Selecting a Launch Vehicle Design For Maximum Financial Success

By

Karl Gallagher

July 18th, 2002
AE 590 Report
Prof. Axelband

Introduction

Aerospace engineers tend to think of their problems as purely technical, divorced from economic considerations. A professor once told the author's class "Cost is not something we think of as an engineering parameter." While the world of government contracting may encourage this, commercial projects must focus intently on their financial viability. If they cannot promise a greater return than other possible investments no money will be available to build the system. This drives them toward the civil engineering maxim "An engineer is someone who can build for one dollar what any fool can build for two dollars."

The author first encountered this situation while working at a launch vehicle start-up company, Pioneer Rocketplane. Pioneer had developed a technically sound system architecture, but the investors and strategic partners were hesitant to proceed with development because of qualms over how well it would compete with other launchers for the available market. The performance requirement for the system had been keyed to a specific large customer, without considering the market as a whole. That customer had since changed their satellite design to one too heavy for Pioneer's design to lift. Redesigning the system for that one customer could have resulted in a system too costly to be profitable. The question then was "What portion of the market should Pioneer aim for and how successful is it likely to be?" The author conducted a trade study to answer it.

This study revisits that trade, using the data available at the time and developing a new methodology to perform a more rigorous analysis of the issue.

Description of Sizing Problem

The driving requirement for Pioneer's launch system design was the amount of payload mass to be placed in orbit. Increasing the requirement would require a larger system or higher performance subsystems, with corresponding cost. Our goal was to capture a large portion of the satellite launch market by offering highly reliable service at reduced cost. The more our system cost to develop the harder it would be to underprice the competitors and still make the profits our investors required.

The satellite systems planned for when Pioneer would be operational (2002 on) covered a wide range of weight classes. The more performance our system had the more sales we could make. The Pioneer architecture could not be scaled up indefinitely, but a more powerful version was definitely possible.

The tradeoff was between lowering expense and increasing revenue, with the goal of finding the highest profit alternative. Pioneer's design philosophy concentrated on using off the shelf components and technology so the system could not be arbitrarily resized. New designs using alternate components would have to be developed, and their performance analyzed. Once this was done the designs could be evaluated for their relative profitability.

Description of the Pioneer Rocketplane System.

The Pioneer launch vehicle is a two-stage system. An airplane-like first stage carries an expendable upper stage out of the atmosphere, which then deploys the satellite payload in orbit. To improve performance the first stage does not carry its full propellant load at takeoff but takes off with empty LOX tanks and meets a tanker at altitude to fill them. This reduces the structural weight of the vehicle, which increases the amount of payload it can lift. After allowing the tanker to maneuver away the rocketplane would boost out of the atmosphere. The upper stage would be ejected and then ascend to orbit. Once the upper stage has separated the “rocketplane” first stage reenters and glides to a landing.

At the time of this trade Pioneer had developed a full architecture for the system, including detailed work on the rocketplane and upper stage designs. Multiple analyses had shown the vehicle would meet the performance requirements and would be safe to operate. The only obstacle to continuing to PDR was the agreement of investors to fund the remaining effort. A key strategic partner raised questions about the performance requirement in light of changes in the satellite market. This led to this trade study to determine the proper payload requirement for Pioneer’s system.

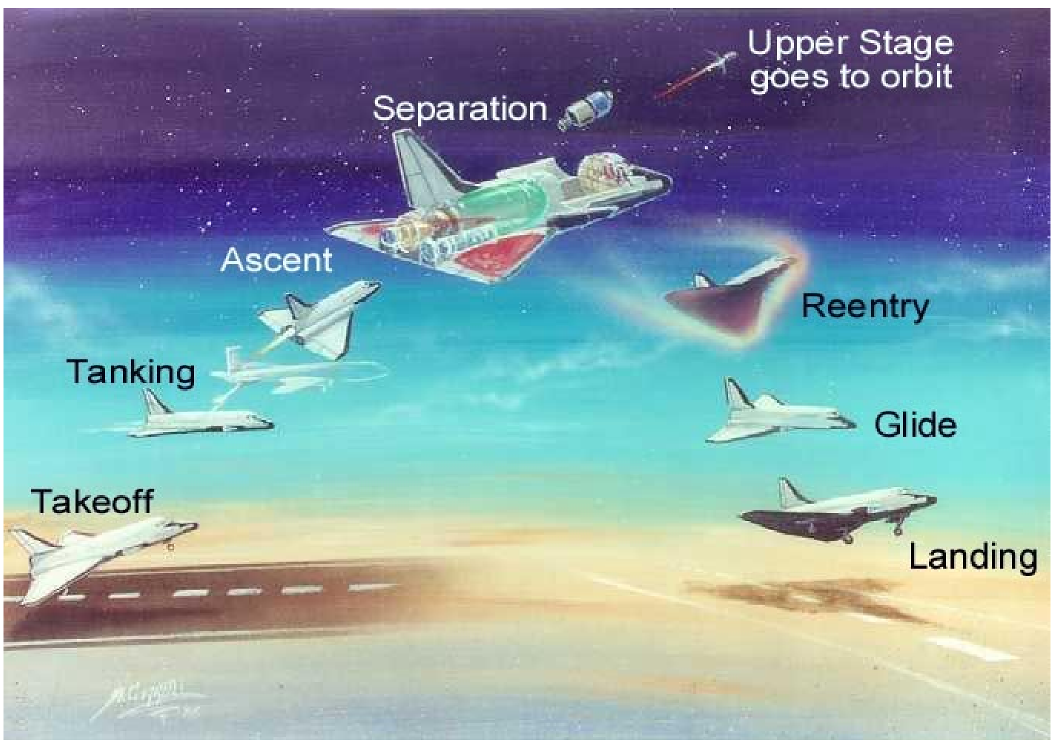

Figure 1. Pioneer Rocketplane concept, with expendable upper stage

Design Inputs

Potential design changes.

A large number of alternate design options were available for the system. Most options were judged by how easily they could be incorporated into the current design. The size of the upper stage could be increased while reducing the rocketplane's propellant load (keeping the total weight constant). The combination of the RD-120 fixed main engine and RD-8 verniers could be replaced by a higher thrust gimbaled NK-33 engine. Careful management of the propellant tanks and lines could reduce the amount of residual propellants. Using higher-tech components could reduce the weight of some subsystems. An upgrade to the RD-120 engine would improve its performance (this had been proposed by Pratt & Whitney). The propellant transfer from the tanker could take place at a faster rate, reducing the altitude loss during the transfer (at a price of more complex equipment and higher risk).

The alternative to small changes to the current baseline was to completely scale-up the size of the system. This would have been slower and more expensive than modification to the baseline but offered greater performance increases than any of the other options.

Other options were ruled out by Pioneer's business strategy. A clean-sheet development of a new engine for the rocketplane was risky and would have added one to three years to the first launch date, which was longer than our investors were willing to wait. Developing light-weight thermal protection materials involved similar risks and delays and was ruled out.

Sensitivities.

The impact of small changes to the design could be easily assessed as the team had analyzed the performance of the system relative to a number of variables. This gave the value of the partial derivative of the payload to orbit relative to mass, thrust, drag, lift, ignition speed or altitude, etc. (for both the rocketplane and upper stage). This had proved very useful during earlier design iterations and allowed a quick first cut of new concepts before using the OTIS trajectory software to determine performance.

Combinations.

Five new system designs were created for the study using combinations of the options above.

Small ΔV Split: The upper stage is enlarged ~30% with a corresponding reduction in the rocketplane propellant load. This changes the ΔV split between the stages.
NK-33 Swap: The RD-120 engine and vernier is replaced by an NK-33 with roughly double the thrust.

Technology Insertion: A number of changes are made to the baseline including an upgraded RD-120 engine, high-rate LOX transfer, and weight savings from using leading-edge components.
Large ΔV Split: The NK-33 engine is used along with a ~60% larger upper stage.
Scale Up: The GLOW of the system is increased 50% with a much larger upper stage.

<u>Assessed cost, performance and risks.</u>

Design	Baseline	Small ΔV Split	NK-33 Swap	Technology Insertion	Large ΔV Split	Scale Up
Rocketplane GLOW (lbs.)	240,000	240,000	240,000	240,000	240,000	360,000
U/S GLOW	20,000	25,000	20,000	20,000	30,000	72,000
Tanker changes	None	None	None	High LOX transfer rate	None	50% more LOX
NRE ($M)	275	282	287	362	296	550
Cost/flt ($M)	2.1	2.6	2.3	2.1	3.6	5
Payload (lbs.)	3667	4034	4217	4327	4620	5498
Risk	Low	Low	Med	High	Med	High
Time to PDR	6 m	8 m	8 m	12 m	10 m	14 m

NRE – Non-Recurring Expenditures. The cost of designing and building the launch system (including rocketplane, upper stage factory, tanker modifications, and ground facilities).
GLOW – Gross Light-Off Weight – mass at ignition of rocket engine (after filling propellant tanks).
Payload is measured to the reference orbit: 85° inclination, 300 km altitude.

Figure 2. Summary of design option characteristics.

The risk for the NK-33 Swap and Large ΔV Split options was not driven by technical issues but by the availability of the NK-33 engines. At the time of the study all existing engines were contractually obligated to a competitor, Kistler Aerospace. Kistler's financial troubles made it very unlikely that they would be able to fulfill their contract, which would make the engines available for other customers. If Kistler found the money to pay for the engines none would be available for Pioneer.

The Technology Insertion option is high risk because the changes require using less proven components and developing a new rocket engine. Any technical hurdles could result in a severe cost and schedule impact to the project.

The Scale-Up option is also considered high risk as the increase in size would require redoing much of the work done to date. Some issues such as thermal protection become much harder with a heavier vehicle. The larger propellant transfer would be much more difficult and require more expensive support equipment. While first cut analysis showed the Scale-Up to be feasible, there was a real danger of finding a show-stopper problem that would prevent the design from meeting the desired performance.

Market Inputs

Market survey.

The federal government's Office of Commercial Space Transportation tracks trends in the satellite industry to guide its policies on launch systems. This data is public and was used to support Pioneer's market model. Additional data was supplied by direct contacts with satellite manufacturers.

At that time (1998) forecasts of the satellite market had consistently underestimated the number of satellites being built in upcoming years, so Pioneer used the optimistic end of the FAA forecast—five "Big LEO", four "Little LEO", and three "Broadband" systems. Many more systems had been proposed in each of those categories but only the likeliest candidates were included in the market model. An additional two launches per year were assumed for science and remote sensing satellites in the Pioneer payload class.

The chosen systems were categorized by weight and deployment schedule. Replenishment launches were also included. If the planned replenishment rate was unknown we assumed between 5 and 10% of the constellation would be replaced every year.

System Name	Number of Satellites in Constellation	Satellite Weight (lb)	Year of Deployment	Replenishment Launches Per Year
KITComm	21	220	2003	3
E-Sat	6	250	2001	0.5
LEO One	48	275	2002	5
FAISat	38	332	2001	4
Globalstar	56	985	1999	6
Sci/RS*	-	1000	-	2
Laredo	72	1500	2006	5
Iridium	66	1516	1998	12
ECCO	54	1550	2001	6
Skybridge	68	1770	2003	10
Ellipso	17	2200	2001	1
Salina	96	3775	2004	5
M-Star	84	4400	2005	5
Celestri**	70	7000	2002	7

*Two science/remote sensing missions per year in this weight class are assumed—historical count has varied widely.
**At the time of the original study this system was being merged with Teledesic.
Fig. 3. Market model for satellite launches.

The market model only covered launches over the years 2002 to 2007, from the initial operation of the rocketplane (assuming development began in 1999) to the end of service life for the first vehicle. The weight of satellites was adjusted if necessary to account for an orbit requiring higher vehicle performance (i.e., higher altitude or inclination). The systems chosen for the model and their attributes are shown in Fig. 3.

The market model has to handle deployment and replenishment launches differently. A replenishment launch is replacing a single failed satellite in a constellation and cannot be easily combined with other satellites. A deployment launch could place all the satellites for a single orbital plane up on a single booster. This made the competition tougher for the deployment launches—Pioneer could find itself competing by launching one or two satellites against a Delta II which can launch 6 or more.

If Pioneer succeeded in lowering the cost of reaching orbit satellite developers would respond. An increase in the market would come from projects that had been marginal now becoming worth undertaking due to the smaller launch costs. With both lower launch costs and faster response times, some satellite developers might reduce their costs by eliminating high reliability components and rapidly replacing satellites that fail. However, the long design cycle times for new satellites make both those responses too slow to have an impact during Pioneer's initial operations. Quantifying the response is also very difficult, so no new customers will be included in the market model.

Competition survey

The prices and capabilities of the existing Western launch vehicles are public information. The most cost-effective launcher for each payload size can be easily identified as shown in Figure 4. For any particular payload size the price difference between Pioneer and its competitor can be easily determined. The ratio between prices was used to calculate the market share Pioneer would obtain.

Other launch vehicles were forecasted for the market but not included in this analysis. A number of non-Western launchers were available but were not expected to be viable competitors in our time frame. Russian launchers relied on the inventory and infrastructure from the Soviet Union, which was no longer being maintained. Chinese systems were unreliable and customers were being pressured by the US government to avoid them. Their share of the market during the 2005 time frame was expected to be minimal so they were not included in the analysis.

New launch systems were also being developed by other start-up companies competing for the same market as Pioneer. Including them in the model would be difficult, as not much data is publicly available and the designs are subject to change. Only a few of the start-ups could be expected to actually deploy their systems. The price per flight and launch rate for them is also hard to estimate. This made it impossible to include them in the list of competitors.

To model the impact of new competitors they must be combined with competitive responses from the existing companies. If Pioneer took many customers from Boeing or Lockheed-Martin they would immediately respond to the challenge. Lower prices, restrictive contracts with customers of other parts of the company, or lobbying for political restrictions on Pioneer's operations would help them retain more of their market share. Accounting for those effects must also include competition from overseas or new entrants to the market.

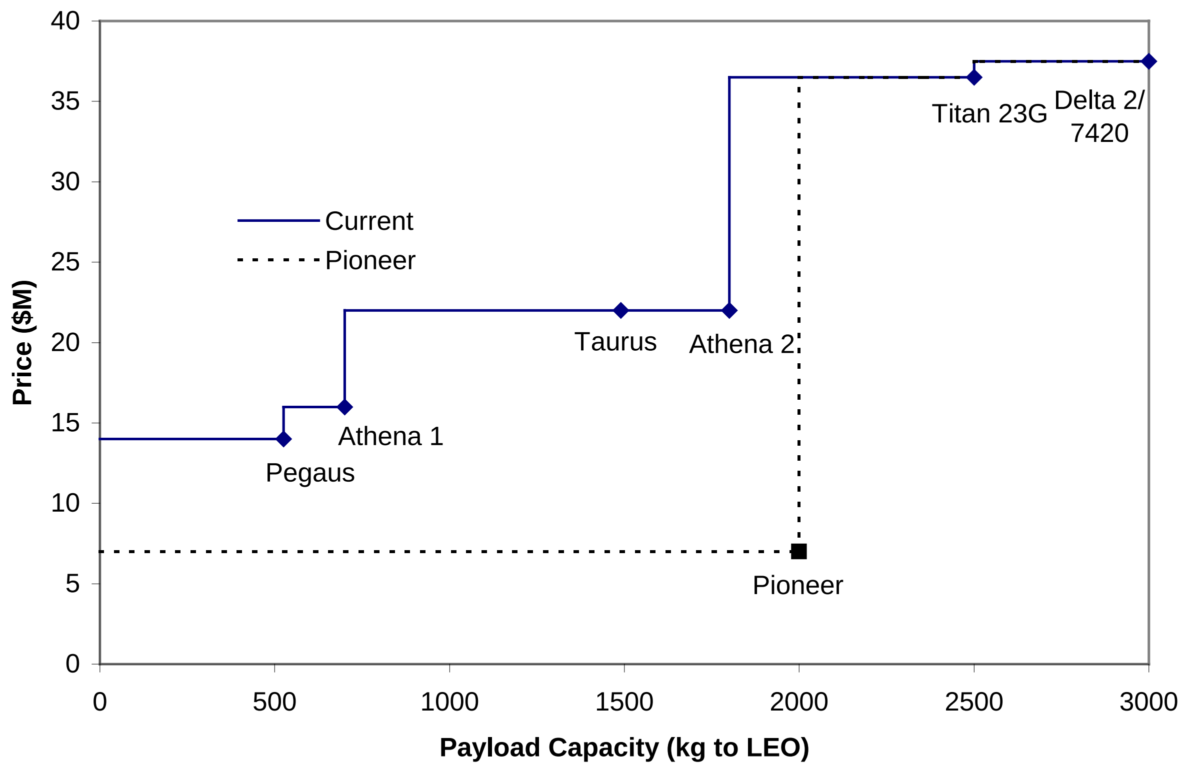

Figure 4. Launch vehicle prices and capabilities. Pioneer price is $7M/launch in this example, using the baseline design.

Market capture curve

Pioneer intended to win market share not just by underpricing its rivals but also by offering advantages in reliability, scheduling, and flexibility in payload integration and launch locations. While those attributes are important to satellite owners they are hard to quantify and Pioneer would not receive credit for them from customers until a track record had been established. To convince customers to bet their satellites on an unknown would require a significant price savings.

Even if Pioneer dramatically underpriced all other launch providers, it could not capture the whole market because the owner of a large constellation such as Teledesic would be unwilling to rely on a single system, risking huge losses if it failed. Instead launches would be shared among several providers, as Iridium did with Boeing and two others. Other customers could be unwilling to use Pioneer because of political considerations or contractual constraints.

To make our model workable, the fraction of the market captured is assumed to be a function of the ratio between Pioneer's price and that of the competitor for that customer. The disadvantages of being a new entrant to the market can be captured by restricting the function. The competitive responses and new competitors can be factored in by further restricting the market share captured.

Once the function is developed it can be applied to the market model. For each customer launching satellites in a particular year the competitor's price is known. For each potential Pioneer price the market share percentage is calculated and the (integer) number of launches done by Pioneer can be found and included in the financial model.

The simplest market capture function is to assume that Pioneer could obtain up to 50% of the market, and would do so if it underpriced the competitors. This is the "Step" curve shown in Fig. 5. In reality, the greater the price difference the more likely customers would be to switch to Pioneer. This produces the "Linear" curve, where an equal price leaves all customers with the established launch systems but a free service can get them all. To include both the constraints on maximum market share and the need for significant discounts the third curve was generated as a quadratic fit.

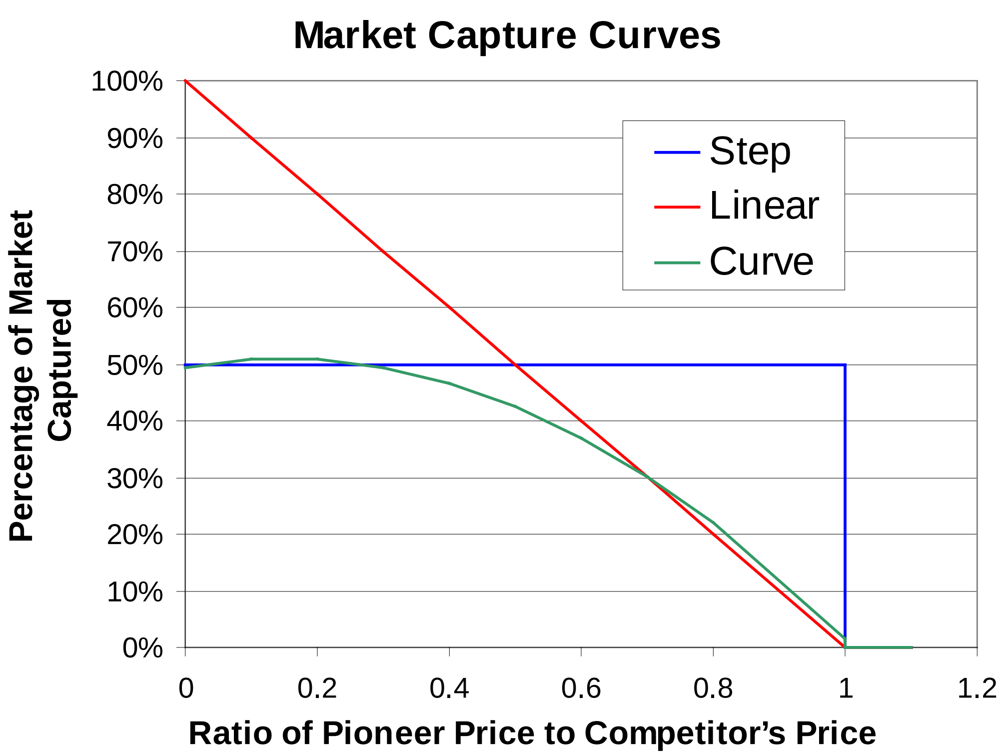

Figure 5. Alternate market capture ratios functions.

The third market capture curve was chosen for the model as the most conservative option. Pioneer's high reliability architecture and other advantages would probably allow it to charge a price premium over existing providers for at least some customers. However, a model assuming that would not account for competitive responses and new competitors, which would drive prices down.

The factors driving Pioneer's market share will change over time. Early in its operational life customers will be reluctant to risk their satellites on an unproven vehicle. Once it is accepted competitors will begin cost-cutting and other actions to preserve their market share. Eventually other new launch systems will be deployed to compete with Pioneer at the same level, displacing the last of the old systems. As Pioneer and the new ventures lower the price of reaching orbit the market will expand but new customers may not arrive until the end of the time period under analysis. The conservative market capture curve should envelope all of the restrictions on Pioneer's market share.

The goal of this analysis is to provide reasonable comparisons between the different design options, not to provide a price forecast of future revenues. The market share function should not favor a design that needs the entire current market forecast or no new competition to be profitable. Conversely, a design which is optimum for very low prices should also be rated low.

Fig. 6 shows that the conservative market capture curve has its optimum between those of the other curves (This shows the final data for the "large ΔV split" design option). This makes it most useful for our purposes by excluding extreme results, as well as considering the constraints on market share above.

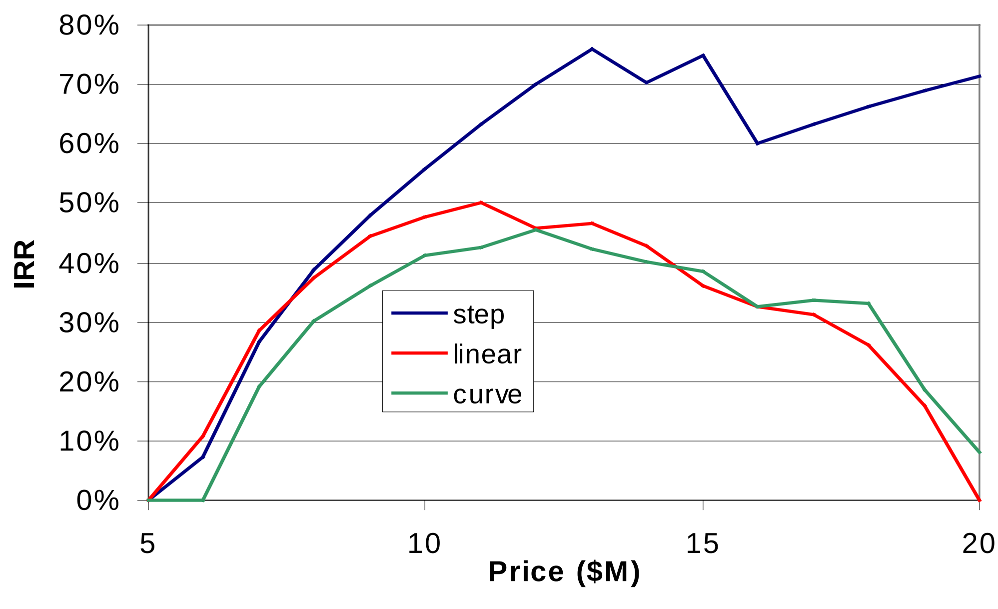

Fig. 6. Effect of market capture assumptions on system IRR.

Financial model.

The actual financial structure of Pioneer Rocketplane, like most startups, was subject to change as directed by new investors and other participants. This trade uses a reference model that assumes the company will begin with $50 million in equity and borrow the remainder of the money needed for development of the system. The loan will be completely repaid as early as possible before returning any money to investors. The rocketplane is depreciated as capital equipment but all other equipment is expensed (the tanker and ground facilities are leased and the upper stages are expended). Interest on the loan is at 10% and corporate income tax is 35%. The rocketplane is fully depreciated at the end of 2007 and all other assets are assumed to have no residual value.

Year	**1999**	**2000**	**2001**	**2002**	**2003**	**2004**	**2005**	**2006**	**2007**
# of Flights				8	25	30	38	36	12
Revenues	-	-	-	96.0	300.0	360.0	456.0	432.0	144.0
Expenses									
Development	53.2	106.4	106.4	-	-	-	-	-	-
Operations	-	-	-	28.8	90.0	108.0	136.8	129.6	43.2
Interest	-	0.3	11.0	24.6	21.6	9.0	-	-	-
Depreciation	-	-	-	6.0	9.6	5.8	3.5	3.5	1.7
Gross Revenue	-	-	-	36.6	178.8	237.2	315.7	298.9	99.1
Income Taxes	-	-	-	12.8	62.6	83.0	110.5	104.6	34.7
Net Revenue	-	-	-	23.8	116.2	154.2	205.2	194.3	64.4
Cash Flow Statement									
Operating Activities									
Net Income	-	-	-	23.8	116.2	154.2	205.2	194.3	64.4
Depreciation	-	-	-	6.0	9.6	5.8	3.5	3.5	1.7
Investment Activities									
Equity Sales	50.0								
Capital Expenses			30.0						
Financing Activities									
Borrowed Funds	3.2	106.4	136.4	-	-	-	-	-	-
Principal Repayment	-	-	-	29.8	125.8	90.4	-	-	-
Net Cash Flow (before loan repayment)				29.8	125.8	159.9	208.7	197.8	66.1
Net Cash Flow (after loan repayment)				-	-	69.5	208.7	197.8	66.1
Total Debt	3.2	109.6	246.0	216.2	90.4	-	-	-	-
Total Cash on Hand	-	-	-	-	-	69.5	278.2	476.0	542.1
Payments to Investors	(50.0)	-	-	-	-	69.5	208.7	197.8	66.1
IRR	46%								

Fig. 7. Financial model's balance sheet for Pioneer Rocketplane. All numbers are in millions of 1998 dollars. Case shown is Large ΔV Split design at $12M/launch price.

Find Optimum

Choose target optimum

Determining the proper figure of merit (FOM) for this trade requires more attention to the business than the technical issues. Technical figures of merit such as payload capacity, mass fraction, or even marginal cost per pound of delivering payload do not address the question of what will be the best choice for Pioneer's stockholders. To be a financial success the company must provide a substantial return to the investors.

To be successful from the shareholders' view the company must not just sell a large number of launches but make a profit on them. The driving factor in profitability is the need to pay back the loan for developing the system before returning money to investors. Therefore using number of launches or gross revenue as the FOM will not reflect how the more expensive options will reduce profits. Total profits is a good measure of success but does not reflect the time value of money—the investors want their return as soon as possible. The internal rate of return (IRR) tells the investors what interest rate they would need to receive from another investment to match the returns from this project. IRR includes weighting for time and is our preferred figure of merit.

Analysis results

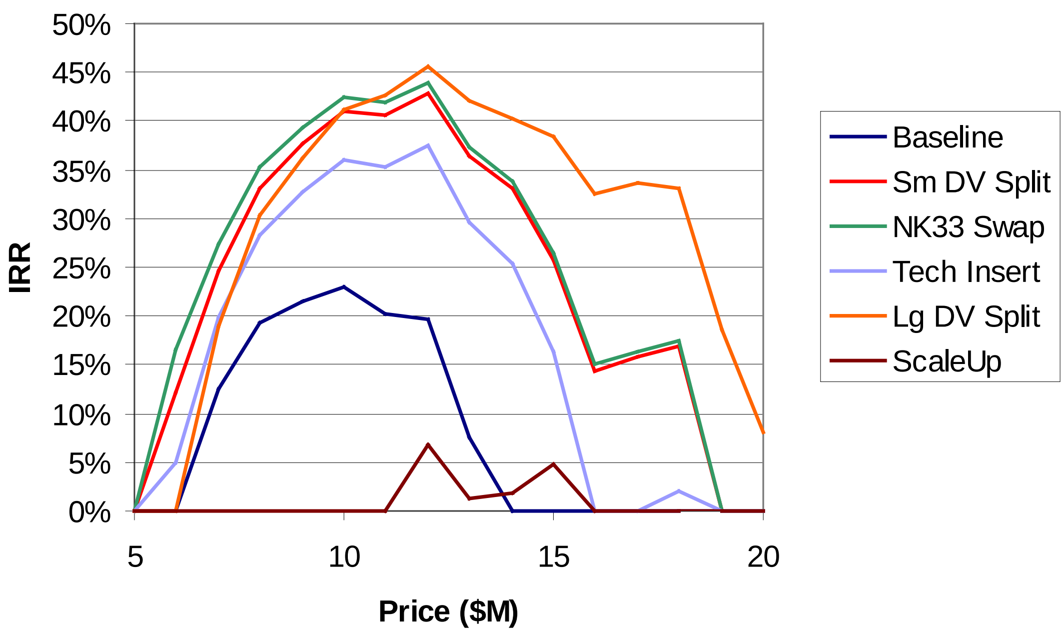

Fig. 8. IRR for each design option.

The IRR for every design option was generated for launch prices in the range of $5M to $20M. The results are shown in Fig. 8. The spreadsheet model takes a design option and launch price input by the user and applies it to the market model. The model compares the price to that of the competitors and calculates a market share for each customer per year, excluding satellites too big to be launched by that design option. The launches per year are totaled and limited to 50 (the maximum flight rate for Pioneer). The launch counts are then copied to the financial balance sheet. The balance sheet calculates the revenue and expenses per launch and applies net income against the loans required to develop the system (note that expenses and development costs are design-dependent). Cash remaining after the debts are repaid is returned to the investors and included in the IRR calculation.

The FOM of each design option is the highest IRR found. The Baseline option was limited by its payload capacity, which did not allow it to serve many possible customers. The Scale-Up option was so expensive to develop that repaying the development loan and interest would consume almost all revenues. The resulting IRR for each option was:

Baseline	23%
Small ΔV Split	43%
NK-33 Swap	44%
Technology Insertion	38%
Large ΔV Split	46%
Scale-Up	7%

<u>Sensitivity to market changes</u>

Since the IRRs of the top designs are close together, the results should be tested to see if they are an artifact of the specific market model chosen. While the design options part of the analysis is based on solid engineering, the market model is speculation intended to provide a reasonable approximation of what the market would look like over the next decade. The final design choice should be one that outcompetes the others in many situations, not just one narrow range.

Looking at the distribution of satellite launches (Fig. 9), some patterns are obvious. Most launches are small enough for all designs to handle them. A large gap in sizes above them exists, with no other systems small enough for the baseline design to support. Three more systems are in the model—one of which can be handled by all of the new designs, one which can be handled by the Large ΔV Split and Scale-Up designs, and one which is beyond the capability of any of them.

This distribution obviously works against the Scale-Up option. It has a nearly half-ton advantage in throwweight over the next biggest option, but this wins no sales because no customers exist in that range. The Baseline similarly has a third of its payload capacity always going to waste because of the gap in the payload sizes. (This may argue for a "Scaled-Down" option but the investors were universally in favor of increasing payload capacity).

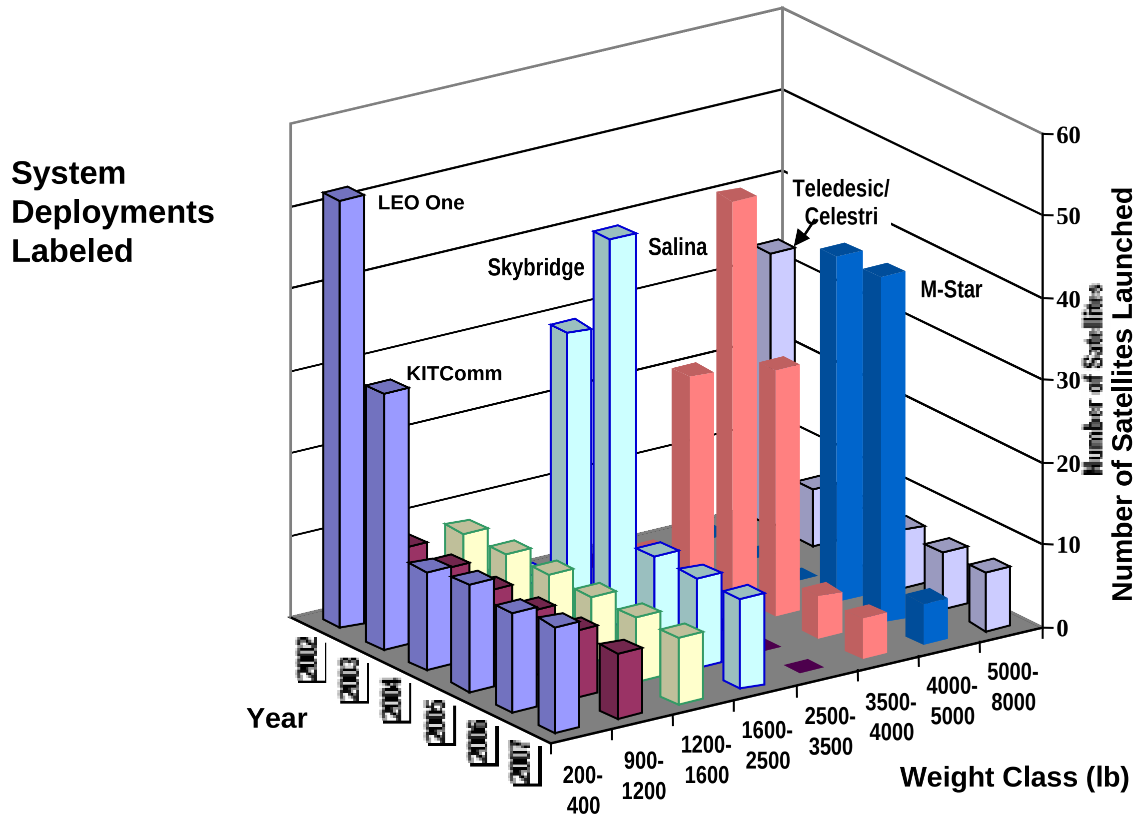

Fig. 9. Satellite market distribution

Four variations on the market model were examined. They were:

1. Doubling the Globalstar replenishment rate (favors Baseline)
2. Reducing the weight of Celestri to 5000 lbs. and doubling the satellite count (favors Scale-up)
3. Increasing the weight of Salina to 4200 lbs. and doubling the satellite count (favors NK-33 Swap and Technology Insertion over Small ΔV Split)
4. Eliminating Skybridge (hurts all)

The analysis described above was repeated for all six designs to find the best IRR possible in each variant of the market. The results are summarized in Fig. 10.

	Reference	Variant 1	Variant 2	Variant 3	Variant 4
Baseline	23%	29%	23%	23%	2%
Small ΔV Split	43%	45%	43%	43%	31%
NK-33 Swap	44%	46%	44%	61%	33%
Technology Insertion	38%	40%	38%	56%	24%
Large ΔV Split	46%	47%	46%	59%	37%
Scale-Up	7%	11%	40%	33%	0%

Fig. 10. Maximum IRR of each design in variant market models.

The ranking of the design options was stable over all the market variants. The same two designs held the top two ranks in every case, with their IRRs within 4% of each other.

The other designs only varied up or down one place in rank. This seems to be a good confirmation of the results given by the original market analysis.

<u>Selected option</u>

The "Large ΔV Split" design has the best return for Pioneer's investors. This design is essentially a combination of the "Small ΔV Split" and "NK-33 Swap" designs, which allows a graceful fallback if development problems are encountered. If the upper stage can't be scaled up, the system would become the NK-33 Swap design, and if the NK-33 engine is unavailable the old engine can be used as in the Small ΔV Split. Had the market conditions of 1998 held, Pioneer would have proceeded with this design.

Description of Trade Methodology for Generic Projects.

Applying this type of analysis to other design decisions requires solid knowledge of the customers for the system. Pioneer's market consisted of customers who were publicizing their plans and absolutely had to buy launch services for their satellites. Doing this for other systems may be much more difficult. If an entirely new market is being targeted (rather than replacing products in an existing one) active research may be needed and the accuracy of the market model will be hard to evaluate.

Once the market has been assessed, the customers must be broken into segments by their needs. For the satellite launch market this was evaluated on a single variable, the weight capacity to orbit. Other markets may require more than one variable, or be broken into qualitative segments (i.e., for customers wanting specific features). A segment's value for a variable may be a threshold (no design below this is acceptable) or a goal (less performance may be acceptable). If the market potential is uncertain, several variations should be developed to compare designs under different circumstances. Developing the variants in advance of the analysis can avoid the problem of "slanting" a variant to favor one design or another.

Each of the designs under consideration must be rated to show if it can satisfy the market segments. This will require values for each variable defining the customers, and noting if they meet the qualitative requirements. The recurring and non-recurring costs for each one must also be defined.

Competitors must be assessed for each market segment. This can usually be done through public information—trade journals track the services being offered for sale. If the system being developed is not an essential service, the biggest competitor may be "nothing". Customers may refuse to buy anyone's system and this must be recognized either as a competing choice or as a part of the market capture function.

The function to calculate market share must be carefully developed. A design's market share is automatically zero for a segment if that segment's requirements aren't met (i.e., the design's performance doesn't meet the threshold for a variable or is missing a feature). The market share should be a function of the design's cost relative to competitors and how its performance compares to the goals for the segment. Other qualitative factors should be considered as well—some segments may be unlikely to adopt a new system or have a ceiling on the possible market share they will allow. If some segments would be willing to pay a premium for the design's new features that may also be reflected in the capture function.

The financial model for the project should reflect the goals of the backers. A start-up company will generally require immediate profits and should rate designs by their IRR. A large company may be willing to take a temporary loss to break into a new market and would evaluate them by the market share achieved. The financing method (equity, loans, internal investment) must also be reflected in the model.

Once the models have been developed each design can be analyzed for its IRR (or other figure of merit) vs. price curves. The best value should be used to judge it against the other candidates. This result may not be the only important factor in choosing which option to develop—risk and schedule may drive the project manager to choose a design which placed 2nd or 3rd in the business analysis.

Conclusion

This analysis found that Pioneer had good launch system designs for the market seen at the time. The investors and strategic partners were interested in pursuing development, but before the deal could be closed Iridium financial situation began to unravel. With the failure of the Iridium service the planned follow-ons were cancelled and companies planning to compete in that market quietly dropped their plans. Within six months of the original study there was no prospect of a satellite launch market large enough to support the development of a new commercial system.

The analysis methodology provides useful data for a head-to-head comparison between design alternatives. The market and financial models aren't accurate enough to provide a high-quality prediction of the project's financial performance. The additional detail needed to do that would include more market analysis, finding the changes over time in the market and capture function, a finer breakdown of expenses, and finding ways to support earlier returns to investors before full loan repayment. Those changes would affect the performance of all the design options equally so there would be no change in the outcome of this trade from the additional effort.

Integrating the technical and business sides of a project is always difficult. This methodology shows one way to bring the technical team into the gap by allowing them to consider the financial impacts of their design decisions. This also shows part of the difference between government and commercial projects. A government contractor, even when working in a "Cost As an Independent Variable" environment, will be dealing with requirements that are pure step functions with no indication of the benefits from meeting them or effect of missing them. A commercial company can boil down the performance differences to whether a customer will pay more for them and find if the extra expense is worth it.