
Our National Space Science Program: Strategies to Maximize Science Return

by Greg S. Davidson

There is a concern at NASA to learn from the lessons of the past and to respond to the concerns of the space science community. We have been grappling with the problems of:

- Big versus small missions
- Multiple simple spacecraft versus spacecraft servicing
- A culture of risk avoidance
- Unrealistic budget planning
- Institutional and political forces
- Linkage to the manned space program

As we are in a very dynamic time for space astrophysics, our discussion must be placed in the context of current events. The missions of the 1980s have been launched and data is beginning to arrive. More new astrophysics missions were expected to be launched between 1991 and the end of 1993 than in the decade of the 1980s. With the greater emphasis on long-term operations designed into several of these recent missions, the supply of new data should continue to grow. Over the last several years, the growth in funding for space astrophysics has exceeded inflation by roughly 15 percent, and the growth in funding for science and data analysis is keeping up with the growth in science data.

There has also been tremendous criticism of the NASA astrophysics program from the media, Congress, public and some members of the science community. Problems with the Hubble mir-

ror have been fodder for comedians and commentators, while other difficulties such as those with hydrogen leaks on the Space Shuttle, or the rash of problems overcome on the Astro mission, have helped to create a perception of serious problems. It is always appropriate to review our programmatic strategies for conducting space astrophysics in light of experience, and to develop our strategy for the future. So let us briefly note the recent scientific output of our space astrophysics program, and then discuss the areas of concern identified above.

■ Status of Astrophysics in January 1991

Cosmic Background Explorer

(cryogenically cooled mission complete
November 1990)

- A smooth big bang
- No later bangs
- Unrivaled data from the infrared background

Hubble Space Telescope

(checkout complete, beginning science verification and operations)

- Potential black hole in nearby galaxy
- Circumstellar ring around Supernova 1987a
- Storms on Saturn
- High-resolution imagery of Pluto and Charon

Roentgen Satellite

(survey complete - U.S. observation time begun February)

- 80,000 new X-ray sources
- 1,000 new extreme ultraviolet sources

Astro/BBXRT

(mission complete December 1990—data being analyzed)

- Results to come . . .

Recent and Upcoming Science Missions

- Gamma Ray Observatory (1991)
- Array of Low-Energy X-ray Imaging Sensors (1991)
- Extreme Ultraviolet Explorer (1991)
- Solar-A/Soft X-ray Telescope/ISAS (1991)
- ASTRO-SPAS/Orbiting and Retrievable Far and Extreme Ultraviolet Spectrometer/ESA (1992)
- Diffuse X-ray Spectrometer (1992)
- KONUS/WIND (1992)
- Solar, Anomalous, and Magnetospheric Particle Explorer (1992)
- Astro-D/Spectroscopic X-ray Observatory/ISAS (1993)
- HST Wide Field and Planetary Camera II (1993)
- Spectrum-X-Gamma/USSR (1993)
 - All-sky monitor
 - Stellar X-ray Polarimeter

During the next several years, some of the other areas of substantial work will include the Advanced X-ray Astrophysical Facility, the X-ray Timing Explorer, the Shuttle test of the Gravity Probe-B instrument, the Submillimeter Wave Astronomical Satellite, additional HST replacement instruments, instruments for the European X-ray Multi-Mirror mission, support for the Russian Radioastron mission and the Japanese VLBI Space Observatory Program, and definition work on the Space Infrared Telescope Facility, the Stratospheric Observatory for Infrared Astronomy, and the Far Ultraviolet Spectroscopic Explorer. Rocket, balloon and airborne activities will continue, as will data analysis from previously flown missions.

■ Big and Small Space Missions

The scientific rationale for a mix of big, moderate and small missions in a balanced program of astrophysics is that there is no optimal mission size to address the incredible variety of scientific phenomena we wish to investigate. In practice, mission size is scaled down to the lowest level required to fulfill the science goals. The U.S. is unique in its capability to conduct science missions that require the largest observatories, but these observatories are balanced by many smaller efforts in our overall program of space astrophysics research. A diversity of missions also helps to develop and maintain our institutional capability to carry out scientific investigations today and in the future. Diversity in mission size supports a variety of implementation strategies, and it helps establish a broad portfolio of missions that can help weather unanticipated and adverse external events.

The terms we use to describe programs may present a misleading perspective to mission size. In one sense, Congress categorizes the entire physics and astronomy budget, which includes astrophysics, space physics, and Shuttle payload mission management, as a single item in the budget.

research and analysis, guest observer and archival research programs round out the currently available astrophysics opportunities. While the diversity in scale of missions is not entirely under our control, we do believe that this combination meets the wide range of needs of the community as they have been expressed to us. Are we correct in this assessment? Some are asking for a greater emphasis on smaller missions in the future. What type of service, product, opportunity or efficiency should be pursued? We welcome further discussion of this issue.

■ Single vs. Multiple Spacecraft in Conducting Missions

There are cost and benefit tradeoffs associated with the strategy in which missions are designed to be conducted by a single complex spacecraft rather than several simpler spacecraft. The most cogent examples of the single complex spacecraft missions are HST and AXAF, and since they represent half of the Great Observatories, questioning the fundamental strategy they embody could provide enormously important insights.

HST and AXAF are not merely large and complex observatories, but they are also serviceable observatories. They reflect a strategy to provide 15 years of on-orbit lifetime through regular replacement of instruments and other hardware. In contrast, the strategy to provide a similar on-orbit lifetime for the Tracking and Data Relay Satellite System (TDRSS) and the Earth Observing System (EOS) involves development of a series of replacement spacecraft on a regular basis. The intent of the servicing strategy is to optimize costs for long-term missions that require an expensive spacecraft (such as HST or AXAF with their large precise optics), and that can also operate in a Shuttle-accessible orbit. TDRSS requires a geosynchronous orbit and EOS requires a polar orbit. Neither TDRSS nor EOS have structural elements analogous to HST or AXAF optics.

The comparison above points to two strategic reasons to use a single, serviceable spacecraft to achieve long life, and to reap the benefits from sharing critical infrastructure. But at what cost? The complexity of a single mission raises costs and increases the required development time, which increases costs further.

While a single spacecraft approach is more vulnerable to an irreparable system single-point failure, servicing provides a programmatic means to regularly repair subsystem failures. The only previous astrophysics missions on a scale even roughly analogous to HST were the Orbiting Astrophysical Observatory (OAO) and the High Energy Astrophysics Observatory (HEAO) series. Two of the four OAO spacecraft failed, one from a launch failure and another on the second day of the mission from a power problem. All three of the HEAO missions were launched and operated successfully for approximately two years. Of course, it was expensive to build serviceability into HST and to purchase the first set of replacement hardware, but the cost to build, launch and operate in the early 1960s would equal \$2.4 billion in 1993 dollars. The life cycle costs of HST, including six Shuttle flights at \$350 million each, is a little over three times that amount. For comparative purposes, the HEAO lifecycle cost through the two years of operations was \$130 million. Is there disagreement that the expected 15-year scientific return of HST will easily surpass that of HEAO and OAO, even in the context of a much advanced technological state-of-the-art?

Multiple missions provide a certain “safety in numbers” for launch vehicle or other flaws, but numbers provide no easy fix to generic failures. If we had built two simple HSTs, both primary mirrors would likely have been distorted by the same faulty null corrector. Since the most cost-effective way to build multiple spacecraft is to have one roughly one to two years ahead of the other, the second Hubble in this example would have been

essentially completed at the time the problem was identified in orbit. In that scenario, repairing the mirror on the ground could easily take as long and be as expensive as the fixes we are implementing today for the real HST.

We must also be careful in using launch rate as a surrogate measure for scientific productivity. Launching five simpler Hubbles with one instrument each might have resulted in a higher flight rate, but the number of launch vehicles by itself does not increase productivity. Implicit in a strategy of flying single multi-instrument spacecraft rather than many single-instrument spacecraft is a lower launch rate. Since a launch itself does not yield science, we have to look for some other measure of scientific productivity that is correlated with flight rate. What are the appropriate ways to measure scientific value: launch rate, number of instruments, weight of instruments, observation time, data returned, refereed publications, new knowledge?

A problem with slow programs is that we do not reap critical information for many years. While this is of most concern in the scientific arena, it also hinders the expansion of our knowledge of how to conduct space science missions. HST is our first experiment with a planned spacecraft servicing strategy, and as a pathfinder it will wind up costing more than programs that can benefit from Hubble's servicing lessons learned. HST was begun in an environment where almost weekly Shuttle flights were anticipated. Throughout the development period, as we have learned about the Shuttle and what it can do, the HST servicing strategy has shifted and adapted. After less than a year on-orbit, we have a small but real database on actual mission events and the programmatic flexibility servicing provides to accommodate them. Servicing will enable key fixes to HST solar arrays and optics, but are these advantages enough to justify the extra expenses? Adopting an empirical approach, let's see the data, let's discuss it, and let's see what we can learn from it.

■ Risk-taking and Risk Avoidance

Can we change the environment to support a level of risk-taking that will increase the long-term efficiency of our space science expenditures? It would be almost impossible to make state-of-the-art spacecraft so reliable that we could be 100 percent certain that there are no technical risks, and even if we could do this, the last bit of reliability would probably cost a lot. A cheaper and more practical approach that NASA has employed is to design our difficult missions to provide additional capability or flexibility that enables us to survive unanticipated problems. If we then can build that mission cheaper, we are getting more science for the dollar.

However, a strategy that includes some risk-taking has one critical implication—sometimes failures happen. The problem with risk-taking strategies is that NASA, Congress, the science community, and the general public are usually unwilling to accept failures. NASA provides a symbol of American technological excellence; thus, NASA successes and failures have a context that exceeds science return for the dollar. NASA receives the budget that it does partly because of this symbolism in the minds of members of Congress and their constituents, but NASA's stature also complicates our simple cost-benefit analysis. Imagine an airline that decided that the strategy to yield the most cost-effective transportation for the dollar would be to reduce safety to the expected fatality level associated with driving a car. Even if this decision could be implemented, what would happen after the first crash?

The Hubble mirror aberration was tragic, but it was also typical of many spacecraft failures in that it was from an utterly unexpected source. But unlike previous missions, the HST program strategy was failure-resilient. On-orbit servicing provides the programmatic flexibility through which even this utterly unexpected technical problem can be addressed and corrected. From a program

perspective, we will have an observatory that will be less than expected for three years, but for the remaining 12 years should live up to its full potential. Nevertheless, the failure in the mirror fabrication has had a profound and fundamental impact on the science community, the public, the Congress and on all of us here.

It is still possible to have a strategy that involves risks, and NASA does designate payloads in one of several categories depending on the level of risk deemed appropriate. Planetary missions are Class A, with requirements to use only the best possible parts and the greatest level of redundancy. Recently, the AXAF spacecraft (with a few exceptions) was deemed to be a Class B mission, in part because of the additional flexibility provided by servicing. The recent Astro flight on the Shuttle Columbia, like many other Shuttle-attached payloads, was developed as a Class C mission.

We agree that failure-resilient strategies should be pursued. So how are we going to change the environment so that there will be support for these strategies? In the abstract, few would disagree. The challenge is to look for ways to enlist and maintain the support for programmatic flexibility and risk-taking even after a failure occurs. If we cannot accept failures of any sort, the cost of missions will inevitably rise.

Budget "Realism" and Strategies to Optimize Financial Risk

How do you estimate the cost of something that has never been done before? NASA starts out with several simultaneous approaches. NASA and contract engineers develop what is known as a "grassroots estimate," in which the working level people estimate their own effort required, and these estimates are aggregated. Although it is critical to have the input from the people who will actually do the job, there are also some problems inherent in a grassroots estimate. Those people do

not yet know how they are going to overcome the unique challenges associated with the missions that yield the state-of-the-art science we are usually pursuing. There is usually some optimism on the part of the engineers, and it is difficult for a grassroots estimate to properly account for the aggregate effect of complex interactions of separate groups working on difficult tasks. So in parallel with a grassroots estimate, a parametric estimate is made using statistical inference based on previous mission experience. By using factors such as subsystem weight or complexity and mission type (such as cryogenically cooled, super lightweight planetary probe, or low Earth orbiting instrument platform), a budget estimate is developed.

The grassroots and parametric estimates are then compared, the information from both estimates is presented, and a single budget estimate is developed for the project. This budget estimate (along with the associated technical and scientific plan for accomplishing the mission) are then reviewed by a team of "non-advocates" who scrutinize the plans and assumptions of the new project, as well as the grassroots and parametric estimates underlying the assumptions.

Given all of this knowledge, some of it contradictory, what budget estimate should be sent to the Office of Management and Budget? If NASA requests a very high budget, we increase the chance that we will look good later, because the chance of overrun is reduced. There is less stress on NASA managers when you have a lot of money for your project. At the same time, there is probably some price at which a program is too expensive to be funded (although it is hard to know what that really is). Another problem is that NASA budgets are a matter of public knowledge, and so all of the contractors know your program's funding. A comfortably large budget can become a tempting target, and so you may find the effort on your mission growing to fit the available budget.

The balance to be struck is to propose a budget which is achievable without being comfortable, and then to monitor and track all changes from this original baseline. Over time, your initial baseline (based on grassroots and parametric estimates) will be modified by contractor bids and negotiated settlements, design reviews, experience with fabrication, assembly and test, and all of the other activities associated with conducting a space science mission. Increases in cost will be of two sorts. Where we have misestimated costs, we request what is known at NASA as a “reprice”—more money to do the same work. Contractor overruns are a subset of repricings, because programs tend to budget for more than the dollar value of the contract to protect against overrun. Only when these reserves are depleted will the program request a repricing to cover a contractor overrun. Sometimes in the development of a mission, we learn that a new activity is required to accomplish the mission, or that spending additional funds to develop a new capability may yield sufficient return to justify the investment. Additional funds to do new work is referred to as an augmentation or as additional scope.

Twice every year NASA formally reviews the budgets of all of our missions to anticipate pending problems, to assess problems that have been identified, and to look for areas where new scope might bring large benefits. It is not entirely a zero-sum competition between these programs for additional funds, but the pressure to make tradeoffs is always there. Within some programs the tradeoffs are internal, and no additional funds are requested. If this is not possible, we must prioritize any annual requests for additional funds. It might look better if there were never any requests for additional funding—hypothetically, NASA would quote a price and come back years later with a spacecraft. But how would we know what price to quote? We could keep eliminating parts of the program to fit within the initial estimate, or we could ask for a high enough budget that we could afford anything. The way the process works now, we make those choices, but we do so incre-

mentally over the life of the missions. Every year our information gets better on what each mission needs and what is possible to accomplish. Most of our effort, and that of our contractors, is dedicated towards learning about the mission and the hardware that can accomplish it—the materials, fabrication and assembly of spacecraft are a minor part of our expenses. If our management at NASA Headquarters is to be based on the science and engineering fundamentals of the missions we are conducting, our management and budgeting must also be a continuous and incremental process. Of course, as our projects and contractors will tell you, this does not mean that we treat budget growth kindly. We must treat an increase in one area as if it were a cut to another, because sometimes that is exactly what we have to do.

We must also counter the tendency towards a focus on the short term, an orientation which is shared by the stock market and indeed with much of our current national character. Congress votes NASA its budget one year at a time. Unlike the private sector, we cannot borrow money from a bank even if it will yield an enormous benefit downstream. The only source of funds in a given year for a new requirement in one mission is to take the money from another. Consequently, everyone’s concern is drawn towards the current year’s budget (which we are spending), and the next year’s budget (which is at OMB or Congress where tradeoffs are being considered). Funding for the next four years beyond that is controlled at OMB, but there is a tendency not to focus on these “outyears.” Unfortunately, our overall scientific productivity depends on choices made throughout the life cycle of our missions. When we develop missions for 15 years of operational life, and 30 years of data analysis to follow, a short-term perspective will not work.

We have a sign on the wall around here: “If everyone keeps saying ‘screw the outyears,’ eventually we will all live in outyears that someone else has screwed.” For long-term missions, this means we must expand our vision even

beyond the five-year horizon typically used for planning. The science return on our missions must be weighed against life cycle costs, and if we only focus on the narrow window of the development period, we are likely to make trade-offs that optimize costs in the short run but are a net loss overall. We have learned the value and insights that come from this perspective in planning the 15-year Hubble operational lifetime, and we plan on implementing the same scope of vision on AXAF. Today's estimates for AXAF operations in the period from 1998 to 2013 must necessarily be soft, but by attempting to track the implications of today's decisions across a mission lifetime, we hope to make decisions that will look as good in retrospect as they do in formulation.

I would also like to raise the controversial premise that part of the "realism" of budget planning depends on an assessment of appropriate financial risk-taking. The problem is analogous to the question of optimizing insurance coverage. If we could reduce the number of programs under development and thereby increase the reserves on each, we could reduce the odds of an overrun on each of our missions. Indeed, the incentive on the individual managers who are responsible for a single project is to be as conservative as possible—to take no financial risks. But at Headquarters our job is different. Our goal is to maximize science return for the dollar. This creates a natural dynamic: the project manager is looking to optimize on behalf of a specific mission, while Headquarters makes tradeoffs between missions and levels of risk. If we are always cutting and delaying every program, the level of risk is too high. If we never have to make a tradeoff between programs, the level of risk we are taking is probably too low. This raises the question: Are we taking the wrong level of financial risk?

Unfortunately, while the downside costs of taking risks are very visible (project cuts or slips), the benefits are not as easily traceable. The quick response to Supernova 1987a was funded by stripping funding flexibility and thus taking

financial risks across the board. In 1988, while AXAF was sent as a new start in the proposed budget to Capitol Hill, we turned down a request from the HST project for \$50 million of additional reserves on development activity. We took what we felt was an appropriate level of risk on HST, independent of concerns for the pending AXAF decision. If we had insisted on having the extra reserves as insurance, it may well have prevented us from starting AXAF that year. We were correct in our assessment that HST development effort did not require the extra reserves to accommodate the problems they were concerned with at that time, but the benefit to astrophysics and space science from this type of risk-taking is usually not as visible as the costs.

One negative aspect of a risk-sharing strategy is that the severe problems in one mission spread across a range of programs. A defining attribute of the Explorer program is that individual projects have reserves that are much lower than usual for other NASA flight programs, and that problems are accommodated within queues. The mission development and launch vehicle problems of the mid-1980s have stretched out the Delta-class Explorer queue to the point where the next mission under development, the X-ray Timing Explorer, was selected 14 years ago. Is this delay acceptable? If not, should we begin to emphasize flight rate more strongly above science performance in Delta-class developments? Should we also wait longer before starting Explorers to increase the likelihood that stable funding will be available?

A particular fear concerning this risk-sharing strategy is that problems in one big project can decimate many other small projects. Put differently, risk-sharing is not appealing to the many if there is one elephant and a lot of mice. However, the existence of several observatories at different stages of their life cycles creates a separate field for elephants, so sensible risk-sharing is now possible. In FY 1991 Congress provided an additional \$30 million for HST, but also levied a similar

reduction that was borne by AXAF. It was deeply disturbing to have to upset the AXAF baseline; we are aware of the inefficiency that funding changes can cause. Nevertheless, we cannot afford to regularly carry reserves to insure against major unanticipated crises such as the HST spherical aberration. If we had carried an extra \$30 million of reserves for Hubble over the past three or four years, these funds would have remained wastefully idle (or worse yet, they would have been spent inefficiently merely because they were available). We had low reserves on Hubble but sufficient funding to support several years of AXAF mirror definition and development work.

Were these the right choices? At the time our assessment of financial risk appeared sound. We also recognize that this is not a clear-cut case because, as events unfolded, we were hit by a more pessimistic scenario than we had planned for. Nevertheless, in retrospect we believe that the overall science productivity has been increased by these choices. In terms of the broader inquiry of operating strategies, the more fundamental questions are whether such financial risk-based strategies are appropriate, and if so, are there further principles or guidance to improve the process?

The problem of financial risk-sharing, as with any risk-based strategy, is that it is difficult to take a broad perspective. There is a cognitive bias in human judgment of risks which has been empirically demonstrated. Negative events resonate far more loudly than positive ones. McCray and Stern ["NASA's Space Science Program: The Vision and the Reality" (1991)] express a concern that the cost of accommodating spherical aberration on HST "may raid small, individual investigator groups of development funds." This fear has a basis in the memory of the so-called "slaughter of the innocents" in 1983 and 1984 when Hubble development problems were solved by cuts primarily from small mission efforts in astrophysics and other science disciplines. These

were truly tragic cuts that caused real damage to individual scientists and teams. Psychologically, the impact of these cuts resonates very deeply. The actual level of reduction was 8 percent of non-Hubble astrophysics in 1983 and 1984, and there has not been a hit on small missions caused by big ones since, but the concern remains because the "slaughter of the innocents" is such a powerful and psychological force in shaping our cognition of risks.

We believe that by establishing one risk pool for large missions and another for small missions, we make it possible to make efficient use of risk-sharing, which yields the maximum amount of science productivity without threatening the "innocents." Inside a single risk pool, a major unanticipated problem such as the HST aberration threatens AXAF, but future AXAF problems may also be weighed against HST funding. It is natural to have a general concern that Hubble will continue to need more and more funding, because that has happened on several occasions. At the same time, the HST Science Institute was specifically created as part of a strategy to counter the institutional tendency of NASA to underinvest in operating missions in favor of new development activity. In general, we want to provide sufficient reserves to our programs so that they can accommodate a nominal range of problems. In our assessment, it is not efficient to provide insurance in the form of reserves to cover very pessimistic scenarios. Since pessimistic scenarios do occur occasionally, we will sometimes be forced to trade off priorities between missions in the same risk pool. We intend to make these tradeoffs in the context of the priorities established (and regularly reiterated) with the science community.

Should we be more risk averse? Holding higher reserves means responding less quickly to opportunities and starting fewer missions. Remember, our risk posture is not the only factor that can influence funding. These have been our choices to date, but we welcome dialogue on this issue.

■ Institutional and Political Constraints/Forces

Institutional constraints sometimes prevent NASA from achieving the best possible science for the dollar. While we do not have the authority to change many of the rules under which we operate, it is worthwhile to discuss the institutional setting of NASA space science because solutions may exist; we may just not see them.

The budget process takes about two years, of which more than half is activity outside NASA by Congress and the Administration. Space science experience has shown that reliable cost estimates are frequently difficult to make until you have already invested about 10 percent of the development cost in definition. However, if we wait until that level of definition and then begin the budget process, we are adding a substantial delay to the program. And once a program is inserted into the budget process, that is no guarantee it will emerge successfully. We have adopted and implemented a strategy embodied in the OSSA Strategic Plan which tries to prioritize and sequence major new starts for a period of several years in order to focus our resources on a few key mission candidates and reduce the complexity of tradeoffs once missions are proposed within the budget process. This strategic planning avoids some types of inefficiency, but inevitable time lags remain in the system.

National political forces sometimes favor highly visible—thereby large—space mission. As public and private individuals pursuing government-supported space science, we are in a bind. The Executive and Congressional process by which the U.S. approves scientific investigations provides the fundamental legitimacy we have to do our jobs. At the same time, institutions and processes can tend towards certain results by virtue of their structure. If in fact there is an institutional predisposition towards large missions in the space mission approval process, then the sources of that structural preference must be

specifically identified and countered. Otherwise, the science administrator who proposes a program of space research that includes fewer large missions and more small missions is likely to lose in the competitive budget arena to others who cater to the existing bureaucratic and political tendencies.

While the current complement of space science missions presents a diversity of large and small science, it is possible that it is not the right mix. The institutional and political process is shaped by the actions and contributions of both public and private space scientists, engineers and managers. The process begins with mission proposals from the scientific community and ends with Congressional approval. If the system has a bias towards bigness, what specific changes can we make or promote to get the system to support the optimal size diversity for space missions? What actions can we take today? How shall we plan to address this issue over time?

■ Linkage to the Manned Space Program

There is a major role in space science for unmanned missions, and we take advantage of the opportunities available. Smaller Explorers have always used expendable launch vehicles, as do even smaller rocket experiments. As experience teaches us more about the capabilities of the Space Shuttle, science mission strategies have been shaped to optimize their mission within the envelope of possibilities. The experience of Challenger and the evolution of the space launch arena since that time have taught us that the Shuttle is generally not an appropriate launcher for larger Explorers. COBE and EUVE had to be redesigned for launch on expendable vehicles, at a significant cost. XTE is planned for Shuttle launch and on-orbit installation on EUVE's Explorer Platform; based on our current investments and options, the servicing strategy embodied in a reusable platform is still the most cost-effective way to pursue the XTE mission. FUSE

is further downstream, and the continuation of this Explorer Platform strategy should be carefully assessed, based on our experience to date and the existence of alternatives. Other Explorers are planned for unmanned vehicles.

Despite the changes in the Shuttle program since the initiation of HST development, that on-orbit observatory is poised to take full advantage of the manned space program in pursuing its 15-year science mission. The Space Infrared Telescope Facility (SIRTF) was originally designed to be a Shuttle attached payload. When Shuttle flight experience indicated the existence of an orbital phenomenon (atomic oxygen glow), the Shuttle Infrared Telescope Facility was changed into a free-flying spacecraft to fulfill the mission needs. Later study has revealed that a 100,000-kilometer orbit optimizes the SIRTF mission return, and so now the current baseline is for launch on a Titan IV unmanned vehicle. There was no institutional resistance to this change, and we intend to continue making launch vehicle choices on the basis of science priorities and cost effectiveness.

Current budgeting policy does not charge differential costs to account for the differences between expendable launch vehicles and the Space Shuttle. Unfortunately, the very fact that the Shuttle is not expendable makes it very hard to calculate the appropriate cost of a single flight. In terms of the expendable fuels used and flight-specific effort, the cost of flying six Shuttle flights in a year instead of five is very small (closer to \$40 million than the \$600 million cost estimate that some members of the science community have used). In economic terms, the variable costs are insignificant when compared to the fixed costs. A greater share of the cost of using a Shuttle flight is due to another economic concept: opportunity cost. If we were to change our minds and launch SIRTF inside a Shuttle, we would have to push some other payload off the manifest. (Of course, some payloads require only partial use of the pay-

load bay and mission timeline.) The value of opportunity cost would depend on the importance of the payload to be replaced. We are also not amortizing the development of the Shuttle in our costs above (nor that of ICBMs upon which our fleet of expendable vehicles is based), because the final economic principle which we are pursuing is that sunk costs should not be considered in making today's choices: we want the most cost-effective way of accomplishing the science missions we are pursuing.

The dialogue should not stop here. We are attempting to pursue better operating strategies to address some of the important problems that have been raised by members of the space astrophysics community. As these efforts progress, we will want to examine their effectiveness, adopt and improve what is successful, and change what is not working. Other key issues remain unanswered, and while some of the institutional problems appear inevitable and unchangeable, we should be wary of complaisance. If we are doing something that can be done better another way, we should try the better way. America's space science program yields benefits to all of us, and it is the duty of those entrusted with conducting this exploration of the universe to grapple with our common challenges and surmount them.

Efforts to understand and to maximize science return have continued. Over the past two and a half years since this was written, the strategy for both XTE and AXAF has been changed from the Shuttle servicing mode, with AXAF split into two smaller, cheaper spacecraft. HST remains on schedule for Shuttle servicing missions in December 1993 and March 1997. Institutional factors now appear to be shifting towards small missions over large ones. Hopefully, this trend will not be simply an exchange of one inappropriate bias for another, but rather an opening of a wider variety of alternatives from which we can pursue the optimum.