procurement rule, which we assume for estimation, with a static procurement rule where the buyer commits to choosing a rocket based only on its current flow payoffs. This comparison not only allows us to understand the effects of allowing for a forward-looking buyer within our model, but it is also relevant given current litigation over U.S. space procurement practices, including the weight that the government should place on future supply considerations (Erwin (2019)).[4] Existing theoretical (Lewis and Yildirim (2002)) and computational (Sweeting et al. (2022)) analyses have shown that, when a buyer faces duopoly suppliers, benefiting from cost-side LBD, and has, at most, a fairly unattractive outside option, a commitment to a static procurement rule can be buyer-optimal because it intensifies supplier competition. In our empirical model, where we consider alternative market structures and the buyer frequently chooses not to launch, we find more nuanced results. Duopolists tend to underinvest in winning launches to increase their experience or know-how, and a forward-looking buyer can offset this underinvestment, so that future buyer surplus increases. On the other hand, we find that a static procurement rule is optimal for the buyer when it faces a ULA monopolist benefiting from the learning synergy, because the commitment to a static procurement rule causes the monopolist to lower prices, as competition with the outside good is strengthened.

[3] Of course, in reality, there is significant government involvement in investment decisions in the launch industry. For example, up to 2018, the U.S. government had committed approximately $1.2 billion to the development of Vulcan Centaur, out of total development costs of around $8 billion (Albon (2024) and `https://en.wikipedia.org/wiki/Vulcan_Centaur`, last accessed on January 19th, 2026).

[4] More broadly, the U.S. government is currently considering potentially "revolutionary" reforms to its procurement policies (Office of Federal Procurement Policy (2025)).

Fourth, we study the effects of competition from SpaceX (SPX) on welfare and on the incentives of ULA to innovate. As we discuss, SPX's success was far from certain and, in many countries, innovative private firms are not allowed to supply government launches. SPX's Falcon 9 rockets introduced a lower launch cost technology, and SPX's use of Falcon 9 to roll out its commercial Starlink system also meant that Falcon 9's launch experience increased rapidly. We show that competition between SPX and ULA's Atlas V and Delta IV created significant gains for the government. Our model also predicts that the probability that ULA will want to innovate, to have its own low-cost rocket family, doubles when SPX enters, which is consistent with ULA's commitment to develop Vulcan Centaur, a rocket designed to be less expensive than ULA's older Atlas V and Delta IV systems, once Falcon 9 had been launched successfully. We also find that, when innovation is possible, buyer control can be even more valuable.

The structure of the rest of the paper is as follows. After a brief review of the related literature, Section 2 provides institutional background on the industry and Section 3 describes our data. Section 4 details our model, and notes its limitations. Section 5 explains how we fix certain parameters (e.g., know-how depreciation rates) and estimate the others, and discusses our parameter estimates and model fit. Section 6 describes our counterfactuals and their policy implications. Section 7 concludes.

Related Literature. Our paper builds on the theoretical and empirical literatures that use models of dynamic competition to study questions in antitrust, industrial policy and international trade (for example, theoretical papers by Fudenberg and Tirole (1983), Cabral and Riordan (1994), Besanko et al. (2010) and Besanko et al. (2014), and empirical papers by Baldwin and Krugman (1988), Irwin and Pavcnik (2004), Benkard (2004), An and Zhao (2019) and Barwick et al. (2025)). The models that we use draw directly from this literature, although we will allow for LBD to improve the buyer's valuations of products, as well as reduce the sellers' costs. Similar to the analysis of the Boeing/McDonnell Douglas merger in An and Zhao (2019), our work provides retrospective analysis of a consummated merger using a dynamic model. In our case, the JV was a true merger-to-monopoly until SPX entered.

Our model differs from the papers considered above in allowing for there to be a monopsonistic and forward-looking buyer, rather than atomistic buyers.[5] Lewis and Yildirim (2002) present a theoretical model where such a buyer spreads purchases between two sellers who benefit from LBD in order to maintain competition, even though this has the effect of raising prices. Sweeting et al. (2022) examine the effects of forward-looking buyer behavior, extending the Besanko et al. (2010) and Besanko et al. (2014) models, allowing for buyers to internalize some proportion of their effects on future buyer surplus, where the monopsonist case corresponds to full internalization.[6]

Our post-2016 counterfactuals also consider the possibility of innovation, and we analyze the effects of competitive pressure on innovation in a dynamic model. Goettler and Gordon (2011), Yang (2020) and Igami and Uetake (2020) provide related analyses in the CPU, smartphone chip and hard disk drive industries respectively.[7] Federico et al. (2020) and Gilbert (2022) provide conceptual arguments for how innovation should factor into competition analysis. We provide a complementary analysis that evaluates innovation incentives under different market structures when there is a forward-looking monopsonistic buyer.

We analyze the effects of a merger that internalizes learning across products within the same firm on market outcomes and welfare. As noted, the FTC allowed the ULA merger to proceed because it recognized enhanced learning economies, including their benefits on reliability, as efficiencies. The extent to which mergers actually generate efficiencies and the extent to which claimed efficiencies should be credited in a merger evaluation is controversial (Rose and Sallet (2019)). Our paper is thus related to the growing empirical literature that provides model-based estimates of different types of merger efficiencies and their effects on

[5]The dynamic demand literature, such as Erdem et al. (2003) and Hendel and Nevo (2006) allows for forward-looking buyers, who ignore any effect that they might have on the evolution of the state. In our model, the buyer takes these effects into account.

[6]As well as showing that a forward-looking buyer can soften competition so much that static procurement is optimal, Sweeting et al. (2022) also show that multiple equilibria, which are common in these models when buyers are atomistic, are eliminated when buyers are even moderately forward-looking. While our model is not identical to the model in Sweeting et al. (2022), and so we cannot be certain that the same results apply, the results in that paper may provide the reader with some suggestive comfort that, because we assume a forward-looking buyer who fully internalizes how their choices affect future buyer welfare, our model, which we solve repeatedly in estimation, is likely to have a unique equilibrium. In practice, we have not found multiple equilibria across different simulations at the estimated parameters.

[7]Liu and Siebert (2022) combine a model of firm competition with cost-side LBD in the 64KB static random access memory chip industry, with firm choices about whether to enter the market, whereas we focus on the choice of an incumbent to innovate by introducing a lower cost product.

market outcomes (e.g., Fan (2013), Miller and Weinberg (2017), Elliott et al. (2025) and Fan and Yang (2025)). Our evaluation suggests that, in our setting, the efficiencies identified by the Commission were large enough so that the merger benefited the buyer.

Finally, we contribute to an emerging literature studying economics related to human activities in space and transportation to space. Groesbeck (2020) simulates a model of learning-by-doing for the space launch industry but studies how international embargoes affect reliability. As noted above, Weinzierl (2018) provides a survey of recent developments of both space launches and the development of other activities in space, as well as emerging problems such as the externalities caused by space debris, an issue also analyzed in Rao and Rondina (2025). We analyze some of the trade-offs between inefficiency and planner control suggested by Weinzierl. Our paper therefore connects to the non-space empirical literatures that have tried to compare the efficiency of state-owned vs. non-state-owned firms along different dimensions, although our exercises are focused on control rather than ownership, and are model-based predictions rather than descriptive empirical analyses. For example, La Porta and López-de Silanes (1999) examine the effects of privatization of many Mexican firms in the 1980s and 1990s, finding privatization increases productivity. Chen et al. (2021) similarly find that previously state-owned firms become more productive after privatization in China. Cao et al. (2020) find that the innovative efficiency of firms, measured by the number of patents per dollar spent on R&D, is higher for firms that are partially state-owned than for fully state-owned or fully private firms. In the space context, Kantor and Whalley (2025) study the spillover from publicly-funded R&D to manufacturing productivity and value-added during the Cold War space race between the U.S. and the Soviet Union.

2 Space Launch Industry and Institutional Background

To put our modeling choices and our counterfactuals in context, we now describe the relevant history and institutional background of the space launch industry.

2.1 Timeline of U.S. Launch Systems After the Space Shuttle

Our focus is on unmanned launches that put payloads, including satellites, into stable Earth orbits or outer space.[8] For example, the International Space Station, Starlink and most spy satellites, and the Hubble Space Telescope are all in low Earth orbit (LEO), usually defined as 160 to $2,000$ kilometers above Earth, whereas many weather satellites stay fixed above a particular location and are at much higher altitudes. Examples of suitable U.S. launch systems are Atlas V, Delta IV and Falcon 9, with Ariane 5, Long March 5 and Soyuz 2 as European, Chinese and Russian systems, respectively. These systems are classified as medium-to-heavy launch vehicles (McConnaughey et al. (2012)), defined as being capable of lifting 2 to 50 tonnes of payload to LEO.

The Apollo and Space Shuttle systems were developed in a highly centralized way by the National Aeronautics and Space Administration (NASA), with commercial vendors providing important components. However, following slow development, large cost overruns and accidents, President Reagan's February 1988 "Directive on National Space Policy" instructed future US policy "to encourage to the maximum extent feasible, the development and use of United States private sector space transportation capabilities without direct Federal subsidy" with the aim "to reduce the costs of space transportation and related services." As noted by Weinzierl (2018), the President's 2004 Commission on Implementation of United States Space Exploration Policy argued that "NASA's role must be limited to only those areas where there is irrefutable demonstration that only government can perform the proposed activity." This emphasis on commercial launch services has been repeated by subsequent administrations.

In 1994, the U.S. government initiated the "Evolved Expendable Launch Vehicle" (EELV) program to make launches more affordable and reliable through the development of lower cost systems based on older technology.[9] This was followed by a sequence of both innovations and changes in market structure.

In 1995, Lockheed and Martin Marietta, which had previously developed the Atlas II and

[8] Stable orbits are typically at altitudes above 160 kilometers to avoid atmospheric drag.

[9] The EELV program became the National Security Space Launch program in March 2019. Within the Department of Defense (DoD), the EELV program was initiated by the Air Force, but it transitioned to the Space Force in December 2019.

Titan systems respectively based on intercontinental ballistic missile technologies, merged and developed the Atlas V system as part of the EELV program.[10] 1996 and 1997 saw mergers of Boeing, Rockwell (an engine developer) and McDonnell-Douglas, with Boeing developing the Delta IV as an alternative EELV.[11] The first post-testing and certification EELV launches took place in 2002, with Boeing and Lockheed Martin subsequently competing for government launches. The competition for government contracts was viewed as fierce, with two Boeing (former Lockheed Martin) employees criminally charged with undertaking industrial espionage (Spaceflight Now (2003)).

In 2006, Lockheed Martin and Boeing proposed to form the United Launch Alliance joint venture. When it did not challenge the deal, the FTC issued a public document, the "Analysis of Agreement Containing Consent Order to Aid Public Comment", explaining that the Department of Defense (DoD) "has informed the Commission that the creation of ULA will advance U.S. national security interests by improving the United States' ability to access space reliably", citing how "[f]irst, the single ULA workforce will benefit from a launch tempo [...] greater than could be expected from the two separate Lockheed and Boeing workforces" and "[i]n addition, integrating the two firms' complementary technologies will infuse each firm's launch vehicles with the technical improvements and innovations of its competitor, further enhancing the reliability". The analysis concluded that "[u]nder these unique circumstances, the increase in reliability can be recognized as an efficiency flowing from the joint venture."[12] Then-Commissioner Kovacic later explained that the "FTC's approval rested on two assumptions: that the claimed efficiencies were significant, and that the DoD and [NASA] would use best efforts to facilitate entry into the launch services sector" (Kovacic (2019)).

The connection to possible entry reflected how the ULA JV was proposed around the same time that SPX was undertaking the first test flights of Falcon 1, which was based on new technological choices with the potential to lower launch costs dramatically, but whose

[10] The Federal Trade Commission agreed to a consent order with the parties that was aimed at allowing other launch system developers access to certain technologies (`https://www.ftc.gov/sites/default/files/documents/commission_decision_volumes/volume-119/ftc_volume_decision_119_january_-_june_1995pages_618-723.pdf`, last accessed on January 19th, 2026).

[11] Boeing and McDonnell Douglas had jointly developed earlier versions of Delta rockets.

[12] `https://www.ftc.gov/sites/default/files/documents/cases/2006/10/0510165analysis.pdf`, last accessed on January 19th, 2026.

success was far from certain.[13] SPX's Falcon 9 rocket was not certified for national security launches until May 2015, and its first such launch took place in 2017. Therefore, between 2007 and 2016, the government had to rely on a monopolist ULA.

After SPX's certification, the launch market was, once again, essentially a duopoly, with anecdotal evidence of fierce price competition. For example, SPX won its first contract by undercutting ULA on an Air Force satellite contract by 40%, leading ULA to lower its Atlas V and Delta IV prices by as much as one-third (Klotz (2016), Klotz (2017)). Simultaneously, ULA announced the development of its Vulcan Centaur system, which was also designed to have reusable first-stage rockets, to lower costs and to replace Atlas and Delta.[14] Delta IV and Atlas V undertook their final government launches in 2024 and 2025, and after various delays, Vulcan was certified in 2025.[15]

Falcon 9 was also heavily used to launch SPX's commercial Starlink satellites, with over 8,811 satellites launched as of October 30, 2025.[16] Figure 1 shows how commercial (with Starlink broken out), civilian (non-military government) and military launches have varied over time, both in the US and globally. Apart from a brief period in the late 1990s, when there were many commercial launches related to the dotcom bubble, civilian government and military launches accounted for the majority of launches until Starlink's development. We will explain how we account for commercial launches when we simulate our model.

2.2 The Role and Conduct of US Government Procurement

We briefly discuss some real-world features of launch procurement and how they relate to the assumptions of our model.

[13]Falcon's initial development was partly funded by a $396 million NASA "Commercial Orbital Transportation Services" development contract, but an extra $1.6 billion contract for deliveries to the ISS was only secured in 2008 when the fourth launch of Falcon 1 was successful. In a 2014 "60 Minutes" interview, Musk described how "[w]e were running on fumes at that point. We had virtually no money... a fourth failure would have been absolutely game over. Done."

[14]As well as competition from SPX, ULA faced the problem that Atlas V used Russian RD-180 engines, which were subject to congressional policy which made them more difficult to use after the Russian invasion of Crimea in 2014 (Smith (2022)).

[15]Vulcan began to win future launch contracts in 2022 but did not receive government certification for national security payload launches until 2025.

[16]The real time count can be tracked at `https://www.space.com/spacex-starlink-satellites.html`, last accessed on November 5th, 2025. A Falcon 9 launch can put between 24 and 48 Starlink satellites into orbit.

Figure 1: US and Global Rocket Launches 1985-2024 by Launch Types

(a) US Launches

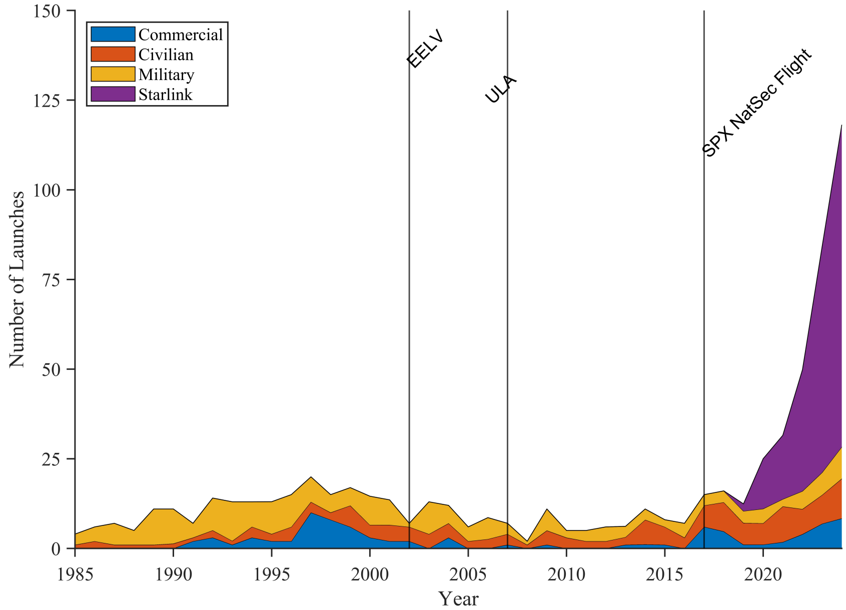

(b) Global Launches

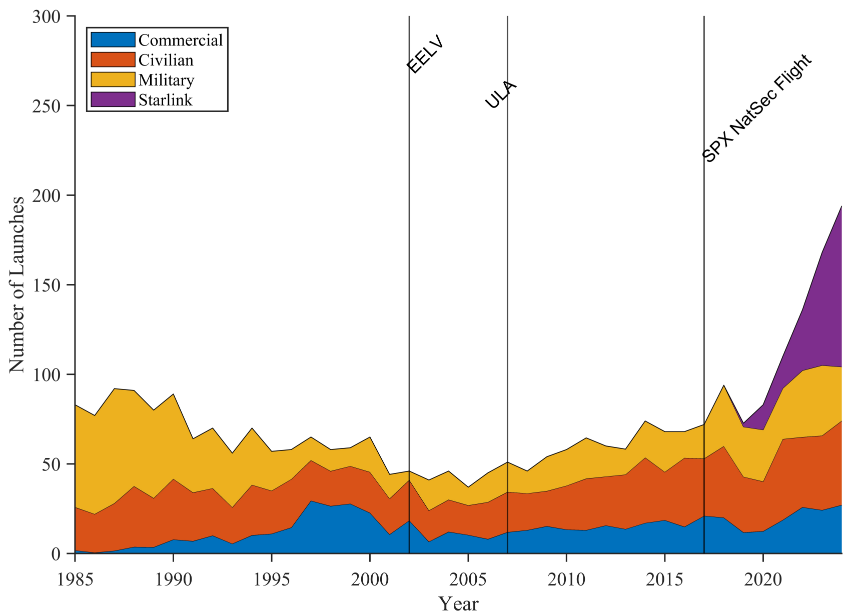

Note: For each payload, we define its type based on the status of its owner. Payloads owned by defense agencies are classified as "Military", while those owned by other government agencies or academic institutions are classified as "Civilian". When a launch carries multiple payloads with different types, we split the launch count in proportion to payload weights.

2.2.1 Long-Lived Forward-Looking Buyer

Our model will assume that the US government acts as a long-lived forward-looking strategic buyer when procuring launches, and we will consider a static procurement rule as a counterfactual. In practice, agencies procure launches under standard federal procurement rules.[17] These rules require agencies to promote "full and open competition through the use of competitive procedures" (Competition in Contracting Act of 1984), although Section 6.202 of the Federal Acquisition Rules allows contract actions to "establish or maintain an alternative source or sources for the supplies or services being acquired if the agency head determines that to do so would (1) [i]ncrease or maintain competition and likely result in reduced overall costs for the acquisition, or for any anticipated acquisition" as well as allowing choices that aim to maintain facilities and capabilities that are in the interests of national defense. As we will discuss below, the expected reliability of launch systems depends on previous experience, and plays an important role in procurement choices. Therefore, the government's rules and practical considerations suggest that the agencies would choose to be forward-looking when purchasing launches, even though litigation has affirmed that they are not legally required to do so.[18]

2.2.2 Procurement Formats

We will treat each potential launch as a separate procurement, and we will be focused on the launch service, not ancillary services (e.g. maintaining launch pads) that support launches.

The reality of contracting is, unsurprisingly, more complicated and has varied over time. For example, from 2002 to 2005 (during what will be our "era 1"), the government used firm fixed-price contracts to purchase launches, but provided Boeing and Lockheed Martin with

[17] Military and intelligence launches are primarily done through Space Force's Space Systems Command, while NASA conducts procurement for civilian science and space exploration launches. We assume that these separate agencies internalize all of the effects on each other's future surplus.

[18] Notably, Blue Origin objected to the Government Accountability Office, and subsequently litigated, the Air Force's decision to only award contracts to two providers in 2019, arguing that this would be "insufficient to incentivize firms to continue to develop launch systems to compete for future awards" (GAO cited in Erwin (2019)). The GAO ruled that procurement rules did not "mandate that the government make multiple contract awards in order to incentivize future private investment necessary to satisfy the government's fulfillment of its future requirements" and that the Blue Origin's claim reflected a policy disagreement with the government rather than an objection to the fairness of the procurement procedure.

extra $500 million contracts to support development and infrastructure, after the bursting of the dotcom bubble reduced commercial demand (U.S. Government Accountability Office (2008)). After 2006, the government signed launch services contracts with ULA for "launch vehicle hardware and labor directly associated with building and assembling launch vehicles" (U.S. Government Accountability Office (2014)) and separate cost-plus launch capability contracts for infrastructure. We will use prices from the launch services contracts, although U.S. Government Accountability Office (2014) criticized the contract split as giving DoD "limited insight into EELV launch costs".

In 2013, DoD made a "major change from past year-to-year contracting" (U.S. Government Accountability Office (2014)) by using a "block buy" contract for the production of 35 launch vehicle booster cores[19] to be produced over the period 2013 to 2017 for launch capability from 2014 to 2019. The contract aimed to allocate costs to launches more clearly, although it continued to include both fixed price and cost-plus elements. The prices we use are from task orders for individual launches covered by these contracts.[20] After SpaceX entered, DoD returned to awarding launch services contracts individually or in small lots.

The purpose of multi-unit purchases is often to "save money and provide predictability for industry" (Congressional Research Service (2025)). One might ask how far our treatment of the time series of prices as all coming from individual procurements is distorted by the fact that some of them come from multi-unit purchases. In Appendix A.2.4, we use the fact that we often see multiple purchases of launches at around the same time under contracts of different sizes, to test whether the differences in prices are substantial. We find that the size of the multi-unit purchase has an effect on launch prices which is statistically insignificant, and the point estimates imply a relatively modest relationship: doubling the purchase size is associated with roughly a 5% reduction in price.[21] This suggests that multi-unit purchases

[19]Delta IV Heavy launches use three cores, whereas regular Delta IV and Atlas V launches use only one (U.S. Government Accountability Office (2014)).

[20]The contract also required ULA to make payments to DoD when ULA performed launches for non-DoD buyers, including NASA.

[21]We find the same pattern when we can make direct comparisons in the data. For example, in 2017 the Air Force procured one launch service for the GPS III-3 satellite and awarded the contract to Falcon 9 for $96.5 million. In 2018, the Air Force procured three launch services for the GPS III-4, -5, and -6 satellites in a single lot and again awarded the contract to Falcon 9, for a total of $290 million, implying a per-launch price of about $96.7 million, i.e., with an almost identical unit price despite the difference in contract size.

do not have large effects on our launch-level price series.

2.2.3 Value of Reliability

We will assume that the government values expected reliability, and that this is reflected by flight experience. This is consistent with how agencies, such as NASA, categorize payloads and launches, although we will not try to capture variation in payload valuations.[22] The government bears the risk if the payload is lost or unable to function. NASA is prohibited from purchasing insurance (Bedell (2011)) and for both commercial and government launches, U.S. law also requires the launch provider and the payload provider to enter in reciprocal waivers of any claims against each other for property damage or loss (FAA (2002)).[23]

3 Data and Descriptive Analysis

In this section, we briefly describe the sources of our data, the construction of our price measure, and present evidence that rocket reliability increases in the number of recent launches, which will motivate us to include a measure of experience, or know-how, in the buyer's expected utility from a launch. Further details of the data are included in Appendix A.

3.1 Launch Data and Reliability

Our first dataset contains information on unmanned orbital rocket launches using medium- to heavy-lift rockets by governments or government agencies in major countries (US, Soviet Union/Russia, China, Japan, India or members of the EU), between 1985 and 2024. It is taken from the General Catalog of Artificial Space Objects (GCAT) (McDowell (2025)), with some augmentation from Gunter's Space Page (Krebs (2025)). The dataset includes

[22]For example, NASA uses a four category payload classification. For example, Class A for payloads is for missions with "the lowest risk tolerance ... a very high priority mission with very high complexity", including, for example, the James Webb telescope (NASA (2023)). Rockets are classified into three risk groups based on their "management systems, flight experience, design and analysis, testing and risk management strategies" (NASA (2024)). The lowest risk flight category (3), which is used for class A payloads, must have a "more robust flight history representing a 95% demonstrated reliability."

[23]The law does require launch providers to have insurance covering damage to government property during launch or re-entry.

information on launch and rocket characteristics (e.g., rocket family and variant, payload capacity, success of the launch outcome), but not prices.

Our focus is on rocket families, such as Lockheed Martin or ULA's Atlas V family or SPX's Falcon 9 family. Families usually have several variants that can be used for particular types of missions or payloads, but they share common components such as engines. We refer to the most frequently used variant of each family as the representative variant and use this variant's characteristics to describe the rocket family.

Table 1: Summary Statistics of Global Governmental Launches, 1985–2024

	Obs	Mean	Std. Dev.	Min	Max
Launch Success	1,934	0.963	0.188	0	1
#{Launches in Past 2 Yrs}	1,934	26.779	28.588	1	157
LEO Payload Capacity (in tonnes)	1,934	9.032	5.638	2.100	25.000
$\mathbb{1}$\{#stages>2\}	1,934	0.658	0.474	0	1
$\mathbb{1}$\{US Launch\}	1,934	0.213	0.409	0	1
$\mathbb{1}$\{Russia/Soviet Union Launch\}	1,934	0.447	0.497	0	1
$\mathbb{1}$\{China Launch\}	1,934	0.202	0.402	0	1
$\mathbb{1}$\{EU Launch\}	1,934	0.063	0.242	0	1
$\mathbb{1}$\{Japan/India Launch\}	1,934	0.075	0.264	0	1

Note: This table presents summary statistics for the global launch data from 1985 to 2024. Each observation corresponds to a governmental orbital launch. For each launch, #{Launches in Past 2 Yrs} denotes the number of launches of the corresponding rocket family within the preceding two years, including both governmental and commercial launches. For Atlas V and Delta IV, this measure also includes launches by the other rocket family occurring after the formation of the ULA.

Table 1 presents the summary statistics from this global launch data. We exclude the initial launches of each rocket family in its entry year, as well as any launch conducted by a family that has been inactive for more than two years, since these launches are typically test flights. There are 1,934 launches meeting our criteria from 1985 to 2024. About 64% of these launches target LEO and these launches account for 70% of the total payload mass in our sample.

In the table, we define a launch as a success if and only if the payload reaches the pre-specified orbit. There are a number of failure scenarios, ranging from serious accidents such as explosions during takeoff, leading to total payload loss, to potentially correctable failures such as malfunctions of upper stage boosters that leave the payload in a lower orbit.[24]

The rockets in our sample have an average carrying capacity of 9 tonnes to a "standard" LEO,[25] though this capacity varies substantially across rocket families.[26] The main U.S. vehicles, such as Atlas V, Delta IV, and Falcon 9, have an average LEO payload capacity of 13 tonnes from 1985 to 2024. The average reported US payload mass to LEO in our data is 5 tonnes.

The number of stages approximates the complexity of the rocket, where 66% of launches use rockets with more than two stages.[27] Of the 1,934 launches, 21% are from the United States. Approximately half of the Soviet/Russian launches took place before the dissolution of the USSR.

3.1.1 Recent Launch Experience and Reliability

Government-mission payloads are usually far more expensive than the rockets that carry them, so that, from the buyer's perspective, reliability, or the expected success rate, is one of the most important characteristics of a rocket.[28] Although some features of a rocket cannot be changed after the design phase, the industry perceives that reliability can be improved significantly by accumulating experience (Office of Commercial Space Transportation (2002), Moore (2019)) through continuous operations, post-flight performance analysis, and iterative design refinements. The FTC's analysis of the ULA JV also recognized that more frequent launches (a higher "launch tempo") could improve operational safety.

We measure a rocket family's experience by its recent launch cadence, defined as the number of launches of the rocket family within the preceding two years. This experience measure includes both governmental and non-government (commercial) launches based on the GCAT data. We choose the two-year window to reflect the fact that, because of orga-

[24]Even if an injection problem can be corrected, this usually diminishes performance or reduces the payload's useful lifespan (NASA Jet Propulsion Laboratory (2022), Carlier (2016)).

[25]The standard LEO is defined here as a 200 km (120 mi) circular orbit at 28.7° inclination.

[26]The capacity also differs across orbits for a given rocket variant, and across variants within a rocket family. Higher target orbits imply lower capacities. Some variants have additional boosters allowing them to launch heavier payloads to a given orbit.

[27]Multiple stages increase fuel efficiency by discarding deadweight during ascent.

[28]Walter Lauderdale, Chief of the Falcon Division and former Technical Director of the EELV Program at the U.S. Air Force, quoted in Hitchens (2021), stated that earlier GPS III satellites had a price tag of about $500 million, while a Falcon 9 launch costs around $65 million.

nizational forgetting (Benkard (2000)), experience is likely to depreciate.[29] We will explain, in the model and estimation sections, how this measure serves as a proxy for intangible know-how. Table 1 shows that the average rocket family experience across launches in our data is 27, reaching a maximum of 157 for SPX's Falcon 9 at the end of our sample.

When calculating experience, we assume that after the ULA joint venture, Atlas V and Delta IV share learning: each family's experience equals its own launches in the preceding two years plus the other family's post-ULA launches in that time window. We merge the experience of the two ULA rocket families for two reasons. First, Kovacic (2019) explains that the ULA JV was expected to generate learning synergies between the two rocket families. Second, Appendix A.3.2 presents additional regression results that are consistent with additional launches of one ULA rocket family increasing the reliability of the other.

We find a robust positive relationship between the number of past launches and mission success. Appendix A.3.1 presents graphical reliability curves illustrating the relationships between accumulated success rates and the number of launches that we see for both the US and non-US launches in our data. The relationships are similar for all major launch systems from different countries. Here, we take a regression approach, and explore whether selection could affect the reliability relationship. We use global launches for this analysis because launch failures are rare.

Our first model is the simple probit model shown in equation (1). The dependent variable, l_{ijt}, is the success of a launch of rocket family j for mission i at time t. On the right-hand side, e_{jt} denotes family j's experience at time t, measured as the number of launches of family j within the preceding two years. The control variables X_{ijt} include rocket carrying capacity, number of stages, and mission country/region fixed effects, and in an alternative specification, X_{ijt} consists solely of rocket family fixed effects.

$$l_{ijt} = \mathbb{1}\left\{\alpha_0 + \alpha_e \ln e_{jt} + \alpha_X X_{ijt} + \psi_{ijt} > 0\right\}, \quad \psi_{ijt} \sim N(0,1) \tag{1}$$

Columns (1)–(4) of Table 2 report the marginal effects for the estimated probit mod-

[29] Office of Commercial Space Transportation (2002) notes that constant monitoring of vehicle performance is necessary to maintain high reliability; even proven systems may lose reliability due to changes in manufacturing or operating procedures.

Table 2: Rocket Reliability Estimation Results

	Probit Model				Probit Model and Selection Equation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$\mathbb{1}$\{Succ.\}	$\mathbb{1}$\{Succ.\}	$\mathbb{1}$\{Succ.\}	$\mathbb{1}$\{Succ.\}	$\mathbb{1}$\{Succ.\}	$\mathbb{1}$\{Succ.\}	$\mathbb{1}$\{Succ.\}	$\mathbb{1}$\{Succ.\}
Rocket Reliability								
$\ln(\#\{$ Launches in Past 2 Yrs$\})$	0.014***	0.015***	0.021***	0.019***	0.015**	0.017***	0.024**	0.020***
	(0.004)	(0.004)	(0.005)	(0.006)	(0.006)	(0.005)	(0.010)	(0.007)
$\ln(\{$LEO capacity$\})$	-0.015**	-0.015**	-0.009		-0.016**	-0.016**	-0.010	
	(0.007)	(0.007)	(0.007)		(0.008)	(0.007)	(0.008)	
$\mathbb{1}\{\#\text{stages} > 2\}$		-0.015	0.003			-0.017	0.003	
		(0.010)	(0.011)			(0.010)	(0.013)	
Error Term Distribution								
σ_ξ					0.065	0.075***	0.121	0.051
					(0.200)	(0.020)	(0.264)	(0.138)
ρ					0.748**	0.790***	0.776**	0.745
					(0.300)	(0.263)	(0.362)	(0.485)
Mission Region FE			Yes				Yes	
Rocket Family FE				Yes				Yes
LR Test Statistic: Demand Shifters					195	183	109	86
Observed Launches	1,934							

Note: *,**,*** denote significance at the 10%, 5% and 1% levels. Standard errors in parentheses. Results in the Rocket Reliability panel are average marginal effects. Columns (1)–(4) correspond to the probit model specified in Equation (1); columns (5)–(8) correspond to the joint model of rocket selection and rocket reliability specified in Equations (1) and (2). The outcome $\mathbb{1}$\{Succ.\} equals 1 if and only if the GCAT data indicates that the launch and the insertion of the payload were completely successful.

els with different control variables based on our global launch data. The coefficient on $\ln(\#\{\text{Launches in Past 2 Yrs}\})$ in Column (4) implies that a doubling of experience increases the probability of success by approximately 1.32 percentage points, which, given the average probability of success is 0.96, means that the probability of some type of failure falls by around one-third. For the 122 ULA Atlas V and Delta IV launches, the model predicts 3.67 failures, declining to 2.56 expected failures if the experience level at the time of each launch is doubled.

There are several possible selection issues with this specification. We would be especially concerned about false positives, where, for example, governments choose rocket families with low experience for more complex missions, which are more likely to fail, because these families have unusual characteristics tailored for these missions.

We address this issue by augmenting the probit model with a selection stage modeled as a multinomial logit model. The model captures the choice of a government with a potential launch opportunity to use one of the rocket families available to it, which we restrict to families that have been active in its country within the past year.[30]

For the logit model, we define the "indirect utility" of assigning family j to opportunity i at time t as

$$u_{ijt} = \beta_0 + \beta_e \ln e_{jt} + \beta_X X_{ijt} + \beta_z Z_{ijt} + \xi_{ijt} + \epsilon_{ijt} \tag{2}$$

$$\begin{bmatrix} \psi_{ijt} \\ \xi_{ijt} \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho\sigma_\xi \\ \rho\sigma_\xi & \sigma_\xi^2 \end{bmatrix} \right) \tag{3}$$

where Z_{ijt} denotes a set of rocket family-specific demand shifters, which do not appear in the reliability equation. For these shifters, we interact two-year-lagged worldwide war casualties and the two-year-lagged GDP of the mission country with rocket family characteristics.[31]

[30]To construct the set of potential launches, we aggregate launches by year, country, sector (civil or military), and orbit, and then, for each country–sector–orbit combination, identify the maximum annual launch count across years. We then define the number of potential launches for each category as 1.1 times this maximum observed number of launches, thus giving the government an outside option of no launch. In our sample of 1,934 launches, 7.5% of government launches use foreign rockets. Of these launches, 53% belong to the EU, India, or Japan; 41% are Russian missions between 1991 and 2006 that used Ukrainian rockets inherited from the Soviet Union; the remaining 6% are joint missions involving the US, which account for 2% of US launches. The Chinese government never uses foreign rockets. Overall, international trade in government launches is very limited.

[31]Specifically, the rocket family characteristics include (i) $\ln(\#\{\text{Launches in Past 2 Yrs}\})$, (ii)

The unobservable ξ_{ijt} represents a demand shock, which we allow to be correlated with the reliability equation unobservable, and ϵ_{ijt} follows a standard type I extreme value distribution. The utility from the outside option is $u_{i0t} = \epsilon_{i0t}$. We estimate the models jointly using maximum likelihood.

Columns (5)–(8) of Table 2 report the marginal effects from the reliability equation in the joint model, using the demand shifters as the Zs.[32] The estimated marginal effects are similar to those in columns (1)-(4), consistent with the space launch engineering literature's evidence that launch experience improves reliability. The demand shifters are jointly significant across all specifications, as shown by the likelihood ratio (LR) test statistics reported in the table, all of which exceed 80 (p-values less than 0.001).[33] The estimated standard deviation of the unobserved demand shock, σ_ξ, is small, and declines further once rocket family fixed effects are included, suggesting that selection on unobservables is unlikely to be a concern for the reliability probit estimates. Finally, in additional sensitivity analysis, we find that the estimated effects of experience are robust to including additional covariates such as the number of years since a rocket family was introduced (vintage).

3.2 Price Data

GCAT does not record either bids or transaction prices. Appendix A describes how we construct a new dataset of launch transaction prices from FAA launch reports, government contract records, public news reports, and industry journals, and how we use these data and

$\ln(\{\text{LEO capacity}\})$, (iii) $\mathbb{1}\{\#\text{stages} > 2\}$. The world casualties from wars are constructed using data from the Polynational War Memorial (https://www.war-memorial.net/wars_all.asp, retrieved on September 10, 2024). For each conflict, we observe its duration and total number of battle deaths. We assume that battle deaths are distributed evenly across the years of the conflict. The GDP data are obtained from the National Accounts Estimates of Main Aggregates published by the United Nations Statistics Division (https://data.un.org/Data.aspx?d=SNAAMA&f=grID%3a102%3bcurrID%3aUSD%3bpcFlag%3a0, retrieved on August 9, 2025). When computing the GDP of the European Union, we include all countries that are ever part of the EU. For Russia, we use the GDP of the Soviet Union before 1991 and the GDP of Russia thereafter. The logic of these demand shifters is that international tensions may, for example, increase demand for spy satellite launches in general, and/or those rocket families that are particularly suited for launching spy satellites. GDP may affect a country's ability to pay for more complex launch systems.

[32] When computing marginal effects, we hold the selection process fixed and focus solely on the reliability equation, computing marginal effects only for selected rockets.

[33] The logit coefficients imply that increases in a country's GDP or in global conflict intensity are associated with greater demand for rocket families with higher payload capacity and more stages, which are better suited for heavier payloads such as spy satellites or large scientific instruments.

hedonic regressions, to construct a measure for the average launch price in each year for the representative variant of each US rocket family.

Figure 2: Imputed U.S. Rocket Family Price Series (2010 dollars)

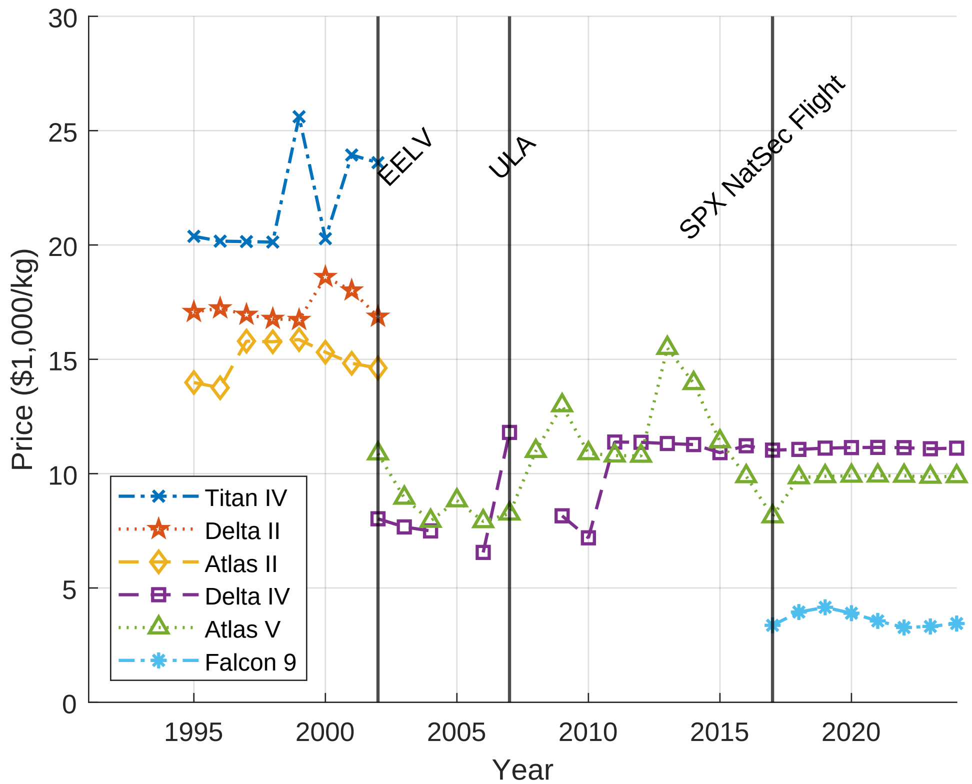

Note: See Appendix A for details of the sources and construction. Each rocket family's price corresponds to the representative variant and is expressed in $1,000 per kilogram of payload capacity to standard LEO in 2010 dollars. The first line indicates the initial launch of the EELV rockets (Atlas V and Delta IV). The second line marks the JV between Boeing (owner of the Delta series) and Lockheed Martin (Atlas series) to form ULA. The third line marks the entry of Falcon 9 rockets into national security launches.

Figure 2 presents the resulting rocket family-level launch prices. Following a standard industry practice (Mathieu et al. (2022)), the price units are \$1,000s (2010 dollars) per kilogram of LEO payload capacity of the family's representative variant.[34] Three vertical lines mark key events in the US launch market. The first line indicates the initial launch of the EELV rockets (Atlas V and Delta IV). The second line marks the JV between Boeing (owner of the Delta series) and Lockheed Martin (Atlas series) to form ULA. The third line marks the entry of Falcon 9 rockets into national security launches. For comparison, Space Shuttle launches cost \$54,500/kg.

Average prices have fallen with each new generation. Vulcan is not shown, but, in 2023, ULA was awarded contracts for 11 future national security missions on Vulcan rockets. Given Vulcan's payload capacity, the implied average price was \$3,293 per kilogram (Sheetz (2023)), similar to Falcon 9's \$3,317/kg price in 2023. Finally, the figure shows that EELV rocket prices, especially for Atlas V, increased after the ULA JV and fell around the time of SPX's entry.[35]

4 Model

This section describes our model, the assumed process for know-how evolution and equilibrium strategies.

4.1 Overview

We solve distinct stationary, infinite horizon, discrete state dynamic games for three eras since 2002: era 1 covers the Boeing and Lockheed Martin single-product duopoly from 2002 to 2006; era 2 covers the ULA multi-product monopoly from 2007 to 2016; and era 3 covers the ULA and SPX duopoly from 2017 to 2023. We assume that each transition to the next era is an unanticipated and exogenous shock.

[34]Within each rocket family, variants differ in carrying capacities and therefore can have substantially different prices when measured per launch. However, these variants often have similar prices per kilogram of capacity to LEO, as illustrated in Appendix Figure A.1, providing a fairly consistent measure of cost-effectiveness.

[35]Due to lead times in contracting over actual launches, it is not surprising that we may see some disciplining effect of SPX on ULA prices before SPX launches begin.

In eras 1 and 2, we assume that there is no product entry or exit, so that the state space is defined by the know-how (depreciated launch experience) of each product. There is a potential U.S. government launch procurement in each period. A product's know-how will affect the government's valuation of using that product, and it may also affect the seller's costs. Both the government's preferences and the seller's production costs will also depend on private information shocks. In era 3, ULA stochastically receives opportunities to innovate and upgrade its products. In each period, timing is as follows:

1. Sellers privately observe their production costs (but not government preferences) and simultaneously submit bids for each product.

2. The government buyer privately observes its preferences and chooses one of the products, or the outside good, to maximize its payoff.

3. Conditional on the procurement outcome and the ownership of the products, know-how evolves stochastically.

4. (Era 3, ULA only) If ULA has not innovated, ULA stochastically receives an opportunity to innovate and observes a cost of innovation. Innovation replaces Atlas V and Delta IV with a new product, Vulcan Centaur, which has a lower production cost given the same know-how level.

In the above, if the buyer uses a static procurement rule, the payoff is its current flow payoff. If the buyer is forward-looking, which we assume in estimation, the buyer's payoff is the sum of its flow payoff and its discounted future value, recognizing that its current choice can affect its future surplus through the evolution of sellers' know-how.

4.2 Stage Game

We begin by explaining payoffs and strategies in the procurement stage game, considering the most complicated era 3 case with two competing sellers, denoted by $i \in \{\text{ULA,SPX}\}$. We use j to denote a product, where $j \in \{\text{Atlas (A), Delta (D), Falcon (F)}\}$ before ULA innovation and $j \in \{\text{Vulcan (V), Falcon (F)}\}$ afterwards. Atlas, Delta and Vulcan are owned by ULA.

Procurements (time periods) are denoted by t, and we assume $t = 1, ..., \infty$. The discount factor is δ, and when we build in commercial launches, we will be assuming that there are 20 procurements per year. Before ULA innovation, the industry state is $\mathbf{e} = (e_A, e_D, e_F)$, where the know-how states (e_j) will be discrete. The post-innovation state is (e_V, e_F). The players (sellers and buyer) are assumed to observe the current know-how state.

4.2.1 Seller Costs

If product j is chosen in a procurement at time t, the production cost of the launch will be

$$c_{jt} = c_j(e_{jt}, \zeta_{jt}) = \exp\left(\gamma_j - \gamma^e \ln e_{jt} + \sigma_\zeta \zeta_{jt}\right) \tag{4}$$

where γ_j is a rocket family intercept term (we will assume $\gamma_A = \gamma_D$, but they may be different from γ_F), and ζ_{jt} is an i.i.d. normally distributed private cost shock, with probability density function $g(\zeta)$, observed by the owner of product j. Therefore, a pre-innovation ULA observes two cost shocks, while a post-innovation ULA, or SPX, observes one. As know-how will tend to increase with past launches, $\gamma^e > 0$ implies the existence of "learning-by-doing" economies in production.

4.2.2 Procurement and Buyer Payoffs

In each period, there is a procurement for a single launch.[36] The sellers simultaneously submit bids, b_{jt}, for each product. The chosen product will be paid its bid. If product j is chosen for the launch, the flow indirect utility of the buyer is

$$u^B_{jt} = \beta_{\text{era}} + \beta_j + \beta^e \ln e_{jt} - \alpha b_{jt} + \varepsilon_{jt}, \tag{5}$$

where for the β_js, we assume $\beta_A = \beta_D$, but they may be different from β_F. If the outside good (no launch) is chosen, the indirect utility is

$$u^B_{0t} = \varepsilon_{0t}. \tag{6}$$

[36]We do not distinguish between national security and civilian government launches.

The εs are type I extreme value i.i.d. payoff shocks that are private information to the buyer when bids are made.

A feature of our model is that know-how, and therefore, learning-by-doing, affect both demand and costs. Of course, it is not necessarily the case that a model where there is learning-by-doing on both sides will give different equilibrium predictions than a model where there is a greater learning-by-doing effect on one side,[37] but in our model, we assume that know-how proportionally reduces the production costs, consistent with NASA's reported estimates of progress ratios for aerospace industries (Section 5.1). Given the different, nonlinear effects of know-how on the demand and cost sides, differences will emerge in equilibrium.

In practice, payload weights and orbits differ across observed missions. However, there is no information for launches that were not procured. We therefore abstract away from these sources of variation and instead think of all procurements as representative missions taking a payload capacity of 15 tonnes to standard LEO.[38]

4.2.3 Equilibrium Procurement Strategies

We assume that the firms play a stationary, symmetric Markov Perfect Nash Equilibrium (MPNE, Maskin and Tirole (2001)) in each era. The Markovian restriction is that strategies can depend only on payoff relevant state variables. Here we describe strategies in the procurement stage game. To do so, we will use notation for intermediate value functions, $W_i^{S,INT}(\mathbf{e}, k)$ for sellers and $W^{B,INT}(\mathbf{e}, k)$ for the buyer, defined as values after the buyer's procurement choice to buy rocket (or no rocket) k has been made, but before know-how transitions or innovation choices (if relevant) are realized.

Given a vector of bids $\mathbf{b}$ and the logit assumption on the preference shocks, the optimal

[37]For example, Nocke and Schutz (2018) show that with static logit demand, Bertrand Nash equilibrium outcomes, such as choice probabilities, depend on indices that vary with the $\delta_i - c_i$ of each product, where δ_i is the expected indirect utility intercept of good i and c_i is the marginal cost of good i. Therefore, if a characteristic x_i linearly increases indirect utility and linearly lowers marginal costs, what is important for the equilibrium outcomes will be the difference between the demand and cost coefficients on x_i. The dependence of outcomes on this difference also holds in logit demand examples we have considered with multiple periods and forward-looking buyers, when we allow for indirect utility and costs to vary linearly with know-how.

[38]The capacity of 15 tonnes is close to the average reported LEO carrying capacity of rockets used in US government LEO missions (14.5 tonnes) and in all US government missions (15.2 tonnes) from 2002 to 2024.

choice probability of a forward-looking buyer for one of the inside products k is

$$D_k(\mathbf{b},\mathbf{e}) = \frac{\exp\left(W^{B,INT}(\mathbf{e},k) + \beta_{\text{era}} + \beta_k + \beta^e \ln e_{kt} - \alpha b_{kt}\right)}{\exp\left(W^{B,INT}(\mathbf{e},0)\right) + \sum_{n=A,D,F} \exp\begin{pmatrix} W^{B,INT}(\mathbf{e},n) + \beta_{\text{era}} + \beta_n \\ + \beta^e \ln e_{nt} - \alpha b_{nt} \end{pmatrix}}. \tag{7}$$

The optimal seller bids will be characterized by first-order conditions where a seller's costs are appropriately adjusted to reflect future values and diversion to commonly-owned products, as well as the realized production costs of the products. For Falcon, before ULA innovates, $b_F^*(\zeta_{Ft},\mathbf{e})$ will satisfy

$$\begin{aligned} \iint \Bigg\{ & D_F(b_F, \mathbf{b}^*{}_{-F}(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e}) + \frac{\partial D_F(b_F, \mathbf{b}^*{}_{-F}(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e})}{\partial b_F}\left[b_F - c_F(e_{Ft},\zeta_{Ft}) + W^{S,INT}_{SPX}(\mathbf{e},F)\right] \\ & + \sum_{n=0,A,D} \frac{\partial D_n(b_F, \mathbf{b}^*{}_{-F}(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e})}{\partial b_F} W^{S,INT}_{SPX}(\mathbf{e},n) \Bigg\} g\left(\zeta_{Dt}\right) g\left(\zeta_{At}\right) d\zeta_{Dt} d\zeta_{At} = 0. \end{aligned} \tag{8}$$

For pre-innovation ULA, the first-order condition for Atlas with respect to $b_A^\star$ will be conditional on ζ_{At} and ζ_{Dt},

$$\begin{aligned} \int \Bigg\{ & D_A(b_A, b_D, b^*{}_F(\zeta_{Ft},\mathbf{e}),\mathbf{e}) + \frac{\partial D_A(b_A, b_D, b^*{}_F(\zeta_{Ft},\mathbf{e}),\mathbf{e})}{\partial b_A}\left[b_A - c_A(e_{At},\zeta_{At}) + W^{S,INT}_{ULA}(\mathbf{e},A)\right] \\ & + \frac{\partial D_D(b_A, b_D, b^*{}_F(\zeta_{Ft},\mathbf{e}),\mathbf{e})}{\partial b_A}\left[b_D - c_D(e_{Dt},\zeta_{Dt}) + W^{S,INT}_{ULA}(\mathbf{e},D)\right] \\ & + \sum_{n=0,F} \frac{\partial D_n(b_A, b_D, b^*{}_F(\zeta_{Ft},\mathbf{e}),\mathbf{e})}{\partial b_A} W^{S,INT}_{ULA}(\mathbf{e},n) \Bigg\} g\left(\zeta_{Ft}\right) d\zeta_{Ft} = 0. \end{aligned} \tag{9}$$

The ULA's first-order condition with respect to b_D is similarly defined. Given logit demand,

$$\frac{\partial D_j(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e})}{\partial b_j} = -\alpha D_j(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e})(1 - D_j(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e})) \tag{10}$$

and

$$\frac{\partial D_j(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e})}{\partial b_k} = \alpha D_j(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e}) D_k(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e}) \text{ if } j \neq k. \tag{11}$$

Computational Approximation. When we solve the model, we assume that the private cost shocks are restricted to be between the 2.5% and 97.5% quantiles of the normal distribution, and that bids are restricted to be from 50% of the production cost (when the cost shock corresponds with the 2.5% quantile) and 3 times the highest cost (when the cost shock corresponds with the 97.5% quantile). For computational reasons, we will approximate the continuous distribution assumption by using 10 discrete values of the possible cost shocks and 10 discrete possible bid levels. Our computational method for solving the procurement game follows Bichler et al. (2025). Its role within our algorithm for solving the dynamic game is explained in detail in Appendix B.

4.2.4 Value Functions, State Transitions and Innovation

We assume that production takes place immediately after the procurement. Conditional on the buyer's choice k, the know-how will transition with probabilities $\Pr(\mathbf{e}_{t+1}|\mathbf{e}_t, k)$. We detail the transition probabilities in Section 4.3.

Without possible innovation, the intermediate values $W_i^{S,INT}$ and $W^{B,INT}$ are weighted averages of the values at the start of the next period:

$$W_i^{S,INT}(\mathbf{e}, k) = \delta \sum_{e'} \Pr(\mathbf{e}'|\mathbf{e}, k) W_i^S(\mathbf{e}') \tag{12}$$

and

$$W^{B,INT}(\mathbf{e}, k) = \delta \sum_{e'} \Pr(\mathbf{e}'|\mathbf{e}, k) W^B(\mathbf{e}'), \tag{13}$$

when the winning rocket is k.

For pre-innovation ULA in era 3, we assume that ULA receives an innovation opportunity with probability $\frac{1}{60}$ each period, implying that it will receive at least one opportunity in a twenty-period year with probability 0.29. The realized innovation cost for a given opportunity is $C + \sigma_C \nu_{nt}$, where ν_{nt} is an i.i.d. logistic random variable. The parameters we set for C and σ_C are described in Section 5.[39] Given these parameters, the intermediate values will

[39] We note that ULA, or Boeing and Lockheed Martin, might have considered introducing a new rocket in earlier eras, as the experience of the Space Shuttle and SPX's development efforts made the theoretical possibility of lowering costs through reusability widely known. While we do not include the possibility of era 1 or era 2 innovation in the model, our counterfactuals will assess how ULA's innovation incentives in

be weighted averages of these functions and the value functions where Atlas and Delta have been replaced by Vulcan Centaur, in the lowest know-how state, in the next period. For ULA, the intermediate value is

$$W_{ULA}^{S,INT}(e,k) = \delta \sum_{e'} \Pr(e'|e,k) \left(\frac{59}{60} W_{ULA}^{S}(e') + \frac{1}{60} \tilde{W}_{ULA}(e') \right), \tag{14}$$

where

$$\tilde{W}_{ULA}(e') = \sigma_C \ln \left(\exp\left(\frac{1}{\sigma_C} W_{ULA}^{S}(e') \right) + \exp\left(-\frac{C}{\sigma_C} + \frac{1}{\sigma_C} W_{ULA}^{S}(e_V = 1, e_F') \right) \right), \tag{15}$$

and, based on the implied ULA innovation probabilities, we also adjust the intermediate values of the buyer and SPX.

If ULA innovates, Vulcan will compete indefinitely with Falcon 9 starting from the next period, when Vulcan will start with the lowest level of know-how. As Vulcan was only certified in 2025, we cannot estimate its demand and cost parameters.[40] We therefore assume that Vulcan will have the same cost intercept as Falcon 9, but the same demand intercepts as Atlas V and Delta IV.[41]

era 3 would have been different absent Falcon 9.

[40] When we simulate model-predicted moments for estimation, we will condition on innovation not having occurred in order to match them to the data.

[41] The new Vulcan rocket may be equipped with new technologies that improve the government buyer's utility for choosing the rocket. Our model does not capture this dimension of improvement. Nonetheless, we believe that demonstrated know-how is still the first-order issue in the government decision. For example, delays in the certification flights of Vulcan led the government to shift missions from Vulcan to Falcon rockets (Erwin (2025)).

4.2.5 Value Functions

Values are defined at the start of each period, before buyer preference shocks or seller cost shocks are known. Therefore, before ULA innovates, the buyer's value is

$$W^B(\mathbf{e}) = \Upsilon + \iiint \ln \Bigg(\exp\left(W^{B,INT}(\mathbf{e},0)\right) + \sum_{n=A,D,F} \exp\begin{pmatrix} W^{B,INT}(\mathbf{e},n) + \beta_{\text{era}} + \beta_n \\ + \beta^e \ln e_{nt} - \alpha b^*_{nt}(\zeta_n,\mathbf{e}) \end{pmatrix} \Bigg) g\left(\zeta_{Ft}\right) g\left(\zeta_{Dt}\right) g\left(\zeta_{At}\right) d\zeta_{Ft} d\zeta_{Dt} d\zeta_{At}, \tag{16}$$

where Υ is the Euler constant. SPX's value, also defined before the realization of the idiosyncratic shocks, is

$$\begin{aligned} W^S_{SPX}(\mathbf{e}) = \iiint \Bigg\{ & D_F(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e}) \left[W^{S,INT}_{SPX}(\mathbf{e},F) + b^*_F\left(\zeta_{Ft},\mathbf{e}\right) - c_F(e_{Ft},\zeta_{Ft})\right] \\ & + \sum_{n=0,A,D} D_n(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e}) W^{S,INT}_{SPX}(\mathbf{e},n) \Bigg\} g\left(\zeta_{Ft}\right) g\left(\zeta_{Dt}\right) g\left(\zeta_{At}\right) d\zeta_{Ft} d\zeta_{Dt} d\zeta_{At} \end{aligned} \tag{17}$$

In the case of pre-innovation ULA,

$$\begin{aligned} W^S_{ULA}(\mathbf{e}) = \iiint \Bigg\{ & D_A(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e}) \left[W^{S,INT}_{ULA}(\mathbf{e},A) + b^*_A(\zeta_{At},\zeta_{Dt},\mathbf{e}) - c_A(e_{At},\zeta_{At})\right] \\ & + D_D(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e}) \left[W^{S,INT}_{ULA}(\mathbf{e},D) + b^*_D(\zeta_{At},\zeta_{Dt},\mathbf{e}) - c_D(e_{Dt},\zeta_{Dt})\right] \\ & + \sum_{n=0,F} D_n\left(\mathbf{b}^*(\boldsymbol{\zeta}_t,\mathbf{e}),\mathbf{e}\right) W^{S,INT}_{ULA}(\mathbf{e},n) \Bigg\} g\left(\zeta_{Ft}\right) g\left(\zeta_{Dt}\right) g\left(\zeta_{At}\right) d\zeta_{Ft} d\zeta_{Dt} d\zeta_{At}. \end{aligned} \tag{18}$$

Values post-innovation are integrated with respect to Vulcan's cost shocks ζ_{Vt}, as opposed to ζ_{Ft} and ζ_{Dt}. ULA's value function post-innovation is defined similarly to SPX's.

4.2.6 Equilibrium: Existence and Uniqueness

The MPNE will be defined by seller values, optimal (expected value-maximizing) seller bidding and innovation strategies, buyer values, and optimal buyer procurement strategies, where expectations are consistent with the state transition processes and other players' strategies.

Assuming that an equilibrium exists in the static procurement games when continuation values are held fixed, existence of an equilibrium in the full dynamic game follows from an application of the logic of Proposition 2 in Doraszelski and Satterthwaite (2010). In practice, we have always been able to find an equilibrium up to numerical tolerances.

We would require additional assumptions to claim that there must be a unique equilibrium (e.g., no know-how depreciation, so the game is directional (Iskhakov et al. (2016))). Besanko et al. (2010) show that multiplicity exists for many technology parameters in a duopoly game where there is LBD and stochastic know-how forgetting on the cost-side, and buyers are short-lived and there are no idiosyncratic cost shocks. Our model differs in several ways, including the existence of a forward-looking buyer, private information idiosyncratic cost shocks and an outside good that is often chosen. These features may tend to eliminate multiplicity (Sweeting et al. (2022)), and we have not identified examples of multiplicity when estimating or solving our model for counterfactuals for any era.

4.2.7 Connecting Eras

As already noted, we solve distinct games for the three different market structure eras in our data. The transitions between eras are therefore treated as exogenous surprises to both the government buyer and the sellers.

4.2.8 Extensions

Static Buyer. We estimate the model assuming a long-lived buyer. However, we will investigate what would happen if the government committed to considering only its flow payoffs. The equilibrium in this model is characterized by the same equations, if we set the δ in equation (13) equal to zero (the δ in equation (12) or (14) would be unaffected).

Social Planner. We find the social planner solution as a benchmark. We follow Deng et al. (2025) in thinking of the planner as a long-lived buyer who faces sellers whose bids are equal to their current production costs. The planner's purchase probabilities can then be found by solving equations (7), (13) and (16).

4.3 Know-How States and Transitions: Model and Simulations

We now discuss our definition of the know-how state space and the transition process. We explain how, when we simulate the model, we adapt the transition process so that it can capture more variation in know-how and commercial launches. The details of solution and estimation are given in Section 5 and Appendix B.

As we solve distinct dynamic games for each era, and estimate parameters using a nested fixed point method, we have to limit the size of the state space to make the computation tractable. The burden is highest in era 3 when there can be three products. We consider four levels of know-how for each product, corresponding to $e_{jt} \in \{1, 5, 10, 15\}$.[42] Following Benkard (2004), who also aggregates know-how, we will use a stochastic process to model transitions between these aggregated states.

We assume that know-how transitions occur in two stages. First, if the product is sold and the state is e_{jt}, the know-how state will increase to the next level with probability $\frac{1}{e^U_{jt} - e_{jt}}$ where e^U_{jt} is the know-how of the next bin. If know-how does not increase, it will remain at e_{jt}. For the learning synergy, we assume that both Atlas V's and Delta IV's know-how may increase independently to the next level if either rocket is chosen. The post-learning know-how is denoted by e'_{jt}. Second, know-how depreciation may occur. Specifically, the know-how state after the learning event will drop down to the next level below e'_{jt}, e^L_{jt}, with probability $\lambda \dfrac{e'_{jt}}{e'_{jt} - e^L_{jt}}$, and otherwise stay the same. The value of λ is chosen to be 0.0341. If know-how were continuous, a fraction λ of know-how would depreciate at every procurement, and a value of $\lambda = 0.0341$ would imply that one-half of know-how would be lost every year without a sale. This choice approximates the measure of know-how in Section 3 based on the number of launches in the past two years, where all know-how would be lost after two years of inactivity.

We accommodate commercial launches by Atlas V or Delta IV by slightly modifying the model when we simulate it for counterfactuals. If a rocket has an observed commercial launch in a year, we assume that the rocket makes a sale in the last procurement of the

[42]Benkard (2004) uses a significantly larger state space reflecting the fact that Benkard can estimate the cost and demand functions without solving the model, whereas we use a nested fixed point procedure for estimating a number of parameters.

year. This modification is simple and reasonable as it is infrequent for these rockets to have multiple commercial launches. Between 2002 and 2023, there are eleven years with one Atlas V commercial launch, and three with two launches. Delta IV launched two commercial missions in two separate years.

In contrast, SPX's Falcon has many commercial and Starlink missions, with 15 such missions before 2017 and then an average of 31 missions per year in 2017–2023. As Falcon's depreciated know-how is therefore almost certain to be well above 15 throughout era 3, we simply fix SPX's know-how at 15 throughout the game, except when considering counterfactuals that assume SPX has no commercial launches.[43]

4.4 Limitations

Aside from our focus on representative missions, which ignores how participants may understand that some rockets are better suited to some missions, we recognize several additional limitations. In particular, we do not account for how contracts may be signed many months before launches take place, and our treatment of innovation is (deliberately) very stylized in that we assume that, when it decides to innovate, ULA can introduce Vulcan immediately.

We also do not account for multi-unit purchases where the government procures several launches in a single contract. In particular, ULA has argued that it can offer launches at lower prices when the government provides quantity guarantees through purchases in a block buy contract (Gruss (2014)).[44] However, while we acknowledge that incorporating contracts would be an interesting extension of our model, we show in Appendix A.2.4 that per-unit price discounts on contracts covering more launches seem limited.

[43]We have considered more complicated alternatives to fixing SPX's know-how without seeing significant changes in the results.

[44]The advantage of a block buy contract to the government is that it allows a seller to set prices on the basis of LBD benefits that it can be confident will be realized. Our model, which assumes prices are set procurement-by-procurement, has sellers set prices based on the expected future LBD benefits of a sale. In practice, when we simulate the market from a low-know-how state under different market structures, the average winning price of a rocket over time (averaged across states and cost shocks) typically increases and then levels off, consistent with the relatively flat prices that would emerge in a block buy contract.

5 Estimation

In this section, we explain how we choose certain cost parameters and estimate the remaining ones. We provide some intuition for identification.

5.1 Choice of Cost LBD and Innovation Parameters

We assume that there are twenty periods per calendar year, and set the discount factor at $\delta = 0.9974$, implying an annual discount factor of around 0.95. Unlike Benkard (2004), we lack direct data on production costs or variable inputs per launch, so that it is difficult in practice to estimate both cost and demand-side LBD. We therefore assume that $\gamma^e = 0.234$ in our preferred specification, and consider $\gamma^e = 0$ as a robustness check. $\gamma^e = 0.234$ implies a progress ratio of 0.85, or that costs decrease by 15% if know-how doubles. This is consistent with NASA's reported estimates of progress ratios for aerospace industries.[45] As mentioned in Section 4, the depreciation parameter, λ, is set to 0.0341.

We will take the entry of SPX as exogenous, and consider ULA's incentives to develop Vulcan only in era 3. We set ULA's innovation cost parameters (C, σ_C) based on the reported development costs of ULA's Vulcan ($6 billion, Gruss (2016)), and projected costs of SPX's Starship ($10 billion, Smith (2024); Space Investments (2024)), Blue Origin's New Glenn ($10 billion, Bogaisky (2025)) and NASA's Space Launch System ($40 billion, Spaceflight News (2025)), so that innovation costs have a mean of $16.5 billion and a standard deviation of $15.8 billion. Given our price units of $1,000 per kilogram capacity to standard LEO in the model, and our assumption of a representative rocket capacity of 15 tonnes to LEO,[46] we convert the mean and the standard deviation to $C = 1,100$ and $\sigma_C = 580$ respectively.[47]

[45] `https://web.archive.org/web/20120830021941/http://cost.jsc.nasa.gov/learn.html`, accessed October 14, 2025.

[46] In practice, the government conducts numerous launch missions carrying payloads of varying masses and targeting different orbits on rockets of similar LEO capacities.

[47] We calculate $C = 1,100$ by dividing $15.8 billion by 15 tonnes and changing the price units to $1,000 per kilogram. We calculate σ_C in the same way but take into account the standard deviation of the logistic shock ($\frac{\pi}{\sqrt{3}}$).

5.2 Simulated Method of Moments Estimation

The remaining parameters are era intercepts in the indirect utility of the buyer, the price coefficient, the (demand-side) reliability coefficient on know-how, the production costs of Atlas V/Delta IV and Falcon 9/Vulcan, and the cost shock scaling parameter. We estimate these parameters by matching the model's predictions of the average values of several outcomes, given the parameters, to similar observed average outcomes in the data.

To calculate both data and model-predicted moments, we treat rocket know-how as variables that we can observe, specifically defining know-how by counting the number of government and commercial launches in the previous two calendar years, and then allocating the state for the year to one of the know-how bins. While this does not match up exactly with the stochastic know-how accumulation and depreciation processes that our model assumes (for example, it means the assumed level of know-how is the same for launches within a year), this approach has two practical advantages. First, it provides a way for us to account for observed commercial launches, which should also increase know-how, as well as government launches. Second, it allows us to match conditional moments, where we can compare, for example, the number of launches observed in the data with those predicted by the model (choice probabilities) conditional on a state, rather than performing a complicated integration over the alternative values of the know-how states given the history of the game. Third, in addition to simplifying the estimation process, we can use know-how as an untargeted moment to check the model fit. Specifically, after estimation, we will compare know-how levels when we simulate our model forward, allowing know-how to evolve stochastically, with the levels calculated from the data using the two-year approach. We match the following moments related to choice probabilities and transaction prices,

1. the average choice probability, and the average of the choice indicator interacted with the era indicators;

2. the probability that the highest know-how rocket wins across all years;[48] and the

[48] The data moment is the number of wins by the highest know-how rocket in each year, divided by the assumed total number of procurements, which is the number of years times twenty. For the model moment, the equilibrium strategy predicts the winning probability of the highest know-how rocket for an observed know-how state in each procurement, and we sum these probabilities across years and then divide the sum by the total number of procurements.

probability that a rocket wins a launch when a rival rocket has the highest know-how;

3. yearly average transaction prices;

4. the standard deviation of winning prices (across years).

The predicted choice probability moments are calculated directly from the solution to the dynamic model. However, the calculation of the price moments is more complicated because of the idiosyncratic cost shocks that happen in each procurement, and the large number of procurements that happen each year within an era. We therefore use 100 simulations to calculate the predicted price moments. When we solve the model, we allow for the possibility of innovation during era 3, although, as noted, the innovation cost parameters are fixed rather than estimated. As Vulcan was only successfully launched after the end of our data, we therefore calculate the model-predicted era 3 moments conditional on innovation not having occurred, even though its possibility affects equilibrium pricing strategies and sale probabilities.

We weight each moment by the inverse of the corresponding data moment's standard deviation. The objective function is minimized by searching across a wide parameter space using a combination of the genetic algorithm and surrogate optimization.

5.3 Identification

We only observe 267 U.S. government launches between 2002 and 2023, so our model is necessarily parametric. However, we can discuss the intuitions for the identification of the demand and cost parameters. The intuitions presume that the model has a unique equilibrium for all parameters.

The key sources of variation are changes in rocket know-how (number of past launches), the change in EELV ownership structure, and the entry of SPX. We are assuming the transition process for know-how, and that rockets start at the lowest level of know-how (i.e., initial conditions are observed).[49] If we observed every bid, as well as the buyer's choices, and we assumed that the buyer and the sellers made static payoff-maximizing choices, then

[49]Even if we assumed that we cannot measure know-how directly from data, we would still be able to assume that each new rocket family begins with the lowest level of know-how as we observe the first launches of the EELVs and Falcon 9.

standard arguments (for example, based on the inversion of the choice probabilities) would imply that these three types of variation would be sufficient to identify the cost intercepts and the price, know-how and intercept coefficients in the demand system, once we normalize the value of the outside good to zero. Variation in bids across procurements where the sellers have the same know-how would identify the scaling of the cost shocks.

We are assuming that the buyer and the sellers make forward-looking choices, but, given the assumed know-how transitions and the assumed discount factor (Chow (1994), Abbring (2010)), the buyer's and the sellers' values are simply the discounted sums of the components of static flow payoffs so that the same identification arguments apply.

A practical challenge is that we only observe, at best, transaction prices, and more precisely, given how we construct prices of a representative variant of each rocket family (see Appendix A), the average winning bid across all transactions for a given rocket in a given year, rather than all bids. Of course, it is not unusual for a researcher to observe only winning bids, or transaction prices, in auction-like settings, and cost and value distributions can still be identified under suitable assumptions about unobserved auction and bidder heterogeneity (e.g., Athey and Haile (2002), Cosconati et al. (2025)). Critically, we are assuming that there is no common heterogeneity, observed by the bidders, but not the researcher, on the demand-side or the supply-side which we need to account for.

5.4 Estimation Results

Table 3 lists the estimated parameters. The "baseline" column provides the estimates when we assume $\gamma^e = 0.234$, and the second column gives the estimates when $\gamma^e = 0$.

The baseline price and know-how coefficients imply that, for a typical 15 tonne LEO capacity launch, a doubling of know-how would increase the buyer's willingness to pay by $\left(\ln 2 \times \frac{1.550}{0.630} \times 15 \approx\right)$ \$25 million. The demand intercept for era 3 is lower than the intercepts for eras 1 and 2, and that ULA's rockets are preferred to SPX's. The higher rate of launches in era 3, and SPX's higher market share, are therefore explained by lower prices/costs and higher levels of reliability. The estimated σ_ζ implies that there is substantial heterogeneity in costs across procurements, so that a one standard deviation shock increases a bidder's costs by approximately 21%.

Table 3: Model Parameter Estimates

	Baseline with Cost LBD	Alternative with No Cost LBD
Demand Parameters		
Price ($1,000 per kg)	0.630***	0.657***
	(0.064)	(0.045)
ln(Know-How)	1.550***	1.546***
	(0.132)	(0.136)
SPX	-3.014***	-3.929***
	(0.467)	(0.349)
Era 2	0.036	0.596***
	(0.185)	(0.024)
Era 3	-1.165***	-0.343
	(0.079)	(0.376)
Constant	2.188***	2.321***
	(0.248)	(0.230)
Cost Parameters		
Cost LBD (γ^e)	0.234	0
	fixed	fixed
Intercepts		
Atlas/Delta	2.784***	2.320***
	(0.086)	(0.048)
New Generation (Falcon/Vulcan)	0.373***	0.227***
	(0.023)	(0.037)
Cost Shock Scale Parameter σ_ζ	0.210***	0.263***
	(0.075)	(0.054)

Note: *,**,*** denote significance at the 10%, 5% and 1% levels. Standard errors, calculated using a bootstrap with 200 repetitions, are in parentheses. The cost LBD coefficients, the discount factor, and the parameters that determine the evolution of know-how are assumed.

If we assume that there is no cost LBD, the reliability and price demand coefficients are almost unchanged, while the cost intercepts fall. The main changes, however, are to the era-specific and the SPX demand intercepts: without cost-side LBD, the estimated model explains rising sales later in the data by demand increases/smaller demand decreases.

Table 4 provides an assessment of model fit, based on both the targeted moments and some moments that are not targeted. For the purposes of this table, we simulate bids and buyer choices, from 2002 through 2023, 1,000 times following the method described in Appendix B, and report the averages of those simulations where there was no innovation in era 3. In particular, know-how endogenously and stochastically evolves, reflecting buyer choices and random shocks in each simulation. Therefore, differences between model and data

Table 4: Fit of Model Moments

	Data	Baseline with Cost LBD	Alternative with No Cost LBD
Targeted Moments			
Choice Probabilities			
Overall	0.389	0.406	0.380
× Era 2	0.186	0.195	0.185
× Era 3, ULA	0.080	0.044	0.062
× Era 3, SPX	0.107	0.136	0.116
Highest Know-How	0.339	0.344	0.300
Highest Know-How Rival	0.152	0.187	0.157
Price			
Era 1	8.244	9.706	8.689
Era 2	11.043	10.504	10.302
Era 3, ULA	10.349	9.437	9.520
Era 3, SPX	3.645	3.491	3.671
Standard Deviation	1.119	1.234	1.522
Untargeted Moments			
Prob. Choose Highest Know-How			
× Era 1	0.016	0.030	0.016
× Era 2	0.186	0.172	0.154
× Era 3	0.136	0.147	0.130
Average EELV Know-How			
Era 1	1.800	2.216	1.738
Era 2	9.700	8.917	8.643
Era 3	12.143	8.877	10.040

Note: The upper panel reports average outcomes, based on the moments targeted in estimation, for the two specifications. We report the average choice probability, the average of the choice indicator interacted with era indicators, the average of the choice indicator interacted with the indicator that the chosen rocket has the highest know-how at the procurement, and the average of the interaction with the indicator that the chosen rocket's rival has the highest know-how. The bottom panel considers outcomes that are not targeted in estimation. In the first set, we report the average of the choice indicator interacted with an era indicator and that the chosen rocket has the highest know-how in a procurement. The sum of these moments matches the probability of choosing the highest know-how rocket in the top panel. The second set reports the average level of ULA rocket know-how within each era when we simulate the model forward, and compares it to the know-how levels that we take from the data (i.e., the number of observed launches in the last two years).

moments can emerge because of how our assumed know-how accumulation and depreciation processes do not quite match the data, as well as a lack of perfect fit of the conditional moments matched in estimation.

The upper panel reports average outcomes, based on the moments targeted in estimation, for the two specifications. Both models match choice probabilities accurately, in each era. Prices are matched a little less well: we over-predict era 1 prices and under-predict ULA prices in eras 2 and 3. Overall, there is no clear reason to prefer one specification based on fit alone, so that our preference for the baseline specification comes from the received industry wisdom that cost-side LBD is important.

The bottom panel considers outcomes that are not targeted in estimation. For the first set, we report the average of the choice indicator interacted with an era indicator and that the chosen rocket has the highest know-how in a procurement. The sum of these moments matches the probability of choosing the highest know-how rocket in the top panel. The second set reports the average level of ULA rocket know-how within each era when we simulate the model forward, and compare it to the know-how levels that we take from the data (i.e., the number of observed launches in the last two years). We find that the know-how levels are matched well in eras 1 and 2, but we under-predict the know-how in era 3. This is partly due to high levels of launches towards the end of era 2, when there were, on average, more than twelve EELV rocket launches per year. The know-how levels of ULA rockets were thus at the highest level at the beginning of era 3 when we apply our two year approach to the data. If we instead simulate the era 3 using the know-how levels in 2017 from the data, the average ULA know-how would be 11.6 in the baseline specification, closer to the data moments.

Additional Specifications

Synergy Parameter. The specifications above assume that each ULA rocket family benefits from cross-family LBD after the JV. While ULA did move the production of both families to a single plant in Decatur, AL, it is possible that the spillover was incomplete. We have therefore experimented with specifications of the ULA transition function where a Delta IV rocket's know-how increases to the next level $e^{U}_{\text{Delta IV},t}$ with probability

$\rho \times \dfrac{1}{e^U_{\text{Delta IV},t} - e_{\text{Delta IV},t}}$ when Atlas V is chosen in a procurement, and vice versa. Our estimates assume $\rho = 1$. We find that the estimated parameters are not sensitive to different values of ρ, although, when we simulate the model forward as in Table 4 (i.e., not treating know-how as observed; instead, know-how evolves endogenously), era 2 choice probabilities and know-how levels are lower than we observe (for example, the probability of a launch falls by 32% if we assume $\rho = 0$). This suggests that, conditional on our other assumptions, allowing for a large JV learning synergy is probably necessary to match the data.

Nested Logit Demand. We have also experimented with allowing for nested logit demand, putting all of the available launch vehicles in a single nest, separate from the outside good, to allow for more flexible substitution patterns. Specifically, we replace the logit taste shock in (5), with a composite shock $\eta_t + (1 - \sigma^{\text{buyer}})\tilde{\varepsilon}_{jt}$ that allows for the buyer in procurement t to have a common preference, η_t, for all launch options. The correlation parameter σ^{buyer} is identified from variations in market structure (e.g., the JV and SPX entry). To facilitate estimation, we therefore set the era 2 demand parameter, which is small in our baseline estimates, equal to zero, and use the same moments as before, as they differentiate across market structure eras. We estimate a small and insignificant σ^{buyer} of 0.02, suggesting that, in this setting, our simple logit demand model is appropriate.

5.4.1 Elasticities

Given the forward-looking buyer behavior, there are likely to be differences between short-run and long-run demand elasticities. In the context of storable goods, Hendel and Nevo (2006) show that consumers respond much more to short-run price changes than to permanent price changes, because the former lead to intertemporal substitution. We do not allow intertemporal substitution, but differences can still arise because future price increases may reduce a buyer's incentives to invest, by making current purchases, to lower sellers' future costs.

The upper panel of Table 5 shows the short-run own-price elasticities of Atlas V (which is symmetric to Delta IV up to know-how states) in each state based on the era 2 demand intercept, but the results would be almost unchanged if we were to use the nearly identical

Table 5: Short-Run Elasticities with Respect to Atlas V Price Change

	Duopoly				Monopoly				
					Atlas V Own-Price Elasticity				
Atlas V \ Delta IV	1	5	10	15		1	5	10	15
1	-2.605	-2.686	-2.610	-2.568	1	-2.229	-2.691	-3.144	-3.479
5	-1.955	-1.971	-1.968	-1.960	5	-1.836	-1.976	-2.096	-2.206
10	-1.845	-1.873	-1.877	-1.874	10	-1.776	-1.841	-1.913	-1.973
15	-1.837	-1.876	-1.887	-1.885	15	-1.756	-1.790	-1.839	-1.885
					Cross-Price Elasticity				
Atlas V \ Delta IV	1	5	10	15		1	5	10	15
1	0.193	0.120	0.116	0.121	1	0.238	0.109	0.064	0.042
5	0.369	0.287	0.269	0.273	5	0.433	0.286	0.216	0.177
10	0.423	0.332	0.312	0.316	10	0.500	0.368	0.302	0.260
15	0.425	0.329	0.306	0.310	15	0.527	0.415	0.351	0.310

Note: Know-how of Atlas V is in rows, and know-how of Delta IV is in columns. Statistics are calculated for era 2, and are averaged across realized cost shocks.

Table 6: Long-Run Elasticities with Respect to Atlas V Price Change

	Duopoly				Monopoly				
Atlas V \ Delta IV	1	5	10	15		1	5	10	15
1	-12.867	-14.398	-13.762	-13.333	1	-3.280	-4.323	-5.119	-6.232
5	-4.001	-4.281	-4.230	-4.133	5	-1.761	-2.196	-2.648	-3.207
10	-2.517	-2.621	-2.566	-2.494	10	-1.451	-1.670	-1.984	-2.410
15	-1.964	-2.028	-1.953	-1.885	15	-1.279	-1.335	-1.534	-1.885

Note: Know-how of Atlas V in rows, know-how of Delta IV in columns. The long run elasticities are calculated as the relative decrease in Atlas V demand at each state when firm bids are reduced by 5% for each realized cost shock and the buyer re-solves the optimal forward-looking procurement problem. Statistics calculated for era 2, and are averages across realized cost shocks.

era 1 intercept. The elasticities are short-run in that they reflect changes in price that happen in the current period, with prices in future periods unchanged (even if the current state is visited again), so that buyer continuation values are held fixed. Table 6 reports the "long-run" own-price elasticities of Atlas V in each state, reflecting the change in demand in the state when Atlas V's prices are increased in both the current state, and all future states, by the same percentage amount. In this case, buyer continuation values will change. We report the elasticities both for the case of ULA monopoly, and for the case of ULA duopoly. The statistics are weighted averages across the idiosyncratic cost realizations in each state. The reader can find the associated equilibrium average prices and sale probabilities in Table 7.

The comparisons between short-run and long-run own-price elasticities are economically interesting and informative about how the model works. Under duopoly, the long-run elasticities are substantially larger (in absolute value) than the short-run elasticities, except in the highest know-how (15,15) state. The difference reflects how, when future Atlas V prices increase, a forward-looking buyer will want to substitute away from Atlas V as it knows it will make fewer future Atlas V purchases because of the price increases, and also the expectation that this will lead to higher Atlas V costs (and therefore even higher Atlas V prices), and potentially lower Delta IV costs. In the highest know-how state, the effects coming from future cost reductions will be more limited (as Atlas V is likely to make enough sales to maintain its low costs). The cost effects will also be muted under ULA monopoly because the assumed learning synergy means that Atlas V's future costs are just as likely to fall as those of Delta IV when Delta IV launches are sold. Therefore the differences between short-run and long-run elasticities are smaller under seller monopoly than under duopoly, although they are still non-trivial, when considered on a proportional basis, in several states.[50]

The second panel of Table 5 reports short-run cross-price elasticities (i.e., the change in Delta IV demand when the Atlas V price increases). The patterns are consistent with Delta IV and Atlas V being close substitutes when Delta IV has more know-how, although this

[50] It is noticeable that, under monopoly, the Atlas V long-run elasticity is smaller (absolute value) than the short-run elasticity in states where Atlas V has higher know-how. From Table 7, we can see that Delta IV has much higher prices in these states, and almost no probability of being sold. The logic behind the difference is therefore that the forward-looking buyer is less likely to substitute away from buying Atlas V, because it recognizes that its reduction in demand in future states will increase the likelihood that both rockets lose know-how, which is undesirable for the buyer.

also means that Delta IV's sale probability is higher, which tends to lower the value of the calculated cross-price elasticity. The table also illustrates how demand for the outside good plays a quantitatively important role in our model.

6 Counterfactuals

We now discuss our counterfactuals. We first consider the period 2002–2016, where we assume no innovation, and then the period 2017–2023, when SPX entered and ULA chose to develop Vulcan. We will primarily be interested in buyer and total surplus, but we will discuss changes in average transaction prices, the number of government launches, the average Lerner indices $\left(\frac{p-c}{p}\right)$, and average know-how to explain the predicted changes. All counterfactuals assume the baseline estimates (i.e., with LBD in costs as well as reliability).[51]

6.1 No Innovation Counterfactuals: 2002–2016

We begin by considering the effects of market structure, before considering the organization of supply (private versus government ownership) and the effects of forward-looking vs. static procurement rules.

6.1.1 Industry Structure and the ULA JV

We compare the simulated equilibrium outcomes under the observed, "factual", market structure (duopoly in 2002–2006, monopoly in 2007–2016, therefore "duopoly-then-monopoly") with two alternatives: (1) no JV, so that Boeing and Lockheed Martin remain separate ("always duopoly"); (2) ULA is formed at the start of 2002 ("always monopoly"). In all three cases, we allow for the estimated small change in the "era" demand intercept after 2006. In each case, we simulate the model forward 1,000 times.

To understand the comparisons we make below, it is useful to review the average equilibrium prices and buyer choice probabilities in each state during the two eras under each

[51] We have also conducted the counterfactuals using the parameters that presume no cost-side LBD. The qualitative results are similar, although magnitudes can change partly because of the differences in the intercept estimates.

Table 7: Average Bid Prices, Costs and Winning Probabilities: Era 2

	Duopoly				Monopoly			
				Average Bid Prices				
Atlas V \ Delta IV	1	5	10	15	1	5	10	15
1	14.674, 14.674	15.154, 11.266	15.523, 11.237	15.545, 11.514	12.769, 12.769	16.276, 11.044	18.121, 10.990	19.318, 11.514
5	11.266, 15.154	11.522, 11.522	11.523, 10.972	11.525, 11.235	11.044, 16.276	11.763, 11.763	13.061, 11.308	14.164, 11.515
10	11.237, 15.523	10.972, 11.523	10.963, 10.963	10.969, 11.235	10.990, 18.121	11.308, 13.061	11.844, 11.844	12.573, 11.791
15	11.514, 15.545	11.235, 11.525	11.235, 10.969	11.235, 11.235	11.514, 19.318	11.515, 14.164	11.791, 12.573	12.297, 12.297
				Average Costs				
Atlas V \ Delta IV	1	5	10	15	1	5	10	15
1	16.603, 16.603	16.603, 11.384	16.603, 9.676	16.603, 8.798	16.603, 16.603	16.603, 11.384	16.603, 9.676	16.603, 8.798
5	11.384, 16.603	11.384, 11.384	11.384, 9.676	11.384, 8.798	11.384, 16.603	11.384, 11.384	11.384, 9.676	11.384, 8.798
10	9.676, 16.603	9.676, 11.384	9.676, 9.676	9.676, 8.798	9.676, 16.603	9.676, 11.384	9.676, 9.676	9.676, 8.798
15	8.798, 16.603	8.798, 11.384	8.798, 9.676	8.798, 8.798	8.798, 16.603	8.798, 11.384	8.798, 9.676	8.798, 8.798
				Winning Probabilities				
Atlas V \ Delta IV	1	5	10	15	1	5	10	15
1	0.034, 0.034	0.021, 0.277	0.009, 0.345	0.010, 0.315	0.126, 0.126	0.010, 0.386	0.001, 0.444	0.000, 0.407
5	0.277, 0.021	0.229, 0.229	0.213, 0.327	0.229, 0.284	0.386, 0.010	0.219, 0.219	0.090, 0.356	0.037, 0.391
10	0.345, 0.009	0.327, 0.213	0.301, 0.301	0.321, 0.253	0.444, 0.001	0.356, 0.090	0.234, 0.234	0.143, 0.303
15	0.315, 0.010	0.284, 0.229	0.253, 0.321	0.271, 0.271	0.407, 0.000	0.391, 0.037	0.303, 0.143	0.212, 0.212

Note: Know-how of Atlas V in rows, know-how of Delta IV in columns. We report the average prices bid in each procurement, the average costs and buyer's choice probabilities for Atlas V and Delta IV rockets in each state. The demand intercept is based on the era 2 estimates, which is nearly identical to the era 1 estimates.

market structure. The era 2 values are shown in Table 7, and, because the difference in the demand intercepts is very small, the era 1 values are almost identical. The first point to notice is that, when we average across the idiosyncratic cost realizations, the probability of the buyer choosing one of the products is no more than 0.6 in any state (it is highest in state (10,10) under duopoly), so that the outside good is often chosen. As we will discuss below, this leads to some differences from the existing dynamic competition literature.

Second, a monopolist charges higher average prices and margins than the duopolist when both products are in the highest know-how state. This is consistent with a monopolist exercising market power, but, as a consequence, the monopolist has stronger incentives to accumulate know-how in lower know-how states, which explains why state (1,1) average prices are lower under monopoly than duopoly (both are below the average cost), and the probability that the monopolist makes a sale is several times higher than under duopoly.

Third, in some asymmetric states, such as (1,10), the monopoly price of the leading product is also lower than the duopoly price. This pattern is explained by the fact that, given the assumed learning synergy, the most profitable way for the monopolist to improve its laggard product is to make a sale of the leading one (for example, in (1,10) the laggard product is only sold with equilibrium probability 0.001). The effect of the synergy also leads to the rockets' know-how tending to be fairly symmetric under monopoly.

Panels (a) and (b) of Figure 3 show the changes in average prices and the number of launches each year under "always monopoly" and "always duopoly", relative to the averages of the factual simulations. These figures capture how the differences in strategies across states detailed in Table 7 translate into differences in outcomes over time, as states evolve. Panels (c) and (d) present the average levels of know-how (unweighted averages across the products) and Lerner indices under the three different scenarios. While the different era 1 market structures create large differences between the factual and "always monopoly" outcomes in 2007 and 2008, these differences are almost eliminated after 2010.

Know-how starts at the lowest level in 2002. A monopoly would have lowered prices so that know-how would have accumulated more quickly, with 4 more government launches per year in 2002 and 2003, an increase of more than 100%. Between 2007 and 2009, "always monopoly" is associated with higher prices than the factual "duopoly-then-monopoly", not

Figure 3: Effects of Alternative Market Structures on Expected Market Outcomes 2002–2016

(a) % Price Difference

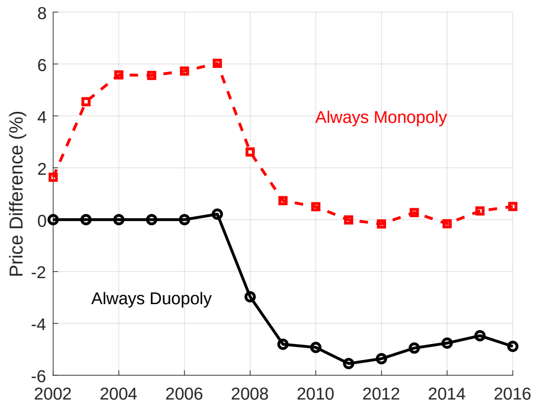

(b) Quantity Difference

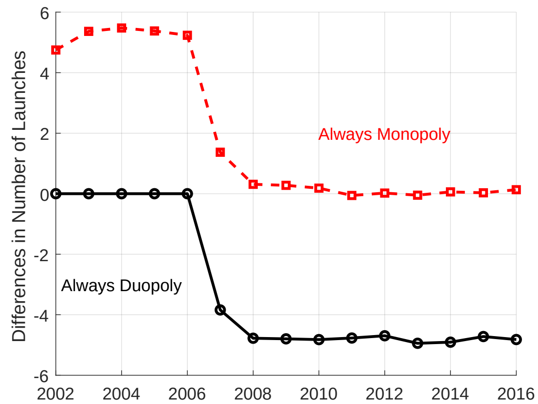

(c) Average Know-How

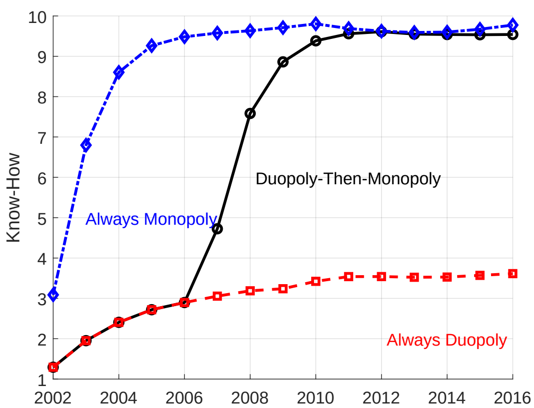

(d) Lerner Index

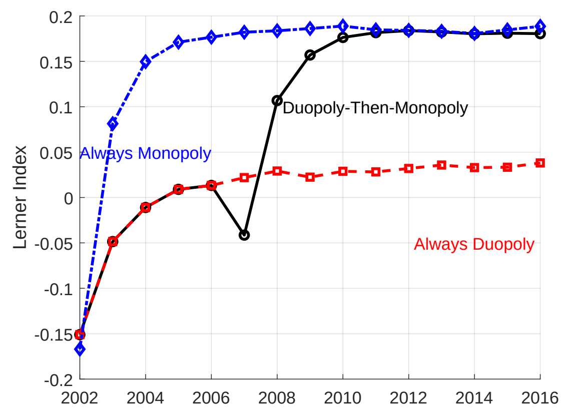

Note: "Always monopoly" assumes that ULA was formed as a multiproduct monopoly in 2002 (start of the EELV era). "Always duopoly" assumes that the ULA JV was never consummated. Panel (a) shows the % difference in average transaction prices for each year in 2002–2016 between simulations of "always monopoly" and "always duopoly" relative to the factual "duopoly-then-monopoly" market structure. Panel (b) shows changes in the average number of launches in each year. Panel (c) shows the evolution of the average know-how of the sellers in each market structure, and panel (d) shows the average evolution of the Lerner index. Statistics are calculated by averaging across 1,000 equilibrium paths in each case.

because of any differences in strategies (the same monopoly strategies apply in both cases), but when there is monopoly in earlier years, both products will tend to have accumulated higher know-how, increasing buyer valuations and seller margins.

We therefore predict that if the joint venture had been blocked, prices would have been lower. But, the buyer's valuation of launches would also have been lower, and the number of launches would have fallen.

We can compare welfare over different time frames. For the alternative market structures, Table 8 compares (i) average outcomes in 2016 (the last year of era 2) and (ii) average annual outcomes across all years in era 1 and era 2 (2002–2016). We also report the buyer and seller values (PDVs) calculated when solving our model averaged across states based on the mix of states reached in 2016 across our 1,000 simulation paths.[52] The total value is the sum of these buyer and seller values.

Comparing the "always monopoly" (column 2) and "always duopoly" (4) outcomes to the factual (1) outcomes, we see that continued duopoly would have lowered both buyer and total surplus (whether measured in terms of flow surpluses or PDVs). The product-level HHI is lower under monopoly than under duopoly, reflecting how ULA's learning synergy tends to keep Delta IV and Atlas V fairly symmetric even when Atlas V has more commercial launches. Under duopoly, commercial launches tend to lead Atlas V to have higher know-how and hence a greater probability of securing government sales.

6.1.2 Social Planner Comparison

Although monopoly market outcomes are preferred to duopoly outcomes in our setting, monopoly outcomes can still be inefficient. We therefore also simulate outcomes under the social planner solution with two products. The differences between the planner outcome and the "always monopoly with a long-lived buyer equilibrium" outcome come from how the social planner purchases at production cost, whereas the long-lived buyer faces prices set by a strategic monopolist. We assume that the social planner enjoys the same learning synergy as the ULA monopolist, and that, when we simulate outcomes, the pattern of commercial

[52] The know-how states reached in 2016 reflect commercial launches in prior years. However, because the buyer and seller values are taken from the model, they do not account for future commercial launches.

Table 8: Market Outcomes With No Innovation: Eras 1 and 2

	Factual	Always Monopoly		Always Duopoly		Planner		Down-select Atlas V		Always Monopoly No Synergy	
			Δ		Δ		Δ		Δ		Δ
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Price ($1,000/kg)											
2016	10.63	10.68	0.05	10.11	-0.52			10.18	-0.45	10.17	-0.46
2002-2016	10.47	10.57	0.11	10.05	-0.41			10.11	-0.35	10.13	-0.34
Number of Launches/Year											
2016	8.60	8.73	0.13	3.78	-4.82	16.76	8.16	2.65	-5.95	2.94	-5.65
2002-2016	6.52	8.42	1.90	3.39	-3.14	16.63	10.11	2.34	-4.18	2.72	-3.81
Know-How											
2016	9.54	9.77	0.23	3.43	-6.11	12.93	3.39	4.33	-5.21	2.80	-6.74
2002-2016	6.34	8.87	2.53	2.58	-3.76	12.46	6.12	3.34	-3.00	2.07	-4.27
Lerner Index											
2016	0.18	0.19	0.01	0.04	-0.14			0.04	-0.14	0.03	-0.15
2002-2016	0.13	0.16	0.03	0.02	-0.11			0.02	-0.10	0.02	-0.11
Product-Level HHI											
2016	6,771	6,759	-12	8,889	2,119	5,658	-1,113	10,000	3,229	9,320	2,549
2002-2016	5,381	5,218	-163	6,415	1,034	5,051	-330	10,000	4,619	6,825	1,444
Buyer Surplus ($bn/Year)											
2016	0.42	0.43	0.01	0.30	-0.12	0.97	0.54	0.29	-0.13	0.29	-0.13
2002-2016	0.36	0.40	0.04	0.29	-0.06	0.87	0.51	0.28	-0.07	0.29	-0.07
Variable Profits ($bn/Year)											
2016	0.24	0.25	0.01	0.02	-0.22			0.02	-0.22	0.02	-0.22
2002-2016	0.13	0.21	0.08	0.01	-0.12			0.01	-0.12	0.01	-0.12
Avg PDV (2016 States, $bn)											
Buyer	8.65	8.65	0.01	6.20	-2.44	19.16	10.52	5.89	-2.76	6.06	-2.59
Seller	5.24	5.25	0.01	0.66	-4.58			0.54	-4.70	0.65	-4.59
Total	13.89	13.90	0.02	6.86	-7.02	19.16	5.28	6.43	-7.46	6.71	-7.18

Note: The price panel reports the average winning prices. The quantity panel reports the average number of launches per year. The know-how panel reports the unweighted average know-how accumulated from launches, accounting for commercial launches and stochastic depreciation. The Lerner index is defined as the difference in winning price and production cost relative to the winning price, and we report the average value across successful sales. The product-level HHI is defined as the sum of the squares of each product's number of launches divided by the total number of launches, times 10,000. The buyer surplus is the buyer's annual flow surplus, which is the expected surplus given the firm bids in each procurement, summed to the year level (with no discounting). The variable profit is the seller expected flow surplus similarly aggregated across procurements to the year level. The reported PDVs are the average buyer, or seller, values from the model, averaged across states, with weights that reflect the states reached in 2016 in our simulations.

launches is unaffected.

Column (6) of Table 8 reports the average annual outcomes for our three different time horizons under the planner. Appendix Table D.1 lists costs and the planner's choice probabilities in each state. The planner purchases almost twice as many launches as the monopolist, so that both rocket families are likely to reach the highest know-how quickly. The resulting increases in buyer and total values are substantial (for example, the reported PDVs increase by over $10 billion and $5 billion respectively). In Section 6.1.4, we will evaluate how much production inefficiency a government would be willing to tolerate to achieve these gains from planner control.

6.1.3 A Down-Select Option and the Role of The Know-How Synergy

So far, we have assumed that both Atlas V and Delta IV are available to the government in each procurement during era 2. Kovacic (2019) notes that, when the ULA joint venture was being evaluated, the FTC believed that, if the JV was not consummated, the procuring government agencies might commit to buying from only one of the companies in order to speed up the chosen system's accumulation of know-how, a policy known as a permanent "down-select" (where the down-selected seller is the one that is chosen).

Column (8) of Table 8 shows outcomes if Atlas V, which has more commercial sales, is down-selected.[53] Compared to continued duopoly, the down-select increases the know-how of purchased launches, but slightly reduces buyer and total surplus, reflecting a loss of competition and variety. However, the factual (two-product monopoly after 2007) outcomes are significantly better than the down-select outcomes, so we continue to interpret the FTC's decision not to challenge the JV as the correct one.[54]

One reason why monopoly is preferred to duopoly, or the down-select, is that the learning synergy, and the variety offered by two systems, combine to speed the accumulation of know-

[53]We assume the commercial launches on Atlas V are the same as observed in the data, i.e., it does not take additional commercial launches from Delta IV.

[54]The reductions in buyer surplus under the down-select partly reflect the standard "loss-of-variety" effect when a product is removed in a demand system with logit preference shocks (Ackerberg and Rysman (2005)). In this setting, maintaining alternative rocket families may be especially important as, if a system is rendered inoperative after a failed launch, the government might be left without the ability to undertake critical missions. Our surplus calculations do not account for the fixed costs of maintaining a rocket family, which are likely to be significant in practice.

how. Column (10) of Table 8 shows average simulated outcomes under the two-product monopoly when there is no synergy. Relative to monopoly with the learning synergy, know-how and the average number of launches are lower, prices are higher, and the product-level market structure tends to be much more asymmetric. Buyer and total surpluses are slightly higher than under the down-select, but they are now lower than under duopoly. Therefore, our results also support the FTC's stated view that the existence of the learning synergy was critical to its enforcement decision. It is noticeable that the no synergy monopoly and duopoly welfare outcomes are quite similar, showing that, in this context, and absent the synergy, the market power of the monopolist would have been roughly offset, from the buyer's perspective, by the greater incentive of the monopolist to lower prices to reduce its future costs, a type of efficiency that a dynamic model of competition can quantify.[55]

6.1.4 The Tradeoffs of State Control: Planner Solution with Cost Inefficiency

As noted in Section 2, governments own or more tightly control launch providers in other countries. This was also true in the U.S. historically. One might assume that the benefit of state control is that planner outcomes would be implemented. On the other hand, as noted by Weinzierl (2018), state control could lead to various types of inefficiencies. We therefore evaluate the level of cost inefficiency that would offset the gains from the social planner's outcome.

Specifically, we assume that, under state control, the buyer pays the costs of the launches as in the planner scenario, but, in each state, each system's realized production cost would be multiplied by a scalar $\mu \geq 1$.[56] We resolve the planner problem for μ values between 1 and 1.60 (1.60 corresponding to 60% cost inefficiency). We assume that, given know-how, rocket reliability is unchanged and that the planner still enjoys the learning synergy from

[55] The discussion above based on Table 8 compares the values of different market structures in 2016. Of separate interest is how a merger or down-select affects welfare at the time the ULA merger was proposed. We therefore compute the 2007 buyer PDV under the alternative market structures using the states reached in the factual simulations at the end of era 1, taking into account the synergy from the merger. We find that the present discounted buyer value under ULA monopoly with the learning synergy is still the highest, \$2.48 billion higher than under a duopoly market, and \$2.80 billion higher than when the buyer down-selects Lockheed Martin's Atlas V.

[56] One might view inefficiencies as resulting from the use of less suitable technologies rather than inefficient management of the production process for a given rocket. As an extreme example, NASA's Space Shuttle cost about \$450 million per launch (Adler (2023)), almost 4 times the price of an Atlas V launch for comparable lift capacity.

Figure 4: Differences Between Planner and Factual Outcomes: Production Cost Inefficiency

(a) Buyer Value Difference (b) Quantity Difference

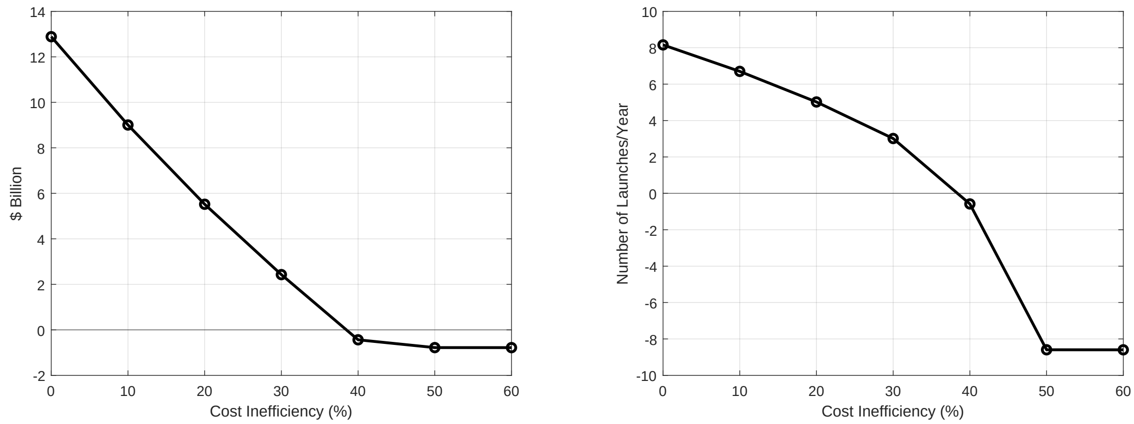

Note: Figure (a) plots the planner's 2016 PDV minus the buyer value in the market equilibrium at different levels of cost inefficiencies. The right figure (b) plots the average number of planner launches per year minus the average number in the market equilibrium. Our analysis uses the era 2 demand intercept.

managing both families.[57] We compare these planner-with-inefficiency outcomes to those under the monopoly equilibrium outcome using the era 2 coefficients.

Panel (a) of Figure 4 shows the difference in the buyer's 2016 PDV between the planner and market outcomes, while panel (b) shows the difference in the average number of launches. For cost inefficiencies of 30% or less (based on interpolation, the threshold would be just over 38%), the buyer value is higher under the planner. However, cost inefficiencies of 50% or more would lead to the planner choosing to undertake almost no launches, and almost all value being eliminated.

6.1.5 Static vs. Dynamic Procurement Strategy

Our model assumes that, in each procurement, the government makes a forward-looking decision about which option to choose, recognizing how its choice will affect the evolution of suppliers' state variables and, therefore, future buyer surplus. As highlighted above, it is not clear, as a matter of economic theory, that it is optimal for the government to use a

[57]We note that Europe's Ariane 5 and Russia's Soyuz-2 rockets are similar to ULA's rockets in terms of reliability.

Figure 5: Difference Between Static Buyer and Dynamic Buyer

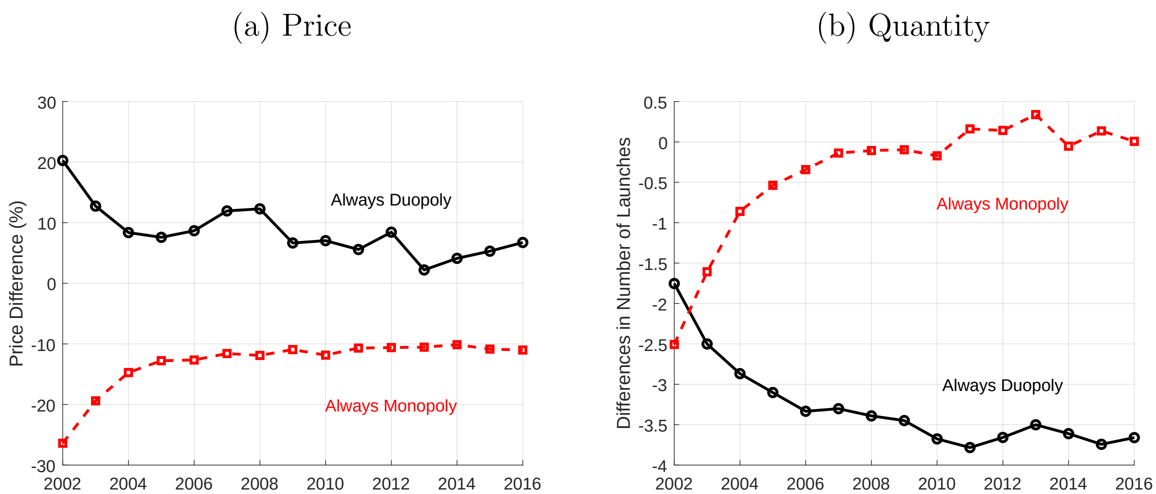

Note: The red line in panel (a) shows the % difference in average transaction prices for each year in 2002–2016 between simulations of "always monopoly" under a dynamic and static buyer (positive numbers mean prices are higher with static procurement). The red line in panel (b) shows changes in the average number of launches in each year. The black lines compare a static and a dynamic buyer under the "always duopoly" market structure. We simulate and average 1,000 equilibrium paths in each case.

forward-looking rule, even when it cares about future surplus. In reality, the government's approach to creating future competition in the space launch industry has been subject to litigation, as we described in Section 2.2.1.

Figure 5 shows the differences, over time, in prices and quantities when the buyer is committed to using a static procurement rule (i.e., makes choices based only on its current flow payoffs) under the "always monopoly" and "always duopoly" market structures. Table 9 reports the average market outcomes. The detailed buyer choice probabilities, bids, costs and winning probabilities at each state are given in Appendix Table D.2.

Under "always duopoly", static procurement causes prices, and Lerner indices, to rise, especially initially, while the number of launches falls. By 2016, both products are likely to be in the lowest know-how state, and both buyer and seller surpluses are reduced. The result that a forward-looking procurement policy is better is different from the findings emphasized by Lewis and Yildirim (2002) and Sweeting et al. (2022). Those studies, also considering dynamic duopoly, show that a static policy can benefit the buyer by promoting

Table 9: Market Outcomes With No Innovation: Dynamic vs. Static Procurement

	Always Duopoly			Always Monopoly		
	Dynamic Buyer	Static Buyer	Δ	Dynamic Buyer	Static Buyer	Δ
	(1)	(2)	(3)	(4)	(5)	(6)
Price ($1,000/kg)						
2016	10.11	10.79	0.68	10.68	9.51	-1.18
2002-2016	10.05	10.84	0.79	10.57	9.28	-1.29
Number of Launches/Year						
2016	3.78	0.12	-3.66	8.73	8.74	0.01
2002-2016	3.39	0.10	-3.29	8.42	8.05	-0.38
Know-How						
2016	3.43	0.17	-3.26	9.77	9.62	-0.16
2002-2016	2.58	0.10	-2.47	8.87	8.36	-0.51
Lerner Index						
2016	0.04	0.08	0.04	0.19	0.11	-0.08
2002-2016	0.02	0.08	0.06	0.16	0.06	-0.10
Product-Level HHI						
2016	8,889	9,925	1,036	6,759	6,782	23
2002-2016	6,415	9,064	2,650	5,218	5,263	45
Buyer Surplus ($bn/Year)						
2016	0.30	0.27	-0.03	0.43	0.57	0.14
2002-2016	0.29	0.27	-0.02	0.40	0.54	0.14
Variable Profits ($bn/Year)						
2016	0.02	0.00	-0.02	0.25	0.12	-0.13
2002-2016	0.01	0.00	-0.01	0.21	0.06	-0.14
Avg PDV (2016 States, $bn)						
Buyer	6.20	5.44	-0.76	8.65	11.40	2.75
Seller	0.66	0.02	-0.64	5.25	2.75	-2.50
Total	6.86	5.45	-1.41	13.90	14.15	0.25

Note: The price panel reports the average winning prices. The quantity panel reports the average number of launches per year. The know-how panel reports the unweighted average know-how that takes into account commercial launches and stochastic depreciation. The Lerner index is defined as the difference in winning price and production cost relative to the winning price. The product-level HHI is defined as the sum of the squares of each product's number of launches divided by the total number of launches, times 10,000. The buyer surplus is the buyer's annual flow surplus, which is the expected surplus given the potential bids from firms in each procurement aggregated to a year. The variable profit is the firm expected flow surplus similarly aggregated from each procurement to a year. The seller PDV is the average present discounted value from the equilibrium model solution averaged across the states reached in 2016 across simulation paths. The PDV of a static buyer is the present value of its future flow surpluses based on equilibrium strategies and state transitions.

supplier competition to gain a cost advantage. The main difference to the current setting is that those studies also assume that the buyer always, or almost always, makes some purchase, so that the significant differences between the static and dynamic buyer outcomes are in the degree of seller know-how asymmetry along the equilibrium path, not the level of accumulated know-how. In our setting, and especially under duopoly, the buyer usually chooses not to purchase, which can result in little know-how accumulation by either seller. A forward-looking buyer recognizes that not purchasing a launch tends to increase future costs and lower future reliability, so that a forward-looking procurement leads to more efficient outcomes with higher quantities sold.

On the other hand, under "always monopoly", which is not a case considered by the earlier studies, the buyer does better when it is committed to a static policy because the monopolist, with stronger incentives to accumulate know-how, and a better learning technology due to the synergy, lowers its prices, especially in low know-how states, to offset the static buyer's reduced propensity to purchase (i.e., the commitment makes the monopolist compete harder against the outside good). Relative to the dynamic buyer case, the loss in seller surplus is offset by the gains in buyer surplus, resulting in a net increase in total surplus.[58]

6.2 Era of Innovation: 2017–2023

We now study the role of SPX entry and ULA innovation over the period of 2017-2023. We explore (i) how SPX's entry affected outcomes; (ii) how SPX's entry affected ULA's innovation incentives; (iii) whether ULA's innovation incentives are efficient. We assume that ULA was formed in 2006, as observed in the data, and that ULA's EELVs benefit from the learning synergy.

6.2.1 Entry and Innovation

Column (1) of Table 10 shows the simulated era 3 market outcomes under the factual multi-product ULA and SPX duopoly. We start the simulations at the highest know-how levels for all rockets, given the high number of observed EELV launches at the end of era 2, and SPX's

[58] Given that we do not consider sellers' fixed costs, we thus ignore that sellers may exit when the buyer adopts a static procurement strategy.

Table 10: Market Outcomes With Innovation: Era 3

	Duopoly EELVs, Falcon	Duopoly Vulcan, Falcon	Δ	Monopoly EELVs	Δ	Monopoly Vulcan	Δ	Duopoly EELVs, Falcon*	Δ
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Price ($1,000/kg)									
ULA, 2017	9.38	2.83	-6.55	9.69	0.30	3.52	-5.86	9.30	-0.09
2017-2023	9.37	4.48	-4.89	9.52	0.15	5.38	-3.99	9.27	-0.10
SPX, 2017	3.38	2.49	-0.89					1.25	-2.13
2017-2023	3.41	1.74	-1.67					2.46	-0.96
Number of Launches/Year									
ULA, 2017	4.17	13.49	9.33	6.13	1.97	14.74	10.58	5.04	0.87
2017-2023	3.58	14.05	10.47	4.91	1.33	14.94	11.36	2.21	-1.36
SPX, 2017	8.20	3.91	-4.29					5.80	-2.41
2017-2023	8.38	3.45	-4.93					8.08	-0.29
Know-How									
ULA, 2017-2023	11.60	10.53	-1.07	8.22	-3.38	10.72	-0.88	6.27	-5.33
SPX, 2017-2023	15.00	7.31	-7.69					7.93	-7.07
Lerner Index									
ULA, 2017	0.23	0.58	0.34	0.22	-0.01	0.66	0.43	0.20	-0.03
2017-2023	0.21	0.79	0.59	0.14	-0.06	0.83	0.62	0.16	-0.05
SPX, 2017	0.77	0.66	-0.11					-0.09	-0.87
2017-2023	0.77	0.41	-0.36					0.58	-0.19
Firm-Level HHI									
2017	5,909	6,725	816	10,000	4,091	10,000	4,091	5,689	-220
2017-2023	6,437	7,031	594	10,000	3,563	10,000	3,563	7,756	1,320
2017 Surplus ($bn/Year)									
Buyer	0.71	0.96	0.25	0.42	-0.30	0.74	0.02	0.44	-0.27
ULA	0.12	0.33	0.20	0.18	0.05	0.50	0.38	0.12	0.00
SPX	0.31	0.10	-0.21					-0.00	-0.31
2017 PDV ($bn)									
Buyer	20.66	23.47	2.82	12.84	-7.82	17.22	-3.43	17.98	-2.68
ULA	0.27†	15.80	15.53	15.97†	15.70	21.39	21.12	0.24†	-0.02
SPX	3.42	2.16	-1.26					1.50	-1.92
Total	24.34	41.43	17.09	28.81	4.46	38.62	14.27	19.72	-4.62

Note: We report the average winning prices, number of launches per year, average rocket know-how, Lerner indices, firm-level HHI, average flow surpluses of buyers and firms, based on 1,000 simulations. The year 2017 is the first year of era 3. The 2017 present discounted values of the buyer and firms are taken from the model solution, and therefore ignore commercial launches. We start at the maximum know-how level of 15 for all rockets as observed in the data. We include commercial launches in the know-how at the end of each year. We start at the know-how level of 1 for Vulcan and assume that it has no commercial launches. We fix Falcon 9 know-how at 15, except in column (8), indicated by *, where initial Falcon 9 know-how is set to the lowest level (1) and we allow its know-how to evolve endogenously and do not include its commercial launches.

†: We report the results where ULA has the option to innovate but does not innovate in era 3. The PDVs take into account the potential benefits and costs of the innovation.

commercial launches. We calculate strategies and the PDVs allowing for possible ULA innovation, based on the innovation opportunity arrival rate and innovation cost assumptions explained in 4, but, when running the simulations, we choose draws such that ULA does not replace the EELVs with Vulcan before the end of 2023 in order to match what we actually observe in the data.

SPX's entry creates significant gains to the buyer, and reductions in ULA's surplus and the know-how of ULA's rockets. Although, everything else equal, the government prefers ULA launches, Falcon 9 wins two-thirds of government launches even in 2017 because of its lower costs and prices.

Column (2) summarizes outcomes if ULA's innovation occurs in 2017, and the EELV rockets (Atlas V and Delta IV) are replaced by Vulcan at the start of era 3 (although, of course, this ignores the time required to develop and test Vulcan). Recall that Vulcan has the same cost intercept as Falcon 9, and that the government prefers Vulcan to Falcon 9 when they have equal know-how, but that Vulcan begins with the lowest level of know-how, while Falcon 9 has, by assumption due to its many commercial launches, the highest. In PDV terms, the government is $2.8 billion better off when its options are Vulcan and Falcon 9, compared to Atlas V, Delta IV and Falcon 9, with ULA capturing the vast majority of the producer surplus.[59]

Column (6) shows outcomes when Vulcan is the only option. Comparing columns (4) (EELVs only) and (6) shows that, because Vulcan has much lower costs and a monopolist ULA has a strong incentive to accumulate know-how for Vulcan, buyer welfare is greater when only Vulcan is available than when ULA has an Atlas V and Delta IV monopoly, despite the reduction in variety.[60] ULA's expected flow profits are also higher with Vulcan. Therefore, from a welfare perspective, ignoring the large sunk costs of either innovation, both the replacement of EELV rockets by Vulcan and the entry of SPX are desirable.

Table 11 reports the probabilities that ULA introduces Vulcan, based on our assumptions on the arrival rate of opportunities to innovate and the costs of innovation. The first row shows the probability of innovation in 2017, assuming that innovation results in the immedi-

[59] Our estimate of SPX's surplus and value does not include Starlink.

[60] One can reach this conclusion by comparing either the expected 2017 buyer surpluses (0.74 vs. 0.42) or the buyer PDVs.

Table 11: Predicted ULA Innovation Probabilities

	Factual	Monopoly		Duopoly		Planner	
	EELV, Falcon	EELV	Δ	EELV, Falcon*	Δ	EELV	Δ
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
2017	0.14	0.06	-0.08	0.18	0.04	0.10	-0.05
2017-2023	0.15	0.07	-0.08	0.17	0.02	0.10	-0.05

Note: We report the average annual innovation probability in 2017, and over the years 2017 to 2023, conditional on no ULA innovation in prior years, based on 1,000 simulations. For EELV (Atlas V and Delta IV) and Falcon rockets, we start at the maximum know-how level of 15 as observed in the data. We fix Falcon 9 know-how at 15, except in column (4), indicated by *, where initial Falcon 9 know-how is set to the lowest level (1) and we allow its know-how to evolve endogenously and do not include its commercial launches.

ate introduction of Vulcan. The second row shows the average annual probabilities over the years 2017-2023 conditional on innovation not occurring in prior years. Comparing columns (1) and (2), we see that ULA is more likely to introduce Vulcan if it is competing with SPX. This is consistent with the fact that, when SPX is in the market, ULA's EELV rockets struggle to compete so that there are almost no existing profits for Vulcan to cannibalize, whereas the EELVs are quite profitable without SPX.[61]

One might wonder if these results reflect Falcon 9's high level of know-how due to its commercial launches. Therefore, the final columns of Table 10 and column (4) of Table 11 report outcomes and innovation probabilities if we assume counterfactually that SPX only provides government launches. Falcon 9's know-how starts at the lowest level and evolves as it competes with ULA's EELV rockets for government sales. The probability of ULA innovation actually rises, as now that SPX can benefit from LBD, Falcon is priced more aggressively, which lowers ULA's EELV profits.

6.2.2 Planner and Innovation

We can compare market innovation choices with those which might arise under planner control. In this context, we think a sensible comparison is with a state-controlled ULA, where SPX is not used for government missions. This would reflect what happens in, for example, Russia, where the government does not use companies in the private sector for its

[61] We also explore the effects of a static procurement strategy in Appendix C.1. We find that static procurement generally lowers ULA values, and slightly lowers the probability of ULA innovation in a ULA-SPX duopoly market.

Table 12: Market Outcomes with Innovation in Era 3: Planner Solution

	Factual Market	Planner		Planner	
	EELVs, Falcon	EELVs	Δ	Vulcan	Δ
	(1)	(2)	(3)	(4)	(5)
Number of Launches/Year					
ULA, 2017	4.17	13.00	8.83	19.83	15.66
2017-2023	3.58	12.93	9.36	19.83	16.25
SPX, 2017	8.20				
2017-2023	8.38				
Know-how					
ULA, 2017-2023	11.60	12.38	0.78	12.54	0.95
SPX, 2017-2023	15.00				
2017 Surplus ($bn/Year)					
Buyer	0.71	0.66	-0.05	1.38	0.66
ULA	0.12				
SPX	0.31				
2017 PDV ($bn)					
Buyer	20.66	34.38	13.72	44.78	24.13
ULA	0.27				
SPX	3.42				
Total	24.34	34.38	10.04	44.78	20.44

Note: Column (1) corresponds with column (1) in Table 10. In columns (2)–(5), we report the average number of launches per year, average rocket know-how, average buyer flow surpluses, and 2017 present discounted values of the buyer in the planner outcomes. We start at the maximum know-how level of 15 as observed in the data. We start Vulcan know-how at 1.

launches.

Table 12 compares outcomes when the ULA EELV rockets and Falcon 9 compete for government launches in a market (column (1)), when the planner uses only the EELV rockets (column (2)) and the planner only uses Vulcan (column (4)). The planner would launch the EELV rockets nearly 13 times a year, slightly higher than the EELVs and Falcon 9 combined in the market outcome, but the flow surplus is slightly lower without SPX, because of the loss of variety. If the planner introduces Vulcan in 2017, the number of launches is even higher, although the buyer has lower 2017 (i.e., initial) surplus in the planner solution, because the new rocket starts with higher costs than Falcon 9, and lower know-how than both EELVs and Falcon rockets in the factual market. The planner's PDV when it operates the EELVs and can innovate is \$13.7 billion greater than the buyer's in the ULA-SPX duopoly market, and \$24 billion greater when it operates Vulcan.

Column (6) of Table 11 shows the probabilities of planner innovation. The probabilities are higher than the market ULA monopolist probabilities in column (2), reflecting the fact that the planner internalizes some of the benefits of Vulcan which ULA does not capture. However, the planner is less likely to innovate than ULA would be if it were to face competition from SPX. The differences in the probabilities are consistent with ULA and planner values in Table 10 and Table 12. When competing with SPX using the older EELV rockets, ULA's gain from innovation (\$15.8-\$0.27 billion in Table 10) is greater than the planner's gain from operating Vulcan rather than the EELVs (\$44.8-\$34.4 billion).[62]

Appendix C.2 examines how much cost inefficiency would be needed to offset the gains from the planner solution when innovation is possible. We find that the required inefficiencies have to be much larger than the production cost inefficiencies identified in our era 1 and era 2 analysis.

[62] One interpretation of ULA being more likely to innovate than the planner is that it reflects ULA taking business from SPX, i.e., a classic excess entry (Mankiw and Whinston (1986)) result. However, SPX's PDVs in Table 10 are small even when it is competing against the EELVs. This suggests that the low ULA PDV in the market outcome is also driven by the EELV rockets losing their know-how over time, whereas, because the planner procures more launches, they tend to keep their high reliability and low costs.

7 Conclusion

We have estimated a dynamic model of competition and procurement for the space launch industry in order to try to provide model-based answers to several important economic and policy questions arising in this strategically important industry.

We find that, in the presence of dynamic learning economies, the FTC's decision not to challenge the United Launch Alliance joint venture was likely the correct one, even though it was, in essence, a merger-to-monopoly. Second, we show that the certification of SpaceX for national security launches and the subsequent development of Vulcan are likely beneficial, and that competition from SpaceX was likely important in ULA deciding to develop the Vulcan system. In terms of procurement design, we also find that it may be optimal for the government to follow a static procurement rule when facing a monopolist supplier, but that a forward-looking rule can be important in advancing supplier know-how when there is competition. We also quantify the trade-offs that may exist between the benefits of government control of launch suppliers (marginal cost pricing and the internalization of the benefits that innovation creates for the buyer) and the types of cost inefficiency that government control is likely to generate (Weinzierl (2018)). The fact that we find that government control can be advantageous unless inefficiencies are quite large may help to explain why both models of control persist in the real-world.

Our results highlight that in industries where learning-by-doing is important, government procurements not only add demand, but can also act as industrial policy. Each award is an investment that could reduce procurement costs or improve product quality in the future. Although we focus on competition between firms whose presence in the market we treat as exogenous, this lesson could apply more broadly to thinking about entry. For example, recognizing the catalytic nature of government procurement, NASA financed a significant portion of the development cost of Falcon 9 through contracts to supply the International Space Station, even though NASA had the option to use reliable and relatively cheap Russian Soyuz rockets. Such examples suggest that forward-looking procurement policies may be especially valuable when there is scope to shape both future market structure and the technological trajectory of an industry.

References

Abbring, J. H. (2010): "Identification of Dynamic Discrete Choice Models," *Annual Review of Economics*, 2(1), 367–394.

Ackerberg, D. A., and M. Rysman (2005): "Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects," *RAND Journal of Economics*, 36(4), 771–788.

Adler, D. (2023): "Why Did NASA Retire the Space Shuttle?," *Astronomy Magazine*. `https://www.astronomy.com/space-exploration/why-did-nasa-retire-the-space-shuttle/` (accessed January 8, 2026), Last updated May 18, 2023.

Albon, C. (2024): "United Launch Alliance's Vulcan Flies Second Certification Mission," *Defense News*. `https://www.defensenews.com/space/2024/10/04/united-launch-alliances-vulcan-flies-second-certification-mission/` (accessed January 8, 2026), Last updated October 04, 2024.

An, Y., and W. Zhao (2019): "Dynamic Efficiencies of the 1997 Boeing-McDonnell Douglas Merger," *RAND Journal of Economics*, 50(3), 666–694.

Athey, S., and P. A. Haile (2002): "Identification of Standard Auction Models," *Econometrica*, 70(6), 2107–2140.

Baldwin, R., and P. Krugman (1988): "Industrial Policy and International Competition in Wide-Bodied Jet Aircraft," in *Trade Policy Issues and Empirical Analysis*, pp. 45–78. University of Chicago Press.

Barwick, P. J., H.-S. Kwon, S. Li, and N. B. Zahur (2025): "Drive Down The Cost: Learning-by-Doing and Government Policies in the Global EV Battery Industry," Discussion paper, National Bureau of Economic Research.

Bedell, D. (2011): "NASA Launch Services Program Role in Mission Assurance," *National Aeronautics and Space Administration, Launch Services Program*. `https://ntrs.nasa.gov/citations/20110014253` (accessed January 8, 2026), Last updated May 3, 2011.

Benkard, C. L. (2000): "Learning and Forgetting: The Dynamics of Aircraft Production," *American Economic Review*, 90(4), 1034–1054.

——— (2004): "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft," *Review of Economic Studies*, 71(3), 581–611.

Besanko, D., U. Doraszelski, and Y. Kryukov (2014): "The Economics of Predation: What Drives Pricing When There is Learning-by-Doing?," *American Economic Review*, 104(3), 868–97.

——— (2019): "How Efficient is Dynamic Competition? The Case of Price as Investment," *American Economic Review*, 109(9), 3339–64.

Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010): "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics," *Econometrica*, 78(2), 453–508.

Bichler, M., M. Fichtl, and M. Oberlechner (2025): "Computing Bayes–Nash Equilibrium Strategies in Auction Games Via Simultaneous Online Dual Averaging," *Operations Research*, 73(2), 1102–1127.

Bogaisky, J. (2025): "With New Glenn Launch, Bezos Looks To Break Musk's Stranglehold On Space," *Forbes*. https://www.forbes.com/sites/jeremybogaisky/2025/01/11/new-glenn-bezos-blue-origin-musk-spacex/ (accessed January 12, 2025), Last updated January 11, 2025.

Cabral, L. M., and M. H. Riordan (1994): "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica*, 62(5), 1115–1140.

Cao, X., D. Cumming, and S. Zhou (2020): "State Ownership and Corporate Innovative Efficiency," *Emerging Markets Review*, 44, 100699.

Carlier, N. (2016): "Spacecraft Recovery Operations Conducted to the Galileo FOC-1," in *14th International Conference on Space Operations*, p. 2561.

Chen, Y., M. Igami, M. Sawada, and M. Xiao (2021): "Privatization and Productivity in China," *RAND Journal of Economics*, 52(4), 884–916.

Chow, G. C. (1994): "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics, Volume 4*, ed. by J. Rust, vol. 4 of *Handbook of Econometrics*, pp. 3081–3143. Elsevier, Amsterdam and New York.

Congressional Research Service (2025): "Defense Primer: National Security Space Launch Program," https://crsreports.congress.gov/product/details?prodcode=IF12900 (accessed January 8, 2026), Last updated April 28, 2025.

Cosconati, M., Y. Xin, F. Wu, and Y. Jin (2025): "Competing Under Information Heterogeneity: Evidence From Auto Insurance," Discussion paper, California Institute of Technology.

Deng, S., D. Jia, M. Leccese, and A. Sweeting (2025): "Dynamic Competition with Bargaining: Comparative Statics and Implications for Subsidy and Competition Policies," Discussion paper, University of Maryland.

Doraszelski, U., and M. Satterthwaite (2010): "Computable Markov-Perfect Industry Dynamics," *RAND Journal of Economics*, 41(2), 215–243.

Elliott, J. T., G. V. Houngbonon, M. Ivaldi, and P. T. Scott (2025): "Market Structure, Investment, and Technical Efficiencies in Mobile Telecommunications," *Journal of Political Economy*, 133(5), 1401–1459.

Erdem, T., S. Imai, and M. P. Keane (2003): "Brand and Quantity Choice Dynamics under Price Uncertainty," *Quantitative Marketing and Economics*, 1(1), 5–64.

Erwin, S. (2019): "In Protest Decision, GAO Negates Blue Origin's Claim that Air Force Launch Procurement Favors Incumbents," *SpaceNews*. https://spacenews.com/in-protest-decision-gao-negates-blue-origins-claim-that-air-force-launch-procurement-favors-incumbents/ (accessed January 8, 2026), Last updated November 23, 2019.

——— (2025): "Space Force Reassigns GPS Satellite Launch from ULA to SpaceX," *SpaceNews*. https://spacenews.com/space-force-reassigns-gps-satellite-launch-from-ula-to-spacex/ (accessed January 8, 2026), Last updated April 7, 2025.

FAA (2002): "Liability Risk-Sharing Regime for U.S. Commercial Space Transportation: Study and Analysis," *U.S. Department of Transportation, Federal Aviation Administration*. https://rosap.ntl.bts.gov/view/dot/15754 (accessed January 8, 2026), Last updated April 01, 2002.

Fan, Y. (2013): "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market," *American Economic Review*, 103(5), 1598–1628.

Fan, Y., and C. Yang (2025): "Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety," *Journal of Political Economy*, 133(6), 1886–1931.

Federico, G., F. Scott Morton, and C. Shapiro (2020): "Antitrust and Innovation: Welcoming and Protecting Disruption," *Innovation Policy and the Economy*, 20(1), 125–190.

Fudenberg, D., and J. Tirole (1983): "Learning-by-Doing and Market Performance," *Bell Journal of Economics*, 14(2), 522–530.

Gilbert, R. J. (2022): *Innovation Matters: Competition Policy for the High-Technology Economy*. MIT Press.

Goettler, R. L., and B. R. Gordon (2011): "Does AMD Spur Intel to Innovate More?," *Journal of Political Economy*, 119(6), 1141–1200.

Groesbeck, T. (2020): "Boosting the Competition: Reliability, Learning, and Trade Barriers in the Space Launch Industry," Working paper.

Gruss, M. (2014): "U.S. Air Force Claims Big Savings on EELV Block Buy," *SpaceNews*. https://spacenews.com/39348us-air-force-claims-big-savings-on-eelv-block-buy/ (accessed January 8, 2026), Last updated January 31, 2014.

——— (2016): "ULA's Parent Companies Still Support Vulcan ... With Caution," *SpaceNews*. https://spacenews.com/ulas-parent-companies-still-support-vulcan-with-caution/ (accessed January 8, 2026), Last updated March 10, 2016.

GUTHRIE, J. (2017): *How to Make a Spaceship: A Band of Renegades, an Epic Race, and the Birth of Private Spaceflight*. Penguin.

HENDEL, I., AND A. NEVO (2006): "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, 74(6), 1637–1673.

HITCHENS, T. (2021): "GPS III Launch Will Provide Global M-Code," *Breaking Defense*. https://breakingdefense.com/2021/06/gps-iii-launch-will-provide-global-m-code/ (accessed January 8, 2026), Last updated June 14, 2021.

IGAMI, M., AND K. UETAKE (2020): "Mergers, Innovation, and Entry-Exit Dynamics: Consolidation of the Hard Disk Drive Industry, 1996–2016," *Review of Economic Studies*, 87(6), 2672–2702.

IRWIN, D. A., AND N. PAVCNIK (2004): "Airbus Versus Boeing Revisited: International Competition in the Aircraft Market," *Journal of International Economics*, 64(2), 223–245.

ISKHAKOV, F., J. RUST, AND B. SCHJERNING (2016): "Recursive Lexicographical Search: Finding All Markov Perfect Equilibria of Finite State Directional Dynamic Games," *Review of Economic Studies*, 83(2), 658–703.

JONES, A. (2019): "Private Space Launch Firms in China Race to Orbit: Four companies set the pace with scheduled launches over the next two years," *IEEE Spectrum, Aerospace News*. https://spectrum.ieee.org/private-space-launch-firms-in-china-race-to-orbit (accessed January 14, 2026), Last updated April 26, 2019.

KANTOR, S., AND A. WHALLEY (2025): "Moonshot: Public R&D and Growth," *American Economic Review*, 115(9), 2891–2925.

KLOTZ, I. (2016): "SpaceX Undercut ULA Rocket Launch Pricing by 40 percent: U.S. Air Force," *Reuters*. https://www.reuters.com/article/lifestyle/science/idUSKCN0XP2T0 (accessed January 8, 2026), Last updated April 28, 2016.

KLOTZ, I. (2017): "United Launch Alliance cuts Atlas Rocket Price Amid Competition," *Reuters*. https://www.reuters.com/article/lifestyle/science/idUSKBN17706L (accessed January 8, 2026), Last updated April 4, 2017.

KOVACIC, W. E. (2019): "Competition Policy Retrospective: The Formation of the United Launch Alliance and the Ascent of SpaceX," *George Mason Law Rev.*, 27, 863.

KREBS, G. D. (2025): "Gunter's Space Page," https://space.skyrocket.de/index.html (accessed April 3, 2025).

LA PORTA, R., AND F. LÓPEZ-DE SILANES (1999): "The Benefits of Privatization: Evidence from Mexico," *Quarterly Journal of Economics*, 114(4), 1193–1242.

LEWIS, T. R., AND H. YILDIRIM (2002): "Managing Dynamic Competition," *American Economic Review*, 92(4), 779–797.

Liu, A.-H., and R. B. Siebert (2022): "The Competitive Effects of Declining Entry Costs over Time: Evidence from the Static Random Access Memory Market," *International Journal of Industrial Organization*, 80, 102797.

Mankiw, N. G., and M. D. Whinston (1986): "Free Entry and Social Inefficiency," *RAND Journal of Economics*, 17(1), 48–58.

Maskin, E., and J. Tirole (2001): "Markov Perfect Equilibrium: I. Observable Actions," *Journal of Economic Theory*, 100(2), 191–219.

Mathieu, E., P. Rosado, and M. Roser (2022): "Space Exploration and Satellites," *Our World in Data*. https://ourworldindata.org/space-exploration-satellites (accessed January 8, 2026).

McConnaughey, P. K., M. G. Femminineo, S. J. Koelfgen, R. A. Lepsch, R. M. Ryan, and S. A. Taylor (2012): "NASA's Launch Propulsion Systems Technology Roadmap," Discussion paper, NASA Office of the Chief Technologist.

McDowell, J. C. (2025): "General Catalog of Artificial Space Objects, Release 1.7.3," https://planet4589.org/space/gcat (accessed April 03, 2025).

Miller, N. H., and M. C. Weinberg (2017): "Understanding the Price Effects of the MillerCoors Joint Venture," *Econometrica*, 85(6), 1763–1791.

Mishra, R. (2025): "Trump Administration Now Holds Stakes In 5 Public Companies: Here's A List—INTC, MP, LAC And More," *Yahoo Finance*. https://finance.yahoo.com/news/trump-administration-now-holds-stakes-023008085.html (accessed January 8, 2026), Last updated October 7, 2025.

Moore, M. (2019): "Launch Vehicle Mission Success," *The Aerospace Corporation*. https://aerospace.org/getting-it-right/jun-2019/launch-mission-success (accessed January 8, 2026), Last updated June 24, 2019.

NASA (2023): "NASA Launch Services Risk Classification Fact Sheet," https://www.nasa.gov/wp-content/uploads/2023/12/risk-classification-fact-sheet.pdf (accessed January 8, 2026).

——— (2024): "NPR 8705.4B: Risk Classification for NASA Payloads," https://nodis3.gsfc.nasa.gov/npg_img/N_PR_8705_004B_/N_PR_8705_004B_.pdf (accessed January 8, 2026).

NASA Jet Propulsion Laboratory (2022): "Basics of Space Flight — Chapter 13: Navigation," https://science.nasa.gov/learn/basics-of-space-flight/chapter13-1/ (accessed January 8, 2026).

Nocke, V., and N. Schutz (2018): "Multiproduct-Firm Oligopoly: An Aggregative Games Approach," *Econometrica*, 86(2), 523–557.

Office of Commercial Space Transportation (2002): "Commercial Space Transportation: First Quarter 2002 Quarterly Launch Report," https://www.faa.gov/about/office_org/headquarters_offices/ast/media/quarter0201.pdf (accessed January 8, 2026).

Office of Federal Procurement Policy (2025): "Revolutionary FAR Overhaul," https://www.acquisition.gov/far-overhaul (accessed January 8, 2026).

Rao, A., and G. Rondina (2025): "The Economics of Orbit Use: Open Access, External Costs, and Runaway Debris Growth," *Journal of the Association of Environmental and Resource Economists*, 12(2), 353–388.

Rees, M. J. (2003): *Our Final Hour: A Scientist's Warning: How Terror, Error, and Environmental Disaster Threaten Humankind's Future in this Century–On Earth and Beyond.* Basic Books (AZ).

Robinson-Smith, W. (2025): "U.S. Space Force Awards $13.7 Billion in New National Security Launch Contracts to Blue Origin, SpaceX and ULA," *Spaceflight Now*. https://spaceflightnow.com/2025/04/05/u-s-space-force-awards-13-7-billion-in-new-national-security-launch-contracts-to-blue-origin-spacex-and-ula/ (accessed January 8, 2026), Last updated April 5, 2025.

Rose, N. L., and J. Sallet (2019): "The Dichotomous Treatment of Efficiencies in Horizontal Mergers: Too Much? Too Little? Getting it Right," *University of Pennsylvania Law Review*, 168, 1941–1984.

Sheetz, M. (2023): "Space Force Awards $2.5 Billion in Rocket Contracts to SpaceX and ULA for 21 launches," *CNBC*. https://www.cnbc.com/2023/11/01/space-force-awards-spacex-ula-with-2point5-billion-for-21-launches.html (accessed January 8, 2026), Last updated November 1, 2023.

Smith, M. (2022): "Senate Reaches Agreement on Russian RD-180 Engines," *SpacePolicyOnline.com*. https://spacepolicyonline.com/news/senate-agreement-reaches-on-russian-rd-180-engines (accessed January 8, 2026), Last updated February 19, 2022.

Smith, R. (2024): "The Secret to SpaceX's $10 Million Starship, and How SpaceX Will Dominate Space for Years to Come," *The Motley Fool*. https://www.nasdaq.com/articles/the-secret-to-spacexs-%2410-million-starship-and-how-spacex-will-dominate-space-for-years-to (accessed January 8, 2026), Last updated February 11, 2024.

Space Investments (2024): "Starship Basic Economics: Unlocking the Space Economy," https://www.spaceinvestments.io/space-economy-market-intelligence/starship-economics (accessed January 8, 2026).

Spaceflight News (2025): "NASA Budget for 2026 Phases Out SLS and Orion, Slashes Earth Science Programs," https://www.spaceflight-news.com/single-post/nasa-budget-for-2026-phases-out-sls-and-orion-slashes-earth-science-programs (accessed September 18, 2025).

Spaceflight Now (2003): "Pentagon Strips 7 Launches from Boeing Delta 4 rocket," https://spaceflightnow.com/news/n0307/24eelv/ (accessed January 10, 2026), Last updated July 24, 2003.

Sweeting, A., D. Jia, S. Hui, and X. Yao (2022): "Dynamic Price Competition, Learning-by-Doing, and Strategic Buyers," *American Economic Review*, 112(4), 1311–33.

Troutman, P. (2020): "A Future of Humans in Space," *NASA Langley Research Center*. https://ntrs.nasa.gov/api/citations/20205009062/downloads/A_Future_of_Humans_In_Space_v3.pdf (accessed January 8, 2026), Last updated October 26, 2020.

U.S. Government Accountability Office (2008): "Space Acquisitions: Uncertainties in the Evolved Expendable Launch Vehicle Program Pose Management and Oversight Challenges," https://www.gao.gov/products/gao-08-1039 (accessed January 8, 2026), Last updated September 26, 2008.

——— (2014): "The Air Force's Evolved Expendable Launch Vehicle Competitive Procurement," https://www.gao.gov/products/gao-14-377r (accessed January 8, 2026), Last updated March 4, 2014.

Vance, A. (2015): *Elon Musk: Tesla, SpaceX, and the Quest for a Fantastic Future*. HarperCollins, New York.

Wagenblast, B. N., and R. A. Bettinger (2024): "Statistical Reliability Estimation of Space Launch Vehicles: 2000–2022," *Journal of Space Safety Engineering*, 11(4), 573–589.

Weinzierl, M. (2018): "Space, The Final Economic Frontier," *Journal of Economic Perspectives*, 32(2), 173–192.

Yang, C. (2020): "Vertical Structure and Innovation: A Study of the SoC and Smartphone Industries," *RAND Journal of Economics*, 51(3), 739–785.

A Data Appendix

This appendix provides additional information on the data, and, in particular, on the estimation of the annual transaction prices for different US rocket families, including Atlas V, Delta IV and Falcon 9.

A.1 Launch Records

Our primary data source for launches and rocket characteristics is the General Catalog of Artificial Space Objects (GCAT) (McDowell (2025)). The dataset is organized by launches and records the launch date, payload and its owner, the associated rocket (our focus), and a measure of the launch success ranging from 0 to 100%. For each rocket, GCAT reports the manufacturer, number of stages, and carrying capacity. We validate and, where necessary, augment GCAT with records from Gunter's Space Page (Krebs (2025)).

We construct a dataset with the following coverage:

- Years 1985-2024
- Civil (e.g. scientific) or military launches on behalf of government agencies, although, in our regressions and when constructing moments, we will measure rocket family experience or know-how using counts that include commercial launches.
- Unmanned launches to put objects into stable Earth orbit using medium- to heavy-lift rockets (LEO payload capacity greater than 2,000 kilograms).
- Launches where the primary payload is owned by United States, Soviet Union/Russia, countries that have ever been in the European Union, China, Japan, and India. These countries account for 95% of government launches between 1985 and 2024.

Our analysis is at the rocket-family or, equivalently, rocket-system level (for example, the Atlas V rocket family/system that was developed by Lockheed Martin and then operated by ULA). A rocket family typically has a number of variants for use on missions with different payload weights or shapes, or with different orbital heights (for example, `https://www.ulalaunch.com/rockets/atlas-v`, accessed October 17, 2025, lists ten variants of Atlas V), although the variants share core components such as engines. As detailed below, when defining prices, we try to create a price that tracks the price of the most commonly launched, or "representative" variant, using hedonic regression to control for cross-variant price differences.

A.2 Rocket Prices

GCAT does not report launch prices. We therefore construct a new dataset of launch prices, drawing on three sources: FAA launch reports, government contract records, and public news reports, which we then use to create a measure of the average government transaction price for the "representative" variant in each rocket family for each year. Here we discuss the data and the steps in constructing the measures. Throughout we deflate prices to 2010 U.S. dollars using the Aerospace Producer Price Index.[63]

A.2.1 Sources

FAA reports: The Federal Aviation Administration's Commercial Space Transportation launch reports cover 1996–2009 (quarterly) and 2010-2017 (annual). From 2000 to 2017 prices are reported at the launch level, but there are a significant proportion of missing values after 2009, and we drop a few values as the prices seem to include some additional services.[64] We use the midpoint when the reported price is a range.

Government contracts: We also use data from government rocket purchase contracts obtained from the U.S. Department of Defense/War (DoD) contract announcements (`https://www.war.gov/News/Contracts/`). These contracts fall under the Evolved Expendable Launch Vehicle (EELV) program, renamed the National Security Space Launch (NSSL) program in 2019. This program covers DoD national security missions using Atlas V, Delta IV, and Falcon families. We focus on launch services contracts, which correspond directly to rocket purchases, and exclude launch support contracts, which are related to launch infrastructure. For each contract, we observe the rockets, the quantity of each rocket purchased, and the total contract amount. The contract data span 2008–2023.

News reports: Finally, we supplement our dataset with prices, where we can find them, from internet news reports and industry journals between 1995 and 2024. These observations are recorded at the rocket–year level.[65] When prices appear as ranges, we again use the midpoint as the reference value.

These sources cover both US and non-US rockets. Using hedonic regressions that we describe below, we use these prices to form an annual time series for the prices of US rockets.

A.2.2 Government Contract Decomposition

The government contracts that we observe often report a single contract price for several launches, often using different variants. We use the following procedure to try to decompose the price to the launch-level.

1. For contracts containing only multiple units of the same rocket variant, we calculate an average price by dividing the total contract amount by the number of rockets in the contract.

[63]Source: U.S. Bureau of Labor Statistics, Producer Price Index by Industry: Aerospace Product and Parts Manufacturing [PCU33643364], retrieved on May 20, 2025 from FRED, Federal Reserve Bank of St. Louis; `https://fred.stlouisfed.org/series/PCU33643364`.

[64]For example, an Atlas V version 401 launch on April 18, 2017 has a price of $150 million, which is greater than ULA's 2017 list price of $109 million for that variant.

[65]The internet reports typically mention the price of a rocket without specifying the associated mission or exact time. Therefore, we record these data at the rocket–year level, using the reported price and the time of the report as the observation.

2. We use information from single-rocket contracts within the same year, or the single variant contracts from Step 1, to decompose bundled contracts. When all rockets in a contract can be matched to observed prices, we allocate the total contract amount proportionally based on the relative price ratios. If exactly one rocket cannot be mapped, we assign the residual value to that rocket after deducting the average prices of the mapped ones. If more than one rocket cannot be matched, we proceed to the next step.

3. We estimate price ratios across rocket variants and use these ratios to decompose more complex bundled contracts. For Atlas V, we use data from ULA's RocketBuilder pricing tool to compute variant-level price ratios, assuming these ratios remain stable over time.[66] For Delta IV, we derive their relative price ratios from single-rocket contracts and nearby-year FAA reports. For Falcon rockets (Falcon 9 and Falcon Heavy), we use the list prices published on SpaceX (SPX)'s official website to estimate their relative price ratio. When a bundle includes rockets from different rocket families (e.g. Atlas V and Delta IV), we adjust the ratios using price information from single-rocket contracts and the preceding steps. These ratios are then applied to impute the rocket-level prices for the bundled contracts that cannot be decomposed in Steps 1 and 2.

A.2.3 Imputation through Hedonic Regression

Based on the data, we do not observe transaction prices for the representative variant of each rocket family in each year that the family is active. To fill these gaps, we estimate a weighted hedonic regression to capture the relationship between rocket prices and observable characteristics, and use the fitted model to impute, where necessary, missing prices for representative variants.

An observation in the regression would ideally be the price for one particular launch, although some of our price data is more aggregated. We weight our observed prices to reflect how many launches, possibly happening at different times, the observed average price is based on.

- For FAA launch-level prices, the weight is set to one.
- For FAA quarterly-level prices, the weight equals the number of launches of that rocket in the quarter.
- For government contracts, each observation corresponds to a procurement contract. When multiple contracts exist for the same rocket–year, we compute a quantity-weighted mean price based on the number of units purchased. Because government contracts report only the award date and there is generally a two-year lag between contract award and rocket launch (Robinson-Smith (2025)), each aggregated rocket–year

[66] ULA released the interactive pricing tool, RocketBuilder, in 2016 and discontinued it in 2017. We accessed the tool via the Wayback Machine and collected prices for all Atlas V variants. The archived version is available at: `https://web.archive.org/web/20161203124622/https:/www.rocketbuilder.com/start/configure`.

price is given a weight equal to the number of government launches of that rocket in the award year and the following year.

- For online report prices, each observation corresponds to a report. When multiple reports exist for the same rocket–year, we compute the unweighted average price. Following the approach for the contract data, we assign each rocket–year price a weight equal to the number of commercial launches of that rocket in the report year and the subsequent year.

Our hedonic regression model is then specified as follows: the price of rocket j from rocket family f at time t, observed in data source s, is specified as:

$$\begin{aligned}\ln\left(p_{jfst}\right) =& \theta_1 \ln\left(\{\text{LEO Capacity}\}\right)_{jf} + \theta_2 \ln\left(\{\text{Rocket Length}\}\right)_{jf} \\ &+ \theta_3 \mathbb{1}\left\{\text{Heavy}\right\}_{jf} + \theta_4 \mathbb{1}\left\{\text{Reuse}\right\}_{jf} + \theta_5 \#\left\{\text{Yrs Since Initial Launch}\right\}_{ft} \\ &+ FE_f + FE_f \times \mathbb{1}\left\{t \geq \text{ULA Merger}\right\}_t + FE_f \times \mathbb{1}\left\{t \geq \text{SPX Entry}\right\}_t + \omega_{jfst}\end{aligned} \tag{19}$$

where $\{\text{LEO Capacity}\}_{jf}$ denotes rocket j's payload capacity, in kilograms, to the low Earth orbit (LEO). $\{\text{Rocket Length}\}_{jf}$ represents the overall length of the rocket in meters. $\mathbb{1}\left\{\text{Heavy}\right\}_{jf}$ is an indicator for heavy-lift variants (e.g., Delta IV Heavy within the Delta IV family). $\mathbb{1}\left\{\text{Reuse}\right\}_{jf}$ is an indicator for reusable rockets.[67] $\#\{\text{Yrs Since Initial Launch}\}$ measures the number of years since the initial launch of rocket family f. $\mathbb{1}\left\{t \geq \text{ULA Merger}\right\}_t$ equals one for years from 2006, after the formation of ULA. $\mathbb{1}\left\{t \geq \text{SPX Entry}\right\}_t$ takes the value one for years from 2017, when SPX entered the national security launch market. By interacting rocket-family fixed effects with the indicator for the period following the ULA JV and with the indicator for the period following SPX's entry, we allow the effects of these two events on rocket prices to vary flexibly across rocket families. In one specification, we extend Equation (19) by including rocket family recent launches, $\#\{\text{Launches in Past 2 Yrs}\}$, which is defined as the number of launches of the rocket's family in the previous two years.

Table A.1 presents the regression results based on rockets from all countries. Column (1) excludes rocket-variant-specific characteristics. The estimate on $\#\left\{\text{Yrs Since Initial Launch}\right\}$ indicates that, on average, rocket prices decrease by about 0.5% per year. The overall R^2 is 0.86, and the within rocket family R^2 is 0.07. Column (2) adds rocket characteristics. The estimated coefficients align with expectations: the coefficient on $\ln\left(\{\text{LEO Capacity}\}\right)$ is positive, while the coefficient on $\mathbb{1}\left\{\text{Reuse}\right\}$ is negative, suggesting that rockets with greater payload capacity are more expensive, whereas launches using reusable rockets tend to be cheaper. The within rocket family R^2 increases to 0.16, so that variant characteristics explain a non-negligible part of the within-family variation in rocket prices.

[67] Not all Falcon 9 and Falcon Heavy rockets are reusable. Early Falcon 9 variants lacked recovery capability, and in some missions, operators use expendable versions to increase payload capacity by not reserving fuel for first-stage recovery. We obtain prices for both expendable and reusable configurations from online reports. We construct regression weights for prices of reusable and expendable configurations separately. To do so, we identify whether each Falcon launch uses a reusable or expendable configuration based on launch level information on first stage recovery attempts from the Falcon launch history page on Wikipedia (`https://en.wikipedia.org/wiki/List_of_Falcon_9_and_Falcon_Heavy_launches_(2010%E2%80%932019)`, last accessed on January 19th, 2026). Regression weights for prices of reusable and expendable configurations are then constructed separately following the weighting rule described above.

Table A.1: Hedonic Regression Results for Rocket Prices

	(1)	(2)	(3)	(4)	(5)
	$\ln(p)$	$\ln(p)$	$\ln(p)$	$\ln(p)$	$\ln(p)$
$\#\{$Yrs Since Initial Launch$\}$	-0.005**	-0.004*	0.006**	0.003	-0.001
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
$\ln(\{$LEO Capacity$\})$		0.324***	0.270***	0.282***	0.320***
		(0.038)	(0.045)	(0.042)	(0.039)
$\ln(\{$Rocket Length$\})$		-0.555***	-0.549***	-0.527***	-0.537***
		(0.162)	(0.183)	(0.178)	(0.171)
$\mathbb{1}\{$Heavy$\}$		0.033	0.008	0.019	0.032
		(0.048)	(0.058)	(0.053)	(0.049)
$\mathbb{1}\{$Reuse$\}$		-0.136***	-0.105**	-0.095**	-0.126***
		(0.036)	(0.046)	(0.039)	(0.036)
$\ln(\#\{$Launches in Past 2 Yrs$\})$					-0.020**
					(0.010)
R^2	0.864	0.877	0.884	0.871	0.877
Within R^2	0.071	0.158	0.177	0.180	0.159
Observations	684	684	671	684	676
Weight	Launch-Adjusted	Launch-Adjusted	Current-Year	Non-Starlink	Launch-Adjusted
Rocket Family FE	Yes	Yes	Yes	Yes	Yes
Rocket Family FE $\times \mathbb{1}\{t \geq$ ULA Merger$\}$	Yes	Yes	Yes	Yes	Yes
Rocket Family FE $\times \mathbb{1}\{t \geq$ SpaceX Entry$\}$	Yes	Yes	Yes	Yes	Yes

Note: *,**,*** denote significance at the 10%, 5% and 1% levels. Standard errors in parentheses. Because some price observations are aggregated, we assign regression weights to reflect the number of launches each observation may represent. FAA launch-level prices receive a weight of one. FAA quarterly-level prices are weighted by the number of launches of that rocket in the corresponding quarter. Annual-level government contract prices and online report prices are weighted by the number of government and commercial launches of that rocket in the award/report year and the following year, respectively. Column (3) applies an alternative weighting scheme in which annual-level prices are matched only to launches occurring in the award/report year. Column (4) excludes all Starlink missions when constructing the regression weights.

Column (3) applies an alternative weighting scheme in which government contract prices and internet report prices are matched only to launches occurring in the report year, excluding launches in the subsequent year. The estimated coefficients are broadly similar to those in column (2). In column (4), we exclude all Starlink missions when constructing the regression weights, and the results remain robust.

Column (5) further includes rocket family recent launches, $\ln(\#\{\text{Launches in Past 2 Yrs}\})$, and excludes observations for rocket families with no launches in the preceding two years. The coefficient on $\ln(\#\{\text{Launches in Past 2 Yrs}\})$ is negative and statistically significant. Given that we find that rocket family recent launches increase reliability and the buyer's value of a launch, this finding is consistent with rocket family recent launches also lowering launch costs (Benkard (2000), Benkard (2004)). Among the five specifications presented in Table A.1, we select the specification in column (5) as our preferred model for imputing representative rocket variant prices.

To help the reader to understand how the imputed representative prices relate to the underlying price observations, Figure A.1 shows both series for the major U.S. rocket families.[68] For the observed price, each marker represents a rocket-year level weighted average of prices from the indicated data source, with weights defined in the hedonic regression. "Rep." refers to representative variants; "Non-rep." refers to non-representative variants. Imputed prices are predicted using the preferred hedonic regression model. For the observed prices, marker size is proportional to the square root of the weight defined in the hedonic regression, which reflects the number of launches associated with that price observation. Vertical lines mark the initial launch of the EELV program rockets, the merger of Boeing and Lockheed Martin to form ULA, and the entry of Falcon 9 rockets into national security launches. Prices are deflated to 2010 U.S. dollars and expressed per kilogram of LEO payload capacity.

Based on the data sources and the hedonic model described above, we construct the rocket family–year prices used in the estimation of the structural model. For each family–year, when the representative variant's price is observed, we compute a weighted average price based on the regression weights defined earlier; when it is not observed, we impute the representative-variant price using the preferred hedonic specification. All prices are expressed in 2010 U.S. dollars per kilogram of LEO payload capacity.

[68]The figure displays Falcon 9 prices only from 2017 onward, corresponding to the period in which Falcon 9 competes in the national security launch market. We also use pre-2017 prices for (non-Starlink) commercial or government civilian launches in the regressions, which are similar in levels.

Figure A.1: Observed and Imputed Rocket Prices for Major U.S. Rocket Families

(a) Delta IV

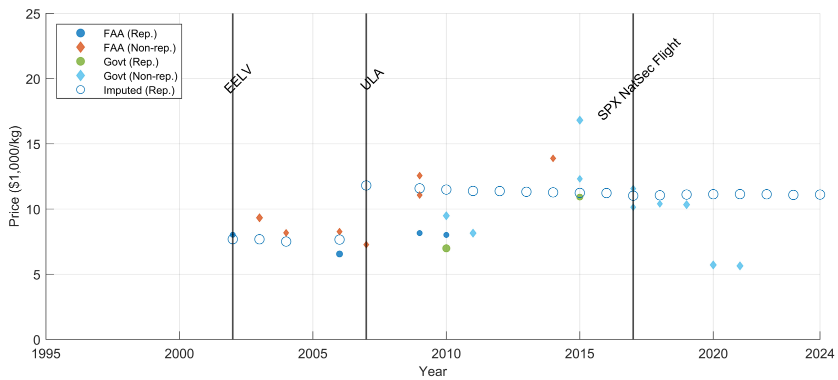

(b) Atlas V

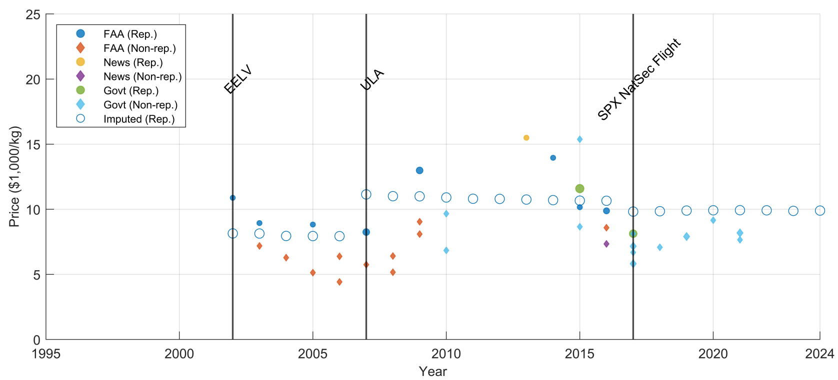

(c) Falcon 9

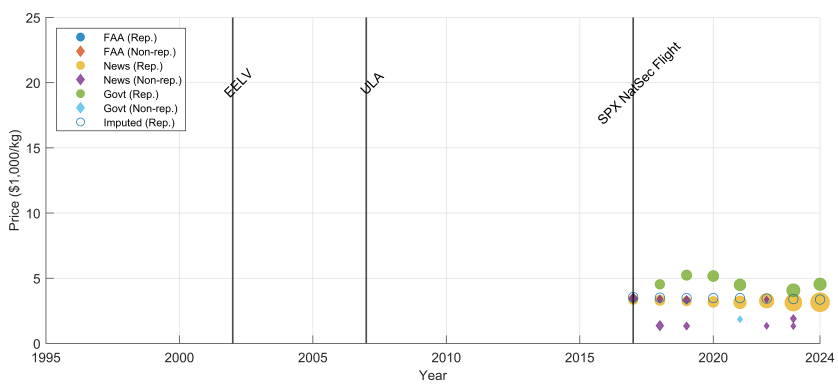

Note: Each marker represents a weighted average of rocket–year prices derived from the indicated data source, with weights defined in the hedonic regression. "Rep." refers to representative variants; "Non-rep." refers to non-representative variants. Imputed prices are predicted from the hedonic regression model (column (5) of Table A.1). For the observed prices, marker size is proportional to the square root of the weight defined in the hedonic regression, which reflects the number of launches associated with that price observation. Vertical lines mark the initial launch of the EELV program rockets, the 2006 merger of Boeing and Lockheed Martin to form ULA, and the entry of Falcon 9 rockets into national security launches. Prices are deflated to 2010 U.S. dollars and expressed per kilogram of LEO payload capacity.

A.2.4 Effect of Procurement Lot Size on Launch Prices

In our analysis, we treat each launch price as if it comes from an individual procurement. Since multi-unit purchases are often justified as cost-saving (Congressional Research Service (2025)), one concern is that our assumption may distort the time series of prices if some contracts involve multiple units and contracting practices have changed over time.

We test the extent to which launch prices vary with procurement lot size using EELV rocket purchase contracts posted on the DoD contract release website. In total, we observe 23 EELV procurement contracts covering 78 individual rocket launches.[69] For each contract, we observe the rocket configurations, the quantity of each configuration, and the total contract value. Launch-level prices are obtained using the decomposition procedure described in Appendix A.2.2.

We define the government procurement lot size as the number of booster cores purchased in a given contract. The booster core is typically the first stage of a rocket and is shared within a rocket family. For example, each Delta IV Medium vehicle uses one booster core, while a Delta IV Heavy uses three. In practice, government procurements typically specify the number of booster cores (U.S. Government Accountability Office (2014)). For example, in 2013 the Air Force procured 35 booster cores from ULA in a single contract.

Table A.2: Effect of Procurement Lot Size on Launch Prices

	(1) ln (Price per kg)	(2) ln (Price per kg)	(3) ln (Price per kg)	(4) ln (Price per kg)
ln ({Procurement Lot Size})	-0.098	0.042	-0.052	-0.048
	(0.070)	(0.055)	(0.042)	(0.061)
{Award Year - 2002}	-0.016			
	(0.026)			
ln (#{Launches in Past 2 Yrs})				0.008
				(0.280)
R^2	0.585	0.655	0.794	0.710
Within R^2	0.130	0.007	0.015	0.011
Observations	78	76	76	58
Era FE	Yes			
Award Year FE		Yes	Yes	Yes
Rocket Family FE			Yes	Yes

Note: *,**,*** denote significance at the 10%, 5% and 1% levels. Robust standard errors in parentheses. The sample consists of EELV rocket purchase contracts, and launch-level prices are obtained from the decomposition procedure described in Appendix A.2.2. Each observation is an individual rocket launch. All prices are deflated to 2010 dollars and expressed as the launch price per kilogram of the rocket's LEO payload capacity. The sample size decreases across columns because of the changing set of controls. In the full sample, the average lot size across procurements is 5.26, and the 5%, 25%, 75% and 95% quantiles are 1, 2, 5 and 23. Columns (2) and (3) include 76 observations because in two award years, there is only a single unit contract; the two "singleton" contracts are absorbed by the award year fixed effects. The dropped observations in Column (4) correspond to Vulcan contracts. Vulcan rocket family had no launches prior to these award years.

Table A.2 reports regressions of log launch price on procurement lot size. All prices are

[69]The majority of observed contracts are for launches after 2013, and the sample is smaller than the universe of all EELV launches.

deflated to 2010 dollars and expressed as the launch price per kilogram of the rocket's LEO payload capacity. The average lot size across procurements is 5.26.

Column (1) includes a linear time trend (defined as {Award Year - 2002}) and era fixed effects. The estimated coefficient on $\ln(\{\text{Procurement Lot Size}\})$ is -0.10 and is statistically insignificant. In column (2), we replace the linear time trend and era fixed effects with award year fixed effects; the estimated effect of procurement lot size remains small and statistically insignificant. Column (3) adds rocket family fixed effects to address the concern that the government might systematically purchase larger lots from particular rocket families at higher/lower prices. After controlling for both award year and rocket family fixed effects, the estimated coefficient on procurement lot size is -0.05 and remains statistically insignificant. Column (4) further controls for rocket family launch experience, #{Launches in Past 2 Yrs}, defined as the number of launches of the rocket family in the two calendar years prior to the award year.[70] This control addresses the possibility that the government procures larger lots from more experienced rocket families, with experience also correlated with the prices those families bid. The coefficient on procurement lot size is again -0.05 and statistically insignificant.

Although the lack of statistical significance may reflect the limited number of contracts, the point estimates are stable across columns (3) and (4), and they imply that doubling the procurement lot size reduces launch prices by roughly 5%, a modest effect.

A.3 Recent Launch Experience and Reliability

In the text, we describe some evidence for accumulated launches being associated with increased reliability, justifying why we include experience, measured as the number of launches within the last two years, in the buyer's indirect utility function. Here we provide some more detail on the analysis.

A.3.1 Reliability Curves

The existing academic and industry literatures (e.g. Moore (2019), Wagenblast and Bettinger (2024)) discuss the role of reliability and its relation to accumulated launches.[71] A common

[70]The sample size decreases across columns because of the changing set of controls. Column (1) includes 78 observations. Columns (2) and (3) include 76 observations because in two award years, there is only a single unit contract; the two "singleton" contracts are absorbed by the award year fixed effects. Column (4) includes 58 observations. The dropped observations correspond to Vulcan contracts. Vulcan rocket family had no launches prior to these award years, so the experience measure, #{Launches in Past 2 Yrs}, is zero; these observations are therefore omitted.

[71]For example, from Moore (2019), "The Aerospace Corporation (Aerospace) periodically generates predictions of the probability of mission success (aka reliability) for upcoming national security space launches, using reliability models based on the success and failure history of over 800 U.S. European launch missions. These predictions are a vital input to forward-looking studies such as functional availability analyses and constellation risk assessments, tools that mission planners utilize to ensure high confidence in enduring constellation success. The predictions are based on the reliability growth principle, which is the continuous improvement in reliability as a system is operated or tested and as/or process defects are discovered and corrected. Analysis of historical launch data ... shows that reliability growth is one of the most significant factors affecting launch reliability—the more experience behind a launch vehicle family, the more reliable future launches are expected to be."

way to represent these relationships is by using "reliability curves". These curves plot, on the y-axis, the cumulated probability of success based on launches up to x, which is the cumulative number of launches. While the curve will drop down after any failure, a path that is upward sloping overall is used to illustrate how reliability tends to increase with experience.

Figure A.2 shows these plots for major rocket families in the US (panel (a)) and the rest of the world (panel (b)). The GCAT data defines a % success of each launch, reflecting how some launches may have partial success (e.g. a satellite was not launched into exactly the correct orbit, but the operator was able to maneuver it into the correct position, at the cost of using fuel which may reduce the satellite's useful life). However, for our purposes, we define a launch as a failure whenever it does not have 100% success. Panel (a) presents the accumulated reliability of U.S. rocket families. As shown in the figure, the reliability of Atlas V, Delta IV, and Falcon increases with the number of launches. Panel (b) displays the accumulated reliability of non-U.S. rocket families, showing that Russian, Chinese, and European rockets also exhibit higher reliability as they accumulate launches.

A.3.2 Learning Synergy

Kovacic (2019) argues that the formation of ULA was expected to strengthen learning processes by increasing launch tempo and enabling technical improvements from each rocket family to diffuse to the other, thereby enhancing overall reliability. We use a regression approach to examine whether we can see evidence for the synergy in the data.

Our regression includes two experience measures: $\ln(\#\{\text{Own Launches in Past 2 Yrs}\})$, which denotes the number of launches conducted by a rocket family in the two prior calendar years, and $\ln(\#\{\text{Partner Launches in Past 2 Yrs}\})$, which captures launch experience acquired from its partner rocket family. The latter variable applies only to ULA rocket families. For each ULA rocket family, it counts launches that (i) were conducted by the other family, (ii) occurred after the formation of ULA, and (iii) took place within the prior two calendar years. The coefficient on $\ln(\#\{\text{Partner Launches in Past 2 Yrs}\})$ is therefore intended to measure learning synergy across rocket families within ULA.

Table A.3 reports the marginal effects based on a linear probability model, where each observation corresponds to a launch. We set $\ln(\#\{\text{Partner Launches in Past 2 Yrs}\})$ to zero for all non-ULA rocket families and for launches that occurred in or before 2007. Columns (1)–(4) include alternative sets of control variables. Across specifications, the regression results are consistent with the presence of learning synergy: launches of one ULA rocket family appear to enhance the reliability of the other. In particular, column (4) indicates that launches by the partner rocket family are as effective as the rocket family's own launches in enhancing launch reliability.

Figure A.2: Accumulated Reliability of Major Rocket Families

(a) US Rocket Families

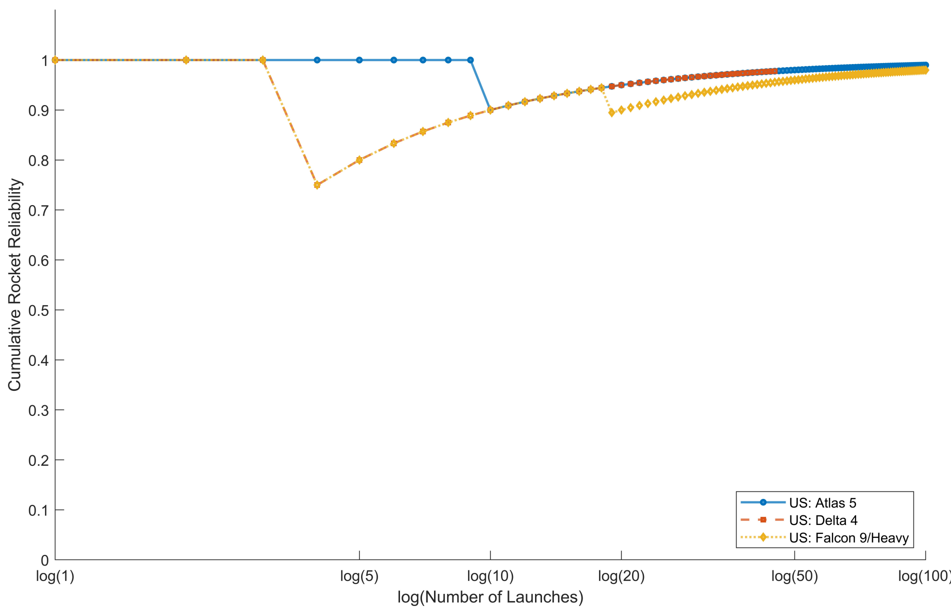

(b) Non-US Rocket Families

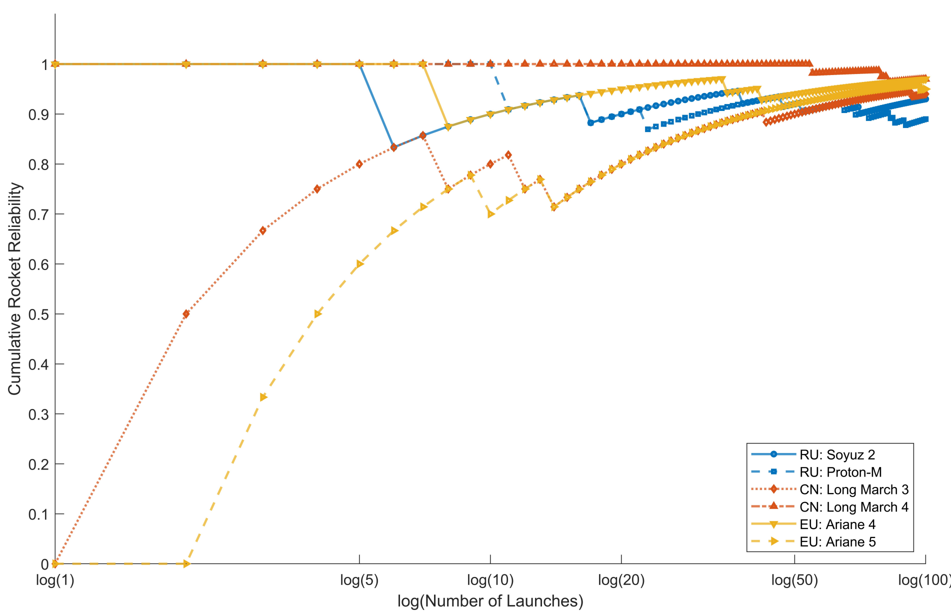

Note: Accumulated reliability is defined as the total number of successful launches of a rocket family divided by its total number of launches. Both governmental and commercial launches are included in the calculation.

Table A.3: Learning Synergy Estimation Results

	(1) $\mathbb{1}\{\text{Success}\}$	(2) $\mathbb{1}\{\text{Success}\}$	(3) $\mathbb{1}\{\text{Success}\}$	(4) $\mathbb{1}\{\text{Success}\}$
$\ln(\#\{\text{Partner Launches in Past 2 Yrs}\})$	0.024***	0.028***	0.026***	0.029***
	(0.003)	(0.004)	(0.005)	(0.007)
$\ln(\#\{\text{Own Launches in Past 2 Yrs}\})$	0.014***	0.015***	0.015***	0.022***
	(0.005)	(0.005)	(0.005)	(0.005)
$\ln(\{\text{LEO capacity}\})$		-0.019**	-0.019**	-0.012
		(0.008)	(0.008)	(0.008)
$\mathbb{1}\{\#\text{stages} > 2\}$			-0.006	0.016
			(0.010)	(0.012)
Mission Region FE				Yes
Observed Launches		1,934		

Note: *,**,*** denote significance at the 10%, 5% and 1% levels. Robust standard errors in parentheses. #{Own Launches in Past 2 Yrs} denotes the number of launches by the rocket family in the preceding two years. #{Partner Launches in Past 2 Yrs} applies only to ULA rocket families. For each ULA rocket family, it counts launches that (i) were conducted by the other family, (ii) occurred after the formation of ULA, and (iii) took place within the two prior calendar years. We set ln (# Partner Launches in Past 2 Yrs) to zero for all other rocket families and for launches that occurred in or before 2007.

B Solving and Simulating the Dynamic Model

In this appendix, we detail how we solve the dynamic games, and how, given a solution, we simulate outcomes to assess fit and to perform counterfactuals.

B.1 Numerical Solution of Dynamic Game

As described in the text, we solve distinct games for each era, based on the assumption that the transitions between eras are unanticipated surprises, in the sense that, in each era, players form strategies assuming that the game played in that era has an infinite horizon.

The state variables of the game are discrete (levels of know-how for the available rocket families), but in each period, the sellers receive i.i.d. private shocks to their launch costs for each rocket, while the buyer receives private preference shocks for procuring a launch from one of the rockets and the outside good.

B.1.1 Continuation Game with Innovation

The most complicated game follows the entry of SpaceX (SPX), as there are three products and we allow for ULA to innovate by replacing Atlas V and Delta IV with Vulcan. In this case, we first solve the dynamic game which follows after the innovation, where no further innovation is possible, and then use the computed values from this game as continuation values for ULA's innovation option in the pre-innovation game. For other eras, or when no innovation is allowed, there is a single dynamic game to solve.

B.1.2 Solution Method

We use an iterative method to solve the game. Given a vector of continuation values for each state, we solve the bidding strategies of the sellers and procurement strategy of the buyer. Given the strategies and the current continuation values, we compute the new values using equations (12), (13), (16), (17), and (18).

The solution of bidding strategies is based on discretizing costs and bids. Bichler et al. (2025) propose a method to efficiently approximate the solution to an incomplete information auction with continuous cost/valuations and continuous bids using discrete costs and discrete possible bids. The method then solves for best-response distributions over the possible bids for each bidder. The updating rule uses simultaneous online dual averaging. Bichler et al. (2025) show that the method has significant speed advantages (i.e. quick convergence to a solution), and, in cases where an analytical Bayesian Nash equilibrium can be computed, that the approximation errors are small.

To apply this method, we discretize a rocket's costs into $K^c = 10$ levels, $c_1 < c_2 < \ldots < c_{K^c}$, based on the quantiles of the normally distributed ν_{it} from the 2.5% quantile to 97.5% quantile. We also discretize possible bids into $K^a = 10$ levels. We start the algorithm using a possible bid set of equally spaced K^a points, $a_1 < \ldots < a_{K^a}$ between $0.5c_1$ and $2.5c_{K^c}$. However, if when we run the algorithm, we predict that sellers would bid at the upper bound, a_{K^a}, we extend the upper bound of the support.[72] We use $\pi_n(k^a, k^c, \boldsymbol{e})$ to denote the

[72]Specifically we increase the upper bound by 10% and re-space the bid levels.

Algorithm 1 Solve for the No-Innovation Dynamic Optimal Pricing Strategy

Input: initial pricing strategy $\pi_{n,1}$

Note: $\eta_t = \dfrac{2}{t+1}$ is the Frank-Wolfe step size.*

Initialize $t \leftarrow 1$, a_{K^a} binding$\leftarrow$true
while a_{K^a} binding
    difference$\leftarrow \infty$,
    while difference$>10^{-5}$
        new buyer choice prob $d_{t+1}(\boldsymbol{e}) \leftarrow$ buyer problem
        update buyer's value function taking into account know-how depreciation
        given $d_{t+1}(\boldsymbol{e})$ and $\pi_{n,t}$, compute each seller's best response $\tilde{\pi}_{n,t}$
        update $\pi_{n,t+1} \leftarrow \eta_t \tilde{\pi}_{n,t} + (1-\eta_t)\,\pi_{n,t}$
        update seller's value function taking into account know-how depreciation
        difference$\leftarrow \max\left\{\|d_t - d_{t+1}\|, \sum_n \|\pi_{n,t} - \pi_{n,t+1}\|\right\}$, $t \leftarrow t+1$
    end while
    if a_{K^a} binds†
        $a_{K^a} \leftarrow 1.1 \times a_{K^a}$
        $a_k \leftarrow \dfrac{k}{10}(a_{K^a} - a_1) + a_1$
    end if
end while

$\star$: We usually achieve convergence with fewer than 50 iterations.
$\dagger$: We use the criterion that the bids bind at the upper bound if for any seller and know-how level, there are at least two cost levels at which the probability of bidding at a_{K^a} exceeds 0.1.

probability that the bid of rocket n is at level k^a given a cost level k^c and the current state of know-how $\boldsymbol{e}$. When ULA has two rockets, it is setting bids for both of them, as a function of both of their cost shocks. While the method allows for equilibrium strategies to be pure or mixed, our solutions are almost always approximations to a pure strategy, with over 99% of the probability weight concentrated at one bid level for every possible cost realization.

The algorithm is summarized in Algorithm 1 for the case when there is no innovation. The extension to the algorithm when innovation is possible is given in Algorithm 2, where we are also solving for the probability of innovation in each state. Our method extends the static case in Bichler et al. (2025) to a dynamic game with both LBD and innovation. Although Bichler et al. (2025) establishes convergence only in static cases, we do not encounter non-convergence in the estimation process, where we search across a large parameter space using a combination of genetic algorithm and surrogate optimization.

Algorithm 2 Solve for the Dynamic Optimal Innovation and Pricing Strategy

Input: initial pricing strategy $\pi_{ULA,1}, \pi_{SPX,1}$, initial innovation probability o_1, post-ULA innovation buyer value V^{inv}, ULA value W^{inv}_{ULA}, SpaceX value W^{inv}_{SPX}.

Note: $\eta_t = \dfrac{2}{t+1}$ is the Frank-Wolfe step size.

1. Initialize $t \leftarrow 1$, inv_difference$\leftarrow \infty$,
2. **while** inv_difference$> 10^{-2}$
3. Run Algorithm 1 given the post innovation value functions, π and o_t
4. Solve for optimal innovation prob $\tilde{o}_t$
5. $o_{t+1} = \eta_t \tilde{o}_t + (1 - \eta_t)\, o_t$
6. inv_difference$\leftarrow \|o_t - o_{t+1}\|, t \leftarrow t + 1.$
7. **end while**

B.2 Simulation

B.2.1 Simulation for Estimation

As explained in Section 5, we treat know-how, measured using the number of launches in the previous two years[73], as observed in estimation so that we can form conditional moments. There are two types of moments. The first type are based on conditional choice probabilities given the know-how states observed in the data, where the know-how of a rocket is the number of launches in the past two calendar years. These moments do not require simulation: we average the winning probabilities based on the equilibrium strategies in each era to form the moments. The second type are based on the average winning prices. We use 100 simulations to calculate winning prices by randomly drawing from the equilibrium bid

[73] SPX's know-how is fixed at 15 because of the large number of Starlink launches.

distributions and drawing a winner using the buyer's choice probabilities, conditional on the know-how states.

B.2.2 Simulation for Assessing Fit and Counterfactual Experiments

To assess model fit and compute counterfactuals, we simulate bids, procurement outcomes and state transitions based on the equilibrium strategies of the sellers and the buyer. All of our simulations will use 1,000 simulated paths.

Era 1. Atlas V and Delta IV are owned by separate sellers, and there is no innovation. We first solve for the optimal pricing strategies and buyer choice probabilities using Algorithm 1.

All simulations begin in the year 2002, the start of the EELV era, with know-how levels equal to 1 for both rockets. In each procurement, given the current levels of rocket know-how, we simulate cost realizations for each procurement opportunity and, use the equilibrium bid strategies and buyer choice probabilities, to simulate the outcome (i.e., whether any rocket is procured, and if so which).

We then simulate the evolution of know-how. To calculate the new know-how stock for the next procurement, we independently simulate a jump to the next know-how level (up to 15) with probability $\frac{1}{e_{jt}^U - e_{jt}}$, where e_{jt} is the current know-how level and $e_{jt}^U \in \{1, 5, 10, 15\}$ is the next higher level. Next, we independently simulate each rocket's drop to the next lower know-how level with probability $\lambda \frac{e'_{jt}}{e'_{jt} - e_{jt}^L}$ until reaching the level of 1, where e_{jt}^L is the next lower level of know-how. We repeat this for the $N = 20$ procurement opportunities we assume for each year.

At the end of the year, if a rocket undertakes commercial launches, we add the number of launches to the know-how stock in the last procurement before the depreciation stage.

Era 2. Atlas V and Delta IV are owned by the same seller. We modify the simulation procedure to account for the synergy in know-how accumulation. We re-solve the equilibrium pricing strategies that account for (1) the era 2 demand intercept and (2) buyer and seller expectations that a new launch adds to the know-how of both rockets. The know-how increases are probabilistic and independent: for example, if either rocket is chosen at state $\boldsymbol{e} = (1, 1)$, the know-how moves to $(1, 5)$ or $(5, 1)$ with probability $\frac{1}{4}\left(1 - \frac{1}{4}\right)$, and it moves to $(5, 5)$ with probability $\frac{1}{16}$. Given the strategies and transitions, we repeat the procedure above to simulate era 2 launches 1,000 times, starting at the simulated know-how levels reached on each path at the end of era 1. We still assume the know-how levels of both rockets depreciate independently, but whenever one of the rockets is chosen, the know-how of both rockets stochastically and independently increases.

Era 3. SPX and ULA compete. We solve for their strategies using Algorithm 2. The know-how accumulation for ULA is the same as in Era 2 for pre-innovation simulations. Post-innovation simulations are similar to the procedure in Era 1 based on solutions using

Algorithm 1 and the low production cost parameter. In all but one of the counterfactual simulations, the SPX know-how level is fixed at 15. In the simulation where the SPX know-how level upon entry is set to 1, its know-how level evolves endogenously in the same way as a ULA rocket. For ULA's Atlas V and Delta IV rockets, we assume that their initial know-how is 15 based on the observed count of launches in the previous two years, and our assumptions about the ULA learning synergy.

C Extensions to Innovation Counterfactuals

C.1 Static Procurement Strategy and Innovation

Commitment to a static procurement strategy can also affect innovation. In particular, a static buyer would favor the current highest know-how and lowest cost seller. This strategy may decrease ULA's value before innovation by reducing the demand for ULA's high-cost EELV rockets. ULA's post-innovation value may also be lower because ULA needs to discount Vulcan rockets more aggressively to accumulate and maintain know-how. In Table C.1, we compute the market outcomes with a static buyer, pre- and post-innovation, and with and without SPX.

Consistent with the intuition above, the ULA values (2017 PDV) are lower than those in Table 10, where forward-looking procurement is assumed, under the same market structures. Furthermore, columns (1) and (2) in Table C.1 show that innovation would increase ULA's value by $13.70 - 0.14 = 13.56$ billion dollars when ULA competes with SPX, which is smaller than the value increase in column (3) of Table 10. Consequently, the implied probability of innovation under a static-buying strategy is 0.13 in 2017, slightly lower than the 0.14 with the dynamic buyer in Table 11.

C.2 Planner Solution: Innovation and Cost Inefficiency

We calculate outcomes and values under planner control with two potential cost inefficiencies: (1) an innovation cost inefficiency where the mean innovation cost is $\left(1+\mu^I\right)\cdot C$, and (2) a production cost inefficiency where the cost of the new rocket post-innovation is $(1+\mu)\,c\,(e_t)$. In Figure C.1(a), we plot indifference curves which show the gains in the planner's value (2017 PDV, \$ bn.) over the buyer's value in the ULA-SPX duopoly market outcome (column (1) of Table 10).[74] In panel (b), we plot the gains in the planner's value over combined buyer and seller values in the market outcome.

We predict that the government should still prefer a planner takeover of ULA even if this increases production costs by as much as 360% (innovation costs held fixed) or if it increases innovation costs by 140% (production costs held fixed), although total surplus would be lower. We also see that the planner would still see significant gains in value (over \$6 billion) even if planner control is associated with, for example, production cost increases of 50% and mean innovation cost increases of 30%. The high tolerance of production cost inefficiency is consistent with the high likelihood of innovation by the planner and the high markup post innovation charged by a monopoly ULA (column (6) of Table 10).

[74] We assume that SPX no longer enters if the government takes over ULA.

Table C.1: Market Outcomes with Innovation: Static Buyer

	Duopoly		Monopoly	
	EELVs, Falcon	Vulcan, Falcon	EELVs	Vulcan
	(1)	(2)	(3)	(4)
Price ($1,000/kg)				
ULA, 2017	9.34	2.21	9.51	2.94
2017-2023	9.30	4.38	9.35	5.16
SPX, 2017	3.41	2.21		
2017-2023	3.52	1.64		
Number of Launches/Year				
ULA, 2017	3.92	11.28	5.43	14.10
2017-2023	2.65	14.42	2.32	14.59
SPX, 2017	8.26	5.34		
2017-2023	8.55	2.34		
Know-How				
ULA, 2017-2023	10.57	10.35	6.23	10.55
SPX, 2017-2023	15.00	6.18		
Lerner Index				
ULA, 2017	0.23	0.38	0.22	0.59
2017-2023	0.21	0.78	0.17	0.82
SPX, 2017	0.77	0.59		
2017-2023	0.78	0.41		
Firm-Level HHI				
2017	6,024	6,036	10,000	10,000
2017-2023	7,193	8,058	10,000	10,000
2017 Surplus ($bn/Year)				
Buyer	0.71	1.11	0.43	0.84
ULA	0.12	0.17	0.15	0.37
SPX	0.32	0.11		
2017 PDV ($bn)				
Buyer	21.29	24.84	13.01	17.76
ULA	0.14	13.70	15.01	20.34
SPX	3.76	2.25		
Total	25.19	40.79	28.02	38.10

Note: We report the average winning prices, number of launches per year, average rocket know-how, Lerner indices, firm-level HHI, average flow surpluses of buyers and firms, and 2017 present discounted values of the buyer and firms in this table. We start at the maximum know-how level of 15 as observed in the data. We fix Falcon 9 know-how at 15.

Figure C.1: Planner Indifference Curves

(a) Planner-Buyer Surplus (b) Planner-Total Surplus

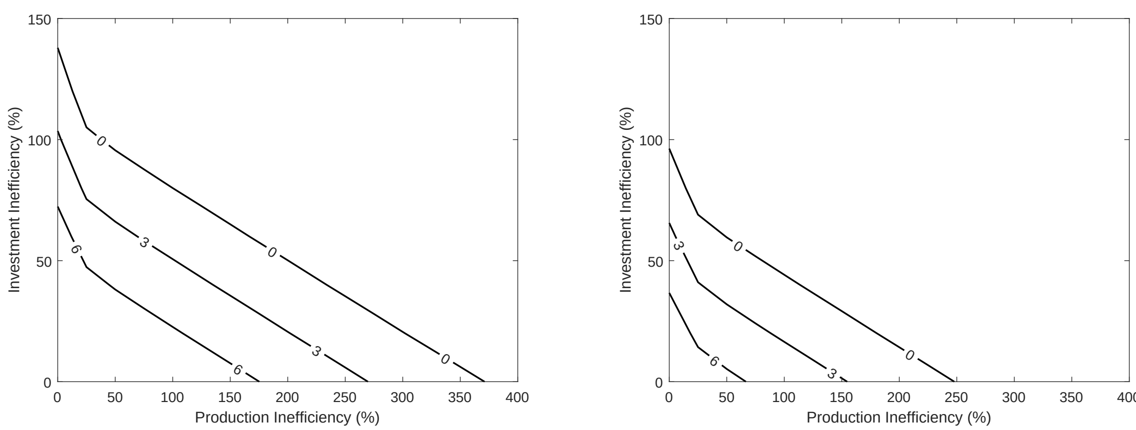

Note: The numbers on the indifference curves are the gains in the planner's value (2017 PDV, $ bn.) over either the buyer's value (panel (a)) or the combined buyer and seller values (panel (b)) in the ULA-SPX duopoly market based on era 3 demand, where the planner takes over ULA and SPX no longer enters.

D Additional Tables

Table D.1: Planner Costs and Winning Probabilities: Era 2

	Average Costs			
Atlas V \ Delta IV	1	5	10	15
1	16.603, 16.603	16.603, 11.384	16.603, 9.676	16.603, 8.798
5	11.384, 16.603	11.384, 11.384	11.384, 9.676	11.384, 8.798
10	9.676, 16.603	9.676, 11.384	9.676, 9.676	9.676, 8.798
15	8.798, 16.603	8.798, 11.384	8.798, 9.676	8.798, 8.798

	Planner Choice Probabilities			
Atlas V \ Delta IV	1	5	10	15
1	0.314, 0.314	0.021, 0.762	0.003, 0.815	0.001, 0.798
5	0.762, 0.021	0.414, 0.414	0.173, 0.669	0.078, 0.743
10	0.815, 0.003	0.669, 0.173	0.428, 0.428	0.264, 0.578
15	0.798, 0.001	0.743, 0.078	0.578, 0.264	0.417, 0.417

Note: Know-how of Atlas V in rows, know-how of Delta IV in columns. We report the average costs and buyer's choice probabilities for Atlas V and Delta IV rockets in each state. The demand intercept is based on the era 2 estimate, which is nearly identical to the era 1 estimate.

Table D.2: Average Bid Prices, Costs and Winning Probabilities: Static Buyer, Era 2

	Duopoly				Monopoly			
	Average Bid Prices							
Atlas V \ Delta IV	1	5	10	15	1	5	10	15
1	17.394, 17.394	17.395, 11.532	17.400, 10.754	17.405, 11.514	10.205, 10.205	13.959, 8.697	16.721, 9.615	18.660, 10.718
5	11.532, 17.395	11.536, 11.536	11.551, 10.702	11.582, 11.235	8.697, 13.959	9.863, 9.863	11.787, 9.928	13.473, 10.765
10	10.754, 17.400	10.702, 11.551	10.692, 10.692	10.694, 11.235	9.615, 16.721	9.928, 11.787	10.929, 10.929	12.251, 11.317
15	11.514, 17.405	11.235, 11.582	11.235, 10.694	11.235, 11.235	10.718, 18.660	10.765, 13.473	11.317, 12.251	12.297, 12.297
	Average Costs							
Atlas V \ Delta IV	1	5	10	15	1	5	10	15
1	16.603, 16.603	16.603, 11.384	16.603, 9.676	16.603, 8.798	16.603, 16.603	16.603, 11.384	16.603, 9.676	16.603, 8.798
5	11.384, 16.603	11.384, 11.384	11.384, 9.676	11.384, 8.798	11.384, 16.603	11.384, 11.384	11.384, 9.676	11.384, 8.798
10	9.676, 16.603	9.676, 11.384	9.676, 9.676	9.676, 8.798	9.676, 16.603	9.676, 11.384	9.676, 9.676	9.676, 8.798
15	8.798, 16.603	8.798, 11.384	8.798, 9.676	8.798, 8.798	8.798, 16.603	8.798, 11.384	8.798, 9.676	8.798, 8.798
	Winning Probabilities							
Atlas V \ Delta IV	1	5	10	15	1	5	10	15
1	0.001, 0.001	0.001, 0.130	0.001, 0.315	0.001, 0.318	0.065, 0.065	0.007, 0.372	0.001, 0.459	0.000, 0.421
5	0.130, 0.001	0.117, 0.117	0.094, 0.298	0.090, 0.330	0.372, 0.007	0.216, 0.216	0.092, 0.392	0.041, 0.403
10	0.315, 0.001	0.298, 0.094	0.254, 0.254	0.248, 0.278	0.459, 0.001	0.392, 0.092	0.250, 0.250	0.141, 0.313
15	0.318, 0.001	0.330, 0.090	0.278, 0.248	0.271, 0.271	0.421, 0.000	0.403, 0.041	0.313, 0.141	0.212, 0.212

Note: Know-how of Atlas V in rows, know-how of Delta IV in columns. We report the average prices bid in each procurement, the average costs and buyer's choice probabilities for Atlas V and Delta IV rockets in each state. The demand intercept is based on the era 2 estimate, which is nearly identical to the era 1 estimate.