Reason
Foundation

JANUARY 1994 POLICY STUDY NO. 173

MEETING SPACE LAUNCH NEEDS— ECONOMICALLY

by Joseph P. Martino

Reason Foundation
3415 S. Sepulveda Blvd., Suite 400 • Los Angeles, CA 90034
(310) 391-2245 • Fax: (310) 391-4395

Policy Study No. 173
January 1994

MEETING SPACE LAUNCH NEEDS—ECONOMICALLY

by
Joseph P. Martino

EXECUTIVE SUMMARY

Getting into space is expensive. The high cost of space launches results from design practices which were inherited from the ballistic missiles of the 1950s, when high performance was demanded almost regardless of cost.

The high cost has been perpetuated by NASAs practice of developing new technology for each new mission, thus piling up development costs which must be amortized over only a few launches. Demand for space launches—like demand for anything else—should be price-elastic: the lower the cost, the greater the demand. Lowering launch costs requires a break from the practices of the past three decades in space.

Three new approaches each promise to lower launch costs, and the three can work synergistically to reduce costs even further. The three are: use of commercial practices for vehicle construction; long production runs of single designs; and designing launchers to be reusable like aircraft instead of one-shot devices like missiles.

Commercial design means using proven technology where it is adequate, and developing new technology only when it is necessary, in contrast with NASA's practice of pushing new technology in each vehicle. Long production runs allow users to amortize development costs over many launches and to gain the economies of mass production. They also increase reliability as design defects are caught and eliminated.

Reusability means that a launcher should need little more than refueling before another launch, instead of the major refurbishing required by the Space Shuttle. Reusability also lowers costs dramatically, since the vehicle is not thrown away after each flight, nor does it need to be remanufactured. Design for reusability can also increase safety, since reusable vehicles can be designed for "intact abort" in the event of a failure. Space need not be as expensive as it now is. Designing launchers according to aircraft design practices rather than missile design practices can lower today's costs significantly.

TABLE OF CONTENTS

I. BACKGROUND

Space is expensive. The total cost of putting a pound of payload into orbit using the space shuttle is about $3500.[1] There are very few commercial products in or out of the world which are worth that price. Communication satellites, pharmaceuticals, and high-purity chemicals just about exhaust the list.[2] Moreover, only the first of these has actually become commercially viable.

How did space turn out to be so expensive? Walter McDougall points out that the roots of the problem go back to the ballistic missiles of the 1950s.[3] High performance was critical to the success of these missiles. Long range with large payloads was the dominant design criterion. From the military perspective, accomplishing the mission was necessary and high cost was acceptable; doing only half the mission was unacceptable even if it was done cheaply. Thus, ballistic missiles were designed for maximum performance. They were expensive to build, expensive to transport, and expensive to repair and maintain.

The Atlas missile, as an example, was designed with minimum internal stiffening, to reduce weight. The missile was essentially nothing but a huge metal balloon. The fuel tanks were thinner than a dime at the base, and tapered to about a hundredth of an inch thickness at the top. When not filled with fuel, they had to be pressurized with compressed nitrogen, or they would have collapsed of their own weight. In flight, the tanks were given only enough additional pressure to permit them to withstand flight loads (this practice is followed for almost all large rockets). Rocket motors, by contrast, were designed for high operating pressure, in order to increase their performance. Thus, to transfer fuel from the low-pressure tank to the high-pressure engine required a pump, which added weight, cost and complexity. A so-called "pressure-fed" engine would have involved heavier fuel tanks (to withstand higher pressure) and a lower-performance engine.

When the U.S. reacted to the Sputnik challenge in 1957, the only launch vehicles available were military missiles. The first successful American satellite was launched by a Jupiter intermediate range ballistic missile. America's (and the world's) first communications satellite was launched by an Atlas booster. Alan Shephard's suborbital flight was launched by a Redstone missile. John Glenn, the first American into orbit, rode an Atlas.

It was not until the Apollo program, with its Saturn vehicles, that the use of converted missiles was abandoned. However, the Saturn was developed by the same team, headed by Werner von Braun, which had earlier developed the Redstone and the Jupiter. It followed the same pattern of high performance at high cost which

---

[1] Gary Stephenson and Greg Freiherr, "How to Beat the High Cost of Space," *Final Frontier*, September/October 1989, pp. 18–23, 52–56).

The recurring cost of a shuttle flight (i.e., ignoring R&D costs, capital costs of launch facilities, etc., but including fuel, crew, and launch personnel) is about $100 million. Payload is about 50,000 pounds. Recurring cost to place a pound in orbit is thus about $2000. NASA estimates shuttle cost at $3500 per pound if the sunk costs of development and the fixed costs of operating shuttle ground facilities are included. However, this figure is based on 14 launches per year. The peak year for shuttle launches was 1985, with 8. Hence, true costs for the shuttle are higher than the NASA estimates. The true figures are probably comparable with the cost of expendable launchers: about $3800/lb on Delta 2, $3600/lb on Titan 4.

[2] Current satellites themselves cost from $5000 to $250,000 per pound, depending on their complexity. Hence the launch cost may be only a small fraction of the cost per pound in orbit. The point is simply that only something which is *worth* $5000 or more per pound is going to be launched at today's launch costs.

[3] Walter A. McDougall, *...the Heavens and the Earth* (New York: Basic Books, 1985).

had marked the military vehicles. This "tradition" was amplified by the need to achieve high reliability in order to meet manned safety requirements.

The possibility of an alternative design philosophy did occur to some rocket designers. As early as 1966, the Aerospace Corporation conducted a study which concluded that launch vehicle cost could be reduced by accepting greater weight, looser design tolerances, and lesser technical sophistication, at least in some components. This concept eventually came to be called the Big Dumb Booster (BDBR).[4] The analogy of a Mack truck versus a sports car has been used to contrast the low-cost, low-sophistication launch vehicle with the current style of launch vehicles.[5]

Why has the "sports car" philosophy dominated, to the complete rejection of the "Mack truck" philosophy? Largely because the U.S. space program has been government-dominated from the beginning. The federal government has not only conducted R&D on space vehicles (as it also conducts R&D on aeronautical vehicles), but has operated what amounts to the nation's only "spaceline." This is in direct contrast with the way NASA's predecessor, the National Advisory Committee for Aeronautics, conducted aeronautical R&D, which was then adopted by commercial manufacturers whose customers were privately-operated airlines. By 1935, some 32 years after the Wright Brothers' first flight, there were profitable commercial airlines already in existence. In 1993, more than 32 years after Sputnik, space launches are still government-dominated and money-losing activities.

James Bennett and Philip Salin point out that this nationalization of space has had a series of unfortunate effects.[6] Two of these are directly responsible for the high costs of space launches.

First, NASA was under no competitive pressure to be efficient or profitable. It had no incentive to minimize costs. On the contrary, it had the usual bureaucratic incentives to avoid embarrassing failures which could be traced to it; to avoid rocking the boat; to reward (or at least avoid upsetting) important constituencies. Thus, there was no incentive to change from the design philosophy which had produced the successes which brought public acclaim. High sophistication at high cost was seen as the only safe way to go.

Second, NASA had strong bureaucratic incentives to keep its engineers and scientists busy by calling for new (but not necessarily better) vehicle designs for each new objective. This meant each design was used for only a few missions. All the nonrecurring costs—development, manufacturing tooling, specialized test equipment, training of equipment operators—had to be charged against those few missions. Using the same vehicle for a

---

[4] "Big Dumb Boosters," Washington, D.C., Office of Technology Assessment, February 1989 (referred to as BDBR).

[5] Historically, the "fly-away cost" of both aircraft and space vehicles has been highly correlated with their empty weight. As a general rule, heavier vehicles cost more. However, at the design margins, when the designer is trying to shave off ounces here and ounces there to meet some performance requirement, cutting out a pound of weight may raise the cost of the vehicle by several thousand dollars. It is in these cases that the designer can save money by accepting greater weight. The problem is that many missiles were designed right at the limit of technological capability, where weight-shaving was needed in order to meet a range or payload requirement. Space vehicles have inherited this design approach.

[6] James Bennett and Phillip Salin, "Privatizing Space Transportation," Issue Paper No. 102, Federal Privatization Project, Santa Monica, Reason Foundation, 1987.

large number of missions would have allowed the nonrecurring costs to be spread over many missions. However, this would have meant less work for NASA's own engineers.[7]

Moreover, there was very little pressure from industry to offset these bureaucratic incentives at NASA. Since launch vehicle contractors were bidding to government specifications, they had no incentive to propose more efficient or less costly designs. In particular, they had no incentive to propose the use of existing or proven vehicles and systems. Moreover, although NASA demanded new systems of its contractors, it often resisted the incorporation of newer technology. The structure of the Shuttle, for instance, is over 90% aluminum, which is 1950s aircraft technology. By the time the Shuttle was designed, various military aircraft had already demonstrated the successful use of titanium as a structural material, yet it was not used in the Shuttle. The computers in the Shuttle have less processing memory than the student calculators sold in college bookstores, and certainly less capability than today's laptop computers.[8]

The end result of all this is that space not only was, but continues to be unaffordable. The situation is the exact opposite of the commercial environment, where each competitor has an incentive to seek cost-reducing innovations, and to copy the cost-reducing innovations of its competitors.[9]

This outcome should have been expected. Richard Nelson and Richard Langlois, in their study of government support for new technology, found the same patterns in almost all the projects they examined.[10] Government-fostered technology tends to reflect government objectives rather than commercial objectives, even when the ultimate intention is commercial viability. This remains true even when the technology is developed for the government by private contractors. The contractors place more emphasis on satisfying the government project managers than they do on meeting commercial requirements.

If space is to be affordable, then, development of launch vehicles must be carried out in a commercial rather than a governmental environment. It must respond to commercial rather than governmental incentives. In particular, the development of launch vehicles must emphasize the use of the cheapest technology which will do the job, whether this is a "tried and true" technology of assured reliability, or a newer technology which offers cost-reducing possibilities. The remainder of this paper will examine how the cost of space can be reduced, by using design philosophies which depart from the NASA tradition.

---

[7] Harry Stine argues that NASA can be compared to a technological army. After the Sputnik scare of 1957, missile experts were brought in from the military and from defense contractors, and given the task of beating the Russians. We won that technological war in 1969. We didn't demobilize the army when the war was over. It has stayed on ever since, and has continued to think up projects to keep itself busy.

[8] In some cases, NASA contractors were able to force the use of newer technology, despite NASA's objections, because the old technology was so obsolete they could no longer obtain spare parts for it.

[9] An anonymous reviewer has pointed out that private contractors "have reduced the cost/lb to geotransfer orbit to *one-third* what it was under NASA" (emphasis in original). This is the result of competition among firms which provide the engines which transfer satellites from low earth orbit to geostationary orbit. Since these engines are considered to be part of the satellite rather than part of the launcher, they are the responsibility of the satellite builder, not of NASA. Satellite builders compete with one another, and therefore have incentives to reduce costs.

[10] Richard R. Nelson and Richard N. Langlois, "Industrial Innovation Policy: Lessons from American History," *Science*, Vol. 219, February 18, 1983, pp. 814–818.

One important issue which will not be covered in this paper is that of government subsidy for foreign launchers (e.g., Ariane, Long March, etc.). This issue, while important, is outside the scope of the paper, which will restrict itself to domestic U.S. considerations. While foreign launchers will be considered to the extent that they impinge on the demand for launch vehicles and launch services, the issue of foreign government subsidies will not be considered.[11]

II. DEMAND FOR SPACE LAUNCHES, 1993–2010

One critical issue is the level of demand for space launches. If there are to be only a relatively few launches, it does not matter much what the design philosophy is. While there are ways to reduce the nonrecurring costs if the work is done by private firms, in a competitive environment, those nonrecurring costs can still be spread over only the few launches, and these launches will inevitably be expensive.

Related to the issue of number of launches is the issue of elasticity, or the sensitivity of demand to cost. If the cost of launches is reduced, will there be additional demand for launches, or will the same number of launches be made almost regardless of cost?

As might be expected, there are contrasting views on the issue of price sensitivity, or elasticity of demand. Some argue that the demand for launches is relatively insensitive to price (inelastic). Others, however, argue that like many other economic goods, the demand for space launches is highly elastic with regard to price. Reductions in price will increase the demand.

John Pike, of the Federation of American Scientists, holds that demand is not sensitive to price. He has been quoted as saying: "There are too many rockets chasing too few launches. If you make a list of satellites to be launched and rockets that might be used, you run out of payloads real fast."[12] In essence, Pike is saying that reduced price will not induce other people to find profitable uses of space, and start purchasing launches. The only launches which will ever take place are those already planned at today's prices.

Similar views were expressed by Elizabeth Corcoran and Tim Beardsley in *Scientific America*[13]. They argue that while existing and planned launchers can provide about 40 launches per year, the market demand will be about 15, primarily for communications satellites. They conclude that the result will be cutthroat competition among launch providers, with each provider seeking to attract customers by offering launches at prices below actual cost. In their view, the only launch providers which can survive in that environment are those which are heavily subsidized.

Dr. Shove, director of the State of Florida's commercial space programs, holds the view that demand is sensitive to price. He told journalist William Broad, "I'm a firm believer that when you bring down those launch costs significantly, you're going to have all kinds of stuff coming out of the woodwork."

---

[11] This issue is discussed in James Bennett, "Creating Competition Space Trade: A Common Market for Space Enterprise," Policy Study No. 123 (Los Angeles: Reason Foundation, 1990).

[12] William J. Broad, "Private Commerce Stands Poised For a Modest Debut in Space," *New York Times*, August 30, 1988.

[13] Elizabeth Corcoran and Tim Beardsley, "The New Space Race," *Scientific American*, July 1990, pp. 72-85.

Both Shove and Pike have vested interests which may color their views on price elasticity of demand for space launches. Shove's interest is clearly visible—commercial payoff for the state of Florida. Pike's vested interest—the status-quo orientation of the Federation of American Scientists—is less visible but no less real. Hence, we need to look at more than just the opinions of protagonists.

In 1992, the total number of launches worldwide for communications satellites (nonmilitary) and navigation satellites was 35, which is already significantly more than the 15 forecast by Corcoran and Beardsley.[14]

However, this number of launches is for communications satellites, which are already profitable at today's launch costs. We need to look at the issue of whether reduced launch costs would increase the market for launches.

Another market is remote sensing satellites. France already sells photos from its SPOT earth resources satellite. In addition, the French space agency CNES plans to launch a radar satellite to complement the SPOT optical imaging satellite.[15] This radar satellite would use an X-band synthetic aperture radar, for a resolution of from 20 to 50 meters. While this resolution is inferior to the 10-meter resolution obtained from SPOT, the radar satellite could obtain images despite cloud cover over the target area. In addition, comparison of microwave and optical reflectivity of the same site on the Earth could provide additional information about the nature of the surface (type of vegetation, mineral and moisture content, etc.). CNES is seeking international partners to share the cost of this satellite. If this satellite were being developed by CNES alone, profit might not be a consideration. However, if international partners are to share in the funding, they will most likely expect a positive return on their investment. Hence, a reduction in launch costs will make this satellite more likely to be launched.

The potential for remote sensing satellites is shown by the fact that in 1992, there were 9 such satellites launched. Again, these are expected to be valuable at current launch costs. Reduced launch costs are likely to increase the demand for photos and radar images from space; thus, increasing the demand for remote sensing satellites.

In 1992 there were 12 launches of scientific payloads, both orbital and interplanetary. The current level of activity in space science is not limited by the potential for research, but by the cost of research. A major portion of the cost of space science is the cost of launching the payloads. Reducing that cost would make a great many more experiments possible.

In examining the need to reduce launch costs, we should not look only at commercial and scientific missions. In an era of declining defense budgets, reduced launch costs can definitely increase the number of military missions demanding launches.

For the U.S., photoreconnaissance satellites are an important military mission. Their importance was reemphasized during the Gulf War, when they played an important role in locating Iraqi forces and spotting SCUD launches. In addition, satellites play an important role for the U.S. in monitoring the nuclear test ban treaty, and in providing early warning of ICBM launches.

---

[14] Tina Thompson, *TRW Space Log*, Redondo Beach, Calif., 1992.

[15] Anonymous, "France Defines Satellites to Complement Spot Series," *Aviation Week and Space Technology*, October 23, 1989, p. 48; and James W. Canan, "How to Make Space Science a Frequent Flier," *Aerospace America*, October 1989, pp. 54–59.

However, another possible military mission is space-based radars for detection of bombers and cruise missiles in flight.[16] Currently the defense of the continental U.S. against air-breathing vehicles requires detection of attackers by ground-based radars. These are limited in range, and their coverage at low levels is full of holes. The former Soviet Union had already deployed what are known as Radar Ocean Reconnaissance Satellites (RORSATs) to track ships at sea. Hence the technology for space-based radars is already within the state of the art. Part of the problem for the U.S. military services is to find room for these satellites in the budget. Reduced launch costs would definitely increase the likelihood of these satellites being deployed.

Another proposed military mission is space-based interceptors for missile defense. Strategic defense against ballistic missiles has been highly controversial, for many reasons which we need not explore here. One reason, however, has been cost. One issue which has been raised is "cost-effectiveness at the margin." That is, an increment to enemy offensive missiles should be more costly than the defensive increment required to offset it. Clearly, launch costs enter into this calculation, since lowered launch costs increase the cost-effectiveness of the defenses.

Current thinking regarding missile defense emphasizes theater defense rather than defense of the continental U.S., and use of ground-based launchers. However, if the cost of launches could be brought down, space-based defense of the continental U.S. would be much more feasible.

These examples suggest that there are many missions waiting "in the wings" for lower launch cost. Reduced launch costs would be expected to increase significantly the market for more launches.

However, the question of elasticity of demand for space launches does not depend solely on whether these possible missions are realized. We can draw an analogy from a related field, air travel.

Figure 1 shows the revenue-passenger-miles flown by domestic U.S. airlines, versus average cost per passenger-mile (in constant 1982 dollars). Historically, lower travel cost has been accompanied by greater passenger travel. Lower fares did indeed attract more passengers.[17]

It is worth noting that in the 1950s, before the advent of air coach fares, the same arguments were made about air travel as are today being made about space launches. It was taken for granted that the demand for air travel was inelastic, that only small numbers of businessmen would ever travel by air and they would do it almost regardless of cost, and that reductions in air fares would only reduce the revenues of the airlines. The introduction of "coach" fares was a radical innovation, resisted by both the existing airlines and the regulatory

---

[16] James W. Canan, "The Big Hole in NORAD," *Air Force Magazine*, October 1989, pp. 54–59; and James J. Coghlan, Jr., "High Altitude Satellites to Support the Lower Echelons," *Defense Electronics*, October 1989, pp. 61-67.

[17] The evidence in Figure 1 is not as solid as might be desired. Since 1971 the trend in real prices for air travel has been downward, and the trend in air travel has been upward. There have been no reversals in either one. From an analytical standpoint, we would like to have data in which some years saw increases and some decreases in real air travel prices, to see if they were accompanied by contrary fluctuations in air travel. We have only one real world to work with, and we cannot conduct controlled experiments in it. Thus, we have to make do with monotonic trend data rather than data with movements in both directions. However, the historical data does exhibit the behavior which would be expected from economic theory. Hence, it can be taken as support for the thesis that, to the extent that air travel and space travel are analogous, lower prices will create demand for additional launches.

authorities. As it turned out, all these assumptions were wrong. It appears highly likely that they will be wrong about space launches as well.

Figure 1.
Airline Revenue Passenger Miles
As a Function of Price, 1977-1990

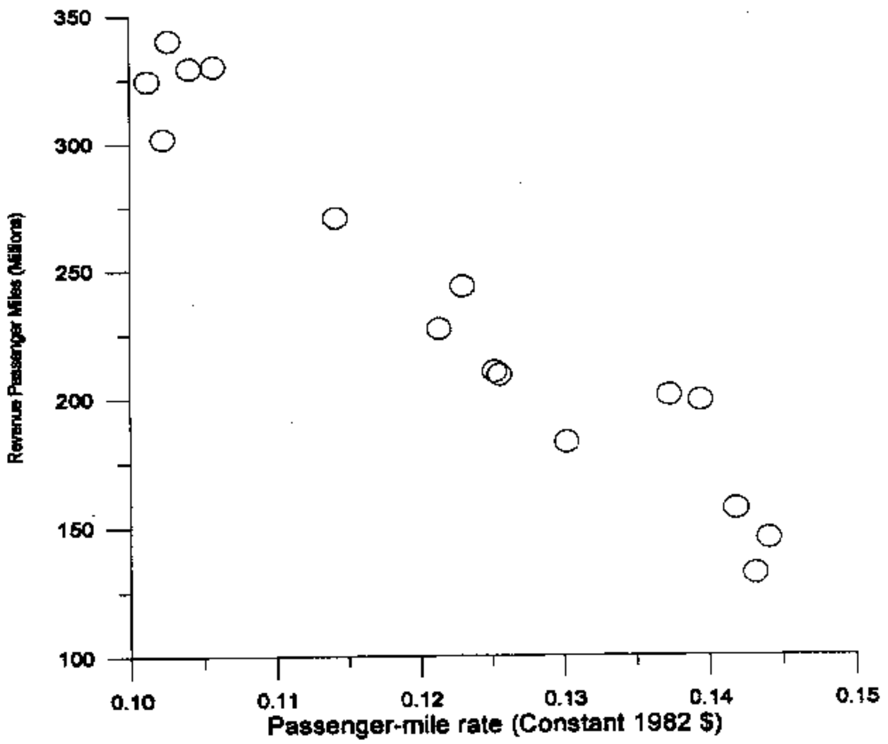

In 1990 the Commerce Department issued a report which stated that the chief barrier to commercial space ideas was the government itself—the government is the "doorkeeper to space."[18] Wider availability of nongovernment launch providers should help overcome this barrier, and lead to greater commercial use of space.

III. HIGH RELIABILITY vs. HIGH COST

Given that lower launch costs will create additional demand for launches, what are the prospects for reducing the cost of getting into space? Is this possible, or are we already doing the best that's possible?

As we have already seen, the history of the U.S. space program has been a series of vehicles which pushed the state of the art, were used for a few missions, then were literally thrown away, to be replaced by another design. This inevitably resulted in high costs per launch. First, the development costs for the vehicles could be amortized over only a few flights. Second, there was no opportunity for long production runs to bring costs down through operation of the aircraft industry's well-known "learning curve." In addition to these cost-raising factors, there has been the problem that the government really had no incentive to reduce costs.

There have been three approaches proposed for getting away from this "high-tech, high-cost" syndrome. These can be somewhat crudely categorized as "commercial practice," "long production runs," and "airplanes vs. bullets". These approaches are not mutually exclusive. It would be possible to implement all three together. Nevertheless, each arises from a different perspective on space launches, and we will look at each individually.

A. Commercial Practice

The practice of developing one-use exotic vehicles is in direct contrast with commercial practices, which emphasize low cost, and which utilize new technology only when it is needed to achieve a commercial goal. The problem is finally being recognized. The Pegasus is a commercial attempt to get away from current practices in space launcher design.

Pegasus. The Pegasus vehicle was developed by Orbital Sciences Corp., in a joint venture with Hercules Aerospace. While privately financed, it is intended to launch payloads for the Department of Defense as well as for private firms. Even before development was complete, the Defense Advanced Research Projects Agency (DARPA) had taken options for space on the first six launches of Pegasus.

The specific reason for developing Pegasus was to reduce launch costs. Company founder and chairman David Thompson was quoted as saying, "The basic concept behind OSC was to start a new company with a real emphasis on lowering the cost of launching satellites into space."[19]

Pegasus is launched from a B-52 bomber at an altitude of 40,000 feet. Air launch has several advantages over ground launch. The launcher can be fired on any desired azimuth, from equatorial to polar. In addition, the aircraft speed itself provides an initial boost to the launcher. Finally, launch at 40,000 feet means the launcher starts off above most of the Earth's atmosphere; thus, losing less energy to drag.

---

[18] Editorial, "Unlock Strategic Space Investments," *Aviation Week & Space Technology*, June 4, 1990 p., 7.

[19] Anonymous, "Rising Stars," *Defense Electronics*, July 1989, p. 63.

A key element in keeping the cost of Pegasus down is the use of existing technology. The guidance system was originally developed for the U.S. Navy's Mark 48 torpedo, and the flight computer was originally developed for the Army's M-1 main battle tank. However, some new technology went into Pegasus. The external surfaces were designed entirely using high-speed computers to simulate the air flow over the vehicle. This eliminated the more conventional wind tunnel tests of subscale models, and the launch of test vehicles. Avoiding these more conventional approaches saved considerable money during development.

Pegasus achieved its first successful launch on April 5, 1990, carrying a NASA experimental payload and a Navy communications satellite. On the basis of that success, Orbital Sciences Corp. contracted for five more launches for DARPA. In the third quarter of 1990, it posted its first profit of $502,000 after previously losing $5.4 million on revenues of $56.8 million.[20][21]

Pegasus is not, of course, really a low-cost launcher. The cost to put a pound in orbit using Pegasus is about $6000—much higher than the Shuttle, and about the same as the existing Scout rocket. The main advantages of Pegasus are that it is a flexible launch system which can respond quickly to unexpected launch demands. Nevertheless, its quick development and low development cost, based on use of subsystems developed for other purposes, is a major departure from conventional launcher design.

The "Barbarian Booster." In 1987 the Strategic Defense Initiative Organization (SDIO) planned a space-based test of a high-powered laser, under the name Zenith Star. The intent was to launch some 100,000 pounds of experimental equipment into orbit in 1990. No existing launch vehicle could carry that weight of payload. However, a preliminary design by McDonnell-Douglas pointed the way to a solution.[22] McDonnell-Douglas engineers had prepared a design for a heavy-lift vehicle consisting of a Delta rocket with a Titan payload fairing. The Delta's thrust was to be augmented by two Shuttle Solid Rocket Boosters. It would have been able to put 87,000 pounds in orbit. The obvious "brute force" approach of the design led to the nickname "barbarian booster." SDIO contracted with McDonnell-Douglas to scale up this design to meet the 100,000-pound payload requirement. SDIO also contracted with Martin-Marietta for a design based on the Titan to achieve the same payload.

The 36-month deadline meant that no new technology, particularly nothing requiring extensive testing or qualification, could be used. Instead, the design goal had to be achieved by combining existing, proven components. Both firms completed designs based on their existing rockets, and other available components. The Martin design is particularly notable in that the engineers deliberately overdesigned certain parts of the structure (safety factor of 1.6 instead of the customary 1.25) in order to avoid the need for testing those components. Designing the structure to be sturdier (and heavier) than absolutely necessary led to savings in both money and time.

The McDonnell-Douglas design would have achieved the 100,000-pounds-to-orbit goal; the Martin design would have been capable of putting 118,000 pounds in orbit. Estimated cost for the first vehicle of each kind was about $500 million; second and later vehicles would have cost about $200 million. That is, nonrecurring costs (design,

---

[20] Bob Davis, "Startup Firm Faces Big Risks," *Wall Street Journal*, April 6, 1990.

[21] Industry Observer, "Business at OSC Looking UP," *Aviation Week & Space Technology*, November 5, 1990, p. 13.

[22] Larry Stafford, and Michael J. Rendine, "Zenith Star Launch System," *Aerospace America*, September 1990, pp. 41–43.

fabrication of launch equipment, etc.) accounted for about $300 million of the cost of the first vehicles. The second vehicle could thus have delivered its payload to orbit at an unsubsidized cost of $2000 per pound—not yet "low cost," but lower cost than anything else available.

As it turned out, neither design was ever built. Cuts in SDI funding delayed the launch of Zenith Star. Furthermore, the barbarian booster ran into political difficulties. To some, it was viewed as a competitor for the then-planned Advanced Launch System (ALS) which, however, has since been canceled. To others, it looked like an attempt at early deployment of SDI, before Congress had approved deployment. The combination of these difficulties led to termination of both projects. Nevertheless, they demonstrated that using commercial practice—proven components, overdesign, "worst case" design—can significantly reduce the cost to put a payload in orbit.

These examples show that the message is getting through: the way to reduce launch costs is to use existing technology when it will do the job, instead of continually pushing the state of the art. Both government and industry have recognized this fact. It is only a matter of time until this recognition is implemented in lower-cost hardware.

B. Long Production Runs

Even if high technology is necessary in a system, this still doesn't mean the system has to be costly. Commercial suppliers have learned that long production runs allow costs to come down. The idea is that "learning" takes place, in the sense that people learn how to make the product more cheaply. There are at least two aspects to this learning.

First, a design may turn out to have manufacturing problems built into it. Specific components or features may be difficult to manufacture or to assemble. If only a few items are to be built, it usually isn't worth redesigning the system to eliminate the built-in cost problem. However, if a long production run is planned, the redesign cost can be amortized over many items, and has more chance to pay for itself.

Second, the manufacturing process itself may have some high-cost elements. Reducing the manufacturing cost may require special jigs, fixtures, or tools. On a short production run, it would not be possible to recover the costs of these special items. On a long production run, however, small savings on each unit produced can add up, and justify the cost of special fixtures or tools.

This effect of declining cost with cumulative production has been given the name "learning curve." The manufacture of commercial transport aircraft is frequently considered to follow what is called an "80-percent learning curve." This means that every time cumulative production doubles, the unit production cost declines to 80 percent of the previous cost.[23] Electronic equipment is frequently found to follow a 70-percent learning curve, while mechanical equipment follows a 90-percent learning curve.

Figure 2 illustrates the learning curve, and the effect of long production runs on costs. The cost of the first unit is arbitrarily taken to be 100. Thus the chart can be read as the percent of first-unit cost for any subsequent unit. For a 90-percent learning curve, the cost of the thousandth unit is down to 35 percent of the first unit cost. For an 80-percent learning curve, cost is reduced to about 12 percent of the first unit cost. For a 70-percent learning curve, cost at the thousandth unit is down to about 3 percent of the first unit cost. These dramatic reductions

---

[23] That is, the second unit costs 80 percent of the first unit's cost, the fourth unit costs 80 percent of the second unit's cost, the eighth unit costs 80 percent of the fourth unit's cost, etc.

Figure 2.
Hypothetical learning curves.

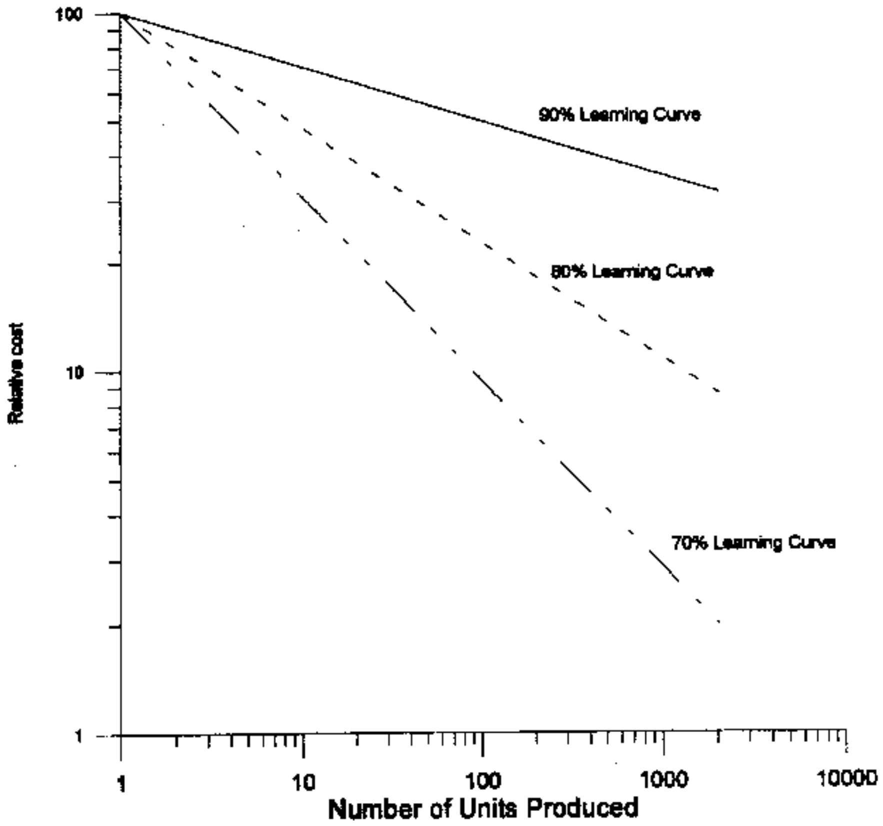

are not simply a mathematical exercise. This behavior has been found time and again, in many different kinds of manufactured products. Long production runs of space launchers should be expected to result in reductions in cost similar to those in the figure. Even if space launchers were found to follow only a 90-percent learning curve, as opposed to the 80-percent typical of aircraft, long production runs would still result in sensational reductions in unit costs of launchers.

Of course, long production runs do not automatically result in reduced costs. The Titan, Atlas and Delta vehicles, for instance, have been in production for 30 years, yet have not achieved low cost. The reason is that they are ordered irregularly, and are never produced at the optimum rate to reduce costs. Although these vehicles have been in use for three decades, they have never really been in "mass production" in the same sense that commercial aircraft such as the Boeing 747 are in mass production. This illustrates the point that market size and production cost exist in synergism. Low production cost is needed to increase the market; promise of a large market is needed to justify the up-front investment in low-cost mass production.

The real-life benefits of mass production can be seen in a comparison between the RL-10 rocket engine used in the Centaur upper stage and a gas turbine helicopter engine. Both have about the same number of parts and are of about the same degree of complexity. However, the RL-10 costs about $2.5 million while the helicopter engine costs about $80,000. The difference is that the production rate for RL-10 engines is 20 to 30 per year, while the production rate for helicopter engines is several thousand per year.[24] An equivalent reduction in launcher costs would have a major impact on the economics of space launches.

While long production runs can reduce costs directly, they have another cost-reduction benefit which is often overlooked in discussions of reducing launch costs. One of the costs which the owners of spacecraft face is insurance. Current premiums for launch insurance run about 20 percent of the cost of the payload. That is, for a $500-million satellite, launch insurance runs about $100 million, or about the cost of the launch itself. It should be recognized, of course, that the insurance premium is intended to be equal to the expected cost of a launch failure. Thus, whether a spacecraft owner buys insurance or instead self-insures (i.e., bears the risk of failure himself), over the long-run the two costs should be the same. Spacecraft owners buy insurance because for them there often is no "long-run." They are making a single launch, or at most a few launches. The insurance company, by contrast, is in the business for the long-run. It charges a premium on each launch which approximates the expected cost of a failure (plus normal profit, of course).[25]

The insurance premium, as a fraction of payload cost, is thus an inverse measure of the likelihood of loss of the payload due to launch failure. As John Higginbotham and Peter Stark note, "Insurance is not what is expensive; it is space itself which is expensive. The risks are high, and those wishing to take the risks must realistically evaluate the associated costs."[26] Decreasing the likelihood of failure will directly reduce the cost of launch insurance, and thereby reduce total cost of deploying a spacecraft.

The Office of Technology Assessment has recognized that launch failures are inevitable.[27] They note that even if NASA is correct in its estimate that the shuttle is 98-percent reliable, this amounts to a 50-50 chance of

---

[24] "Big Dumb Boosters," p. 25.

[25] In reality, space insurers have undercharged space launch operators in the past. The space insurance industry has suffered cumulative losses (premiums less payments for losses) of something like $500 million since 1968 (Higginbotham & Stark).

[26] John B. Higginbotham, and Peter M. Stark, "Insuring Space Ventures," *Space Commerce*, Vol. 1, No. 1, Summer 1990, pp. 19–28.

[27] U.S. Congress, Office of Technology Assessment, *Round Trip to Orbit*, OTA-ISC-419 (Washington, D.C.: U.S. Government Printing Office, August 1989).

failure in 34 flights. At an average of 10 flights per year, this comes to a 50-50 chance of a failure every three years.[28]

As Peter Leonard and William Kisko show, however, launch failure rates decrease as a given vehicle design is used over and over. They categorize causes of failure into design failures and processing failures. Their data show that most design failures are caught early in the life history of a launcher. Subsequent launches of the redesigned launcher therefore have a lower likelihood that they are harboring a design error which will cause a failure under previously untried conditions.[29] Similarly, processing failures become less likely as the same design is launched repeatedly. Mistakes in processing, once identified, can be prevented by establishing and following procedures designed to avoid them.

Leonard and Kisko provide failure data based on 134 Titan launches. No design errors occurred after about 45 launches. The rate at which process errors occurred decreased as more launches were made. After about 80 launches, process errors occurred randomly at a constant rate of about one in 30 launches (e.g., slightly less reliability than the 98 percent claimed for the shuttle).

Aviation Week & Space Technology[30] reported that during 1989 and 1990, there had been two communications satellites lost due to human error in launching. Out of some two dozen launches worldwide during that period, this comes to a loss rate of about 8 percent due to processing failure alone, not counting design failures. Actual loss rates for insured payloads were actually about 21 percent over the period 1983–1988.[31]

An important implication of this is that bringing down insurance costs requires a long sequence of launches of the same vehicle. As design errors are detected and corrected, and as launch procedures are improved to eliminate processing failures, launch reliability increases and insurance rates will decrease; thus, decreasing the insurance costs through long production runs of a given launcher works in tandem with the learning curve, to reduce launch costs. These reductions are over and above those coming from amortizing nonrecurring costs over a greater number of launches.

In summary, long production runs can significantly reduce launch costs. The key is to use the same vehicle repeatedly, instead of switching to a new model, with the latest technology, after only a few launches.

---

[28] This points to the fact that before 1995, the United States will have to decide what to do about replacing the current shuttle fleet. One would hope that it will be replaced with something embodying up-to-date, cost-reducing technology.

[29] Peter B. Leonard and William A. Kisko, "Predicting Launch Vehicle Failure," *Aerospace America*, September 1989, pp. 36–38.

The American Rocket Company suffered a failure on its first launch attempt at Vandenberg Air Force Base, Calif. on October 5, 1989. The failure was traced to a stuck oxygen valve which reduced the thrust of the engine. This valve had been tested extensively in the New Mexico desert without any problems. The cause of failure apparently was ice which formed in the valve under the higher humidity conditions found near the ocean at Vandenberg. This is a design error which was uncovered only under conditions not previously encountered. Once fixed, the problem should not recur.

[30] Washington Roundup, "Oops!," *Aviation Week & Space Technology*, November 26, 1990, p. 23.

[31] Higginbotham and Stark, "Insuring Space Ventures."

C. Airplanes vs. Bullets

As noted above, the earliest space launchers were derived from missiles. They inherited the high-performance, high-cost tradition of missile design. They also inherited something more subtle: what might be called a "round of ammunition" or "bullet" mindset on the part of designers.

A round of ammunition is designed as, literally, a one-shot device. One can reuse a cartridge case by reloading it, but this amounts to nearly complete remanufacture. The case must be resized and reshaped, then reloaded with powder. A bullet, to be reused, must actually be melted and molded again.

A missile is likewise a one-shot device. Even if parts of it are recovered (boosters, first stages, etc.), they must be remanufactured in order to be reused.

If a round of ammunition fails during the attempt to fire it, there is no way to recover it and repair it. Any failure is a complete failure. It is simply not designed to be repaired.

Likewise, any missile failure is a complete failure. Not only can the failed part not be repaired, the missile itself is lost completely.

Testing a round of ammunition means consuming it, using it up, destroying it. There is no way to verify that it will work, then put it back in the armory. It is impossible to conduct 100 percent testing of a production lot of ammunition; there would be none left to fire at the enemy. Instead, a sample must be tested, and the statistics of the sample projected to the rest of the production lot.

Similarly, testing a missile means using it up. If it fails, it is destroyed; if it works, it is destroyed. Testing a missile means testing the *design* by repeatedly launching identically manufactured missiles.

How reliable does a missile need to be? Historically, military services have accepted a launch success rate of about 95 percent as being adequate. This means that 19 out of 20 missiles will hit their targets. Any targets still surviving can be hit in a follow-up attack. If the target is really critical, two missiles can be launched at it simultaneously, with fewer than 3 chances in 1000 that *both* missiles will fail. Hence for ballistic missiles, typical launch reliability figures of about 95 percent are quite reasonable.

Now consider adapting an existing missile for use as a manned space launcher. Reliability of 95 percent means that if an astronaut plans to make 4 flights to orbit, he has only an 81 percent chance that all four will be successful. Put another way, he has a 19 percent chance that he will be killed before the 4 flights are completed. "Russian Roulette" offers better odds. If the missile is put through a program to increase reliability, in order to "man-rate" it, the figures are still discouraging. Even raising launch reliability to 98 percent means that a "career" of 10 launches gives the astronaut about the same odds as Russian Roulette. The basic problem is that a missile has little or no redundancy. If any part fails, the whole thing fails. Moreover, it has no "engine-out" capability. Even on a multi-engine rocket, the failure of a single engine is catastrophic.

The end result is that man-rating a missile means launching missiles of the same basic design repeatedly until all causes of failure appear to have been discovered and fixed through redesign. When space launchers are designed like missiles, the same holds true.[32]

An airplane is designed from a totally different perspective. An airplane can be reused simply by refueling it rather than remanufacturing it. While some airplane failures are catastrophic, most are survivable because the airplane can still be flown to an emergency landing. An airplane is designed to have redundancy (multiple load paths) in the structure; it has multiple fuel tanks; it has "back-up" instruments which can be used if the primary instruments fail; if it has multiple engines, it is usually designed to maintain altitude even with one engine out. Airplane designers don't even use the term "man-rated" because an airplane is inherently man-rated. Once a test pilot has verified that the structure will not fail catastrophically in normal use, it is taken for granted that the airplane can carry a crew and passengers. Airplane builders don't throw away a new airplane for each test flight, trying to gather statistics about the design. They build one test item, and fly it repeatedly. Each test flight goes only slightly beyond the bounds of safe performance already verified on previous flights. In most cases, design failures do not result in loss of the test airplane. Instead, the test airplane is landed, the problem is corrected, and the airplane is flown again.

The third approach to reducing launch costs is based on replacing the missile design perspective with the airplane design perspective. The basic idea is to design the vehicle so that failure is not catastrophic; so that the vehicle can return safely even in the event of a failure. This is referred to as an "intact abort" capability.

This approach is now being demonstrated in the Delta Clipper Experimental (DC-X) vehicle. This vehicle has been developed by McDonnell-Douglas under the sponsorship of the Defense Department's Ballistic Missile Defense Organization. The DC-X flew for the first time on August 18, 1993. It flew to a height of 100 feet, moved sideways 300 feet, and landed on a second prepared landing pad. On September 11, 1993, it performed the first *second* flight of a reusable space vehicle, rising to an altitude of 300 feet and again displacing sideways for 300 feet before landing. The launch crew required was only 10 people.[33] DC-X has thus successfully demonstrated that a reusable vehicle does not need to be returned to the factory before reuse. It needs only to be refueled, like a commercial transport aircraft.

The DC-X is only an experimental vehicle. Its maximum design altitude is 18,000 feet. Thus, it can never reach orbit itself. However, it has already demonstrated everything needed for the airplane design concept as opposed to the bullet design concept. Moreover, it has demonstrated the soundness of the intact abort concept.

Plans for a follow-on DC-Y vehicle call for the ability to deliver 25,000 pounds into low earth orbit at a direct cost of $10 million, or $400 per pound. The cost of each vehicle would have to be amortized over its lifetime, adding to this cost. In addition, nonrecurring costs such as development would also have to be amortized over the lifetime of the entire fleet. However, the point is that over a reasonable lifetime for each vehicle, such as

---

[32] The alternative to man-rating is escape and rescue systems. Often these are cheaper than pursuing reliability improvement. However, escape systems for manned spacecraft would be useful only in the vicinity of the Earth. They would have little value, for instance, on a manned lunar trip.

[33] About 9000 people are required to carry out all the operations necessary to get a space shuttle into orbit and back. By contrast, all U.S. domestic airlines together have an average of 149 employees per airplane. Even experimental aircraft, which are at least as complex as space launchers, typically require less than 400 people to get them off the ground and back. The difference is that manned aircraft are expected to survive even a fairly serious failure, while ammunition-styled space launchers (including the space shuttle) *must not fail* because it is assumed they can't survive a failure.

that exhibited by a 747 aircraft, and a reasonable fleet size, development costs and vehicle costs come to less than the direct costs, and total cost would be less than twice the direct cost.

Whether or not the plans for the DC-Y itself ever come to fruition, the "bottom line" is that a given amount of payload can be placed in orbit far more cheaply by genuinely reusable vehicles than by throwaway vehicles, particularly if the turnaround time is short (i.e., many launches per year by each vehicle). Simply compare the weight of the vehicle fleets required to put a given amount of payload into orbit. One-shot vehicles inevitably mean that more weight is thrown away than is placed in orbit. With reusable vehicles, the weight placed in orbit can be many times greater than the total weight of the vehicle fleet. This translates directly into lower cost to do the job. A reusable vehicle can be more costly than a throwaway vehicle and still be more economical in use.

It may be objected that we have heard the arguments for reusability before. They were made when the shuttle was proposed, and they turned out to be false. The shuttle is actually more expensive than an expendable vehicle.

This is true but irrelevant. Despite the original claims, the shuttle is not really "reusable." Instead it is "refurbishable." About 50 percent of the launch cost of a shuttle is expendable hardware. Reusing the shuttle is more like reusing a round of ammunition than like refueling an airplane. It can be done, but only at great expense and difficulty, since safety of the crew depends upon the complete absence of failure. Despite the original promises, the shuttle was designed like a round of ammunition, not like an airplane.

The airplane concept, as exemplified by the DC-X, means designing for recovery despite failure rather than for complete absence of failure. It should therefore be able to fulfill the promises which the shuttle could not fulfill.

In this discussion of DC-X we have concentrated on the advantages deriving from intact abort capability and the inherent reusability it provides. We have said nothing about reducing the cost of the vehicle itself by commercial design practice and long production runs. However, it should be clear that reusability adds leverage to low production costs and low operating costs. The two in combination can significantly reduce the cost of placing a payload in orbit. Thus the airplane concept should not be seen as competitive with the idea of long production runs of spacecraft designed to commercial standards. On the contrary, the ideas are synergistic. Every dollar that can be taken out of the cost of a single flight comes back every time the vehicle is reused, instead of once only for a throwaway vehicle.

IV. BUYING LAUNCHES vs. BUYING LAUNCHERS

The two preceding sections have argued that space need not be as costly as it has been for the past 30 years. By adopting practices well known in the commercial world, it would be possible to reduce the cost of launch vehicles and of launches. The important question is, how do we get there from here? How do we commercialize space?

One of the most important issues facing the government is whether it should buy launchers, or should instead buy launch services from private firms. This amounts to a decision as to whether the government should operate its own spaceline, or should instead buy tickets for occasional flights on a commercial spaceline. The issue is somewhat similar to the issue of whether the government should run its own airline, or instead buy tickets for its employees on commercial airlines.

Clearly the military services must retain their own space launch capabilities. In wartime they must have access to space, and must exercise this capability in peacetime, just as they exercise their flying capabilities in peacetime.

The case of NASA is at least arguable. So long as NASA is charged with conducting R&D on space vehicles, there may be some need for it to contract for the construction and testing of experimental vehicles.

In any case, there is no need for NASA to operate a spaceline. In fact, NASA's continued involvement raises costs for everyone. Launch of commercial satellites can be done using commercial launch vehicles. Even launches of scientific spacecraft can be done using commercial vehicles. Thus NASA funding of space science should go for construction of scientific spacecraft, and for purchase of launches from commercial sources rather than for purchase of launchers and operation of a "spaceport."

There is precedent for purchase of launches for scientific purposes. On March 29, 1989, the Consort I vehicle was launched by Space Service from White Sands Missile Range in New Mexico. It was a suborbital rather than orbital flight. It carried several experiments requiring short periods of weightlessness. The University of Alabama at Huntsville paid $1 million for the launch. Much of the funding came from a NASA program to encourage commercial development of space. Thus NASA was indirectly purchasing a launch rather than a launcher. A Consort-3 vehicle, launched on May 17, 1990, carried 12 materials-processing and biotechnology experiments in a suborbital flight. Ten of the experiments were successful; partial results were obtained from the other two. Funding again came in part from NASA support of space science; NASA was indirectly purchasing a launch rather than a launcher.

One way to start reducing the cost of space launches, then, is to encourage space scientists to purchase commercial launches. Whether their support comes from NASA or from some other public or private funding agency, they should be free to purchase launch services from the lowest responsible bidder. They should not be a captive market for NASA. This isn't a big market, but it is a start.

Ironically, this would actually benefit space science and space scientists.[34] NASA's commitment to the space shuttle meant that scientists' experimental packages had to be designed for the shuttle even if manned capability wasn't required. In particular, this meant that human safety often became a major consideration in the design of an experimental package, even though the package had no requirement to be man-tended in space. Man-rating space experiments has raised the cost significantly. Moreover, as science payloads are bumped from the shuttle or rescheduled, the scientists themselves waste a lot of time "analyzing and reanalyzing revised mission scenarios, in budgeting and rebudgeting exercises, and in planning and replanning research programs for students, colleagues, and themselves."[35] Clearly these problems could be reduced or eliminated if scientists were free to purchase launch services from providers other than NASA.

Likewise, NASA itself should be encouraged to purchase launch services from the lowest bidder, when the NASA endeavor is the payload itself, not R&D on launchers. NASA should not look upon its own space R&D activities as requiring it to perform the launches as well. NASA should purchase cargo capacity into space just as it purchases air freight, from private suppliers. Again, this is not a large market, but it should be open to competing providers of launch services.

---

[34] M. Mitchell Waldrop, "A Crisis in Space Research", *Science*, Vol. 235, January 23, 1987, pp. 426–429.

[35] *Ibid*, p. 426.

A start has been made in this direction. On July 25, 1990, a commercial Atlas Centaur was launched by General Dynamics for NASA. The payload was a joint NASA-Defense Department satellite intended to measure near-Earth electric and magnetic fields, and the effects of space radiation on microelectronics.[36][37]

It is worth noting that much of NASA's budget has gone towards high-cost launch systems. This has resulted in space science being put on short rations. Giving space science a higher priority (and a bigger share of the budget) would also increase the number of launches, further increasing the market for low-cost launch vehicles.

In short,then, we must change the way government agencies currently think. To the maximum extent possible, they should buy launches rather than launchers. Moreover, when they contract with someone to perform some activity in space, the contractor should be free to purchase launch services from competing sources, rather than be required to utilize NASA launch services. This is especially true when the payload doesn't require the manned capability of the shuttle.

V. THE ROAD TO AFFORDABLE SPACE

For airlines, the direct cost of operation is about one-third fuel, two-thirds everything else.[38] Thus as an ultimate goal, airline-style operation should mean that the cost of placing a payload in orbit is about three times the cost of the fuel. The rocket fuel to place a pound in Low Earth Orbit costs about $5. Thus ultimately we should be able to place a pound of payload in orbit for a direct cost of about $15 in today's dollars.

This goal is still in the future. Nonetheless, it is a realistic one. We should be focusing on how to reach it, rather than continuing to think of space as inherently expensive. Given an emphasis on making space affordable to a wider range of users, it should be possible to reduce launch costs to $50 per pound by the end of the century.

There are three mutually reinforcing ways to achieve that ultimate goal of $15 per pound. These are use of commercial design practices, use of long production runs (uninterrupted and at efficient rates), and shifting the design philosophy from one appropriate for bullets to one appropriate for airplanes. While these approaches can ultimately achieve the goal of making space affordable, space cannot be made cheap in one big step. The design of new launchers, and their mass production, must go hand in hand with development of markets which require large numbers of new vehicles.

The biggest current market, of course, is communications satellites. These can be profitable even without the subsidized launch costs of the shuttle. Under current U.S. policy, NASA will not launch these satellites. Hence, they already provide a market for private launch services suppliers. Moreover, several suppliers are already competing for this market. Since reduced launch costs provide a competitive edge in this situation, private

---

[36] Anonymous, "CRRES Launch Gives NASA Something to Cheer About," *Countdown*, September 1990, p. 11.

[37] Anonymous, "First Commercial Atlas Places Scientific Satellite Into Orbit," *Aviation Week & Space Technology*, July 30, 1990, p. 24.

[38] Direct cost does not include capital cost of the aircraft or fixed facilities. It does include fuel, salaries, and consumable supplies such as in-flight meals.

launch services suppliers already have an incentive to reduce the costs of their services. Thus, we may expect to see a shift to commercial design practices as a means of reducing costs.

The communications satellite market alone is not large enough to provide the long production runs needed to get the cost down to the ultimate level. Nor can it alone justify the shift to a totally new concept like the Delta Clipper. Additional markets will have to be developed.

The historical evidence, however, is that reduced transportation costs attract additional passengers and freight. Thus we can expect to see costs coming down and markets expanding in tandem. Reduced costs will expand the market; expanded markets will allow further cost reductions.

Making space affordable does not require new technology *per se*. Instead, it requires a shift in our ideas about the nature of the round trip to orbit. We must stop thinking of space travel as somehow "different" from things we do here on Earth. Instead, we must apply to space travel the same thinking which brought the cost of airline tickets, automobiles, TV sets, and even computers within the reach of the average person. That kind of thinking will ultimately bring space travel within the reach of the average person.

ABOUT THE AUTHOR

Dr. Joseph P. Martino is a Senior Research Scientist at the University of Dayton Research Institute, where he prepares technological forecasts for industry and government. He retired from the Air Force as a Colonel and also developed advanced technologies for the Air Force Laboratories. Dr. Martino is an Associate Fellow of the American Institute of Aeronautics and Astronautics and a Fellow of the Institute of Electrical and Electronics Engineers.

ACKNOWLEDGEMENTS

I am grateful to the Reason Foundation for providing the support which made this paper possible. I wish to acknowledge the valuable comments of G. Harry Stine, of the Enterprise Institute, and of Major Jess Sponable, of the National Aerospace Plane (NASP) Project Office, as well as those of an anonymous reviewer, all of whom reviewed an earlier version of this paper. All remaining errors of omission or commission remain my responsibility.