

UNIVERSITY OF CALIFORNIA
Los Angeles

Essays on the Market for Prototypes

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Management

by

Thomas Jorgenson Groesbeck

2021

© Copyright by
Thomas Jorgenson Groesbeck
2021

ABSTRACT OF THE DISSERTATION

Essays on the Market for Prototypes

by

Thomas Jorgenson Groesbeck

Doctor of Philosophy in Management

University of California, Los Angeles, 2021

Professor Nico Voigtlaender, Chair

Why do so many firms successfully bring innovative new products to market, only to fail shortly thereafter? I argue that early sales are critical in improving a product’s characteristics, leading to more sales: the result is a “bottle rocket” economy where a minor difference in initial trajectory leads a firm to accelerate to the heavens or just as rapidly crash to earth.

Chapter 1 concerns the role of reliability growth in dynamics. A firm can observe problems encountered by early adopters and modify their blueprint accordingly. Trial-and-error thus allows successful products to become more reliable, and thus even more successful, posing a challenging “competitive moat” to entrants. I apply this model to the market for space launches, showing that observable reliability growth explains the longevity of rockets from the 1950s, the expensive failure of the majority of new orbital launch vehicles, and the apparent backfiring of US sanctions meant to stem the growth of the Chinese space sector.

Chapter 2 extends this model of reliability growth to software: crucially, automatic updates allow early adopters to benefit from continued product improvement. This results in a dynamic game of imperfect information: consumers adopt a product based on both current value and expected future improvements, while firms must decide if delivering on

expectations is profitable. I find that firms release patches at a relatively steady rate until support is ended completely, and present an equilibrium where protracted steady patching serves as both a direct product improvement and a credible signal for long-term support.

Chapter 3 concerns social influence in adopting a new product: users can delay adoption until network effects strengthen or social consensus is positive. I establish a micro-foundation for the canonical “social influence” model of adoption, and develop a novel nonparametric estimator the heterogeneity of user preferences underlying an observed adoption curve. Using adoption curves from the market for video games, I show that this algorithm produces estimates consistent with known product characteristics, i.e. the relative importance of social features versus independent early adopters.

The dissertation of Thomas Jorgenson Groesbeck is approved.

Romain Wacziarg

Elisabeth Honka

Ariel Burstein

Nico Voigtlaender, Committee Chair

University of California, Los Angeles

2021

TABLE OF CONTENTS

1 Boosting the Competition:

Reliability, Learning, and Trade Barriers

in the Space Launch Industry	1
1.1 Introduction	1
1.1.1 Previous Literature	2
1.1.2 Rockets and Economics	4
1.2 Institutional Details	6
1.2.1 Data	6
1.2.2 A Note on Taxonomy	7
1.2.3 Trade Friction Shock	8
1.3 Demand Side	11
1.3.1 Capability	11
1.3.2 Data Details	12
1.3.3 Commercial Availability	13
1.3.4 Logit Model	14
1.3.5 Results	17
1.4 Duplex Launches	21
1.5 Supply Side	25
1.5.1 Reliability Growth	25
1.5.2 Production	29
1.6 Profit	34

1.7	Firm Dynamics with Debugging	35
1.7.1	Steady State	35
1.7.2	Dynamic game	36
1.7.3	Solution Concept	37
1.8	Vehicle Family Entry and Unanticipated Exit	41
1.8.1	Satellite Construction	45
1.8.2	Policy Implications	47
1.9	Conclusion	49
2	Testing for Fun and Profit	51
2.1	Introduction	51
2.2	Setting	52
2.3	Data and Trends	54
2.4	Sketch of the Model	59
2.5	Consumer's Problem	61
2.6	Developer's Problem	61
2.7	Fully-Informed Equilibrium	64
2.7.1	Timing and Trust	67
2.8	Equilibrium with Asymmetric Information	68
2.8.1	Incentive Compatibility Constraints	69
2.9	Conclusion	70
2.10	Assumptions of PM-2 Discrete	71
3	Nonparametric Heterogeneity in Social Influence	72

3.1	Introduction	72
3.2	Literature Review	73
3.3	Setting and Data	74
3.4	Social Influence	78
3.5	Empirical Discussion	83
3.6	A Microeconomic Foundation	86
3.6.1	Myopic Consumers	87
3.6.2	Forward-Looking Consumers	90
3.7	Conclusion	92
.1	Crow-AMSAA MLE	95
.1.1	Pooled Estimator	96
.2	Interpolating Orbital Capability	96
.3	Demand Trends	100

LIST OF FIGURES

1.1	Recent evolution of market shares	10
1.2	Data from NASA Launch Performance Calculator	11
1.3	Fitted Crow-AMSAA results	28
1.4	Evolution of reliability	29
1.5	Recent evolution of reliability	30
1.6	Production time	33
1.7	Production time	33
1.8	Simulated sales paths	39
1.9	Simulated sales paths with an ITAR-like shock	40
2.1	Reliability growth (scatter) and patching (stepped line)	57
2.2	Reliability growth (scatter) and patching (stepped line)	58
3.1	MAU evolution of a “blockbuster” versus a “sleeper hit”. Blue lines indicate new content released for <i>The Witcher 3</i>	77
3.2	Smoothed adoption of <i>Counter-Strike: Global Offensive</i>	81
3.3	Results for <i>CS: GO</i> including 1-month pre-release beta test	85

LIST OF TABLES

1.1	Data Overview	7
1.2	Satellite Operators/Customers	7
1.3	Demand results, disaggregated	18
1.4	Demand model with “distance” measures	19
1.5	Demand with pro-rated prices	24

ACKNOWLEDGMENTS

I would like to thank my fiancée and pandemic office-mate Wenwen Ni, both for her constant support and for providing, day after day, an inspiring example of the curiosity, diligence, industry, and professionalism of an exemplary social science researcher.

I would like to thank my professors at UCLA, especially my committee members Nico Voigtlaender, Ariel Burstein, Elisabeth Honka, and Romain Wacziarg, for their advice and patience throughout my time at UCLA. I am grateful to the professors and students who attended my presentations at the Global Economics and Management, International Trade, and Industrial Organization seminars at UCLA, as well as the attendees and organizers of the Los Angeles Conference in Applied Economics and the Wharton Innovation Doctoral Symposium. I would also like to thank my classmates Bruno Pellegrino, Sebastian Ottinger, Maria Lucia Yanguas, and everyone else who attended our prestigious “Manning Seminar Series”.

I would also like to thank the Center for Global Management, the Harold and Pauline Price Center for Entrepreneurship & Innovation, and the Smith Richardson Foundation for their financial support. I am grateful to all of the academic, government, and industry experts who took the time to speak with me on these topics; I am especially grateful to Dr. Tom Boone IV and the Federal Aviation Administration for graciously sharing their data and estimates on costs and capabilities of modern space launch vehicles.

Finally, I would like to dedicate this paper to my grandmother, Lois Groesbeck, who always read my essays and recommended more books for me to read.

VITA

- 2010–2014 B.A. Economics, B.S. Mathematics, Washington and Lee University
- 2014–2016 Research Assistant, Federal Reserve Board of Governors
- 2016–2019 Anderson Fellowship, UCLA Anderson School of Management
- 2017–2021 Teaching and Research Assistant, UCLA Anderson School of Management

CHAPTER 1

Boosting the Competition: Reliability, Learning, and Trade Barriers in the Space Launch Industry

1.1 Introduction

Can export controls backfire, helping foreign firms excel in strategic industries? I examine one striking case: in 1998, US firms handled 42% of civilian space launches worldwide. The US government, concerned with alleged technological espionage by China, placed space launch services under the stringent International Traffic in Arms Regulations (ITAR). By 2012, the civilian market share for all US firms put together was just 9%, while Chinese firms had grown from 6% to 26% of the global market. Concerned with losses to US rocket and satellite manufacturers, Congress allowed many space goods and services to return to the more lenient Commerce Control List (CCL); US market share recovered somewhat by 2018, but now trailed Chinese firms 28% to 36%.

Was the US's self-imposed trade restriction responsible for the loss of US leadership in the space industry? And if so, why did the 2013 legislation fail to return American firms to dominance? To answer this question, I introduce a model incorporating exogenous demand shocks, endogenous supply decisions, and an approach to learning-by-doing new to economics.

Around 5% of space launches result in catastrophic failure, with the cargo, rocket, and

often nearby buildings destroyed, and less than a third of satellites are insured (Manikowski and Weiss, 2013). As a result, industry analysts and marketing focus on safety records as the main axis of differentiation. I use comprehensive data on launch failures since 1957 and the reliability growth analysis method from Crow (1975), popularized in engineering and computer science but new to economics, to develop a measure of rocket safety over time, quantifying the effects of trial-and-error on the desirability of different launch systems.

The effects of learning and trade barriers are likely mediated by entry, exit, and strategic pricing. I model strategic behavior using the moment-based Markov equilibrium from Ifrach and Weintraub (2012), extended to reflect an international market with learning-by-doing. Simulations with estimated real-world coefficients suggest that ITAR drove demand towards to foreign competitors, learning-by-doing explains the persistence of the shift, and ITAR has additional lasting effects through the extensive margin.

My framework for analyzing competition in reliability generalizes to any industry where firms are differentiated in reliability and product “bugs” are discovered and ameliorated during full-scale production. I find that a single value, the “demand returns to experience”, estimatable from sales and accident data, determines the long-term fate of the market.

1.1.1 Previous Literature

1.1.1.1 Natural Experiments in Infant Industry Protection

The idea that temporary government intervention can help a region develop a new industry dates back to Hamilton (1791), and has more recently been supported by stylized models in Krugman (1987) and Bartelme et al. (2018). Quantifying the actual efficacy of such a policy is more difficult:

1. An industry may simultaneously receive multiple types of government support.
2. Infant industry protection may be implemented at a time of broad-based economic

development: the sector may have developed regardless of intervention, as seen in Irwin (2000)

3. Infant industry protection may be targeted at a sector with exogenously high potential

One way to address this is to consider natural experiments, cases where a trade barrier was imposed for reasons other than industrial development. For example, Juhász (2018) shows evidence that Napoleon’s punitive sanctions on British goods inadvertently supported a new cotton weaving industry in the French Empire, one which outlived these sanctions. Hanlon (2015) looks at a foreign policy, namely the US Navy’s blockade on the Confederacy, which forced UK manufacturers to substitute high-grade American cotton for lower-grade inputs from India. The inventions and investments developed to do so remained after the Civil War, allowing India to retain its improved prices and market share afterwards.

A second class of natural experiments consider domestic interventions as accidental “subsidies” for infant industries. For example, Hansen, Jensen, and Madsen (2003) show that Danish environmental laws inadvertently created a burgeoning export-oriented windmill industry. Liu [WP] [CITE] shows that an abortive attempt to create an indigenous weapon industry in 19th-century China helped some regions establish manufacturing centers for other goods. Ferraz, Finan, and Szerman (2015) use a random-assignment mechanism in government procurement bidding to show that government contracts help small businesses expand into new regions.

I build on this work in two ways. First, the infant industry support I consider is not just inadvertent, but backfiring. The primary goal of the US export restriction was to prevent countries like China from acquiring improved rocket technology, so evidence that it encouraged the Chinese space program is especially compelling. Second, instead of national-level data, I use a global census of suppliers, so that the policy’s effects can be observed on both the entrants and incumbents.

1.1.1.2 Subsidies

By contrast, Kalouptsi (2018) considers an export promotion program of unknown scope but clear timing, namely a broad-based push by the Chinese government to expand ship production. By estimating the dynamic decision-making of both shipbuilders and shipowners, she estimates the total impact of the various subsidies and regulatory changes.

1.1.1.3 Dynamic Equilibrium

This paper also fits into a larger literature on computing dynamic competitive equilibria, starting with Ericson and Pakes (1995) and summarized by Doraszelski and Pakes (2007). The archetypal model considers a firm choosing entry/exit, price, and paying for an investment to improve quality for the future. By contrast, I consider firms “investing” in future profits by choosing prices away from those that optimize static profits. Mahone and Rebessi (2019) has firms raising prices to affect consumer beliefs about quality. This paper is closer to Benkard (2004), which considers learning-by-doing to reduce costs in the wide-body aircraft industry. I build on this framework by adding learning-by-doing in product *quality* and introducing observed product differentiation. To maintain tractability, I set aside forgetting.

1.1.2 Rockets and Economics

There is a small but growing literature addressing the business of spaceflight. Perhaps the closest to this paper is Boone (2017) and its companion piece Boone IV and Miller (2017). Boone considers launch vehicle selection from the perspective of a social planner: given launch vehicle capabilities and prices, how can payloads be assigned to vehicles to minimize cost? Between 2000 and 2013, he calculates that 654 tons of excess capacity on orbital launches was left unused, worth \$8.72 billion (2014). Furthermore, roughly 50% of missions use a vehicle which is not ranked in the top 4 cheapest capable vehicles, leading to losses of up to \$19.1 billion, or 43.8% of total launch costs. By implementing a bin-packing algorithm

for sharing payloads between more cost-effective rockets, the cost of orbital launches can be reduced by 19% holding supply constant, or 53% if launches can be delayed and additional units of cost-effective vehicles can be supplied. Boone suggests that future work could incorporate non-cost drivers of vehicle choice, such as reliability, home bias, and trade sanctions, which serve as a focal point of this paper.

Xu, Hollingsworth, and Smith (2019) takes the additional step of adding reliability into a model of vehicle choice. They consider the problem of an operator looking to launch one or more satellites, allowing the planned system to have redundancy (multiple interchangeable satellites) and/or fractionalization (splitting essential components across multiple satellites), and develop an algorithm to choose the optimal combination of launch vehicles. Their measure of reliability uses a 2-step Bayesian estimator, assuming that launch failures are drawn from a stationary Bernoulli distribution. Cost and reliability are related via an intuitive but restrictive form: the total cost of a launch is $c_l \times n_l(\alpha)$, where n_l is the expected number of launches necessary for a single success with $\alpha\%$ confidence, $1 - (1 - r)^{n_l} \geq 1 - \alpha^1$. Based on these input parameters, they use mixed integer nonlinear programming (MILNP) to assign three historical satellite programs to vehicles with lower expected costs. By contrast, I treat reliability as a non-stationary Poisson process to explicitly model trial-and-error improvements, and estimate the relative preference for reliability and cost from consumer choices. This follows from my focus on the producers of launches, rather than the customers.

Satellite insurance can reach 20% of the cost of a mission, and is well-studied by actuarial literature. Gould and Linden give a background on pricing satellite insurance and estimating liabilities. Manikowski and Weiss (2013), like this paper, takes advantage of the size of the satellite market, which is both large enough for interesting variation and small enough that the global census of transactions can be feasibly collected. Using data from 1968 to 2010 on rates-on-line, capacity, and loss ratios, they find evidence that insurers set premiums based

¹This definition of “reliability-inclusive price” can be reconstructed in my own data $n_l c_l = \frac{\ln(\alpha) PRICE_l}{\ln(MTBF_l)}$

on data from previous losses (the “rational expectations hypothesis”) and charge higher premiums as the market approaches maximum capacity (the “capacity constraint hypothesis”).

Finally, I found several works on space policy instructive for this project. Kutter (2006) suggests that scale, experience, and competition are critical for decreasing space launch costs, and suggests that NASA should auction contracts for future missions to “commercial” launch vehicles like the *Atlas*, rather than relying on large, expensive government-operated vehicles like the *Space Shuttle*. López, Pascuini, and Ramos (2018) and Pekkanen and Kallender-Umezu (2010) describe the incremental advance of the space industries in Argentina and Japan. They outline a “space technology ladder”, where countries first purchase foreign equipment for space services, then develop indigenous satellites, and finally create their own indigenous launch vehicles, with each step of the ladder providing technological spillovers and home-market demand for the next step. Most recently, Triezenberg et al. (2020) argues that, given US demand and a global commercial demand steady at around 20 launches per year, the US can only hope to sustain 2 heavy lift launch vehicles in the long run. One extension of this project will be to directly test this hypothesis using my own learning-by-doing dynamic model.

1.2 Institutional Details

1.2.1 Data

My primary source of data is Gunter’s Space Page, a website which compiles data from a variety of press releases and historical sources to create a digital census of all satellite launches since *Sputnik* in 1957, with satellite name, mass, destination orbit, operator, manufacturer, launch vehicle, launch site, and often additional details like power supply and propulsion system. This is matched to vehicle-variant-level data from Boone (2017), including maximum payload to various orbits. I extend this data to include price variation over time using

Table 1.1: Data Overview

Satellites	8,350
Launches	5,697
Success Rate	91.3%
Partial Success Rate	93.2%
Avg. satellites/launch	1.367

Table 1.2: Satellite Operators/Customers

Operator Category		Number
Corporate	Geospatial	501
	Telecom	704
	Other	222
	Total	1,427
Government	Space Agency	1,164
	Other Civil	246
	Intelligence	237
	Military	1,774
	Total	3,813
University		392
Total	5,240	

Federal Aviation Administration annual reports from 1997-present (FAA, 2019); these compile the best publicly-available data supplemented with estimates by industry experts where necessary.

The data includes military, civil government, and commercial customers. While military and intelligence community customers may choose a launch service provider for geopolitical reasons, we expect multinational corporations to select strictly on cost and quality. On my learning regressions, military and civil government launches provide valuable data, while in the full model they serve as an important “captive market” driver of growth.

1.2.2 A Note on Taxonomy

I abstract the market as a panel of firms (providers), each producing a single vehicle family. A vehicle family, like *Ariane*, shares core components like avionics and engine, as well as

production, assembly, and launch facilities. Given this overlap, I assume that the experience needed to streamline production and prevent accidents accumulates at the vehicle family level.

I will use f to denote the “variant” of the vehicle, including accessories. Variant *Ariane 42L-H10*, for example, denotes a 4th generation vehicle with 2 liquid strap-on boosters and the *H10* upper stage. Variants have a specific set of payloads they can accommodate and a specific price, so I match the demand model at the variant level.

One possible extension is to test for spillovers between different rocket families produced in the same country or marketed by a single firm like United Launch Alliance. In practice, countries with multiple simultaneously-active indigenous rocket families usually have very little overlap in technology or facilities. In the early 2000s, for example, Japan launched the *H-2*, a liquid-fueled medium-lift vehicle from Mitsubishi Heavy Industries launched from Tanegashima, and *Mu-V*, a solid-fueled light-lift vehicle from Nissan and IHI Aerospace launched 97 kilometers away at Uchinoura. By contrast, I treat *Lambda* and *Mu*, with shared technology, facilities, and manufacturer, as a single family F .

1.2.3 Trade Friction Shock

Space technology is closely related to military technology. Space launch vehicles like the *Soyuz* and *Atlas* were developed directly from intercontinental ballistic missiles and frequently used to launch military surveillance and communications payloads (Hill, 2012). At the same time, they have had an expanding role in delivering commercial and civil government satellites: governments have struggled to find a balance between export promotion and arms nonproliferation.

US space technology was historically included on the US Munitions List (USML), subject to International Traffic in Arms Regulations (ITAR). This list includes items like missiles and high-grade night-vision goggles; their export requires a licence from the Department


of State. Following the end of the Cold War, space technology was gradually shifted onto the Commerce Control List (CCL), subject to Export Administration Regulations (EAR). These items are recognized as “dual use”, with both military and nonmilitary applications, such as unarmed drones with a maximum endurance of greater than 1 hour and high-power marine gas turbine engines (BIS, 2020).

In 1996, a US-built satellite, *Intelsat 708*, was launched on a Chinese *Long March (CZ) 3B* rocket in Xichang, China (Zinger, 2015). The vehicle almost immediately veered off course, crashing to the ground in a nearby village and killing six people². Following the accident, a Chinese committee investigated the crash, while a separate investigative committee, including US experts, was assembled by the satellite’s insurance providers (Zinger, 2015). In 1997, US government investigators found that information passing between these committees effectively helped the Chinese rocket industry improve their *Long March* vehicles, and fined Space Systems Loral \$20 million. Congress, however, alleged that the information shared could be used to improve Chinese ballistic missiles; and called for drastic changes. The “Strom Thurmond National Defense Authorization Act for Fiscal Year 1999” shifted all US space-related products, services, and knowledge back onto ITAR.

ITAR compliance represents a substantial burden. First, all US technology is banned from export to China. Second, firms located in other foreign countries are required to apply to the Department of State for an export license, which could take up to 3 months. Business partners from any third country must apply for a separate license to export to that country. In addition to physical products, the USML includes software, blueprints, and expertise. Expertise is loosely defined: one NASA employee interviewed said that a foreign national merely glancing at the the user interface of an unmanned probe might be considered an “export of expertise”. Other aerospace professionals explained that all foreign nationals must wear lanyards to restrict their access to sensitive areas of SpaceX facilities. Restrictions on sharing technical also causes difficulty in securing insurance for a

²This number is disputed, with international groups speculating additional casualties


Figure 1.1: Recent evolution of market shares



launch/satellite. Applying for ITAR and complying with its rules will require a firm to hire a team of lawyers specializing in ITAR or seek outside counsel; failing to comply could lead to fines of up to \$100 million and imprisonment. Taken together, these risks represent a substantial burden on any potential client of the US space industry, foreign or domestic.

Citing massive losses, the aerospace industry successfully lobbied Congress to reverse the Strom Thurmond policy, returning space launch services to Commerce Controls in 2013. Surprisingly, however, US market share has only partially recovered. The new space industry in India and China appears unlikely to go away; a policy meant to curtail Chinese rocket development may have given it an unexpected boost.

Figure 1.2: Data from NASA Launch Performance Calculator



1.3 Demand Side

1.3.1 Capability

Launch vehicles vary dramatically in payload capacity, with the Chinese *Kaitouozhe-1* launching just 100 kg to low earth orbit compared to 25,700 kg on the *Space Shuttle*. I assume a customer exogenously requires a satellite with certain capabilities, hires a satellite manufacturer to design a suitable satellite with as little mass as possible, and then begins “shopping around” for orbital launch vehicles capable of launching their satellite to the desired orbit. Some LSPs advertise their payload capacity for a wide range of orbits (see graph from NASA), but it is more common to show data for just a few common orbits:

- **Low Earth Orbit (LEO)** consists of circular orbits between 200 and 2000 km in alti-

tude, the “edge of space” used by most satellites and the International Space Station.

- **Geosynchronous Orbit (GEO)** is a circular orbit at exactly 35,786 km and directly over the equator. At this altitude, satellites appear to hover above a fixed point on Earth’s surface, a desirable feature for communications. This orbit, however, requires much more fuel per kilogram.
- **Geosynchronous Transfer Orbit (GTO)** consists of an elliptical orbit from LEO to GEO. A satellite will typically use a launch vehicle to reach GTO, then separate and fire its own thrusters to reach GEO.

I focus on the 90% of satellite launches in the data which target either LEO or GTO orbits. One extension of this paper will be utilizing Tsiolkovsky’s Equation to extrapolate capability to other orbits like those used by GPS satellites and interplanetary probes. Note that the relationship between maximum payload to LEO and GTO is nonlinear, depending on the details of the vehicle’s staging and the dry mass of the final stage (see appendix).

Due to incomplete data, I ignore the role of orbital inclination and latitude, which offer an additional source of horizontal differentiation. Tropical spaceports like Arianespace’s Kourou, French Guiana (5° N) are better-suited for reaching equatorial orbits and maximizing mass-to-orbit. By contrast, spaceports like Kosmodrom Plesetsk (63° N) and PSC Alaska (57° N) are more efficient for “high-inclination” orbits over polar regions.

1.3.2 Data Details

I restrict the sample for fitting the demand model to launches between 1995 and 2018. This period is covered by the FAA price estimates, and includes years before (1995-1998), during (1999-2013) and after (2014-2018) space launches were included on ITAR.

I match this to annual price estimates for all vehicles from the FAA Year-in-Review, later renamed the Annual Space Compendium. Where the FAA provides a range of price values, I

choose the midpoint: the 2000 *Ariane 5* is listed at \$150-180 million, which I write as \$165. In the earlier data, prices are only included for vehicles available in the current year; if a vehicle is unlaunched but available, I use the last-listed price.

Industry analysts note that national security launches are often charged a premium for priority in the queue, increased security, etc., compared to the “commercial” prices provided by the FAA data. If all firms charge governments a multiplicative markup x then $\ln(P_{GOV,f}) = \ln(x) + \ln(P_f)$: the customer-level constant term will absorb the premium, so estimates using commercial prices will be unbiased.

I discard 24 observations of vehicles which have only been launched as prototypes, such as the South Korean *Naro*.

I use the measures of bilateral trade frictions from Mayer and Zignago (2011) including language, former colonial relationships, and geographic distance. Note that distance is measured by average distance between cities weighed by population, so that $DIST > 0$ even for domestic transactions. When the customer is listed for multiple countries, I give priority to any country with its own indigenous launch vehicle. I use the headquarters locations to place the European Space Agency and the Intelsat Corporation in France and Luxembourg, respectively.

My initial analysis ignores satellites which constitute less than 50% of the cargo for the mission, since “piggybacking” customers are limited to choosing from vehicles with available space; see below for details on factoring in vehicle sharing into pricing.

1.3.3 Commercial Availability

Previous analysis on this subject carefully documents which rockets are advertised as “commercially available”, and which satellites are “commercially competed”; Triezenberg et al. (2020), for example, concludes that only 20 launches each year are available for international bidding. I take a more liberal approach, assuming that any vehicle with a published



price estimate which is still in production is commercially available, and that each customer considers all rockets capable of their mission and not prohibited by ITAR.

I do this for two reasons. First, I theorize that the choice to commercialize a vehicle or open a satellite launch for commercial bidding is an endogenous decision. For example, while Israel’s *Shavit 2* has exclusively launched Israeli reconnaissance satellites, Israel Aerospace Industries has repeatedly explored commercialization, but struggled with the vehicle’s poor efficiency owing to its unusual westward launch trajectory. I assume that vehicles which are not made available for commercial launches would not have won any contracts if actively marketed. Likewise, I assume that launches not identified as “commercial” prefer the characteristics of their chosen vehicle to all commercial alternatives.

Second, private negotiations, involving vehicles which are not “commercially available” and/or customers who are not conducting a transparent “commercial” bidding process are still important, and the mere possibility of a transaction can have implications for the intensive and extensive margin of the market. For example, in the 1960s the UK government noted that American rockets may become available for launching British satellites: they promptly began negotiating prices with the US and shutting down their own indigenous launch vehicle program (Hill, 2011). By contrast, failed negotiations to acquire an affordable former ICBM in Russia inspired Elon Musk to found SpaceX (Vance, 2015). I argue that decision-makers will assume that any vehicle which proves safe and cost-effective will eventually become commercially available.

1.3.4 Logit Model

Each year t produces a new cohort of satellites $s \in S_t$. The satellites are defined by their mass, destination orbit, region, and operator type (m_s, d_s, r_s, T_s) . The choice set is limited by mass and destination orbit (the rocket must be physically capable) and origin region (the rocket cannot be blocked by trade sanctions): $f \in C_t(m_s, d_s, r_s)$



Given customer choices f_{st}^* , I fit the following model:

$$f_{st}^* = \operatorname{argmax}_{f \in C_t(m_s, d_s, r_s)} u_{sft}$$

$$u_{sft} = \alpha_0 + \alpha_p \ln(p_{ft}) + \alpha_{ITAR} + \alpha_R \ln(R_{ft} ITAR(r_s, F)) + \alpha_{DIST} DIST(r_s, F) + \epsilon_{sft}$$

where

- p_{ft} is the price of the launch vehicle
- R_{Vt} is the updated estimated probability of launch success for the vehicle family
- $DIST(r_s, F)$ is a vector of measurements of “distance” between r_s and F from Mayer and Zignago (2011)
- ϵ_{sft} is an iid extreme value error term

Realized demand for the year is then added to the firm’s experience for next year:

$$E_{F,t+1} = E_{F,t} + \sum_{s \in S_t} \mathbb{1}\{f_{st}^* \in F\}$$

$$\mathbb{E}[E_{F,t+1}] = E_t + \sum_{f_{st}^* \in F} \frac{e^{u_{sf^*t}}}{\sum_{f' \in C_t(m_s, d_s, r_s)} e^{u_{sf't}}}$$

By fitting the model using micro-level data, we control for fluctuations in aggregate demand.

Satellite operator “Types” consist of *MILINT* (military, intelligence, nuclear early warning systems, etc.), *CIVGOV* (civil government agencies, space exploration, weather service, etc.), and *CORP* (commercial telecommunications and earth observation). All coefficients are interacted with *Type*: for example, the US Air Force has a strong preference for a US-launched vehicle, while DirectTV, a multinational corporation headquartered in the US, uses a variety of international vehicles.

1.3.5 Results

Table 1.3 gives results, disaggregated by customer category. All customers have a strong preference for more-affordable vehicles, vehicles from their home country, and vehicles unaffected by ITAR.

Reliability has a significant positive effect on demand, driven primarily by corporate customers. A success rate increasing from 90% to 95% will have the same impact on corporate demand as a 29% decrease in prices³. Interestingly, government demand is not significantly affected by reliability. One possible explanation is that a national government with an untested new indigenous rocket will internalize the benefits of providing more experience to their domestic product. Indeed, column (6) shows that national governments *without* domestic rocket programs are actually *more* sensitive to reliability than their private-sector counterparts. For my remaining analysis, I focus on sample of private-sector customers and government agencies of countries without indigenous launch vehicles.

Table 1.4 looks at a more detailed set of measurements of “distance” between countries from Mayer and Zignago (2011). Home bias is present and significant: producing a rocket domestically is as desirable as reducing its cost by 57%.

The cost of hiring a US-based launch provider under ITAR is estimated to be equivalent to a stunning 303% tariff⁴. Interestingly, there is no significant difference between foreign and US-based companies in this respect. This suggests that the “third country national” elements of ITAR are the most onerous, eg restricting the transfer of technical information from insurance providers, employees, partners and subsidiaries of any nationality besides the US and that of the customer.

Trade is not significantly restricted in the opposite direction: US-based corporations like AT&T use foreign vehicles like *Zenit* (UKR), *Proton* (RUS), and *Ariane* (EU) throughout

$$3 \frac{(1-0.95)^{-1} - (1-0.9)^{-1}}{(1-0.9)^{-1}} * \frac{0.543}{1.872} = .29$$

$$^4 \text{Letting } u(P_1, ITAR) = u(P_2, ITAR), e^{\frac{\alpha_{ITAR}}{\alpha_P}} = \frac{P_1}{P_2} = e^{\frac{2.83154}{2.0304}} = 4.03$$

Table 1.3: Demand results, disaggregated

VARIABLES	Demand, 1995-2020					
	(1) All	(2) Corporate	(3) Civil Gov	(4) Mil/Intel	(5) All Gov.	(6) All Gov + No dom
Log Price (2014 USD, Boone)	-1.646*** (0.0713)	-1.872*** (0.133)	-1.958*** (0.164)	-1.226*** (0.174)	-1.743*** (0.0938)	-1.969*** (0.233)
Reliability (log MTBF)	0.218*** (0.0706)	0.543*** (0.121)	0.197 (0.157)	-0.0529 (0.211)	0.137 (0.0947)	0.674*** (0.210)
Same country	3.134*** (0.113)	1.561*** (0.162)	4.819*** (0.296)	7.246*** (1.347)	4.181*** (0.177)	
US vehicle under	-1.943*** (0.150)	-2.452*** (0.230)	-0.828* (0.429)	-2.271** (1.046)	-1.391*** (0.229)	-1.823*** (0.453)
Observations	24,123	7,498	7,092	3,242	16,695	2,769

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

the ITAR period. It is possible that hiring a foreign rocket for a US payload may pose less of a compliance risk than hiring a US-based rocket for a foreign payload. Alternatively, US-based launch vehicle brokers like International Launch Services may be able to ensure compliance for customers.

Unsurprisingly, all customers prefer to purchase from geographically closer launch service providers, and NATO members prefer to use US and ESA launch vehicles for military and intelligence satellites.

As noted previously, I see little room for unobserved product quality: national origin, cost, capability, and reliability are consistently cited by industry analysts as the only drivers of demand for these delivery services, and prices are relatively stable over time in the data. Nevertheless, it is potentially hazardous to assume that these correlations are unbiased estimates of consumer preferences. As this project continues, I have three plans for addressing this. First, I use a structural model to build intuition on the actual drivers of strategic pricing in this setting. Second, I will develop an IV formulation using drivers of aggregate

Table 1.4: Demand model with “distance” measures

VARIABLES	Demand and trade friction	
	(1) chosen	(2) chosen
Log Price (2014 USD, Boone)	-2.037*** (0.109)	-1.922*** (0.117)
Reliability (log MTBF)	0.873*** (0.106)	0.541*** (0.106)
Same country		1.758*** (0.183)
US vehicle + ITAR		-2.749*** (0.770)
US vehicle + foreign cust. + ITAR		0.847 (0.806)
US cust. + foreign vehicle + ITAR		0.404 (0.849)
Military, within NATO+AUS		1.906*** (0.613)
1 for pairs ever in colonial relationship		
Log distance (km)		
1 for contiguity		
1 if a language is spoken by at least 9% of the population in both countries		
1 for common official of primary language		
Observations	10,197	10,197
Pseudo R2	0.148	0.233

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Excludes governments with indigenous vehicles

VARIABLES		(1) Vehicle chosen	(2) Vehicle chosen
Log Price (2014 USD, Boone)	Corporate	-2.117*** (0.140)	-2.117*** (0.140)
	Civil	0.940** (0.396)	0.940** (0.396)
	Space Agency	-0.420* (0.246)	-0.420* (0.246)
	Mil/Intel	1.106*** (0.215)	1.084*** (0.217)
Reliability (log MTBF)	Corporate	0.392*** (0.123)	0.392*** (0.123)
	Civil	0.612 (0.507)	0.612 (0.507)
	Space Agency	-0.865*** (0.214)	-0.865*** (0.214)
	Mil/Intel	-1.161*** (0.235)	-1.225*** (0.241)
US vehicle under ITAR = 1	Corporate	-3.022*** (0.229)	-3.022*** (0.229)
	Civil	-0.881 (1.983)	-0.881 (1.983)
	Space Agency	3.321*** (0.581)	3.321*** (0.581)
	Mil/Intel	1.206 (1.164)	0.129 (1.369)
Same country = 1	Corporate	0.661*** (0.216)	0.661*** (0.216)
	Civil	6.845*** (1.778)	6.845*** (1.778)
	Space Agency	-0.307 (0.788)	-0.307 (0.788)
	Mil/Intel	3.200** (1.387)	3.938** (1.609)
Log pop-weighted distance (km)	Corporate	-0.381*** (0.0895)	-0.381*** (0.0895)
	Civil	0.394 (0.661)	0.394 (0.661)
	Space Agency	-1.679*** (0.393)	-1.679*** (0.393)
	Mil/Intel	-1.024*** (0.392)	-0.532 (0.503)
	Military within NATO+AUS		1.201 (0.853)
Observations		30,098	30,098

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

demand (tech sector performance, government research spending) and drivers of aggregate supply (prices for kerosene, hydrazine, and aircraft-grade aluminum). Third, I plan to exploit the overlapping capability ranges of different vehicles: the price of a vehicle capable of launching 6-9 tons will induce a price change in one capable of 8-11 tons, which will serve as a plausibly exogenous shock to competitor prices for a vehicle launching 10-12 tons.

1.4 Duplex Launches

Rockets are increasingly launched with multiple payloads on board. As Boone IV and Miller (2017) suggests, this presents an opportunity for substantial cost savings. In 37% of cases this consists of a multi-ton primary payload sharing a fairing with several small CubeSats. These 10 cm cubes are often launched pro-bono for educational institutions, have relatively negligible mass, and are under strict scheduling and design limitations so they do not interfere with the primary customer⁵.

By contrast, roughly 20% of launches in the matched data are split between two large payloads⁶. A “duplex” launch requires a larger and more expensive vehicle variant, but costs less per kilogram than launching the satellites separately on two smaller rockets. In some cases, a larger variant by the same launch service provider appears to be chosen specifically to accommodate a second payload. To cover the cost of a vehicle with higher capacity, compensate the customers for the coordination costs of sharing a single launch date and mission profile, and incentivize choosing a cheaper per-kg option, it makes sense for the rocket provider to offer a discount for shared launches.

This presents two problems. First, the data only includes prices for solo launches. Second, it would be incredibly computationally burdensome to consider all possible combinations of

⁵For example, CubeSats cannot contain volatile fuels, and the CubeSat operator must provide a dummy mass to replace their satellite in case of production delays, so that the primary customer’s launch proceeds on-schedule.

⁶Splits between 90-10 and 50-50

vehicle variant, primary payload, and secondary payload in each year. In fact, finding the optimal social planner’s allocation (a bin packing problem with variable bin capacity and cost) is an NP-complete problem (Boone & Miller, 2017).

In my main analysis I focus on primary payloads ($> 50\%$ of mass). I consider several models for discounting the price of a rocket for the primary customer in a duplex launch. For simplicity of notation, consider primary and secondary payloads A, B of mass m_A, m_B . Assume A is considering a vehicle with solo-launch price p_v and maximum payload to the desired orbit $\bar{m}_v$. The discounts are ordered from most- to least-generous to a customer choosing a shared launch on a larger vehicle:

1. “À la carte model”. The most generous model is to only price primary customers for the payload capacity actually used. Given the small number of duplex launches, this incentivizes waste, but serves as an upper-bound estimate for popularity of large vehicles.

$$p_{Af} = \frac{m_A}{\bar{m}_f} p_f \tag{1.1}$$

2. “Roommate” model. For each pair of satellites A, B which *historically* shared a duplex launch, the price of any rocket which will accommodate both is shared between the two proportionate to mass. These are compared to vehicles which can only carry A .

$$p_{Af} = \begin{cases} p_f * \frac{m_A}{m_A + m_B} & \text{if } m_A + m_B \leq \bar{m}_f \\ p_f & \text{otherwise} \end{cases}$$

This is rather stylized, since the co-payload B is likely a function of the vehicle chosen. On the other hand, it acknowledges that unobserved factors like timing, mission profile, and security considerations likely link A to B .

3. “UberPool” model. The customer pays for their expected portion of the total payload

given the excess capacity of the vehicle. We could assume that this is the actual price paid, as in ridesharing services, or that customers have preferences over *expected* price paid, and are uncertain if a duplex option will be offered. This requires the strong assumption that all customers are equally willing to accommodate a co-payload.

$$p_{Af} = E \left[\frac{m_A}{m_A + m_{B'}} \middle| \frac{m_A}{\bar{m}_f} \right] p_f$$

4. “Complimentary upgrade” model. This assumes that A selects a vehicle for a solo launch, and is subsequently offered a shared launch on a larger vehicle produced by the same firm for the same price.

$$p_{Af} = \min_{f' \in f} \{p_{f'} \text{ such that } m_A \leq \bar{m}_{f'}\}$$

5. No discount

$$p_{Af} = p_f$$

Table 1.5 shows implementations of all five pricing models for shared launches, which all give fairly consistent results, with the exception of Column 2, which restricts the sample to launches where $\frac{m_A}{m_A + m_B} > 0.95$.

Table 1.5: Demand with pro-rated prices

Demand with hypothetical duplex discounts					
VARIABLES	(1) chosen	(2) chosen	(3) chosen	(4) chosen	(5) chosen
Reliability (log MTBF)	0.512*** (0.106)	0.121 (0.121)	0.497*** (0.107)	0.549*** (0.108)	0.549*** (0.108)
Same country	1.162*** (0.198)	1.420*** (0.226)	1.213*** (0.198)	1.084*** (0.202)	1.084*** (0.202)
Military, within NATO+AUS	1.880*** (0.649)	1.153 (1.156)	1.943*** (0.666)	2.046*** (0.704)	2.046*** (0.704)
Log distance (km)	-0.203*** (0.0672)	-0.0769 (0.0804)	-0.193*** (0.0678)	-0.232*** (0.0676)	-0.232*** (0.0676)
US vehicle + ITAR	-2.319*** (0.211)	-2.091*** (0.227)	-2.384*** (0.214)	-2.325*** (0.213)	-2.325*** (0.213)
Log Price (2014 USD, Boone)	-2.007*** (0.122)	-2.424*** (0.150)			
Log price, prorated by expected capacity utilized			-2.018*** (0.128)		
Log price, prorated by capacity utilized if possible				-2.402*** (0.133)	
Log price, assuming complimentary upgrade					-2.402*** (0.133)
Observations	10,130	7,943	9,658	10,130	10,130
Pseudo R2	0.234	0.278	0.228	0.270	0.270

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Excludes governments with indigenous vehicles

1.5 Supply Side

1.5.1 Reliability Growth


At the beginning of each year t launch providers and customers are assumed to observe the instantaneous estimate of reliability of each rocket family, R_{Ft} , a function of the probability that the chosen rocket will deliver their satellite to the desired orbit in working condition.

I assume that reliability is determined at the vehicle family level, since there is significant overlap in technology, procedures, and personnel between successive generations of vehicles within each family. For example, all variants of *Soyuz* flown since 1966 share the same kerosene-oxygen fuel, RD-107 engines⁷, staging, and mission profile, and are shipped from one factory in Samara, Russia to just four launch sites around the world.

The data consists of a sequence of launch success dummies for each vehicle family, with more than 8000 observations in total since 1957. I fit a parametric model to the data, rather than a naive kernel, for the following reasons:

1. Reliability should grow monotonically, as later accidents reflect preexisting undiscovered design flaws
2. Relatedly, potential customers should reward “proven” designs with many observed launches
3. Accidents, and ensuing improvements, arrive stochastically; the reliability estimate should reflect a smooth expectation over the distribution of failure events.
4. A smooth reliability function gives us the marginal reliability improvement from additional launches
5. A parametric model allows us to extrapolate reliability for counterfactual scenarios: what will SpaceX’s expected success rate look like with 200 additional launches?

⁷With the exception of the new *Soyuz 2.1v*



The literature on reliability statistics suggests a simple visual test: the “Duane plot” graphs the accumulated experience versus the average time between failures on log-log scale. An upward-sloping linear function indicates trial-and-error reliability improvement (Dummer, Winton, and Tooley, 1997).

Crow (1975) lays out a theoretical model consistent with with relationship. This “Crow-AMSAA” model, used in software testing, quality assurance, and NASA safety analysis (Dawson, 2011), provides a maximum likelihood estimator and confidence intervals, detailed in the appendix. I make the following assumptions to apply this model to this setting:

1. Launch events represent roughly-equal units of time spent testing the system
2. The vehicle begins with an unknown number of faulty components
3. Faulty components fail via a Poisson process; failure of the component results in failure of the launch
(components are “fail-deadly”)
4. After a failed launch, the culprit component is located and replaced with a faultless

component

5. Launches resume, with all future launches incorporating the improved component.
6. All accidents are attributable to flaws in the product
7. Launch vehicles are identical except for experience-based design improvements

Note that the assumption of a discrete set of fail-deadly components determining product quality is similar to the “O-Ring” theory in Kremer (1993), which takes its name from this setting. Kremer allows the firm to control the defect rate by choosing worker skill levels. By contrast, in my model inexperience drives defects, so that a firm can only improve success rates by building more units.

Item 7 is justifiable if we define the “product” to include safety procedures which mitigate “acts-of-god”. For example, while inclement weather contributed to the loss of the Space Shuttle *Challenger*, after thorough investigation NASA changed launch procedures for future launches to avoid such weather (Commission, 1986).


Item 6 is most dubious; the upgraded *Long March 2E*, for example, proved to be more dangerous than earlier versions of the Long March family and was ultimately discontinued. A future extension of this model will explore the firm’s decision to introduce an upgraded model which expands capability and/or reduces costs, but may decrease reliability.

For mathematical simplicity, “reliability” is defined as the expected number of consecutive launches between failure events, equivalent to Mean Time Between Failures (MTBF) in the NHPP literature:

$$R_{Ft} = E_{Ft}^{\theta} R_{F0} \tag{1.2}$$

where θ is shared across all vehicles, while R_{F0} is a vehicle-family-specific intercept drawn

Figure 1.3: Fitted Crow-AMSAA results



when a new vehicle family is introduced and revealed by observing the success rate of launches over time⁸. R_{F0} can reflect a robust design⁹, rigorous ground testing, and spillovers from using proven components and materials from previous space and missile programs.

Note that this model this model also allows for reliability improvements from *successful* missions in addition to failures: Space Shuttle solid rocket boosters, for example, were recovered after each launch and inspected for evidence of degradation (Moore and Phelps, 2011).


Figure 1.3 shows results for 19 vehicle families with sufficient data alongside kernel-smoothed success rates. The Crow model gives a close approximation of the success data, deviating from the kernel when launches exhibit a string of fortune or misfortune at the very beginning or the very end of the series.

Figure 1.4 shows the growth of expected success rate over time. Note that starting levels

⁸Note that $\theta \equiv 1 - \beta$ and $R_{V0} \equiv 1/(\lambda * \beta)$ in the conventional Crow-AMSAA model notation

⁹For example, the *Falcon 9* has proven capable of completing a mission with the loss of 1 of its 9 first-stage engines, while the Soviet *N1* required all 30 first-stage engines to operate without failure in practice.

Figure 1.4: Evolution of reliability




R_{0F} are higher for later vehicles like *Ariane* and *Falcon*, but they still require significant experience to match the reliability of incumbent firms. This suggests that safe rocket designs and procedures are the result of both publicly-available scientific literature and classified, vehicle-specific, or institutional knowledge.

1.5.2 Production

While the reliability model is applicable to a variety of settings where commercial products are used for “debugging”, the production model is more specialized to the setting of space launches noted in Cates et al. (2018), namely:

- Detailed factors of production are unavailable for most vehicles
- By contrast, product delivery dates are publicly-available
- Production is bottlenecked by a small, fixed number of expensive facilities.
- Customers vary in preference for timeliness: while the US military requires payloads

Figure 1.5: Recent evolution of reliability



to be launched “whenever such payloads are needed in space” (Congress), cubesat operators can wait up to 3 years to “hitchhike” on another customer’s rocket.

Production requires access to expensive specialized facilities with limited capacity, such as “clean rooms” for payload integration, massive vehicle assembly buildings, and isolated launchpads, which we treat as fixed capital¹⁰ $\bar{K}$: this puts an upper bound on the amount of labor, L , which can be employed at one time. Outsourced components and materials are represented by M : we assume these can be purchased at will to complete orders. The annual production rate is thus

$$q_{tF} = \min \left\{ \min \left\{ \bar{K}, L_{tF} \right\}^\zeta E_{tF}^\eta, M_{tF} \right\}$$

¹⁰The construction of new space launch facilities is remarkably rare in data. For example, *Falcon 9* is launched from complexes first constructed in the 1960s at Vandenberg AFB and Cape Canaveral, including the Apollo Program’s SLC-39A.

with total cost

$$C(q_t) = r\bar{K} + cM + wL$$

1.5.2.1 Rush orders

Some orders are purchased by the home-country government for military or intelligence satellites: this enters as an exogenous shock which forces the firm to fulfill the order from $\bar{o}s$ at maximum capacity $L = \bar{K}$. The annualized production rate in the months before these orders are fulfilled is thus

$$\bar{q} = q_{\bar{o}s} = \frac{1}{T_{\bar{o}s}} = \bar{K}^\zeta E_{Vt}^\eta$$

This relationship is used to estimate η .

It should be noted that this methodology is consistent with the popularity of launch rate as an indicator of efficiency. For example, the Ansari X Prize required companies to perform two launches in two weeks using the same reusable vehicle, the DARPA 2018 Launch Challenge called for two launches of disposable vehicles quick succession, and SpaceX frequently cites launch rate as a metric for progress (X-Prize, 2020).

1.5.2.2 Below-Capacity Orders

When production capacity is not binding, firms are free to hire just enough labor L to complete orders:

$$\text{for } q \leq \bar{q} : L_{tF}^*(q) = \left(\frac{q}{E_{tF}^\eta} \right)^{\frac{1}{\zeta}}$$

$$C(q_{tF}) = r\bar{K} + cq_{tF} + w \left(\frac{q_{tF}}{E_{tF}^\eta} \right)^{\frac{1}{\zeta}}$$

Experience thus confers two distinct supply-side advantages: firms have larger production capacity $\bar{q}$, potentially allowing them to spread fixed costs $r\bar{K}$ over more units sold, and below-capacity firms can produce a given quantity at lower marginal cost.

Kalouptsi (2018) discusses the problem of the optimal number of backlog orders and the impact of time-to-build on demand at greater length in discussing shipyard production. This cannot be replicated here as I lack precise data on the backlog for each vehicle, and the value function of diverse satellite operators cannot be generalized as easily as that of bulk cargo ship operators. Instead, I assume that non-rush orders are smoothed over time for one or more of the following reasons:

- Reducing forgetting effects at the employee level by keeping a smaller, continuously-employed staff
- Reducing the chance of shortages from upstream production bottlenecks
- Time to maintain launch facilities between launches
- A launch delayed due to inclement weather will not delay subsequent launches

Price schedule estimates for the *Atlas* and *Delta* programs from Greenberg and Hertzfeld (1992) suggest that $\zeta = 1$.

Figure 1.6: Production time





Figure 1.7: Production time



1.6 Profit

Given the current state of demand S_t , competitors $\mathbb{V}_t$ with experience levels E_t annual profit at optimal labor and capital levels is thus

$$\Pi_{Vt}(q_t, S_t, \mathbb{V}_t) = p(q_t, S_t, \mathbb{V}_t, E_t)q_t - r\bar{K} - cq_t - w_t \left(\frac{q_t}{E_{Vt}^\eta} \right)^{\frac{1}{\zeta}}$$

Each year t the firm chooses whether to exit, $\chi_V = 1$, or continue, after observing a scrap value Φ_V . The Bellman equation for

$$V_V(S, \mathbb{V}) = \max_{\chi_V} \left\{ \chi_V \Phi_V + (1 - \chi_V) \max_{q_t} \left[\Pi_V(q, S, \mathbb{V}, E_t) + \beta \sum_{S', \mathbb{V}'} [P(S', \mathbb{V}', E' | q, \chi, S, \mathbb{V}, E) V_V(S', \mathbb{V}', E)] \right] \right\}$$

1.7 Firm Dynamics with Debugging

1.7.1 Steady State

Suppose we have a set of firms $i \in \{1, 2, \dots\}$, each with a single vehicle capable launching 1000 kg to LEO, and facing total global demand of m 1000 kg satellites per year. Setting aside distance, trade regulations, and home bias, demand is given by

$$q_{it} = m \frac{P_i^{\alpha_P} (R_{i0} E_{it}^\theta)_R^\alpha}{1 + \sum_j P_j^{\alpha_P} (R_{j0} E_{jt}^\theta)_R^\alpha}$$

$$\frac{q_{it}}{q_{jt}} = \left(\frac{P_{it}}{P_{jt}} \right)^{\alpha_P} \left(\frac{R_{i0}}{R_{j0}} \right)^{\alpha_R} \left(\frac{E_{it}}{E_{jt}} \right)^{\alpha_R \theta}$$

Note that $\frac{E_{it}}{E_{jt}}$ is constant in t if and only if $\frac{E_{it}}{E_{jt}} = \frac{q_{it}}{q_{jt}}$, ie new launches are distributed exactly in proportion to previously-accumulated experience. This implies the steady-state price ratio

$$\left(\left(\frac{E_{it}}{E_{jt}} \right)^{\theta \alpha_R - 1} \left(\frac{R_{i0}}{R_{j0}} \right)^{\alpha_R} \right)^{-1/\alpha_P} = \frac{P_{it}}{P_{jt}}$$

Interestingly, the long-term trajectory of experience depends on coefficients: if $\theta \alpha_R > 1$, the more experienced rocket commands a higher price, but if $\theta \alpha_R < 1$, a less-experienced vehicle actually prices higher in the steady state.

Equivalently, if $\theta \alpha_R > 1$, then experience gaps tend to widen over time unless entrants engage in aggressive penetration pricing. Instead, my estimates suggest that $\theta \alpha_R < 1$ for the commercial space industry. Thus vehicle families with considerable experience but high cost, such as the US's *Atlas*, will inevitably lose market share over time.

If this is the case, then why have the US and Russia been launching the same *Atlas* and *R-7* rockets, respectively, since 1957? I argue that entry and exit dynamics are the real drivers of persistence in this market: given entry costs on the order of \$100 million to \$12 billion and fixed costs of \$200 million, government and private investors may not be able to wait several decades for a new vehicle to reach profitable experience levels. Entry may be deterred, while new products may be abandoned almost immediately after entry when they realize cost and quality insufficient for meteoric success. Finally, a vehicle like *Atlas* or *Soyuz* which is still profitable because of considerable experience may be kept in production even as dozens of new vehicles with superior R_{0i} are abandoned. I thus move on to consider the firm's decision-making process for prices, entry, and exit.

1.7.2 Dynamic game

The timing of the dynamic game is as follows:

1. Observe

- Potential entrants privately observe a random entry cost
- Incumbent firms privately observe a random scrap value
- All firms observe the state of all other firms (γ, E_{it})

2. Decisions

- Potential entrants decide whether to enter
- Incumbent firms choose a price and whether to exit at the end of the period

3. Sales are realized, experience levels are increased

4. Realization

- Entering firms draw a starting reliability level γ_F and are charged their entry cost
- Exiting firms receive their scrap value

1.7.3 Solution Concept

Computational tractability is arguably the most challenging barrier to computing an Ericson-Pakes-style dynamic competitive equilibrium. This is exacerbated by the number of independent firms in a market: while Benkard (2004) limits computational experiments to a maximum of 4 independent firms for tractability, I observe an average of 17 separate vehicle families per year.

One solution proposed by Benkard, Jeziorski, and Weintraub (2015) is to consider a “partially oblivious equilibrium” (POE), where firms observe the average long-run state of the industry and the states of a small number of dominant firms. With a C2 ratio ranging from 0.34 to 0.48¹¹, the space industry has an appropriate level of concentration for this method. The information constraint also makes sense: leading firms like SpaceX publicly advertise their prices and product characteristics and livestream launches online. By contrast, the capabilities and sales plans of small firms are often subject to speculation. For example, reports of launches and failures of the Iranian *Safir* rocket are often revealed to the press from US analysis of commercial and intelligence satellite photos (Brumfiel, 2019).

However, POE assumes that the set of dominant firms is fixed over time. This is inappropriate to my setting: of the top 4 vehicle families from 1995, only one remains in the top 4 in 2018. I instead approach the problem from the perspective of “Aggregate Equilibrium”, as laid out in Ifrach and Weintraub (2012). Like POE, this equilibrium concept allows for two tiers of firms: dominant firms whose individual states are observed by the entire industry, and fringe firms whose states are observed only in aggregates. Importantly, fringe firms are allowed to become dominant firms and vice-versa in this model.

I specifically use the utility of all outside goods as an aggregator for competitor states. Note that our demand function is

¹¹Weighted by mass-to-orbit

$$u_{sft} = \alpha_0 + \alpha_p \ln(p_{sft}) + \alpha_{ITAR} ITAR(r_s, F) + \alpha_R \ln(R_{ft}) + \alpha_{DIST} DIST(r_s, F) + \epsilon_{sft}$$

Let $\alpha_E \equiv \alpha_R \theta$

$$\gamma_f \equiv \exp(\alpha_0 + \alpha_{ITAR} ITAR(r_s, F) + \alpha_R \ln(R_{F0}) + \alpha_{DIST} DIST(r_s, F))$$

Setting aside variants, we can then rewrite demand as a function of the firm's own state, the firm's action (price) and an aggregate Y , similar to the demand model in Nocke and Schutz (2018)

$$\begin{aligned} q(P_F, \vec{P}_{-F}, E_F, \vec{E}_{-F}) &= M \frac{P_F^{\alpha_P} E_F^{\alpha_E} \gamma_f}{1 + \sum (P_j^{\alpha_P} E_j^{\alpha_E} \gamma_j)} \\ q(P_F, E_F, \gamma_F, Y) &= M \frac{P_F^{\alpha_P} E_F^{\alpha_E} \gamma_F}{1 + Y} \\ Y &\equiv \sum P_j^{\alpha_P} E_j^{\alpha_E} \gamma_j \end{aligned}$$


The aggregate equilibrium consists of the following:

- An entry strategy $\lambda(Y)$, which gives the maximum price a firm would be willing to pay to enter after observing aggregate market condition Y
- A pricing strategy $\mu(E_F, \gamma_F, Y)$, such that the firm charges a price of $P_F = c(1 + \mu)$ at each experience level, quality, and competition level specified.
- An exit strategy $\nu(E_f, \gamma_F, Y)$, giving the minimum scrap value a firm would accept to exit

To calculate the aggregate equilibrium, I follow the following steps:

1. Simulate 30 years of sales, profit, experience, entry, and exit when all firms follow strategies λ, μ, ν

Figure 1.8: Simulated sales paths



2. Record the average realized continuation values $V(E_f, \gamma_F, Y)$ and entry value $V(Y)$ from the simulation
3. Update $\lambda(Y), \mu(E_F, \gamma_F, Y), \nu(E_f, \gamma_F, Y)$
4. Repeat

My version of Step 2 stores continuation values, rather than per-period profits as in Ifrach and Weintraub (2012), because a single action (price) influences both payoff in the stage game and the firm's future state. Thus, at each possible state visited by the firm, we consider the firm deviating from the existing pricing strategy for one period to move onto a new trajectory.

Figure 1.8 depicts a simulated market history after 1000 iterations of strategy refinement. This replicates several salient features of the historical market. Firms which enter earlier and have higher starting quality rapidly increase in sales. As Y increases in later periods, entry attempts become less frequent and are held to higher standards to avoid cancellation.

Figure 1.9: Simulated sales paths with an ITAR-like shock




Figure 1.9 divides firms into two countries, Red and Blue. Blue firms receive an unanticipated negative shock to quality of e^{-1} in period 10, which is reversed in period 20. Red firms immediately gain more sales, Blue firms immediately lose sales, and several experienced Blue firms respond by exiting¹². Over the course of the decade, Red firms gain a considerable amount of experience, and there is a slight uptick in new firm entry due to the decreased value of Y . After $t = 20$, Blue firms recover in market share, but Red firms now command a solid lead.

¹²Note that the US *Titan* family exited in 2005 after 368 launches.

1.8 Vehicle Family Entry and Unanticipated Exit

Exit events are fairly common throughout the data. The cost of entry into the market ranges from hundreds of millions to billions of dollars. Post-entry, rocket families have no clear “expiration date”; today’s *Soyuz* and *Atlas* rockets are descended from missiles which first flew in 1953 and 1957, respectively. This supports a story where production efficiency and quality are nondecreasing in experience; furthermore, incorporating new innovations into an existing family is generally less costly than starting a new family from scratch.

Given these conditions, the numerous examples of rocket families cancelled or put on indefinite hiatus are striking, and must be the result of severe unanticipated negative shocks to the firm’s continuation value, and are thus valuable in estimating a firm’s value function.

1.8.0.1 Exogenous Exit

PA Yuzhmash, a former Soviet manufacturing firm in Ukraine, produced a variety of vehicles for the global market including *Tsyklon*, *Dnepr*, *Kosmos*, and *Zenit*. Production was abruptly halted after the 2014 Russian military intervention in Ukraine, as Yuzhmash lost access to Russian partners and Russian-leased Baikonur Cosmodrome; the firm has struggled to develop new vehicle variants with alternative launch sites since then. This sort of geopolitical shutdown is fortunately rare.

1.8.0.2 Single-customer products

Relatedly, some heavy-lift vehicles are purpose-built for a single (government) customer and a specific mission profile, with no plans to expand for international commercial competition: waning political support for the mission, due to launch failures, cost overruns, or exogenous factors, can lead to exit. Production of the US *Saturn* and Soviet *N1* rockets ended with

their respective manned lunar programs¹³, while the *Energia* and its payload, the *Buran* space shuttle, were both cancelled with the fall of the Soviet Union. An edge case is the *Space Shuttle*: while initially conceived for a variety of government and private payloads, commercial missions were suspended after the *Challenger* disaster in 1986, after which it was used exclusively for government super-heavy-lift and manned missions like assembling and supplying the International Space Station. These examples are beyond the scope of this paper, namely vehicles competing for commercial customers from 1995-present. However, they can be considered as an extreme case of the patterns described below: the *Saturn V* was the best value for delivering an incredible 140 tons to LEO or 48.6 tons to translunar injection orbit, but exited when the market for delivering a satellite of that incredible size disappeared.

1.8.0.3 Aggregate demand shocks

Commercial launch services are used to install physical capital for high-tech services: as such, aggregate demand can be uncertain. The late 1990s saw significant investment in new orbital communications services, such as satellite phones. The three largest players, OrbComm, Iridium, and GlobalStar, filed for bankruptcy between 1999 and 2002, cancelling planned launches and doing significant harm to the commercial launch industry. One casualty of this crash appears to have been the Lockheed-Martin *Athena*, a low-cost launcher based on Space Shuttle solid boosters last launched in 2001.

A number of new vehicles, such as the US/NZ *Electron* and the Chinese *OneSpace M*, are currently in development to serve hypothetical future demand for numerous on-demand microsatellite launches; should this demand fail to materialize, we could see a similar rash of exits.

To account for this uncertainty, I model aggregate demand as a Markov process, with

¹³One remaining *Saturn V* was used to launch the *Skylab* space station

the total market demand as well as distribution with respect to mass, destination orbit, and country of origin changing over time.

1.8.0.4 Unanticipated competition

The availability and quality of competing vehicles is perhaps the biggest *ex ante* risk for an entering launch vehicle. For example, when the UK began developing their indigenous *Black Arrow* launch vehicle in 1964, it was unclear if contemporary US launch vehicles would ever be made available for foreign customers. After 2 failed launches, the US proposed an export arrangement, and the UK government recommended that “all work on [Black Arrow] should be stopped as soon as satisfactory arrangements have been made for the supply of [a] US launcher” (Hill, 2011).

In the present day, Mitsubishi Heavy Industries expected significant foreign sales of the *H-2* and *H-2A*; these were not realized, likely due to the meteoric rise of *Falcon* and *Long March*, and MHI is increasingly reliant on Japanese military and intelligence launches (Pekkanen and Kallender-Umezu, 2010).

1.8.0.5 Poor reliability

I assume that firms must pay the startup cost to design a vehicle and prepare production facilities before observing the vehicle’s starting reliability R_{V0} ; this value is then revealed through Bayesian learning with additional launches. Due to the stochastic nature of failure events, it may be difficult to discriminate between poor luck and a fundamentally-flawed design until a certain threshold of data is reached. This threshold appears fairly consistent in the data: the Soviet *N1* was cancelled after 4 failed launches, the Brazilian *VLS* after 2 failures, the British *Black Arrow* after 2 failures (with a successful launch after cancellation); the South Korean *Naro-1* after 2 failures and 1 success.

It should be noted that the threshold R_{V0} for exit likely varies with market conditions.

One of the first US vehicles, *Atlas-LV3 Agena-A*, suffered 2 failures in the first 4 launches, but with few alternatives in 1961 the US continued to order more units of upgraded versions. *Falcon 1* failed in 3 of 5 launches, but successfully demonstrated technology which would be incorporated into the cost-effective *Falcon 9*. Most strikingly, Astra Inc. formally named their first *Rocket 3.0* mission “1 in 3” because they “believed that it would probably take three launches before we could successfully deliver a satellite into Earth orbit...but understood how strategically important responsive launch was to the government”.

As this suggests, political will is likely a confounding factor. Skeptical voters may be quick to condemn an expensive indigenous launch vehicle program after one or more highly-publicized accidents. It is also possible that a launch vehicle program with little hope of competing globally may be kept alive by subsidies and government orders as a point of national pride and security (Pekkanen 2010).

I account for the unpredictable future of competition by representing the future experience levels of all vehicles in the market E_{t+1} as a random variable conditional on the current experience levels, E_t , and the firm’s decision q_t . The “political will” behind indigenous rocket development is reflected by the home bias dummies in the demand function, $\alpha_H, \alpha_{H,GOV}, \alpha_{H,MILINT}$, and the total demand for military and intelligence launches from a country over time. By observing the predicted value functions of rockets just before exit, we may find evidence that indigenous launch capability provides intrinsic value beyond the expected discounted future profits of the launch service provider.

1.8.1 Satellite Construction

The firm which assembled each satellite (contractor) is also included in the data. 474 contractors connected to more than 1 satellite are hand-matched to their home country and public/private status. Around 60% of satellites in the full dataset are assembled by a firm in the customer's country. Demand estimates show that the customer is more likely to choose a launch vehicle from the same country as the contractor; controlling for contractor reduces the effect of home bias. This could be justified by

1. **One firm bundling in-house satellite construction and launch services**

This seems likely: CGWIC, for example, advertises bundled satellite manufacturing and launch services.

2. **Reduced cost of transporting the satellite from the manufacturer to the launch site**

This appears less likely: Arianespace, after all, transports entire *Soyuz* vehicles from Russia to a launch site in French Guiana rather than developing a locally-manufactured alternative.

3. **Synergies in travel and regulatory compliance**

I find that firms which buy ITAR satellites are no more likely to hire ITAR vehicles, suggesting that there is no cost-saving by applying for ITAR for both the satellite and launch vehicle.

4. **Synergies in hardware technology**

Besides miniscule CubeSats, there is no standard size and attachment standard comparable to the ISO container, so a satellite bus manufactured by Nippon Electric Company may be built with a payload adapter specific to the Mitsubishi Heavy Industries *H-IIB* in mind. To test for this, I include a dummy which equals 1 if the vehicle option uses RUAG's payload adapter, used by *Ariane*, *Atlas*, *Delta*, *Land Launch*, *Proton*, *Sea*

Launch, *Soyuz* and *Antares*; this is not significant. A second dummy variable tests for a RUAG payload adaptor but excludes satellites manufactured in countries with a non-RUAG indigenous vehicle (India, China, and Japan); this coefficient is significantly positive in more parsimonious models, suggesting that payload adapter compatibility is non-trivial.

This still leaves the question of how to incorporate satellite contractors into the model, since we observe neither prices nor quality differentiation between satellite manufacturers. A model frequently cited by industry analysts is “space technology ladder” (Pekkanen and Kallender-Umezu, 2010), showing how a developing country develops space capability. Each stage gives the country more skilled workers and institutional knowledge for the next stage of development, and provides a larger home market for products in the next stage.

1. Local government agencies and companies launch foreign-built satellites on foreign rockets
2. Domestically-produced satellites are launched on foreign rockets
3. Indigenous satellites are launched on experimental indigenous rockets (possibly using foreign parts¹⁴)
4. Indigenous rockets are marketed commercially for foreign customers
5. Indigenous rockets are used for manned space missions

The focus of this paper is the shift from stage 3 to stage 4: a locally-produced design becomes cheap and reliable enough to compete internationally. Stage 3 is considered for entry and exit conditions. Much of this decision is exogenous, however: the startup costs of most orbital rockets are funded by one or more national governments for national security

¹⁴For example, India’s PSLV and GSLV used the Vikram engine, a licensed version of the Viking engine from Europe, and Japan initially used US-built Thor 1st-stage boosters.

reasons, as an offshoot of a domestic ICBM program and/or to enable a spy satellite program; both conditions have more to do with geopolitics than economics.

For the purpose of the demand model, I treat stages 1 and 2 as exogenous. For the dynamic simulation of pricing strategies, I assume that a country's demand for space services and supply of satellites evolve over time as an increasing function of GDP, serving as demand shifters for launch services. For example, increasing local demand for telecom services led the Angolan government chose to purchase a geostationary satellite (AngoSat 1) in 2009; the scale of the Russian satellite industry contributed to their choice of RKK Energia to construct the satellite, which in turn increased the probability that the Russian/Ukrainian *Zenit-3F* was chosen over Chinese or American alternatives. Angosat 3, by contrast, is being built by Airbus for an unknown rocket, suggesting that the Angolan government is simply sampling from the pool of established satellite manufacturers, then choosing a compatible rocket with a bias towards providers local to the satellite manufacturer.

1.8.2 Policy Implications

Will arms control always backfire? Unlikely. Rather, I argue that policymakers considering regulating US exports of national-security-sensitive technologies must consider a more nuanced approach, using information on each product under consideration for export controls.

1. What is valuable about the product?
 - (a) **An innovative blueprint.** Innovations are visible to the end-user and easily reverse-engineered by a foreign actor after purchasing a few samples. Export controls may prevent spillovers.
 - (b) **Experience-driven innovations hidden from the end-user**, such as production techniques, safety procedures, and a specialized workforce. It is impossible to reverse-engineer this institutional knowledge from the final product. Export controls can have an “infant industry protection” effect which *encourages* foreign

firms to enter and built their own experience.¹⁵

2. What is the size of the market on each side of the proposed barrier?

- Targeted, multilateral sanctions like a NATO sanctions on Iran may prevent the targeted country from acquiring capability while denying them an export market for indigenous alternatives.
- Broad and unilateral sanctions, like the US requiring all foreign firms to acquire ITAR licenses, may backfire by incentivizing foreign firms to buy foreign products.

3. What is the role of civilian users?

- A primarily military product can be sold to allies through government-facilitated transactions.
- Civilian customers, including domestic corporations, may be deterred by expensive licensing and regulatory compliance. The key to a high-quality product may be capturing international sales to civilian customers by streamlining export procedures.

Another key implication of this project is the importance of the size of the home market. I find that governments are willing to buy a subpar new indigenous products, likely internalizing the benefits of helping their “national champion” reach competitive scale and experience. No single national government, however, has the experience to support a product as its only customer until it reaches competitive experience levels. In such an industry, the solution may be to build the product through a multinational organization, where it can be the preferred “domestic” choice for several different governments. Organizations like Airbus, Arianespace, and Eurofighter leveraged a wider home market to help their products reach internationally-competitive quality and price.

¹⁵Note that exports could pose other risks to national security: imported weapons could be used to threaten the exporting country (like Iran’s US-built F-14 fighters) or analyzed for security vulnerabilities (like the US’s study on a captured Mil Mi-25 helicopter).

1.9 Conclusion

How can an advanced economy maintain a competitive edge in national-security-sensitive technology? Counterintuitively, the answer may be to maximize exports. I look at a high-tech industry where the US attempted to restrict the development of its geopolitical rivals by placing its own products under restrictive export controls. This bolstered demand for foreign rivals, especially new, low-cost entrants from Japan, China, and India, serving as “inadvertent infant industry protection”.

I show four key relationships:

1. Consumers are sensitive to reliability and cost
2. Experience allows firms to reduce cost and accident rates, bolstering sales and increasing experience further
3. ITAR served as a massive barrier to US launch service exports from 1999 to 2013, leading the US to lose market share to foreign entrants
4. 1 and 2 resulted in this shift in market share persisting since 2013, even after accounting for shifting customer characteristics
5. Structural models show that the shift is on both intensive and extensive margins.

To complete this model, I endogenize firm pricing, entry, and exit using the Aggregate Equilibrium model of dynamic firm behavior from Ifrach and Weintraub (2012). I modify this model to incorporate learning-by-doing, which ties the pricing behavior of the stage game to the “investment” action of the dynamic game.

The result is a model which can be extended to a variety of applications critical to modern public policy:

- **Infant Industries**, where developing countries may consider implementing trade barriers to enter new sectors
- **Bleeding-edge high tech products**, where early adopters help the producer identify and correct product defects, such as novel software and internet-of-things products
- **Dual-Use Products**, where technology is shared between civilian and military products, and the profitability of private firms must be weighed alongside national security concerns

Canonical economic models and media coverage tend to focus on paying development costs and drawing an excellent blueprint as the driving force of innovation. This paper finds evidence that a second stage of innovation is equally important: selling early models, gathering feedback, and amending designs and production techniques to earn a greater market share in the future. The benefit of these incremental changes forms a significant barrier to entry, which I hope to explore in future work.

CHAPTER 2

Testing for Fun and Profit

2.1 Introduction

In many classical models, firms pay a sunk cost to develop a new product with a new quality level which can then be introduced to the market; information frictions may slow the adoption of a superior product. A parallel literature on “learning by doing” shows that firms can also reduce their costs as they increase experience (cumulative units produced to date). In previous research, I presented an alternative model where an element of a product’s *quality* (reliability) is improving with additional experience, as common problems encountered by customers are remedied through incremental product improvements. This creates a positive feedback effect: firms which sell many units gain more experience, allowing reliability to continue to grow, resulting in even more sales, creating slowly-accelerating adoption even in the absence of information frictions. This “Matthew effect” drives older and more expensive products to retain significant market share while cost-effective new designs are driven out of the market by liquidity constraints or stochastic shocks.

In my research on rockets, I assumed that after each product failure the firm halts production, investigates the accident, and locates and removes the “bug” in the product for all units produced thereafter. This is reasonable for a product like an orbital launch vehicle: the cost of failure is enormous relative to the cost of locating and addressing a bug, and it is little consolation for a customer whose satellite has been destroyed by a recent launch to know that future launches will be safer.

By contrast, consider a consumer computer program which crashes, freezes, or gives incorrect feedback under certain circumstances. The user might spend less than a minute closing and restarting the program, and may expect that the developer may offer a “patch” to remove the glitch from their program in the future. As a result, a developer may choose to ignore a glitch identified by users. Potential future patches offer an incentive for early adoption, provided that users believe that further patching is probable. A signalling game thus emerges, where the developer’s patching behavior to date is used to generate beliefs on the long-term value of adopting the product.

This work can be extended to a variety of software and, perhaps more interestingly, firmware. “Internet of Things” consumer durables are powered by firmware which offers novel features, which could be expanded and improved by downloadable updates. Tesla electric cars, for example, are regularly updated to improve performance and refine their “autopilot” self-driving functionality. By purchasing a product like this, a consumer both buys the existing product and wagers on future software improvements.

2.2 Setting

While patches are *de rigueur* in all software, they are particularly critical for “Early Access” software. Early Access allows consumers to pay developers for unfinished titles, in exchange for immediate access to the current build and automatic free updates as the software is finished (Valve, 2021).

Perhaps the most successful early access title was “Minecraft”. Originally developed by Marcus “Notch” Persson and released in 2009, early access sales proved the appeal of the game’s novel mechanics and unusual art style, and funded the hiring of additional employees. The game was formally released in 2011, and by 2012 the developer had 25 employees and \$240 million in sales (Sarkar, 2013).

Other early access projects are more disappointing. Updates become infrequent: the title

can be formally “released” in an unfinished state or “abandoned” with no further updates. This can be extremely disheartening for early access users: one Steam user in the dataset laments “Almost five years of broken promises and we’re left with a tech demo that doesn’t even resemble the original concept that was sold”.

A recurring allegation is that some developers abuse Early Access programs to “scam” consumers into purchasing a broken title with no intention of delivering the promised updates. While this is likely true in some cases, I propose another possible channel: developers enter early access in good faith, but uncertain of their product’s potential profitability. Early access sales are used to refine beliefs about the sales potential of a flawless version of the product; the firm’s progress in eliminating flaws allows customers to refine their beliefs on the firm’s relative cost of fixes.

For this paper I will focus on the first 12 video games admitted into the Early Access program on Steam, the leading digital marketplace for games on PC, Mac, and Linux. Of these titles, 9 “graduated” into a formal release¹, 2 were abandoned and removed from the store, and 1 appears to have been abandoned and left in Early Access indefinitely (Lane, 2019).

¹One of these was removed after “release” due to extremely negative reviews

Title	Genre	Status
123 Kick It	Rhythm	Last Updated Aug 2014
Arma 3	Shooter	Released Sep 2013
Drunken Robot Pornography	Shooter	Released Feb 2014
Gear Up	Driving	Released Jan 2015
Gnomoria	Citybuilding	Released Feb 2016
Kenshi	Roleplaying	Released Dec 2014
Kerbal Space Program	Space Sim	Released Apr 2015
Kinetic Void	Space Sim	Released Nov 2014
Patterns	Sandbox	Cancelled Oct 2014
Prison Architect	Management Sim	Released Oct 2015
Starforge	Action	Released Sep 2014, Removed 2017
Under the Ocean	Platformer	Cancelled Feb 2014

If abandoned projects came from “bad” firms, we might expect the worst products to come from single-product firms or serial offenders, with “good” developers consistently completing projects to maintain a positive reputation. Even in this small sample there are clear counterexamples to this argument. One developer, Dejobaan Games, produced both the successful *Drunken Robot Pornography* and the abandoned *1...2...3...Kick It!*. More strikingly, the same developer, Linden Labs, abandoned *Patterns*, but has been continuously updating their title *Second Life* since 2003. Instead, I will focus on the idea that consistent patching/updating is used to maintain reputation at the *product* level.

2.3 Data and Trends

Analysis of data from a wide variety of research and development programs, starting with Duane (1964), finds that a prototype system run for E_t hours (with pauses to implement fixes) will an increasing Mean Time Between Failures $MTBF_t$, with relationship

$$\ln(MTBF_t) = \alpha \ln(E_t) + \beta \quad (2.1)$$

Like the “gravity” model of international trade, this empirical regularity can be justified using several theoretical models, but all revolve around the core idea that an unobserved set of product bugs randomly cause failure events, which result in re-designs which eliminate the bug.

My first question is “Does early access software follow the Duane rule?” To answer this, I need measurements of the number of bugs reported by users and the total time that the system has been running across all users.

In my first method, I utilize user posts on the official Steam forum. For each title considered, I collect all posts featuring words related to product failures like “glitch”², words related to less-serious quality issues³, and words critical of the developer’s conduct⁴. I posit that this is a plausible measurement of the true number of failure events (reported and unreported), up to a scalar transformation, for the following reasons:

- Every consumer who accesses the game through Steam⁵ must have the Steam app open and logged in
- Inspection of the data suggests that users will often comment on another user’s bug-related post to describe similar or identical bugs they have also encountered
- Users frustrated by bugs may favor the official Steam product forum for complaints, in

²blue screen, broke, bug, crash, error, freeze, froze, glitch, issue, kick

³nerf, balance, missing

⁴abandon, refund, scam

⁵Note that customers on other sites, such as Kickstarter, will often receive their purchase through the Steam client

the hopes of drawing the attention of the developer and/or marketplace administrators who could remove the product and/or offer refunds

I combine this with an estimate of the total time that all users have spent using the product to date. To do this, I take the monthly average users (MAU_t) values calculated and archived on James Grey’s *SteamCharts* site, and find the cumulative sum to date.

$$E_t = \sum_{t=0}^t MAU_t \qquad MTBF_t = \frac{E_t}{F_t}$$


The results shown are for a slightly more complicated specification utilizing daily data: if the i th failure (bug report) is posted on month $m_i \in \{0, 1, \dots\}$, day d_i since release, with d_{m_i} days in month m_i , the MTBF is calculated as follows:

$$E_t = \sum_{m=0}^{m_i-1} MAU_m + \frac{d_i}{d_{m_i}} MAU_{m_i} \qquad MTBF_i = \frac{E_i}{i}$$

Note that this framework implicitly assumes that failures are random events, independent across individual users (i.e. “my computer freezes when I enter this menu”), as opposed to a correlated event (“the server crashed, booting all active users at 9 AM PST”). I do not find any such “blanket failures” in the current sample, though future work could explore the evolution of “uptime” for internet services as both usage and experience expand.

There are several other viable sources for reliability estimates not included here. First, bug complaints on other forums, such as Reddit, can also be collected as independent estimates of total bug counts. Galyonkin (2021)’s estimates of ownership over time can be used to provide the cumulative months owned over all owners as an alternative measure of E_t . Finally, the fraction of user-written reviews which include bug-related words can serve as an alternative measurement of failure rates; as Steam includes each reviewer’s total playtime,

Figure 2.1: Reliability growth (scatter) and patching (stepped line)




the average could even be weighted by usage.

I combine my reliability estimates with each developer’s official patch notes posted on Steam, and historical price trends from Galyonkin (2021). Data from two of the first titles in Steam Early Access, *Arma 3* and *1...2...3...Kick It!* are shown in Figure 2.1 and 2.2.

Arma 3 was considered a major success, graduating from Early Access to a full release in September 2013 for \$59.99 and selling an estimated 6 million copies. The relationship between cumulative usage and time between failures is piecewise loglinear. The first shift is around July 2014: reliability growth picks up roughly a year after release, corresponding to an increasing cadence of patches and the release of the game’s first paid downloadable content (DLC). Reliability growth tapers off after May 2016; at this point the game goes on sale for \$19.99 for the first time, and patch notes increasingly focus on solving issues introduced by DLCs and detecting cheating, rather than solving any remaining performance issues in the core product. By 2019, patching is reduced, reliability growth is turning negative, and the title is frequently on sale for just \$10.19. Throughout its entire development cycle, the studio consistently offered 12-21 patches per year.

1...2...3...Kick It! is more of a cautionary tale: this rhythm game had just 14.5 users

Figure 2.2: Reliability growth (scatter) and patching (stepped line)



on average at the peak of its popularity, but less than a third of the estimated reliability of *Arma 3*. The game was patched 8 times in 2013: since then, there have been no further patches or updates, and the title has not left Early Access as of May 2021. Interestingly, the patches did not result in a positive reliability growth trend in 2013, and what little positive growth we see in 2014 tapers off immediately. With less than 1 user on average since 2015, there are, unsurprisingly, no further bugs reported on forum.

The relationship between update frequency, failure rate, and users is difficult to quantitatively summarize across many titles, but most examples in the dataset fall on the spectrum between these two examples: the immediately popular, continuously-supported, and increasingly reliable *Arma 3* and the unpopular, troubled, and seemingly abandoned *Kick It*. Successful and abandoned Early Access projects continue to the present day; it should be noted that small and large studios alike also continue to release software using the traditional non-Early-Access model.

From an economic perspective, we have two clear questions:

1. Why do both weak and strong firms initially “pool” into Early Access programs?

2. Following entry into Early Access, why do weak and strong firms rapidly separate into success or abandonment?

In the remainder of this paper, I will attempt to justify this equilibrium with a novel model combining elements of reliability growth analysis and Spence-style signalling.

The result can be thought of as a model of “strategic default”: well-intentioned developers enter the market, find a high cost and low reward for completing their product, and completely abandon the product.

I argue that this “hard cutoff” is enforced by a separating signalling equilibrium: high and steady patching effort serves to both directly improve the product and separate safe investments from soon-to-be abandoned projects. The dual nature of patches is noted by Chris Hunt, lead developer of *Kenshi*: “If we did an update every four months, then everyone would get angry and during that gap would be saying, ‘The game’s abandoned!’...But if we did the same amount of work, but broke it up into an update every day, people felt it was really well maintained.” (Lane, 2019)

2.4 Sketch of the Model

Time is discrete and the horizon is infinite. As the field of early access games is very diverse, I focus on the consumer’s choice between a single game and a generic outside good.

At $t = 0$, a developer pays C_E to enter.

- The developer draws an unobserved time-invariant “intrinsic appeal”⁶ of the product, $a \sim F_A(a)$.
- Each consumer i draws and privately observes their own preference for the product, $a_i \sim F_{a_i}(a_i|a)$, $E[a_i] = a$.

⁶Note that this reflects the value of the “ideal” version of the product, i.e. a perfect version with no observed bugs.

- The developer draws and discretely observes $c \sim F_c(c)$, a cost shifter for fixing all future bugs, reflecting, for example, the quality of existing code and documentation
- The starting probability of using the software for 1 time period with 0 glitches, thereafter “reliability rate”, R_0 is drawn and revealed to both consumers and the developer, e.g. through user and professional reviews.

Each time period t ,

1. R_t is observed by all.
2. Consumers decide whether or not to adopt the product; sales are made, and users observe and report bugs.
3. Developer updates their beliefs on a
4. Developer puts in effort to address bugs $\mu_t \in [0, 1)$
5. Users update their beliefs on a, c
6. R_{t+1} updates

Note that I set aside price-setting for now. In practice, prices appear to be set based on the broad category of the product (\$60 for a photo-realistic game made by a large studio, \$20 for a stylized game by a small team, etc.), with round-valued discounts and regular sales. It may be interesting, however, to explore the relationship between intertemporal price discrimination and strategic patching: a consumer’s adoption decision may hinge on whether they expect the firm to make additional sales through discounts, which reward delayed adoption, or patches, which reward early adoption.

Assuming firms choose to enter into an Early Access program, I outline a signalling game where firms of the strong type (favorable draws of c) will put in more effort over more periods than those of the weak type. Additional effort today thus generates revenue in two ways: the

product is immediately improved, increasing sales, and consumers have a higher expectation of the strength of the firm, thus expecting the product to travel along a higher reliability trajectory.

2.5 Consumer's Problem

Each time period, consumers choose between adopting/continuing usage of the product or consuming the outside good, with the following utility:

$$\text{Don't adopt: } u_{i0t} = u_0$$

$$\text{Adopt: } u_{iAt} = a_i + b \ln(R_t) - P_t$$

$$\text{Continue to use: } u_{iCt} = a_i + b \ln(R_t)$$

where R_t is the current probability of using the product for 1 time period with 0 bugs observed.

Trivially, if R_t is nondecreasing in t , adopters will continue to use the product following adoption. The consumer's value function is thus:

$$V(0, a_i, t) = \max \begin{cases} u_0 + \delta E[V(0, a_i, t+1)] \\ a_i + b \ln(R_t) - P_t + \delta E[V(1, a_i, t+1)] \end{cases} = V(1, a_i, t) - P_t$$

$$V(1, a_i, t) = \sum_{s=t}^{\infty} \delta^{s-t} E[a_i + b \ln(R_s)] = \frac{a_i}{1-\delta} + b \sum_{s=t}^{\infty} \delta^{s-t} E[\ln(R_s)]$$

2.6 Developer's Problem

I use “PM-2-Discrete” from AMSAA (2011) as a reference model. This Department of Defense publication is intended to help the manager of a research and development program

schedule a campaign of alternating testing periods and “corrective action periods” (CAP) to meet a specified reliability goal: for example, a vehicle will receive 1000 orders if it breaks down less than once per 1000 miles driven; it receives 0 orders if it fails to pass this threshold. While this is appropriate for a military contract, I expect the adoption of a consumer product to be a continuously-increasing function of reliability, so that the firm must ask, in each period, whether further testing and CAP are worth the cost.

In this setting, the firm chooses a “management strategy”, which I label $\mu \in [0, 1]$, so that for every n new failure modes discovered, $n\mu$ will be scheduled to be fixed during the next CAP and $n(1 - \mu)$ will be ignored. I assume that, for each period t , the cost of pursuing management strategy μ is

$$C_t(\mu) = \begin{cases} \frac{c}{1-\mu} = c + \frac{c\mu}{1-\mu} & \text{if } \mu > 0 \\ 0 & \text{if } \mu = 0 \end{cases} \quad (2.2)$$

where c is drawn and observed by the firm at $t = 0$, but unobserved by the consumers.

This function requires a few assumptions:

1. The firm pays a lump sum c to “support” the product ($\mu > 0$)
2. As $\mu \rightarrow 1$, $C(\mu) \rightarrow \infty$
3. $C_t(\mu)$ is constant in current failure intensity $\lambda(t)$

Assumption 1 reflects the cost of continuing to catalog potential bugs from customer complaints and administrative data. Such a fixed cost would explain why developers often formally announce the end of support for a legacy product, rather than simply asymptotically reducing patch frequency [cite].

If bugs require different levels of effort to address, then $C_t(\mu)$ would be increasing. There are multiple reasons why it may be expected to increase without bound (Assumption 2):

- A patch to remove a single bug may be ineffective or introduce a new bug, unless an arbitrarily large expenditure is made to test the patch
- Patching more bugs simultaneously has a higher probability of creating new bugs
- Some fixes could address multiple linked bugs, while addressing isolated bugs could be more expensive
- Patching for every possible use case (hardware, operating system, user inputs) could be prohibitively expensive.

Assumption 3 may be considered a convenient knife-edge case, where the decreasing number of new bugs is cancelled out by an equivalent increase in the cost of addressing each new bug. We could imagine that the first bugs reported are relatively simple fixes, like addressing spelling errors in the user interface, while later complaints require more involved fixes⁷. As the most common bugs are addressed, it may also become increasingly difficult to reproduce bugs which only appear under highly unusual circumstances. We might also expect later adopters to use the software in more extreme circumstances, requiring more demanding fixes: for example, the algorithms in statistical software may need to be re-optimized when later users try to process millions of observations; rendering software might require better memory management as it's adopted by big-budget CGI companies who require more complicated scenes.

Note that, for simplicity, a single primitive c drives heterogeneity across firms in both the lump sum and marginal cost of debugging. In future work, I will consider separate primitives affecting the lump-sum and marginal costs of debugging.

Borrowing from AMSAA (2011), provided that management strategy μ is chosen at each time period t , and assuming that all bug fixes implemented are effective, the new reliability

⁷For example, Riot Games was forced to re-write every clock and timer in League of Legends to give referees the ability to “rewind” a game in competitive tournaments

will be:

$$R_{t+1} = R_0^{1-\mu} R_0^{\frac{\nu\mu}{\nu+E_{t+1}-1}}$$

where ν is a universal primitive reflecting the distribution of bug probabilities, and $E_{t+1} = \sum_{s=0}^t (t-s)q_s$ is the total cumulative runtime of the program over all consumers to date.

Consider the effects of continuing to pursue management strategy μ after observing $\Delta_t = E_{t+1} - E_t = \sum_{s=0}^t q_s$

$$\begin{aligned} \frac{R_{t+1}}{R_t} &= \frac{R_0^{1-\mu} R_0^{\frac{\nu\mu}{\nu+E_{t+1}-1}}}{R_0^{1-\mu} R_0^{\frac{\nu\mu}{\nu+E_t-1}}} = R_0^{\frac{\nu\mu}{\nu+E_{t+1}-1} - \frac{\nu\mu}{\nu+E_t-1}} \\ \ln(R_{t+1}) - \ln(R_t) &= \frac{\nu\mu(E_{t+1} - E_t)}{(\nu + E_{t+1} - 1)(\nu + E_t - 1)} (-\ln(R_0)) \\ &= \frac{\nu\mu\Delta_t (-\ln R_0)}{(\nu + E_t - 1 + \Delta_t)(\nu + E_t - 1)} \end{aligned}$$

Thus the returns to individual consumer utility, and thus sales, of continuing to pursue management strategy μ are diminishing towards 0 in E_t ; thus if $c > 0$, then for some $T \geq 0$ a developer will set $\mu = 0$ for all $t \geq T$.

2.7 Fully-Informed Equilibrium

For now, I will set aside the complexity of learning-through-pricing and assume that both the firm and all consumers rapidly converge on an estimate of a based on sales: $E[a|\{P, q_s, R_0\}_{s=0}^t]$. In particular, at $t = 0$, consumers will have no observation of μ or q_s to refine their beliefs on c , so variation in q_0 will reflect only the generic priors of the population and the distribution of private preferences a_i .

Imagining for a moment that patching continues forever, we can combine the evolution of

reliability with the consumer's utility of adoption to create the value function of an adopter as a function of μ and E_t :

$$V(1, a_i, t) = \frac{a_i}{1 - \delta} + b \frac{\ln R_0(1 - \mu)}{1 - \delta} + b \sum_{s=t}^{\infty} \mathbb{E} \left[\frac{\delta^{s-t} \mu \nu}{\nu + E_t - 1} \ln(R_0) \right]$$

This value function has three additive components: the user's idiosyncratic preference for the concept, a_i , the disutility of bugs which will be never addressed, and the expected disutility of undiscovered bugs which will be addressed in the future.

Suppose instead that the consumer expects the producer to halt updates at some time in the future T , as shown above. At $t \geq T$, the product's utility will no longer increase, so we can safely assume that non-adopters will enjoy the outside good ad infinitum. Assuming $\mu_t = \bar{\mu}$ for $t < T$, the value of not owning the product, owning an abandoned product, and owning a continuously-updated product are, respectively,

$$V(0, a_i, t \geq T) = \frac{u_0}{1 - \delta}$$

$$V(1, a_i, t \geq T) = \frac{a_i}{1 - \delta} + \frac{b \ln R_0}{1 - \delta} \left(1 - \mu + \frac{\bar{\mu} \nu}{\nu + E_T - 1} \right)$$

$$V(1, a_i, t < T) = \frac{a_i}{1 - \delta} + b \frac{\ln R_0(1 - \bar{\mu})}{1 - \delta} + b \left[\sum_{s=t}^T \delta^{s-t} \frac{\bar{\mu} \nu \ln(R_0)}{\nu + E_t - 1} + \frac{\ln R_0 \delta^{T-t}}{1 - \delta} \left(1 - \bar{\mu} + \frac{\bar{\mu} \nu}{\nu + E_T - 1} \right) \right]$$

Let $Q_t \equiv \sum_{s=0}^t q_s$, the cumulative sales to date (inclusive).

At $t \geq T$, the number of potential consumers with $V(1, a_i) - P > V(0, a_i)$ will thus be

$$\begin{aligned} Q_{t \geq T} = Q_T = Q_{\infty} &= \int \left[\frac{a_i}{1 - \delta} + \frac{b \ln R_0}{1 - \delta} \left(1 - \mu + \frac{\bar{\mu} \nu}{\nu + E_T - 1} \right) - P \geq \frac{u_0}{1 - \delta} \right] dF(a_i) \\ &= \int \left[a_i \geq u_0 - b \ln R_0 \left(1 - \mu + \frac{\bar{\mu} \nu}{\nu + E_T - 1} \right) + P(1 - \delta) \right] dF(a_i) \\ &= 1 - F_{a_i} \left(u_0 - b \ln R_0 \left(1 - \mu + \frac{\bar{\mu} \nu}{\nu + E_T - 1} \right) + P(1 - \delta) \right) \end{aligned}$$

Given the short time period involved (12-24 months), assume that the firm has no discount rate, and thus tries to maximize the sum of undiscounted profits:

$$\Pi = P * Q_T(\mu) - \frac{Tc}{1 - \mu} \quad (2.3)$$

Under perfect knowledge, the firm chooses $\mu(a, c, R_0)$, $T(a, c, R_0)$ to maximize final sales agnostic of the timing of consumer adoption. The first order condition approximating this (discrete) choice of T is

$$\begin{aligned} \frac{\partial \Pi}{\partial T} &= P \frac{\partial Q_T}{\partial T} - \frac{c}{1 - \mu} \\ &= P \times f_{a_i} \left(u_0 - b \ln R_0 \left(1 - \mu + \frac{\bar{\mu}\nu}{\nu + E_T - 1} \right) + P(1 - \delta) \right) \frac{-b \ln R_0 \mu \nu}{(\nu + E_T - 1)^2} \frac{dE_T}{dT} - \frac{c}{1 - \mu} = 0 \\ c &= \mu[1 - \mu]P \times f_{a_i} \left(u_0 - b \ln R_0 \left(1 - \mu + \frac{\bar{\mu}\nu}{\nu + E_T - 1} \right) + P(1 - \delta) \right) \frac{-b \ln R_0 \nu}{(\nu + E_T - 1)^2} Q_T \end{aligned}$$

Provided that $f_{a_i}(\cdot)$ is bounded ⁸, the RHS of this equation is decreasing towards 0 in T . Thus for higher values of c will require an earlier abandonment time T .

We next turn to the firm's choice of debugging effort μ :

$$\begin{aligned} \frac{\partial \Pi}{\partial \mu} &= P \times f_{a_i} \left(u_0 - b \ln R_0 \left(1 - \mu + \frac{\bar{\mu}\nu}{\nu + E_T - 1} \right) + P(1 - \delta) \right) (-b \ln R_0) \left(-1 + \frac{\nu}{\nu + E_T - 1} \right) - \frac{T}{(1 - \mu)^2} \\ c &= \frac{[1 - \mu]^2}{T} P \times f_{a_i} \left(u_0 - b \ln R_0 \left(1 - \mu + \frac{\bar{\mu}\nu}{\nu + E_T - 1} \right) + P(1 - \delta) \right) (-b \ln R_0) \left(-1 + \frac{\nu}{\nu + E_T - 1} \right) \end{aligned}$$

⁸ $\left(1 - \mu + \frac{\bar{\mu}\nu}{\nu + E_T - 1} \right)$ is decreasing towards $1 - \mu$ in T , so the value of $f_{a_i}(\cdot)$ is increasing towards $f_{a_i}(u_0 - b \ln R_0(1 - \mu) + P(1 - \delta))$ in T , provided we are on the right tail of F_{a_i}

$[1 - \mu]^2$ is decreasing towards 0 on $(0, 1)$; as long as f_{a_i} is bounded⁹, higher values of c will force the firm to chose a lower value of μ .

These first order necessary conditions will therefore pin down the optimal steady-effort strategy, $\sigma(a, c, R_0) = (\bar{\mu}, T)$, with corresponding final reliability $R(\sigma(a, c, R_0))$.

2.7.1 Timing and Trust

Firms facing fully-informed consumers can safely ignore the precise timing of adoptions¹⁰, since Q_∞ is eventually sold regardless, but the exact timing will become crucial to the asymmetric equilibrium below.

Let $Q(T) = Q_\infty$ provided that support is maintained until T .

As a thought experiment, suppose that at time t , consumers have total trust that the firm will support the product until time T^* consistent with the first order conditions outlined above. All $Q(T^*)$ consumers with $\frac{u_0}{1-\delta} - P < V(1, a_i, t, T^*)$ will adopt immediately, and the firm will have no incentive to support the product further, resetting $T = t$. If consumers with $V(1, a_i, t, T = t) < \frac{u_0}{1-\delta} - P < V(1, a_i, t, T = T^*)$ anticipate this, they will decide against adopting this period, so that only $Q(t)$ consumers will adopt during this period. This incentivizes the firm to choose a new value $T > t$, so that some of the $Q(T^*) - Q(t)$ switchers will now be incentivized to switch back.

I propose that an intermediate value, $Q(t, T) \in (Q(t), Q(T))$, exists such

- Adoption is slow enough that maintaining support until time T is incentive compatible for the firm.
- No consumer will wish to deviate to an earlier adoption date.

⁹Similarly, as $\mu \rightarrow 1$ $f_{a_i}(\cdot) \rightarrow f_{a_i}\left(u_0 - \frac{b \ln R_0 \nu}{\nu + E_t - 1} + P(1 - \delta)\right)$

¹⁰Note that faster adoption accelerates the growth of E_t ; this complex relationship is left for future work.

Further work will need to be done to determine the conditions necessary for the existence and uniqueness of an incentive-compatible adoption curve $Q(t, T)$.

2.8 Equilibrium with Asymmetric Information

When c is unobserved, a firm with a high c may wish to use a high initial effort level μ so consumers will impute a low c and high value of T .

Assume that customers observe R_0 , μ_{t-1} , μ_{t-2} , and $E[a|\{P_s, q_s\}_{s=0}^t]$, i.e.

1. R_0 : How “broken” was the title at launch?
2. μ_{t-1} How much effort has the firm put into recent patches?
3. μ_{t-2} Has their effort level changed recently?
4. $E[a|P, R_0\{q_s\}_{s=0}^t]$ Do sales suggest that the underlying concept has great sales potential?

Realistically, we can imagine R_0 being revealed by reviews. $E[a|P, R_0\{q_s\}_{s=0}^t]$ is justifiable by news, social media posts, “best seller” lists, and word-of-mouth. μ_{t-1} and μ_{t-2} could be revealed by official patch notes, as well as word-of-mouth perception.

- Each developer chooses management strategy $\mu = \bar{\mu}(a, c)$ for $t < T(a, c)$, and $\mu = 0$ for $t \geq T(a, c)$. Both $\bar{\mu}$ and T are strictly increasing in (a, c)
- Each consumer holds beliefs $E[c|\mu_t, \mu_{t-1}]$

Suppose that there are only two possible cost levels, $c \in \{c_L, c_H\}$, $c_H > c_L$, unobserved by the consumer. As before, each firm tries to maximize their profit for all sales:

$$\Pi[(a, c_j, R_0), (\mu_j, T_j)] = P \times Q_\infty(a, R_0, \mu_j, T_j) - \frac{cT_j}{1 - \mu_j}$$

yielding a strategy, (μ_L, T_L) for the “strong” firm with c_L . Suppose that c_H is sufficiently high that the “honest” strategy for the weak type is $(0, 0)$. The profits associated with these strategies is as follows:

$$\begin{aligned}\Pi[(a, c_H, R_0), (0, 0)] &= P * Q(T = 0) = P \left[1 - F_{a_i}(u_0 - b \ln R_0 + P(1 - \delta)) \right] \\ \Pi[(a, c_L, R_0), (\mu_L, T_L)] &= P * Q(T_L) - \frac{c_L T_L}{1 - \mu_L} \\ &= P \left[1 - F_{a_i} \left(u_0 - b \ln R_0 \left(1 - \mu + \frac{\bar{\mu} \nu}{\nu + E_T - 1} \right) + P(1 - \delta) \right) \right] - \frac{c_L T_L}{1 - \mu_L}\end{aligned}$$

2.8.1 Incentive Compatibility Constraints

From the fully-informed model, it is easy enough to show that each firm type would not prefer to copy the exact same strategy as the other type:

$$\begin{aligned}\Pi[(a, c_L, R_0), (\mu_L, T_L)] &\geq \Pi[(a, c_L, R_0), (0, 0)] \\ \Pi[(a, c_H, R_0), (0, 0)] &\geq \Pi[(a, c_H, R_0), (\mu_L, T_L)]\end{aligned}$$

A more profitable deviation from the separating equilibrium is for an H -type firm to choose $\mu_t = \bar{\mu}_L$ for $t < \tau < T_L$, i.e. for a weak developer to imitate the strong type for τ periods, resulting in the incentive compatibility constraint

$$P * Q(0, 0) \geq P * Q(\tau, T_L) - \frac{c_H \tau}{1 - \mu_L}$$

Further calculations require stronger assumptions on F_{a_i} to determine the optimal se-

quence of adoptions $Q(t, T), t = 0 \dots T$, but even this broad model offers several interesting properties. For example, The effect of (μ, T) is “scaled” by $b \ln R_0$, so a value of R_0 closer to 1 will reduce (μ_H, T_H) and make a pooling equilibrium more likely.

2.9 Conclusion

In this paper I present a novel dataset combining data on product adoption, product improvements, and

1. As software is patched, reliability improves log-linearly with usage (i.e. the “Duane Rule”)
2. The relationship between reliability and usage changes sharply at certain points, corresponding to changes in patching frequency
3. This can be explained by a repeated game of asymmetric information similar to a Spence model with productive signals

I lay out parameters for one possible separating equilibrium for this model. In future work, I hope to refine this equilibrium using more specific assumptions on the distribution of idiosyncratic preferences a_i .

2.10 Assumptions of PM-2 Discrete

To summarize AMSAA (2011) in my preferred notation, PM2D assumes that the product is subjected to a series of pass/fail tests, with expected probability of success $R(t)$ realized *after* each round of fixes. There are a finite number of bugs in the product, with probability of creating a failure in 1 test $p_1, p_2, \dots, p_k$, and $p_i \sim \text{Beta}(n, x)$, i.e. failure probabilities are drawn independently with pdf

$$f(p_i) = \frac{\Gamma(\nu)}{\Gamma(x)\Gamma(\nu - x)} p_i^{x-1} (1 - p_i)^{\nu-x-1}$$

with ν a known primitive and x an unknown primitive.

CHAPTER 3

Nonparametric Heterogeneity in Social Influence

3.1 Introduction

Since Everett Rogers published “Diffusion of Innovations” in 1962, an ever-growing body of literature has addressed the mechanisms of the gradual spread of new products and industrial processes through a market. One challenge in this field is incorporating heterogeneous actors into these models. Young (2009) proposes a “social influence” model where consumers delay adoption until the product has a certain number of preexisting users (the “threshold”). Young demonstrates how various properties of the distribution of thresholds across the population change the properties of the resulting adoption curve, and shows examples of “S-shaped” adoption curves generated by various parametric distributions.

In previous work, I find that the predictions of the social influence model (e.g. super-exponential growth) are consistent with data showing the adoption of various social video games on Valve Corporation’s *Steam* online marketplace.

1. Why would rational consumers follow a “threshold” adoption rule?
2. What does a realistic distribution of thresholds actually look like?

I address the first question by proposing a theoretical foundation where rational, forward-looking consumers with varying preference for the product and/or the product’s positive network effects but limited information will follow a threshold rule for adoption. In addition to supporting the credibility of social influence as an economic model, this work shows

For the second question, I “invert” Young’s model: given the number of users over time $n(t)$, my nonparametric estimator shows the underlying graph $G(n)$: when there are n users, fraction $G(n)$ of possible users are willing to adopt. From a theoretical standpoint, this allows us to see which parametric distributions are most consistent with real-world data. The results are also useful for firm strategy. How realistic are a growth projections for your product based on graphs of $G(\cdot)$ for earlier comparable products? Have recent promotions or product upgrades shifted the product to a higher growth path, or has adoption merely continued along a typical trajectory?

3.2 Literature Review

A large existing literature exists exploring how innovations and new products spread through a market (Peres, Muller, and Mahajan, 2010). There has been particular recent interest in the effect of the heterogeneity of agents, with characteristics of some agents leading to delayed adoption. Agents may wait to observe the outcomes of early adopters before adopting a product of unknown quality (social learning). A new product’s utility may also depend on a large user-base, so that consumers delay adoption to avoid being “first to the party” (network externalities).

In both social learning and network externalities the structure of the underlying social network is non-trivial. Jackson and Yariv (2007), for example, characterize the equilibrium strategies for a variety of games, including adoption with network externalities, depending on the distribution of adoption cost and interpersonal connections over actors. While I do not observe interpersonal links in my data, the mean-field approximations used in these models is useful for motivating the mechanics of aggregate adoption.

A landmark comparative analysis of diffusion models was provided by Young (2009). Four benchmark models of diffusion (Inertia, Contagion/Bass, Social Influence, and Social Learning) are generalized to allow for heterogeneous characteristics of agents. Young then

identifies unique characteristics of each model that should be observable in aggregate adoption trends. As an example, he compares each of the four models to Ryan and Gross' data on the adoption of hybrid corn.

In previous research I have replicated Young's analysis using the data detailed below, finding evidence in favor of the social learning and social influence models of diffusion in the market for video games. These effects are most pronounced in video games with social elements (multiplayer, user-generated content) and games relying on word-of-mouth advertising (independent studios, non-sequels, early access).

Finally, there is a small but growing literature on the empirical analysis of video games. Dubé, Hitsch, and Chintagunta (2010) and Clements and Ohashi (2005) both consider *indirect* network effects in the adoption of video game *hardware*, where more hardware owners motivate third parties to develop more software for each video game console. By contrast, I look at *direct* network effects in the market for *software*.

3.3 Setting and Data

The video game industry is a massive and growing sector, with US sales estimated at \$30 billion, larger than the combined revenue in recorded and live music (Association). Many of the most successful titles rely on word-of-mouth to advertise novel features: *Minecraft* was produced by a single programmer and never formally advertised, but its originality and robust community has helped it sell 121 million copies since 2009 (Warren)¹. Products may rely on competitive and social features, gradually building a persistent community of active users more akin to a social media network than a traditional media product. *Counter-Strike: Global Offensive*, for example, recorded a peak of 578,933 concurrent active users in August 2019, 71% of its all-time record (Charts, 2019), almost 7 years after release.

¹Early sales revenue was used to hire additional developers and artists to further improve *Minecraft*; this is an increasingly common business model which I hope to explore in future research.

I look at video games for Mac, PC, and Linux sold on Valve Corporation’s *Steam* online marketplace. Many of the most successful titles on this service rely heavily on the presence of other active users: players of *Team Fortress 2* go online to compete against other players, while fans of *Skyrim* have created thousands of “mods” which users can download to improve and expand graphics, art assets, and gameplay. Other successful “indie” titles such as *Hotline Miami* and *Goat Simulator* have benefited from word-of-mouth and press coverage to publicize titles with novel concepts, gameplay, and aesthetics from obscure independent studios.

My data consists of the number of current active users of each video game title, pulled from the Steam API on an hourly basis and archived by James Grey on *SteamCharts.com*, aggregated to the average and peak concurrent users for each month, focusing on monthly average users (MAU)². I used the top 2000 best-selling titles, removing observations for downloadable content, dedicated servers, etc. which show up in sales but not MAU. This leaves 1467 titles with between 1 and 69 months of observations.


An interesting question is whether adoption is driven by the number of *purchasers*, the number of *unique* users, or the number of *average active* users. Many Steam users reportedly buy titles on sale which they never play, and are thus unlikely to convince others to adopt, so purchasers may not be an effective metric. It would be interesting to see if a devout user-base of 1000 players online for 20 hours per month is more or less likely to spur adoption than a community of 20,000 players playing 1 hour per month; this question is left for future work. For the remainder of this paper, I will, with some imprecision, refer to the number of average users as the “number of users”. Note that if users average 1 hour per day on the app, this could be $\frac{1}{24}$ th of the actual user-base.

²This is somewhat of an abuse of notation, as the MAU of a digital product more commonly refers to [Unique] Monthly *Active* Users


Data per Title					Categories and User Tags	
(1)						m
	mean	sd	min	max		
Months	14.75	21.14	1	69	multiplayer	0.
Mean Monthly Avg. Users	4071.3	32965.1	101.0	749414.6	singleplayer	0.
SE Monthly Avg. Users	567.9	6420.4	0.105	170289.6	achievements	0.
<i>Titles</i>	1467				coop	0.
					mmo	0.
					mods	0.
					early_access	0.
					free	0.
					indie	0.

This paper will primarily focus on products with S-shaped adoption curves like *Counter-Strike: Global Offensive* or *Brawlhalla*, titles which started with relatively small user-bases but rapidly grew in popularity, as opposed to titles like *The Witcher 3* which, like blockbuster movies, enjoyed with large openings followed by a long decline as users move on to other titles (Fig. 3.1).

Figure 3.1: MAU evolution of a “blockbuster” versus a “sleeper hit”. Blue lines indicate new content released for *The Witcher 3*.



/2/2



/2/2

3.4 Social Influence

Suppose that we have a unit mass of individuals who could potentially adopt a new title. Consumers observe the adoption choices of others, and delay adoption until they have observed $\bar{n}$ adoption decisions, with two possible motivations. Under a social learning model, customers vary in their expectations of product quality and risk aversion, and each successful observed adoption increases the consumer's *expectation* of the product's quality, as in Ameri, Honka, and Xie (2018). By contrast, in a network effects model the product's *actual* quality is increasing in the number of fellow users, as in Liu, Mai, and Yang (2015). The theoretical section of this paper focuses on this explanation, which is to be expected for

products built around peer-to-peer interactions (like online multiplayer) or user-generated content (like “mods”, fan-made modifications adding or changing content and features in existing software).

Assume that the threshold adoption level $\bar{n}$ will be distributed across the population with cdf $G(\bar{n})$. At time t , we will observe $n(t)$ active users. For the purposes of this exercise, I will assume that active users remain active. $\lambda \in (0, 1)$ will represent an arbitrary delay between decision to adopt and actual adoption. The diffusion process is thus defined by the differential equation

$$\dot{n}(t) = \lambda [G(n(t)) - n(t)], \quad t > 0 \quad n(0) = n_0 \quad (3.1)$$

Taking the derivative with respect to time, we can solve for the PDF $g(n) \equiv G'(n)$:

$$\ddot{n}(t) = \lambda [g(n(t)) \dot{n}(t) - \dot{n}(t)] \quad \frac{\ddot{n}(t)}{\lambda \dot{n}(t)} + 1 = g(n(t)) \quad (3.2)$$

Peyton Young (2009) presents this equation to show the theoretical properties of a “Social Influence” model under a wide variety of parametric distributional assumptions. I wish to recover the threshold distribution $G(n)$ itself. Under ideal conditions, $n(0)$ is observed in the data and the relative acceleration rate $\frac{\ddot{n}}{\dot{n}}$ is a linear transformation of $g(n)$ with a single unknown parameter $\lambda \in (0, 1)$. λ should be chosen such that the implied CDF is between 0 and 1 and nondecreasing.

$$0 \leq G(n(t)) = \int_0^{n(t)} \left[\frac{\ddot{n}(t)}{\lambda \dot{n}(t)} + 1 \right] dn + n_0 \leq 1 \quad \forall(t) g(n(T)) \geq 0 \quad (3.3)$$

The SteamCharts monthly data has several characteristics which prevent a straightforward implementation of this method. The time series are noisy, with minor downward slopes during otherwise steady increases, and we have only a handful of observations during the critical opening months. As we are using first and second derivatives it is imperative that the adoption curve be as smooth as possible, and analytical derivatives would be superior to finite differences. For these reasons I approximate the data with a Gaussian kernel, which I can then differentiate to obtain the first and second derivative.


$$u_s = \frac{s - t}{h} \quad (3.4)$$

$$\text{Let } k(u_s) = \phi(u_s), \text{ thus } \frac{dk(u_s)}{dt} = \frac{k(u_s)u_s}{h} \quad (3.5)$$

$$\hat{n}(t) = \frac{\sum_s k(u_s)n(s)}{\sum_s k(u_s)} \quad (3.6)$$

$$\dot{\hat{n}}(t) = \frac{\left[\sum_s \frac{k(u_s)u_s n(s)}{h} \right] \left[\sum_s k(u_s) \right] - \left[\sum_s k(u_s)n(s) \right] \left[\sum_s \frac{k(u_s)u_s}{h} \right]}{\left[\sum_s k(u_s) \right]^2} \quad (3.7)$$


with $\ddot{\hat{n}}(t)$ defined similarly.




I implemented this solution in R with bandwidths between 1 and 10; visual inspection of

$\hat{n}(t)$ suggests that a bandwidth of $h = 5$ captures interesting variation without over-fitting. One could imagine a theoretical story where the effects of network effects and social learning are delayed by several months, or consumers with satisfied thresholds nevertheless delay until their next vacation or the title goes on sale.

Figure 3.2: Smoothed adoption of *Counter-Strike: Global Offensive*



The main feature of the first derivative of this estimator is bell curve with a peak in month 29 (January 2015) of 19.8%, or 79,658 individuals per month adopting, though the kernel estimator overshoots the first-differences peak considerably. Two peaks in months 54 and 66 are ignored by higher bandwidths.





The graph of the finite-differences-of-finite-differences approximation of $\ddot{n}(t)$ is dominated by noise. By contrast, the kernel estimates show a simpler story of rising acceleration of adoption in the first 25 months followed by deceleration and gradual stabilization at an equilibrium.

With these smoothed derivatives in hand, I manually tried a variety of values for λ to satisfy both equations in 3.3, settling on a value of 0.7. For reasons discussed below, I only show the data to month 37. The resulting PDF is shown.

The distribution of thresholds appears to have two distinct modes: a point mass at 0% and a peak at 28.8%, or 116,055 users. This is consistent with two broad groups of users: players loyal to the series³ who adopt as soon as possible, and players seeking a game with a large active user-base.

³Global Offensive is the fourth game in the *Counter-Strike* series, which started in 2000



3.5 Empirical Discussion

Between month 40 and 50, the raw data shows a substantial *decline* in user-base, followed by a recovery and convergence with the long run equilibrium. While smaller non-monotonicities disappear with higher bandwidths h , this dip is at most reduced to $\dot{\hat{n}} = 0$. This causes our estimate $\hat{g}$ to explode at these values, and if the rate adoption dips negative $\hat{g}$ jumps wildly between very negative and very positive values.

One option to address this is a kernel estimator with a monotonicity constraint, like Hall and Huang (2001). However, this would likely leave us with estimates of $\hat{n}$ close to 0, making it hard to fulfill either restriction in (3). The higher-level issue is that the social influence model does not explain the observed behavior of an adoption process grinding to a halt, then resuming its advance after a delay of half a year. I posit that around this time consumers began to quit the title, but the developer added new features substantial enough that the product was effectively restarted on a new adoption path with a higher $G(n(t))$.

This suggests several directions for further research:

1. Modelling the adoption and abandonment of a product simultaneously, with social learning increasing adoption while network effects both boost adoption and curtail abandonment, similar to Ainslie, Drèze, and Zufryden (2005)
2. Using data on price promotions and product updates, model a product which is shifting to a new $G(n)$ which is first order stochastic dominant over their previous $G(n)$, i.e. each consumer is willing to join at a lower threshold community size. Since we only observe a single time series for each title, this will likely require parametric assumptions on $G(n|\theta)$, where θ includes price, new content, etc.
3. In September 2017 (week 62) *CS: GO* was released in China for the first time. With more substantial data on the geography of active users and release dates, it would be interesting to see how diffusion mechanisms are affected by international borders, with multiple diffusion processes possibly superimposed on our single international aggregate series.


There is some overlap between the practical justifications of our adoption “inertia” term λ and our filter bandwidth h , as both result in a mismatch between a consumer’s threshold-crossing date and their observed adoption date. One could imagine a story like Nair (2007) where consumers waiting for promotions or free time could delay adoption by several months.

An asymmetric kernel might combine the function of both variables, reflecting a data generating process where adoption is driven by the *realized* adoption decisions of their peers, not expectations, and thus the unobserved delays result in later adoption, never earlier.

Relatedly, the first 6 observations of our estimated PDF are high and smoothly decreasing. It is unclear if this is an artifact of a point mass at $n = 0$ erroneously smoothed by the Kernel density estimator (too high h and λ) or a real aspect of the market, with many users willing to join a game at 4% adoption but not 0%.

Finally, it should be noted that we get substantially different results when we include the beta test in July 2012, when a much smaller number of users tried out the incomplete game. Many of the titles in our data have *years* of beta/early access testing before formal release; these “early access” games can develop substantial followings long before they are supposedly “complete”. It would be interesting to see if our estimated diffusion model is impacted by the shift from beta to released product.

Figure 3.3: Results for *CS: GO* including 1-month pre-release beta test



3.6 A Microeconomic Foundation

The social threshold model is typically presented as an approximation without microeconomic foundations⁴. I outline a new model where rational consumers with varying preference for the product’s social and non-social (“intrinsic”) features and limited information will follow a social threshold rule. Inspired by Doraszelski and Judd (2012), I use continuous time to avoid the problem of simultaneous decisions.

Suppose a consumer is considering adopting a network good which is already used by fraction $n(t)$ of all possible users at time t . Users and non-users receive the following utility streams:

$$\text{Use: } u_i(t) = x_i + y_i n(t)$$

$$\text{Do not use: } u_{i0}(t) = u_0$$

where $x_i \in [0, \infty)$, $y_i \in [0, \infty)$, and $u_0 \in [0, \infty)$. Note that the only cost of adopting the product is u_0 , an assumption with several possible justifications:

- The product may be a free ad-supported service like Facebook
- The product may be a subscription service, with the per-period cost included in u_0 , like *World of Warcraft*
- The consumer could be deciding whether to spend time using a product they purchased on sale but have not yet used, a common scenario for Steam users.

⁴See Young (2009): “A notable limitation of the contagion and social influence models is that they provide no clear reason why people would adopt an innovation given that others have adopted it”

Let there be a large number N of total consumers in the market. Let decision opportunities arrive for each user via a Poisson process with rate λ , so that an average of λN consumers are allowed to switch between the adoption and non-adoption each period. λ may reflect fluctuations in the user's free time, weekend sales, or marketing which temporarily brings the product and/or an outside good to an individual consumer's attention.

Trivially, there is an equilibrium where all adopters continue using the product after adoption, so that $n(t)$ is nondecreasing.

3.6.1 Myopic Consumers

A myopic consumer adopts the product if $x_i + y_i n(t) \geq u_0$. Note that the "sample" of $N\lambda$ consumers per period is taken over the entire population, including both users and non-users.

3.6.1.1 Varying Intrinsic Preferences

First, suppose that $x_i \sim F(\cdot)$, $y_i = y > 0$, i.e. consumers vary by their intrinsic value for the product, but all consumers have the same preference for a popular product. The rate of additional adoptions is:

$$\begin{aligned} \frac{dn}{dt} &= \lambda \left[\int_0^\infty \mathbb{I}[x_i + yn(t) \geq u_0] dF - n(t) \right] \\ &= \lambda \left[1 - F(u_0 - yn(t)) - n(t) \right] \\ &= \lambda \left[G(n(t)) - n(t) \right], \quad G(n) \equiv 1 - F(u_0 - yn) \\ \lim_{n \rightarrow -\infty} G(n) &= 0, \quad \lim_{n \rightarrow \infty} G(n) = 1, \quad \forall n G'(n) = F'(u_0 - yn) \geq 0 \end{aligned}$$

Thus $G(n)$ can be called a cumulative distribution function with respect to n .

3.6.1.2 Varying Social Preferences

Now, consider a world where $x_i = x$ and $y_i \sim F_Y(y_i)$, i.e. consumers vary only in the value they place on the “social” aspects of the product. The rate of new product adoption is thus:

$$\begin{aligned}\frac{dn}{dt} &= \lambda \left[\int_0^\infty \mathbb{I}[x + y_i n(t) \geq u_0] dF_Y - n(t) \right] \\ \frac{dn}{dt} &= \lambda \left[1 - F_Y\left(\frac{u_0 - x}{n}\right) - n(t) \right] \\ \frac{dn}{dt} &= \lambda \left[G(n(t)) - n(t) \right], \quad G(n) \equiv 1 - F_Y\left(\frac{u_0 - x}{n}\right) \\ \lim_{n \rightarrow 0} G(n) &= 0, \quad \lim_{n \rightarrow \infty} G(n) = 1, \quad (\forall n > 0) G'(n) = F'_Y\left(\frac{u_0 - x}{n}\right) \frac{u_0 - x}{n^2}\end{aligned}$$

$G(n)$ is thus a CDF with respect to n if and only if $u_0 \geq x$; otherwise, we are in the trivial “inertial” adoption case, $\frac{dn}{dt} = \lambda(1 - n)$. It should also be noted that $n \in [0, 1]$, so there may still be some stationary point $F_Y\left(\frac{u_0 - x}{n}\right) = n$ at $n < 1$ where adoption will grind to a halt.

3.6.1.3 Varying Intrinsic and Social Preferences

Suppose consumers are myopic and $x_i \in (0, \infty)$ and $y_i \in (0, \infty)$ are drawn independently from the distributions $F_X(\cdot)$ and $F_Y(\cdot)$, respectively.

$$\begin{aligned}\frac{dn}{dt} &= \lambda \left[\int_0^\infty \int_0^\infty \mathbb{I}[x_i + y_i n(t) \geq u_0] dF_X dF_Y - n(t) \right] \\ &= \lambda \left[\int_0^\infty \left(1 - \int_0^{u_0 - y_i n(t)} dF_X \right) dF_Y - n(t) \right] \\ &= \lambda \left[\int_0^\infty \left(1 - F_X(u_0 - y_i n(t)) \right) dF_Y - n(t) \right]\end{aligned}$$

Provided that moment generating functions for x_i and y_i $M_X(t), M_Y(t)$ exist, then the variable $u_i = x_i + y_i n$ has moment generating function

$$M_{u_i}(t) = M_X(t)M_Y(nt) = E[e^{tx_i}] E[e^{nty_i}]$$

The distribution of thresholds over consumers, $G(n)$, is not solvable in closed form, but rather identified by finding G such that

$$(\forall t) \int e^{tx} dF_X(x) \int e^{tny} dF_Y(y) = \int e^{tn} dG(n)$$

We can also easily calculate an upper bound of the adoption rate using the Chernoff inequality:

$$\frac{dn}{dt} = \lambda \left[\Pr(x_i + y_i n(t) \geq u_0) - n(t) \right] \leq \min_{t>0} e^{-u_0 t} M_X(t) M_Y(nt)$$

3.6.1.4 Summarizing Myopic Consumers

Assuming consumers adopt the product whenever they get the opportunity and the current flow of utility is positive, I have presented three cases: varying intrinsic value, varying social value, and varying intrinsic and social value independently⁵. In each case, adoption takes the canonical “social influence” form. Each consumer’s characteristics are projected onto a single value, the popularity threshold necessary for adoption, $\bar{n}_i$, such that adoption will occur if and only if $n(t) \geq \bar{n}_i$. The distribution of $\bar{n}$, $G(\bar{n})$, thus dictates the shape of the adoption curve.

⁵A fourth option, jointly varying intrinsic and social value, is left to future work.

3.6.2 Forward-Looking Consumers

Consider consumer i being offered the chance to adopt the product at time t_i . Since adoption is free (save for the ongoing opportunity cost u_0) and reversible, then i only needs to worry about how the value of the product $u_i(t)$ will vary between t_i and the (unknown) date of the next decision opportunity t'_i .

Let $\tau = t'_i - t_i$; note that $\tau \sim \text{Exponential}(\lambda)$, and thus $r\tau \sim \text{Exponential}(\frac{\lambda}{r})$, and finally $e^{-r\tau} \sim \text{Beta}(\frac{\lambda}{r}, 1)$.

Given discount rate r , the consumer will thus adopt if

$$V(x_i, y_i, t_i, n) \equiv \mathbb{E}_{\tau, n(t)} \left[\int_0^\tau e^{-rs} (x_i + y_i n(t_i + s)) ds \right] \geq u_0 \mathbb{E}_\tau \left[\frac{1}{r} - \frac{e^{-r\tau}}{r} \right] = u_0 \left[\frac{1}{r} - \frac{\lambda}{r(\lambda + r)} \right] = \frac{u_0}{\lambda + r}$$

$$y_i \mathbb{E}_{\tau, n(t)} \left[\int_0^\tau e^{-rs} n(t_i + s) ds \right] \geq \frac{u_0 - x_i}{\lambda + r}$$

If $x_i + y_i n(t_i) \geq u_0$ and $n(t)$ is nondecreasing, then $\left(\forall t \in [t_i, t'_i] \right) x_i + y_i n(t) \geq u_0$: thus if a given myopic consumer would adopt at t_i , then a forward-looking consumer with the same preferences would adopt as well.

The complementary case, $x_i + y_i n(t_i) < u_0$, presents a more difficult question: if i expects $n(t)$ to grow sufficiently between t_i and t'_i , they may still try adopting. Importantly, i does not observe the distribution of thresholds $G(n(t))$. Like the myopic consumer, we can assume the forward-looking consumer knows their own preferences x_i, y_i and the current user-base $n(t_i)$, but they do not know their place in either distribution⁶. This is frustrating for the forward-looking consumer, because, as Young (2009) notes, depending on $G(\cdot)$ and $n(t)$, actual growth may be anywhere from nonexistent to superexponential!

Let the consumer take the naive assumption that growth will be approximately expo-

⁶While $n(t)$ is presented as a fraction of total possible consumers, depending on $G(n_t)$, adoption may taper off well short of $n(t) = 1$

nential⁷ for $t \in [t_i, t'_i]$, so that $n(t) = n(t_i)e^{\alpha(t-t_i)}$ for some $\alpha < r$. The value of adoption is thus:

$$\begin{aligned} V(x_i, y_i, t_i, n) &= y_i E_\tau \left[\int_0^\tau e^{-rs} n(t_i) e^{s\alpha} ds \right] = y_i n(t_i) E_\tau \left[\int_0^\tau e^{-s(r-\alpha)} ds \right] \\ &= y_i n(t_i) E_\tau \left[\frac{e^{-\tau(r-\alpha)} - 1}{\alpha - r} \right] = \frac{y_i n(t_i)}{\alpha - r} \left(\frac{\lambda}{\lambda + r - \alpha} - 1 \right) = \frac{y_i n(t_i)}{\alpha - r} \left(\frac{\alpha - r}{\lambda + r - \alpha} \right) = \frac{y_i n(t_i)}{\lambda + r - \alpha} \end{aligned}$$

The same equation also holds when expectations are higher ($\alpha > r$). $e^{\tau(\alpha-r)}$ will be Pareto; provided that $\frac{\lambda}{\alpha-r} > 1$, the value of adopting is:

$$V(x_i, y_i, t_i, n) = y_i n(t_i) E_\tau \left[\frac{e^{\tau(\alpha-r)} - 1}{\alpha - r} \right] = \frac{y_i n(t_i)}{\alpha - r} \left(\frac{\lambda}{\lambda - \alpha + r} - 1 \right) = \frac{y_i n(t_i)}{\lambda + r - \alpha}$$

Thus the forward-looking consumer adopts when

$$\begin{aligned} V(x_i, y_i, t_i, n) &= \frac{y_i n(t_i)}{\lambda - \alpha + r} \geq \frac{u_0 - x_i}{\lambda + r} \\ n(t_i) \frac{(\lambda + r)}{\lambda - \alpha + r} &\geq \frac{u_0 - x_i}{y_i} \end{aligned}$$

This is a simple monotonic transformation of the myopic customer's threshold rule, and thus the forward-looking customer's adoption decisions can also be aggregated to form the Social Influence model $\frac{dn}{dt} = \lambda(G(n) - n)$.

Note that a higher value for α means that a lower threshold $n(t_i)$ is necessary for adoption, while higher discount rates r and decision frequencies λ reduce the adoption threshold towards that of the equivalent myopic consumer.

⁷Note that this is most accurate in a neighborhood where $G(n_0 + \epsilon) \approx \beta(n_0 + \epsilon)$, $\beta > 1$

3.6.2.1 Overzealous Adoption

If an adopter is very disappointed in the realized adoption rate ($n(t'_i) \ll n(t_i)e^{(t'_i-t_i)\alpha}$) they may reverse their adoption decision; if enough users do so, the assumption that $n(t)$ is non-decreasing may be violated. This “retreat” could potentially be used to explain transient non-monotonicities observed in the Steam adoption data. More nuanced consumer expectations, suggested below, may reduce the expected number of de-adoption events.

3.6.2.2 Consumer Expectations

Numerical simulations could be used to find a value for α , the naive expected short-term exponential growth rate, such that consumers using α for adoption decisions will have adoption curve $n(t)$ and $\hat{n}(t) = n(t_0)e^{(t-t_0)\alpha}$ is a relatively reasonable short-run estimate of $n(t)$ across the universe of consumer decision events.

The assumption that all consumers share a single naive expected short-term exponential growth rate α is rather strong. A consumer may observe $\frac{dn(t_i)}{dt}$ and adjust α accordingly, or use their own value x_i, y_i to make inferences about the distribution of G . Varying α is left to future work. Other functional forms of the consumer’s naive projection of $n(t)$ are also omitted for now.

3.7 Conclusion

Using data from the adoption of a new video game with prominent social features, I demonstrated that a “social influence” diffusion model can be inverted to construct a plausible distribution of heterogeneous consumer adoption preferences. Testing the model on data from *Counter-Strike: Global Offensive*, we find two modes of adoption: one group adopting at release, and a second adopting around 28.8% penetration, or 116,055 monthly average users. This is consistent with a wave of loyal fans of the decade-old series adopting the

product on release, followed by a longer drawn-out wave of consumers willing to adopt the product based on substantial positive network effects.

I also provide a novel microeconomic foundation, so that the distribution of primitive preferences over rational, forward-looking consumers with limited information generates a social influence adoption curve.

The next steps will be focused on implementing this model on a larger and more varied set of adoption curves in my data. This will be easier to do with a rule-of-thumb for selecting h , automating the selection of λ to satisfy my CDF constraints, and using the solution in Hall and Huang (2001) to guarantee that the smoothed diffusion process is non-decreasing until adoption reaches 100%.

With hundreds of estimates $(\lambda, \hat{g}(n), n(0))$, we can ask comparative questions like:

1. Is there substantial overlap in appropriate values of λ across titles?
2. Do the derived PDFs $\hat{g}$ look “sane” and relatively simple?
3. Do different titles exhibit similar $\hat{g}$ curves up to a truncation at $n = n(0)$?
4. Which parametric distributions could be a reasonable approximation of $\hat{g}$?
5. Is there any difference between the $\hat{g}$ curves driven by network effects (big-studio multiplayer) and social learning (independent-studio singleplayer)?
6. If a title is continuing to grow, can we use early data to predict the long-run equilibrium user-base of the product?

It should be noted that these findings should be of use far outside the video game industry. Products and services like R and Ebay rely heavily on user-generated content and peer-to-peer interactions, respectively. Physical products like “hoverboards” have been successfully marketed via interactions between early adopters and their peers. Social influence

will be a driving factor in the success of new products in a wide variety of industries, and understanding how consumer heterogeneity affects diffusion will be important for optimizing pricing and marketing efforts in a social consumption environment.

.1 Crow-AMSAA MLE

The Crow-AMSAA model is applicable when the “Duane Rule” holds: the time between failures is decreasing loglinearly in experience. We write this instantaneous failure rate as $\lambda_E = \lambda\beta E^{\beta-1}$.

Let E_i indicate the i th failure occurring during the E_i th launch.

$$\begin{aligned} Pr(E_i > x | E_{i-1} > y) &= \frac{e^{-\lambda E_i^\beta}}{e^{-\lambda E_{i-1}^\beta}} = e^{-\lambda(E_i^\beta - E_{i-1}^\beta)} = 1 - F(E_i | E_{i-1}) \\ \frac{\partial F(E_i | E_{i-1})}{\partial E_i} &= 0 - e^{-\lambda(E_i^\beta - E_{i-1}^\beta)}(-\lambda)E_i^{\beta-1}\beta \\ f(E_i | E_{i-1}) &= \lambda\beta E_i^{\beta-1} e^{-\lambda(E_i^\beta - E_{i-1}^\beta)} \end{aligned}$$

We observe a sequence of failure events $(E_1, \dots, E_n)$ from a sample of E^* total launches with likelihood

$$L = \prod_{i=1}^n \beta \lambda T_i^{\beta-1} \exp\left(-\lambda(E_i^\beta - E_{i-1}^\beta)\right) = \beta^n \lambda^n e^{-\lambda E^{*\beta}} \prod_{i=1}^n E_i^{\beta-1}$$

Taking the derivatives with respect to the primitives λ and β , we have

$$\begin{aligned} \ln(L) &= n \ln(\beta) + n \ln(\lambda) - \lambda E^{*\beta} + \sum_{i=1}^n (\beta - 1) \ln(E_i) \\ \frac{\partial \ln(L)}{\partial \beta} &= \frac{n}{\beta} - \lambda E^{*\beta} \ln(E^*) + \sum_{i=1}^n \ln(E_i) = 0 & \frac{\partial \ln(L)}{\partial \lambda} &= \frac{n}{\lambda} - E^{*\beta} = 0 \\ n &= \lambda E^{*\beta} & \frac{n}{\beta} &= n \ln(E^*) - \sum_{i=1}^n \ln(E_i) \\ \hat{\beta} &\equiv \frac{n}{n \ln E^* - \sum_{i=1}^n \ln E_i} & \hat{\lambda} &\equiv \frac{n}{E^{*\hat{\beta}}} \end{aligned}$$

Note that we must un-bias these estimates by $\frac{N-1}{N}\hat{\beta} = \bar{\beta}$ for time-terminated data,

$\frac{N-2}{N-1}\hat{\beta} = \bar{\beta}$ for failure-terminated data.

.1.1 Pooled Estimator

Assume instead that we observe the history of multiple product families $j \in \{1, \dots, N\}$:

$(E_{1,1}, E_{2,1}, \dots, E_{n_1,1}), (E_{1,2}, E_{2,2}, \dots), \dots (E_{1,j}, \dots, E_{n_j,j})$, and assume that λ_j varies between product families while β is a universal constant.

$$\begin{aligned}
L &= \prod_j \prod_{i=1}^{n_j} \beta \lambda_j T_{i,j}^{\beta-1} \exp \left(-\lambda_j \left(E_{i,j}^\beta - E_{i-1,j}^\beta \right) \right) = \beta^{\sum_j n_j} \prod_j \lambda_j^{n_j} e^{-\lambda_j E_j^{*\beta}} \prod_{i=1}^{n_j} E_{i,j}^{\beta-1} \\
\ln(L) &= \ln(\beta) \sum_j n_j + \sum_j \left[n_j \ln(\lambda_j) - \lambda_j E_j^{*\beta} + \sum_{i=1}^{n_j} (\beta - 1) \ln(E_{i,j}) \right] \\
\frac{\partial \ln(L)}{\partial \lambda} &= \frac{n_j}{\lambda_j} - E_{*j}^\beta = 0 \\
\frac{n_j}{E_{*j}^\beta} &= \hat{\lambda}_j \\
\frac{\partial \ln(L)}{\partial \beta} &= \frac{\sum_j n_j}{\beta} + \sum_j \left[-\lambda_j \ln(E_{*j}) E_{*j}^\beta + \sum_i^{n_j} \ln(E_{i,j}) \right] = 0 \\
\frac{\sum_j n_j}{\beta} &= \sum_j \lambda_j \ln(E_{*j}) E_{*j}^\beta - \sum_i^{n_j} \ln(E_{i,j}) = \sum_j \left[n_j \ln(E_{*j}) - \sum_i^{n_j} \ln(E_{i,j}) \right] \\
\hat{\beta} &= \frac{\sum_j n_j}{\sum_j \left[n_j \ln(E_{*j}) - \sum_i^{n_j} \ln(E_{i,j}) \right]}
\end{aligned}$$

.2 Interpolating Orbital Capability

First, consider the linear velocity of a satellite in orbit. The *vis-viva* equation is derived from the conservation of energy:

$$v^2 = GM \left(\frac{2}{r} - \frac{1}{A} \right)$$

where A is the semi-major axis of the ellipse, r is the distance from earth's center, G is the gravitational constant, and M is the earth's mass. Let $GM = \mu$ for simplicity.

Starting from a circular orbit $r_0 = A$ with a velocity of v_0 , suppose a rocket fires to raise its trajectory into an elliptical orbit. Immediately after firing, the rocket is still at its lowest point in the orbit (periapsis) r_0 above the center of the Earth, but with a higher current velocity v_1 and thus a higher apoapsis of r_1 . We can calculate the necessary change in velocity to achieve this:

$$\begin{aligned} v_0^2 &= \mu \left(\frac{2}{r_0} - \frac{1}{r_0} \right) = \mu \left(\frac{1}{r_0} \right) \\ v_1^2 &= \mu \left(\frac{2}{r_0} - \frac{2}{r_0 + r_1} \right) = \mu \left(\frac{2(r_0 + r_1) - 2r_0}{r_0(r_0 + r_1)} \right) = \mu \left(\frac{2r_1}{r_0(r_0 + r_1)} \right) \\ \Delta v(a_0, a_1) &= v_1 - v_0 = \frac{\sqrt{\mu}}{\sqrt{r + a_0}} \left(\frac{\sqrt{r + a_0}}{\sqrt{2r + a_0 + a_1}} - 1 \right) \end{aligned}$$

Δv sums linearly when a rocket performs multiple maneuvers, serving as a “budget constraint” for the rocket's mission: $\Delta v_T \leq \sum \Delta v_i$ for a sequence of maneuvers i . Δv_T is defined by the characteristics of the rocket and payload:

$$\begin{aligned} \Delta v &= v_e \ln \left(\frac{m_p + m_s + m_f}{m_p + m_s} \right) = -v_e \ln \left(1 - \frac{m_f}{m_p + m_s + m_f} \right) \approx v_e \frac{m_f}{m_p + m_s + m_f} \\ m_p &= \frac{v_e m_f}{\Delta v} - m_s - m_f \end{aligned}$$

where the mass of the payload is m_p , the dry mass of the final stage is m_s , the mass of fuel and oxidizer is m_f , and the velocity of the exhaust is v_e .

Suppose a given rocket can launch payloads m_{p1} , m_{p2} to elliptical orbits with impulse

requirements $\Delta v_1, \Delta v_2$.

$$m_{p1} - m_{p2} = \frac{v_e m_f}{\Delta v_1} - \frac{v_e m_f}{\Delta v_2} = \frac{v_e m_f (\Delta v_2 - \Delta v_1)}{\Delta v_1 \Delta v_2}$$

$$\frac{(m_{p1} - m_{p2}) \Delta v_1 \Delta v_2}{\Delta v_2 - \Delta v_1} = v_e m_f$$

To further simplify, let's assume that the rocket's first stage launches the vehicle to the same trajectory regardless of payload⁸. The final stage contributes Δv_{LEO} to reach a circular orbit, then contributes Δv_{GTO} to reach *GTO*.

If m_f and v_e are known, then we can estimate $\Delta \hat{v}_{LEO}$ such that

$$m_{pGTO} = m_{LEO} - \frac{v_e m_f (v_{pGTO} - \hat{v}_{pLEO})}{\Delta \hat{v}_{pLEO} \Delta v_{pGTO}} \quad (.8)$$

which can then be used to extrapolate

$$m_{px} = m_{LEO} - \frac{v_e m_f (\Delta v_x - \Delta \hat{v}_{LEO})}{\Delta \hat{v}_{LEO} \Delta v_x} \quad (.9)$$

m_f and v_e are available for a variety of rockets, and this method can be used to extrapolate to higher circular orbits like those used by GPS satellites, as well as "earth-escape" missions to other planets.

For now I ignore the role of launch trajectories, latitudes, and orbital planes. The latitude of a launch site limits the range of orbital inclinations that the rocket can initially launch to. Once in orbit, more fuel must be spent to change the inclination. For this reason, the US operates two major launch sites: Cape Canaveral in Florida is more suitable for launches over the equator, while Vandenberg is located at a higher latitude and is thus more suitable


⁸Note that the *Atlas V* payload can vary from 0-1% of its mass at takeoff, and thus has limited effect on first-stage performance

for launches at higher inclinations. Besides customer inclination requirements, launch sites are also non-trivial for rocket efficiency. The velocity of the Earth's surface at the launch site is subtracted from the Δv to reach orbit: a rocket launched from the equator needs 460 m/s less Δv than one launched from the North Pole, or 6% of the cost to reach orbit. This is complicated by trajectory requirements: the Israeli *Shavit* rocket must launch *west* over the Mediterranean to prevent debris from dropping onto neighboring countries, meaning that hundreds of m/s of Δv must be spent to fight against the Earth's rotation.

.3 Demand Trends


What are the determinants for launch vehicle demand? Historically, the most prominent driver of the demand was the Cold War, when NATO and the Soviet Union launched hundreds of satellites for spying, communicating between military commands, and providing navigational data to the armed forces. At the same time, their respective civilian space agencies launched vehicles at a breakneck pace to achieve PR coups like the first lunar landing and test new technology with potential military applications. The end of the Cold War led to a reduction in both national security and civilian launches.

Since the breakup of the Soviet Union corporate satellite launches have closely followed the the high-tech sector, with a bubble in the late 1990s including Intelsat, Globalstar, and Orbcomm (Lim, Klein, and Thatcher, 2005). After a lull in the 2000s, they have once more enjoyed rapid growth in the 2010s. Military satellites have continued to provide a small but consistent segment of demand, with increasing launches from new players like China, Japan, India, Israel, and Iran. Note that demand for launches to geostationary orbit



has been relatively stable over time. Because geostationary orbits are limited to a single

circle of orbits⁹, they risk collision and radio interference. Because of this, the International Telecommunications Union allocates orbit × radio frequency “slots” to governments and corporations on a first-come-first-served basis. The scarcity of these “slots” means that GEO launches are limited to a steady stream of replacements for existing satellites.




Trends in satellite mass differ sharply between GEO and LEO markets. The median mass of both security and civilian satellites in low earth orbit has been decreasing since the 1980s, likely due to miniaturization and other weight-saving solutions like ion engines. The 95th percentile has been decreasing for national security payloads and roughly stable for civilian payloads; note that the latter category includes manned spacecraft like the 7-ton *Soyuz* orbital vehicle which have remained roughly equal in size over time.

By contrast, GEO vehicles steadily increased in size since their introduction in 1961. Consider *Intelsat II F-1*, launched in 1966, compared to *Intelsat 36*, launched in 2016. Since there are few GEO “slots” available for satellites, operators pack more equipment onto each satellite: *Intelsat II* had 2 transponders using 85 watts, while *Intelsat 36* has 44 and requires

⁹0° inclination, 35,786 km x 35,786 km

15,800 watts of power. The later satellites have greater longevity: the *Intelsat II* series was designed to last 3 years, compared to 15 for *Intelsat 36*; this longevity requires more fuel for station-keeping and more redundant systems. As a result, weight has ballooned from 162 kg for *Intelsat II* to 3,253 kg for *Intelsat 36*.



Bibliography

- Ainslie, Andrew, Xavier Drèze, and Fred Zufryden (2005), “Modeling Movie Life Cycles and Market Share,” *Marketing Science*, 24 (3), 508–517, URL <https://doi.org/10.1287/mksc.1040.0106>.
- Ameri, Mina, Elisabeth Honka, and Ying Xie (2018), “Word-of-mouth, observational learning, and product adoption: Evidence from an anime platform,” .
- AMSAA (2011), *Department of Defense Handbook, Reliability Growth Management*, Army Materiel System Analysis Activity, Aberdeen Proving Ground, MD, mIL-HDBK-189C.
- Association, Entertainment Software (????), “U.S. Video Game Industry Generates \$30.4 Billion in Revenue for 2016,” <http://www.theesa.com/article/u-s-video-game-industry-generates-30-4-billion-revenue-2016/>, accessed: 2018-01-04.
- Bartelme, Dominick, Arnaud Costinot, Dave Donaldson, and Andres Rodriguez-Clare (2018), “Economies of Scale and Industrial Policy: A View from Trade,” Technical report, Working paper.
- Benkard, C Lanier (2004), “A dynamic analysis of the market for wide-bodied commercial aircraft,” *The Review of Economic Studies*, 71 (3), 581–611.
- Benkard, C Lanier, Przemyslaw Jeziorski, and Gabriel Y Weintraub (2015), “Oblivious equilibrium for concentrated industries,” *The RAND Journal of Economics*, 46 (4), 671–708.
- BIS (2020), “Commerce Control List,” URL <https://www.bis.doc.gov/index.php/regulations/export-administration-regulations-ear>.
- Boone, Tom (2017), “Financial Operational Losses in Space Launch,” .
- Boone IV, Tom R and David P Miller (2017), “Effects of Multiple Payload Launches on Launch Cost,” *New Space*, 5 (2), 89–98.

- Brumfiel, Geoff (2019), "Iranian Rocket Launch Ends In Failure, Imagery Shows," URL <https://www.npr.org/2019/08/29/755406765/iranian-rocket-launch-ends-in-failure-images-show>.
- Cates, Grant R., Daniel Houston, Douglas Conley, and Karen Jones (2018), "Launch Uncertainty: Implications for Large Constellations," *Center for Space Policy and Strategy*, URL <https://aerospace.org/paper/launch-uncertainty-implications-large-constellations>.
- Charts, Steam (2019), "Counter-Strike: Global Offensive," URL <http://steamcharts.com/app/730>, [Online; posted 7-August-2019].
- Clements, Matthew T and Hiroshi Ohashi (2005), "Indirect network effects and the product cycle: video games in the US, 1994–2002," *The Journal of Industrial Economics*, 53 (4), 515–542.
- Commission, Rogers (1986), "Report of the PRESIDENTIAL COMMISSION on the Space Shuttle Challenger Accident," URL <https://history.nasa.gov/rogersrep/v1ch5.htm>.
- Congress (????), "10 U.S. Code § 2273," URL "<https://www.law.cornell.edu/uscode/text/10/2273>".
- Crow, Larry H (1975), "Reliability analysis for complex, repairable systems," Technical report, ARMY MATERIEL SYSTEMS ANALYSIS ACTIVITY ABERDEEN PROVING GROUND MD.
- Dawson, Jeffrey W (2011), "The use of Crow-AMSAA plots to assess mishap trends," .
- Doraszelski, Ulrich and Kenneth L Judd (2012), "Avoiding the curse of dimensionality in dynamic stochastic games," *Quantitative Economics*, 3 (1), 53–93.
- Doraszelski, Ulrich and Ariel Pakes (2007), "A framework for applied dynamic analysis in IO," *Handbook of industrial organization*, 3, 1887–1966.

- Duane, JT (1964), “Learning curve approach to reliability monitoring,” *IEEE transactions on Aerospace*, 2 (2), 563–566.
- Dubé, Jean-Pierre H, Günter J Hitsch, and Pradeep K Chintagunta (2010), “Tipping and concentration in markets with indirect network effects,” *Marketing Science*, 29 (2), 216–249.
- Dummer, G., R. Winton, and M. Tooley (1997), *An Elementary Guide to Reliability*, Elsevier Science, URL <https://books.google.com/books?id=Yt5-gfqTy-IC>.
- Ericson, Richard and Ariel Pakes (1995), “Markov-perfect industry dynamics: A framework for empirical work,” *The Review of economic studies*, 62 (1), 53–82.
- FAA (2019), “2018 AST Compendium,” *AST*, URL https://www.faa.gov/about/office_org/headquarters_offices/ast/media/2018_AST_Compendium.pdf.
- Ferraz, Claudio, Frederico Finan, and Dimitri Szerman (2015), “Procuring firm growth: the effects of government purchases on firm dynamics,” Technical report, National Bureau of Economic Research.
- Galyonkin, Sergiy (2021), “Steam Spy,” URL <http://https://steamspy.com/>.
- Gould, Allen J and Orin M Linden (????), “Estimating satellite insurance liabilities,” .
- Greenberg, Joel S and Henry Hertzfeld (1992), *Space economics*, volume 144, Aiaa.
- Hall, Peter and Li-Shan Huang (2001), “Nonparametric kernel regression subject to monotonicity constraints,” *Ann. Statist.*, 29 (3), 624–647, URL <https://doi.org/10.1214/aos/1009210683>.
- Hamilton, Alexander (1791), “Report on Manufactures,” .
- Hanlon, W Walker (2015), “Necessity is the mother of invention: Input supplies and Directed Technical Change,” *Econometrica*, 83 (1), 67–100.

- Hansen, Jørgen Drud, Camilla Jensen, and Erik Strøjer Madsen (2003), “The establishment of the Danish windmill industry—Was it worthwhile?” *Review of World Economics*, 139 (2), 324–347.
- Hill, CN (2011), “The Cancellation of Black Arrow,” URL "<http://www.spaceuk.org/ba/blackarrowcancellation.htm>".
- (2012), *A Vertical Empire: History of the British Rocketry Programme*, World Scientific.
- Ifrach, Bar and Gabriel Y Weintraub (2012), “A framework for dynamic oligopoly in concentrated industries,” *Columbia Business School Research Paper*, (12/47).
- Irwin, Douglas A (2000), “Did late-nineteenth-century US tariffs promote infant industries? Evidence from the tinplate industry,” *The Journal of Economic History*, 60 (2), 335–360.
- Jackson, Matthew O and Leeat Yariv (2007), “Diffusion of behavior and equilibrium properties in network games,” *American Economic Review*, 97 (2), 92–98.
- Juhász, Réka (2018), “Temporary protection and technology adoption: Evidence from the napoleonic blockade,” *American Economic Review*, 108 (11), 3339–76.
- Kalouptside, Myrto (2018), “Detection and impact of industrial subsidies: The case of Chinese shipbuilding,” *The Review of Economic Studies*, 85 (2), 1111–1158.
- Kremer, Michael (1993), “The O-Ring Theory of Economic Development,” *The Quarterly Journal of Economics*, 108 (3), 551–575, URL <http://www.jstor.org/stable/2118400>.
- Krugman, Paul (1987), “The narrow moving band, the Dutch disease, and the competitive consequences of Mrs. Thatcher: Notes on trade in the presence of dynamic scale economies,” *Journal of development Economics*, 27 (1-2), 41–55.
- Kutter, Bernard F (2006), “Commercial Launch Services: an Enabler for Launch Vehicle Evolution and Cost Reduction,” .

- Lane, Rick (2019), “Game over, or not: what became of Steam’s first Early Access games?” <https://www.theguardian.com/games/2019/jan/28/steam-early-access-first-wave-games-kenshi-valve>, [Online; accessed 14-May-2021].
- Lim, Jaejoo, Richard Klein, and Jason Thatcher (2005), “Good technology, bad management: A case study of the satellite phone industry,” *Journal of Information Technology Management*, 16 (2), 48–55.
- Liu, Yong, Enping Mai, and Jun Yang (2015), “Network externalities in online video games: an empirical analysis utilizing online product ratings,” *Marketing Letters*, 26 (4), 679–690, URL <https://EconPapers.repec.org/RePEc:kap:mktlet:v:26:y:2015:i:4:p:679-690>.
- López, Andrés, Paulo Pascuini, and Adrián Ramos (2018), “Climbing the Space Technology Ladder in the South: the case of Argentina,” *Space Policy*, 46, 53–63.
- Mahone, Zachary L and Filippo Rebessi (2019), “Consumer Learning and Firm Dynamics,” .
- Manikowski, Piotr and Mary A Weiss (2013), “The satellite insurance market and underwriting cycles,” *The Geneva Risk and Insurance Review*, 38 (2), 148–182.
- Mayer, Thierry and Soledad Zignago (2011), “Notes on CEPII’s distances measures: The GeoDist database,” .
- Moore, Dennis R and Willie J Phelps (2011), “Reusable Solid Rocket Motor-Accomplishments, Lessons, and a Culture of Success,” .
- Nair, Harikesh (2007), “Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games,” *Quantitative Marketing and Economics (QME)*, 5 (3), 239–292, URL <https://EconPapers.repec.org/RePEc:kap:qmktec:v:5:y:2007:i:3:p:239-292>.

- Nocke, Volker and Nicolas Schutz (2018), “Multiproduct-firm oligopoly: An aggregative games approach,” *Econometrica*, 86 (2), 523–557.
- Pekkanen, Saadia and Paul Kallender-Umezu (2010), *In Defense of Japan: From the Market to the Military in Space Policy*, Stanford University Press.
- Peres, Renana, Eitan Muller, and Vijay Mahajan (2010), “Innovation diffusion and new product growth models: A critical review and research directions,” *International Journal of Research in Marketing*, 27 (2), 91–106.
- Sarkar, Samit (2013), “Mojang tallied 2012 revenue of nearly \$240M, looking to expand Minecraft to new markets,” <https://www.polygon.com/2013/2/1/3941948/mojang-2012-revenue-new-markets>, [Online; accessed 23-April-2021].
- Triezenberg, Bonnie L., Colby Peyton Steiner, Grant Johnson, Jonathan Cham, Eder Sousa, Moon Kim, and Mary Kate Adgie (2020), *Assessing the Impact of U.S. Air Force National Security Space Launch Acquisition Decisions: An Independent Analysis of the Global Heavy Lift Launch Market*, Santa Monica, CA: RAND Corporation.
- Valve (2021), “SteamWorks Documentation: Early Access,” <https://partner.steamgames.com/doc/store/earlyaccess>, [Online; accessed 23-April-2021].
- Vance, Ashlee (2015), “Elon Musk’s Space Dream Almost Killed Tesla,” *Bloomberg Businessweek*, URL <https://www.bloomberg.com/graphics/2015-elon-musk-spacex/>.
- Warren, Tom (????), “Minecraft Sales,” <https://www.theverge.com/2016/6/2/11838036/minecraft-sales-100-million>, accessed: 2018-01-04.
- X-Prize, Ansari (2020), “Launching a New Industry,” URL <https://ansari.xprize.org/prizes/ansari>.

- Xu, Qin, Peter Hollingsworth, and Katharine Smith (2019), “Launch Cost Analysis and Optimization Based on Analysis of Space System Characteristics,” *Transactions of the Japan Society for Aeronautical and Space Sciences*, 62 (4), 175–183.
- Young, H. Peyton (2009), “Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning,” *American Economic Review*, 99 (5), 1899–1924, URL <https://ideas.repec.org/a/aea/aecrev/v99y2009i5p1899-1924.html>.
- Zinger, Kurtis J (2015), “An Overreaction That Destroyed an Industry: The Past, Present, and Future and US Satellite Export Controls,” *U. Colo. L. Rev.*, 86, 351.