Article

International Area Studies Review
1–13

DOI: 10.1177/2233865916636834
ias.sagepub.com

Engineering cooperation: How Americans and Russians manage joint operation of the International Space Station

William T. Bianco
Indiana University, USA

Robert Landis
NASA Headquarters, USA

Abstract

The 1990s agreements that created the International Space Station (ISS) described the effort as a partnership of equals, a joint venture between organizations that remained independent in terms of many procedures, norms, goals, and the assumptions underlying these factors. As a result, successful joint ISS operations required the participants, most notably the American and Russian space programs, to reconcile different procedures, norms, and training regimes, as well as the beliefs that underlie these practices. Drawing on a combination of operational experience, first-hand observation, and interviews, this paper focuses on how the two programs reduced conflict and engendered cooperation. It also uses the ISS experience to consider how future joint efforts can be designed to minimize conflict between international partners.

Keywords

Cooperation, prisoners' dilemma, new economics of organization, International Space Station, NASA

Despite many initial warnings (e.g. Oberg, 2002), the International Space Station (ISS) is now often described as an exemplar of effective international cooperation, confirming that self-interested agents can secure mutually beneficial outcomes despite incentives to defect (Axelrod, 2006). Our study highlights an important feature of this apparent on-orbit success: the participants hold

Corresponding author:
William T Bianco, Indiana University, Woodburn Hall 210, Bloomington, Indiana 47401, USA.
Email: wbianco@indiana.edu

		Player 2	
		C	D
Player 1	C	$b_1 - c_1$ $b_2 - c_2$	$-c_1$ b_2
	D	b_1 $-c_2$	0 0

Figure 1. Two-player prisoners' dilemma.

very different ideas of what constitutes cooperation. These disagreements have led to an outcome we label as truncated cooperation, where working together involves maintaining a separation between the two programs, including hardware, procedures, and day-to-day operations, acting in concert only when it is absolutely necessary.[1] While some NASA personnel disparage this practice, the advantages of truncated cooperation constitute a crucial and unappreciated lesson-learned from the ISS program.

Our analysis begins with a two-player model of cooperation, showing how disagreements between the players can lead them to limit the frequency and extent of interactions. Successful cooperation under these conditions will feature high levels of integration and joint effort in some areas but no cooperation in others, reflecting differing levels of consensus and disagreement. The paper than turns to the details of the ISS experience, showing that labeling the ISS as a successful partnership overlooks the extent to which success involved independence rather than interaction – working separately rather than working together. Our point is not that truncated cooperation constitutes a failure. Rather, we argue that a limited form of cooperation is exactly what we should expect from partnerships between established organizations. The final section of the paper draws on the ISS experience to offer some insights into how future joint space exploration efforts can be designed to minimize conflict between international partners.

Prisoners' dilemmas, cooperation, and disagreement

We characterize ISS operations as a kind of iterated prisoners' dilemma or PD (Axelrod, 2006; Miller, 2004). This game is the canonical representation of situations where players can make themselves better off by working together, but where each faces an incentive to deviate from any agreement to work together. This simple game captures crucial aspects of daily ISS operations. The ISS consists of two sections: the Russian Segment (RS) and the US Orbital Segment (USOS).[2] Each segment was designed and built separately, and is commanded by independent groups of flight controllers (individuals with day-to-day responsibilities for station operations), working in NASA's Mission Control in Houston and Russia's TsUP in Moscow. In our analysis, the two teams of flight controllers constitute the players, operating the ISS as a partnership with no central authority exercising overall control. However, there are many situations where the two sides can cooperate to achieve mutually beneficial outcomes. Notably, these outcomes cannot be achieved unilaterally – the dilemma is that cooperation requires both sides to take actions they would not otherwise be inclined to do.

Drawing on the notation developed by Bianco and Bates (1990), Figure 1 describes a single iteration of a simple two-player (1, 2) PD.

In each iteration, each player has two choices: cooperate or defect. By cooperating, player i incurs the cost c_i, and produces a benefit b_j for opponent j. Defection produces no benefits and avoids costs for i. By definition, $c_i > 0$ (cooperation is costly) but $b_i > c_i$ (players prefer mutual

		Player 2		
		C1	C2	D
Player 1	C1	$(p)b_1 - c_1$ $(q)b_2 - c_2$	$- c_1$ $- c_2$	$- c_1$ $(q)b_2$
	C2	$- c_1$ $- c_2$	$(1 - p)b_1 - c_1$ $(1 - q)b_2 - c_2$	$- c_1$ $(1 - q)b_2$
	D	$(p)b_1$ $- c_2$	$(1 - p)b_1$ $- c_2$	0 0

Figure 2. Prisoners' dilemma with uncertainty.

cooperation to mutual defection). Thus, the problem in a single-iteration PD is that each player has a dominant strategy (defect), but if both players use their dominant strategy, their payoff (0) is less than what they would have received from mutual cooperation ($b_i - c_i$).

Many analyses have shown that mutual cooperation can be sustained in PDs if interactions are iterated, meaning players make a series of strategy choices over time, and if the likelihood of future interactions (w) is not too low relative to ($b_i - c_i$). These outcomes are sustained by the use of trigger strategies, where player *i*'s willingness to cooperate is contingent on player *j*'s past behavior – in the most extreme version, *i* begins the game cooperating, but permanently switches to defection following a defection by *j*. Trigger strategies change the incentives for cooperation and defection: player *j* must weigh the short-term benefits from defection against the reduction in their payoff from *i*'s defection in subsequent iterations. When the likelihood of future iteration is high enough, trigger strategies are in subgame-perfect equilibrium – if *i* believes *j* is using a trigger strategy, their best-response is to use a trigger strategy as well (Axelrod, 2006). Moreover, repeated cooperative interactions among a set of actors can drive an expansion of the scope of cooperation, as the players using trigger strategies learn from their interactions about new opportunities to secure mutually beneficial outcomes, and noncooperators see the value of conditional cooperation (Miller, 2004). Moreover, conditional cooperation provides a way to secure cooperative outcomes through self-governance rather than hierarchy (Ostrom, 2005).

Subsequent work shows that incentives for cooperation are also shaped by the larger context. Ostrom (2005) shows how preexisting cooperation in one area can be used as part of a trigger strategy to secure cooperation in new regimes. A second solution analyzed by Bianco and Bates (1990) and Brehm and Gates (1999) involves adding hierarchy to the game in the form of a leader who controls the distribution of benefits from cooperation, and who followers believe will reward cooperation and punish defection.[3] Both of these mechanisms are particularly important for the initiation of cooperative behavior.

Cooperation given disagreement

One critical assumption imbedded in Figure 1 and in the broader literature on cooperation is that the actions needed to produce mutually beneficial outcomes are common knowledge.

Figure 2 relaxes this condition, showing a game where the players have two potentially cooperative strategies – C1 and C2 – and beliefs about which strategy will produce a mutually beneficial outcome. Specifically, player 1 believes that if player 2 plays C1, the probability of player 1 receiving benefits is (p), while if player 2 plays C2 the probability that 1 receives benefits is ($1 - p$). Player 2's corresponding probabilities are (q) and ($1 - q$), respectively.[4]

Adding disagreement creates a new problem: suppose (p) is high while (q) is low – player 1 believes that the players need to choose the C1 strategy in order to secure benefits, while player 2 believes that benefits are more likely to be secured by both choosing C2. Under these conditions, the players agree that cooperation (of some kind) is needed to produce a mutually beneficial outcome, but they disagree on what actions are needed to secure this outcome. Such problems are not considered in the canonical analysis of self-interested cooperation. The implicit assumption is that the players have already agreed on which actions constitute cooperation (or a central authority or leader has told them what to do), and that any remaining uncertainties about consequences are folded into their expected payoffs. As long as $b_i > c_i$ and the likelihood of future interactions is high enough, mutually beneficial outcomes can be secured.

However, what if the players disagree on the specifics of cooperation and there is no leader to make these choices for them? This situation is, in effect, what NASA and the Russians faced when making decisions about how to operate the ISS: while they agreed about how to cooperate in some areas, most of their interactions were beset with strong disagreements. A full equilibrium analysis of the game in Figure 2 is beyond the scope of this paper, but preliminary analysis shows that the impact of disagreements can be divided into three broad cases as follows:

- when players agree on which action is most likely to generate benefits (formally, when p and q are either both less than .5 or both greater than .5), disagreements do not affect their ability to sustain cooperation, as long as the likelihood of future interactions is high enough;
- when players disagree about which action is most likely to generate benefits ($p > 5$ and $q <$ .5, or vice versa), but the disagreement is small enough, the players may disagree on which actions yield the best outcome, but they will agree that any kind of cooperation is preferable to no cooperation at all;
- when disagreement is sufficiently large, cooperation is a dominated outcome: regardless of what is proposed (both playing C1 or both playing C2), one player will have costs that exceed expected benefits.

In words, at one extreme, when beliefs differ but the players nevertheless agree on which actions are most likely to produce benefits, a mutually beneficial cooperative outcome still exists. At the other extreme, where disagreements are large, there is no mutually beneficial outcome, and defection is the expected result. In the intermediate case, cooperation is still sustainable if the players can agree on its form.

Implications for real-world interactions

This analysis has three important implications for real-world cooperation. Firstly, the fact that participants are able to work together in some areas does not imply that they are fully cooperating – there may be other areas where mutually beneficial outcomes cannot be reached because of disagreements. Put another way, disagreements may lead to what we label as truncated cooperation, where participants abandon joint action in some areas while cooperating in others.

These results also suggest a new role for leaders in an iterated PD: rather than allocating rewards, their job may be to identify situations where disagreements do not preclude cooperation – and to move participants away from other areas where disagreements are large. Such advisories may take the form of broad directives that lead players to cooperate in a narrow range of activities, or advisories about specific interactions at some point in time that are exceptions to more general rules. In cases of disagreement, a leader's influence over cooperation may be derived from their ability to provide information rather than their control over benefits.

Finally, the possibility that disagreements may lead to truncated cooperation implies that repeated interactions between PD participants may not lead to the expansion of cooperation, as suggested in Axelrod's work. Limiting the scope of cooperation implies bounds on what participants learn about each other and about additional opportunities to work together.

Returning to the situation posed at the beginning of this section: what kind of cooperation should we expect from two organizations forced to work together in the absence of central authority, given frequent disagreements about joint action? In a word, we should expect truncated cooperation – in some areas, the organizations will work together effectively, taking actions that satisfy both sides, but in others they will operate in parallel, maintaining their independence and minimizing their dependence on each other. Moreover, successful truncated cooperation requires mechanisms that provide information to each side and a means to identify and implement cooperation in exceptional circumstances where the partners happen to be in agreement.

An analogous situation: joint ventures

These insights about truncated cooperation in iterated PDs are consistent with the literature on corporate joint ventures, which share the characteristic of two organizations working together in the absence of central authority (Burke, 2010; Brannen and Salk, 2000; Kotter, 1996; Schein, 2010). While these studies focus on different organizations and levels of joint effort, their conclusions about how to build a successful joint venture are remarkably similar to the arguments made here.

Firstly, success is often associated with an incomplete merger, where the partners maintain separate organizations and do not attempt to build consensus on all relevant issues. The key is to focus attention on the questions that must be resolved for the partnership to succeed – and ignore other matters entirely: even when one side is sure that their approach is the best. Secondly, while some cooperation across the organizational divide is essential for success, high-level mandates about how this outcome must be achieved are often counter-productive. The best strategy often involves some degree of self-governance, where mid-level managers from each organization determine whether and how cooperation will occur within their areas of responsibility. Thirdly, because joint ventures reflect an incomplete merger, success often requires building an interface group composed of people who have deep knowledge of the practices, procedures, and goals of both organizations. The group's job is to keep everyone informed about day-to-day events, as well as unanticipated opportunities for working together.

These findings have the same essential message as the analysis of the game in Figure 2: successful joint ventures may not require a wholesale merger of organizations or widespread cooperation across the organizational divide. Rather, in most situations the optimal partnership will take the form of truncated cooperation: we will see varying levels of interaction and accommodation across different activities, depending on perceived benefits from cooperation and the nature of underlying disagreements.

Cooperation in International Space Station operations

Theoretically, the 1998 Memorandum of Understanding that established the ISS program places NASA in overall control ("NASA will have the responsibility for the overall management and coordination, through the management mechanisms established in this Article, of the operation of the Space Station") and mandates a high level of integration in on-orbit operations ("The Space Station crew will operate as one integrated team with one Commander. Consistent with the principle of an integrated crew, the entire crew will operate under a single timeline for performance of all operations and utilization activities.")[5] However, in practice the ISS is operated as a partnership of equals between the various Russian space agencies and corporations on the one hand and NASA and the other international partners on the other, including the European Space Agency (ESA) and the Japanese Space Exploration Agency (JAXA).

While this organizational restructure reflects the segmented ISS design, the lack of an overarching central authority is notable. While it is common for un-crewed spacecraft to have participation from multiple countries, typically these arrangements involve the construction of a specific instrument (e.g. the ESA provided a camera for the Hubble Space Telescope and a lander for NASA's Cassini mission to Saturn), with overall program management and day-to-day operations being firmly held by a single agency. Several shuttle Spacelab flights had science operations managed by the ESA or JAXA, although NASA was in overall control of the mission. During the earlier Shuttle–Mir program, the rule was that NASA directed all activities inside the US shuttle orbiters, while Russians were in charge of the Mir station.[6]

The impact of the ISS partnership on station operations was magnified by differences in the engineering judgments, experiences, and beliefs that the two sides brought to station design and operations. Discussions with NASA and Russian flight controllers and managers reveal that disagreements persist to the current day, from prosaic items such as planning crew meals to determining the appropriate response to emergency situations.[7]

One example involved a disagreement about what the ISS crew should do if a threat from orbital debris was identified too late to maneuver the ISS out of harm's way (Dempsey, 2009; Matushin, 2010). The two sides agreed that the crew should evacuate into their Soyuz escape crafts until danger had passed. However, the two sides disagreed on critical aspect of the plan: NASA wanted the Soyuz hatches closed and latched, while the Russians argued for leaving hatches ajar. This disagreement reflects several factors. Russian concerns were the desire to preserve limited Soyuz consumables (a closed hatch required Soyuz life support systems to be activated), as well as avoiding two other contingencies: that a latched hatch could jam, and that the debris might hit the Soyuz, which would almost instantly depressurize after an impact if hatches were closed. NASA's position reflected two worries: that unlatched hatches could be damaged by an impact-induced pressure wave, preventing them from being closed for evacuation, and that an impact would cause module walls to tear wide open, decompressing the station quickly enough that the crew would be incapacitated before they could close the Soyuz hatch.

This case neatly reflects the nature of the conflicts that arise in joint ISS operations. Each side has a rationale for their position that is based on their operational experience, but also on untested (and untestable) assumptions. For example, while there has been some research on how ISS modules will fare in a debris impact, these tests have only involved small test articles. Similarly, Russian concerns are based on assessments of the chances of an extended delay in landing – an event that has never occurred.

There are many other examples. The USOS orientation in space must be carefully managed to keep hardware from overheating, freezing, or breaking apart; the RS was built to operate without such constraints. NASA uses Velcro to secure cargo, while Russians prefer to bolt things down. NASA uses iodine to purify water; Russians add ionic silver (as a result, mixing the two water supplies will clog plumbing). Russian crews are given an alcohol ration; the USOS is dry. NASA flight controllers work eight-hour shifts five days a week; Russians work 24 hours on/72 hours off. In all, the RS and the USOS have virtually no hardware or systems in common – the one exception is the toilet, where NASA opted for Russian hardware after several failed attempts to produce a workable US design.[8]

In virtually all of these cases, both sides have detailed rationales for their practice, and are confident that their approach is best.[9] There are multiple costs involved with compromising these differences, including the time and effort (and international travel) required for negotiations, translating between Russian and English, and the need to convince others inside each organization about the virtues of a compromise solution – costs that are substantial even when substantive differences are small.[10]

The result of these differences and the high cost of compromise is a pattern of truncated cooperation in ISS operations. The two segments of the ISS are operated as though they were completely independent entities, with joint efforts limited to areas where there is a consensus on how to work together and clear, substantial benefits from doing so. This is not to say that interactions never happen: during the ISS design phase, the agencies set up working groups to manage the assembly process and determine the joint response to major, time-critical problems that affect the entire station. These efforts continue to the present day. The striking feature of ISS operations, however, is the degree of separation between the two programs. The two sides may know where conflicts exist, but this shared information does not lead to compromise, consensus, and cooperation, except in the broadest sense that the hardware they manage is connected together.

There are two exceptions to this practice. The first involves NASA and Russian controllers who are responsible for the station's attitude and trajectory. Early in the ISS design phase, it was decided that the requisite hardware would be distributed across both segments, with thrusters located on the RS and control movement gyros on the USOS. As a result, Russian and American flight controllers with responsibility for attitude control and trajectory management must share information and interact frequently to maintain the station's attitude, increase the station's altitude, prepare for docking of crew or cargo ships, or other actions. In effect, the station's design forced this subset of flight controllers to arrive at common beliefs about how to manage this aspect of station operations and to cooperate going forward.

The second exception involves controllers who manage life support systems. Both sides have hardware that recycles so-called gray water (waste water and moisture extracted from cabin atmosphere) into oxygen and drinkable water, but these systems have different design philosophies and no parts in common. On occasion, however, one side offers gray water to their counterparts for recycling in their systems. Russians also routinely deliver crew urine to the USOS for recycling because their hardware does not have this capability. Finally, because both systems extract water from the cabin atmosphere, some coordination is needed to keep both systems operating properly. As in the case of attitude control, we argue that cooperation between life-support controllers was the product of constraints that tie the two side together: since maintaining a breathable atmosphere is ultimately a joint responsibility, the two sides had strong incentives to compromise their differences and devise common plans, otherwise neither one could fulfill their responsibilities.

Evidence of segmented operations

Perhaps the most compelling evidence of truncated operations on the ISS comes from first-hand accounts. Interviews of US astronauts with ISS flight experience broke down into two categories. Astronauts who had flown during the early days of ISS construction, when the RS provided living quarters and life-support for the entire station, were very familiar with Russian systems. The same is true for controllers who served during this time. In contrast, astronauts and controllers whose work began after about 2008, when the USOS became the primary location for day-to-day activities of the US crews, had much less knowledge of the RS. Interviews also revealed that other than using astronauts and cosmonauts as subjects in medical experiments, there is virtually no joint research conducted on the ISS. Each side handles maintenance on the systems it controls using its own crew. The standard practice is for the USOS and RS crews to have separate daily planning conferences with their respective control centers. Absent of special events such as spacewalks, a major systems malfunction, or docking of a spacecraft containing cargo or new crew, interactions between RS and USOS crews are usually limited to a joint meal at the end of the working day.

NASA flight controllers are also well aware of the segmented nature of ISS operations. Consider an email sent in 2002 by John Curry, NASA's Lead Flight Director, during the initial phase of ISS crewed operations (quoted by Oberg, 2002, caps in original):

> The Russians ARE DEFINITELY trying to force us into segmented operations, where the Russian crew members obey Moscow [Mission Control] and the Americans are controlled from Houston. This is a direct violation of the fundamental principle of the "unified crew," agreed to by all the international partners. They REFUSE to acknowledge any type of U.S. leadership in the planning world on items such as sleep cycles, priorities, etc., and they proactively maintain a written record of instances when [Houston] makes a planning mistake. This record is then periodically shipped to Houston as justification for why they should still lead planning.

By all accounts, in the years since Curry's email, both sides have accepted segmented operations and truncated cooperation as the defining characteristic of ISS operations.

The Houston Support Group. As discussed earlier, one implication of truncated versus full cooperation is a reduction in opportunities for information transfer and learning. That is, repeated interactions create opportunities for learning and communication, revealing opportunities for cooperation that may have been overlooked. When cooperation is limited, disagreements are likely to persist, and participants may have no way to initiate requests for an exception or changes – they may not even know whom to ask. Thus, successful joint ventures such as the ISS generally create an institution that operates as an interface between the partner organizations, serving as an information provider and communications conduit.

In the ISS program, the interface between NASA and the Russians is the Houston Support Group (HSG), located inside Russia's Mission Control Center in Moscow.[11] The HSG is staffed by five NASA flight controllers who serve three-month tours in Moscow, supported by a number of technicians and interpreters and, in recent years, some Russian systems specialists who were hired away from Roscosmos by NASA. The job of the HSG is to act as an interface between the two agencies. HSG staff monitors the air-to-ground loop on the RS, as well as telemetry from Russian station modules, and gives real-time advice to colleagues in Houston on scheduled and unscheduled events, malfunctions, or any other matters of interest. HSG personnel also act as a communications relay, transmitting requests from NASA to the appropriate Russian controller, system specialist, or manager, and sending the response back to Houston.

If interactions between NASA and the Russians were anything close to full cooperation, there would be no need for a HSG. Rather, repeated interactions across the entire range of ISS operations would give flight controllers a detailed understanding of their counterpart's operations, including motivations, constraints, and information about whom to contact in the event of unanticipated events. In the case of the ISS, even after 15 years of crewed operations, much of this information is held by the HSG. In fact, while NASA and Russian flight controllers are able to communicate directly with each other, in practice most messages (other than the two exceptions mentioned earlier) are sent through the HSG. Thus, while ISS operations involve truncated cooperation, the HSG plays a critical role in sustaining this equilibrium.

HSG interactions data. The quantitative evidence for our claims about the ISS and truncated cooperation come from daily reports prepared by HSG staff during 2008–2012, covering 1340 days of

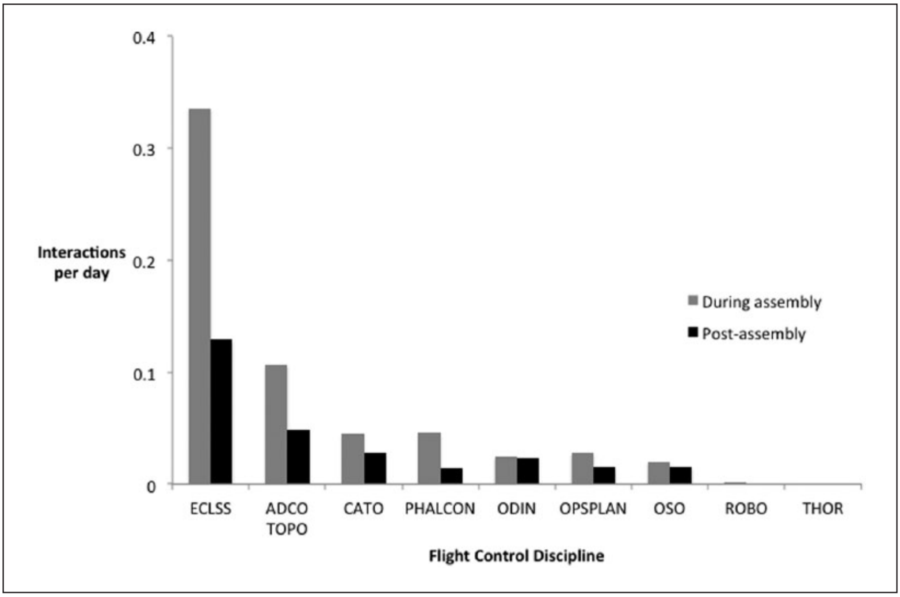

Figure 3. NASA–Russian interactions.

station operations. Among other things, these reports describe interactions between NASA and Russian flight controllers – instances where one side asked for something from their counterparts in the other program. Figure 3 shows interactions per day for different groups of flight controllers, separated into two time periods: ISS assembly (March 2008–March 2011) and ISS assembly complete (April 2011–December 2012).

The abbreviations along the *x*-axis in the figure are labels for NASA flight controllers with different responsibilities for day-to-day ISS operations:

ECLSS: environmental control and life-support
ADCO/TOPO: attitude and trajectory control
CATO: communications and tracking
PHALCON: power, heating, and lighting
ODIN: onboard data networks
OPSPLAN: operations planning
OSO: operations support (stowage)
ROBO: robotics
THOR: thermal control

The figure confirms two elements of our analysis. Firstly, even for the group with the highest probability of interaction, ECLSS, interactions post-assembly are relatively infrequent, only .15 per day. Secondly, for most disciplines, the likelihood of interactions are below .05 per day – and for two, ROBO and THOR, there are essentially no interactions between NASA and Russian flight controllers. The variation between flight control disciplines also makes sense, with ECLSS and ADCO/TOPO having the most interactions for reasons discussed earlier, and ROBO and THOR having the lowest – robotics systems are self-contained on the USOS and RS, and thermal control is not a concern for the RS. The intermediate cases are also as expected – for example, Russian controllers often use US communications hardware, which requires facilitation by CATO controllers.

These findings confirm that interactions between NASA and Russian controllers are best characterized as truncated cooperation. For most controllers, days and weeks go by without any

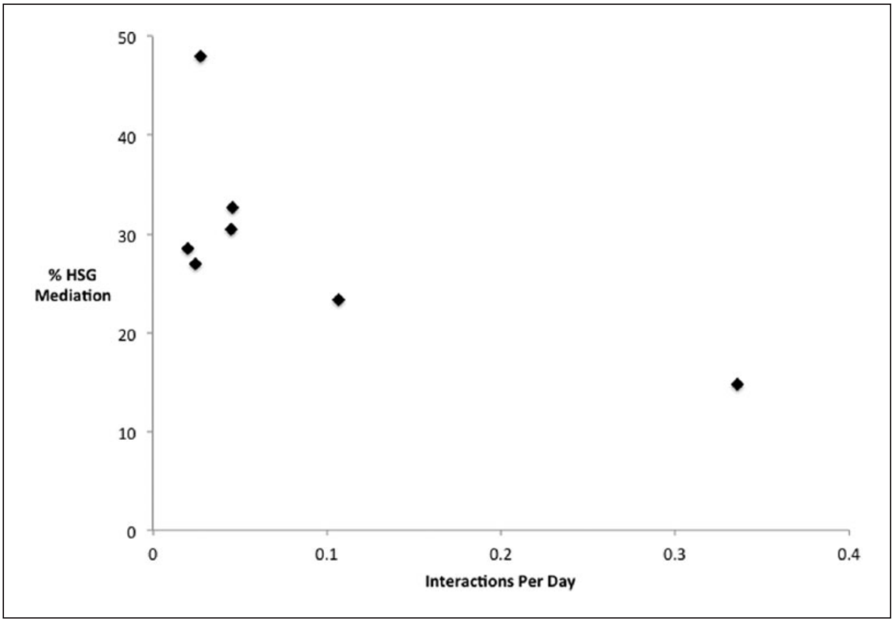

Figure 4. Houston Support Group (HSG) interventions.

communication with their counterparts. Even for ECLSS and ADCO/TOPO, cooperation with the Russians is a relatively uncommon event. Moreover, while interactions were more frequent during ISS assembly, as hardware was being delivered, installed, and de-bugged, they are much less likely during the current period of steady-state operations.

These findings are reinforced by Figure 4, which shows the relationship between the frequency of interaction between controllers and the mode of interaction – whether the interaction involved NASA and Russian controllers communicating directly with each other, or whether interactions went through the HSG. As noted earlier in the discussion of full versus truncated cooperation, frequent interactions should inform each side about their counterparts' motivations and constraints, strengthening communication links and engendering trust. This argument implies there should be a negative relationship between the total number of interactions for a particular flight control discipline and the percentage of these interactions that were mediated by the HSG, a prediction confirmed by Figure 4.

The scatterplot shows that NASA controllers that interact rarely with the Russians (such as ROBO or OSO) tend to do so through the HSG, while controllers whose interactions are more frequent (but still relatively uncommon) are more likely to initiate contact and build consensus without assistance from the HSG – in the latter cases, the HSG reports record that an interaction has taken place, but also note that the HSG did not take part.

Discussion

By all accounts, the ISS program is proceeding well: research goals are being met, conflicts are relatively rare, and unplanned events (such as the failures in 2015 of Russian and US cargo launches) have been accommodated without disrupting station operations. Given these reports, it is no surprise that both sides have characterized the partnership as successful, and argued that the ISS experience can serve as a template for future international space exploration efforts.

This inference is only partially true. The pattern of infrequent conflict does not reflect broad cooperation. Rather, it is the product of segmented operations, where the partners work together only when absolutely necessary. In a narrow range of areas, NASA and Roscosmos personnel interact routinely and effectively, developing and executing plans to achieve common goals. But everywhere else, interactions across the organizational divide are relatively rare – and, when they occur at all, they are likely to involve HSG staff as intermediaries rather than direct communication, even after over 15 years of joint operations. These findings are no surprise. Given the non-hierarchic nature of the ISS partnership, and the significant disagreements between the parties, it makes sense that success would involve truncated cooperation and the development of interface organizations such as the HSG. While we might hope that the two sides would work to a common understanding of how to build and operate a space station, thereby reducing costs and enhancing capabilities, the reality is that achieving consensus on even the smallest issues is extremely time-consuming and often impossible. Under these conditions, truncated cooperation is not an inferior outcome – it is the only feasible way to manage a joint venture.

In this light, the segmented nature of the ISS itself has significant advantages over a more integrated structure. Typically the ISS design is justified on engineering grounds such as differential redundancy, or as a reflection of the political reality that neither side could afford to build its own station. However, given the sharply different operating philosophies of the American and Russian programs, segmented operations reduce conflict by allowing each side to operate systems as they think best without having to coordinate or compromise with their counterparts. The segmented structure also allows each side to retain their staff and mission control facilities, rather than abandoning some of these capabilities as a fully integrated station would require. Rather than being a drawback, the segmented design of the ISS helped to make the partnership a success.

These findings have important implications for the structure of future international space exploration efforts. Preliminary plans for missions to the Moon, asteroids, and Mars assume a significant amount of international cooperation based on the need for cost-sharing, and as a way of ensuring long-term political support for exploration efforts (International Space Exploration Coordination Group, 2013). However, the nature of such cooperation has not been defined. While it is probably true that the partners must arrive at a common understanding of the broad outlines of an exploration effort, the ISS experience suggests allowing them to agree to disagree on some of the specifics. That is, rather than mandating full cooperation across all activities, the venture should be structured in terms of truncated cooperation. One possible approach is to divide responsibilities by function, allowing each agency to build hardware as they see fit and exercise complete control over this hardware during operations. For example, informal discussions of a lunar exploration program have centered on the idea that NASA would be responsible for the craft that takes joint crews to lunar orbit, Russia would build and operate the spacecraft that actually lands crews on the Moon, and the ESA would provide scientific hardware for surface operations. If disagreements make this approach impractical (i.e. agencies are unwilling to defer to the judgments made by others), the ISS experience suggests a strategy of building and operating hardware in parallel. This approach may seem inefficient in that it involves some duplication of effort and additional mass (including interfaces between these modules). These drawbacks must be weighed against the saving in time and effort required for compromise and complete cooperation.

Finally, while the analysis here has focused on the ISS, we believe that this case is an exemplar of a broad class of situations where self-interested actors must cooperate in the absence of hierarchy. Our analysis shows that insofar as participants in these situations disagree about the specifics of cooperation, equilibrium outcomes may feature only limited cooperation – even when interactions are frequent and occur over long time periods. Expectations about how the participants behave in these situations – whether they can “solve the dilemma” – need to consider the extent and magnitude of disagreements along with the costs, benefits, and frequency of cooperation.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Notes

1. Bianco was Fulbright Senior Scholar at the National Research University/Higher School of Economics, Moscow in 2011–2012, where he shadowed NASA personnel working at Mission Control Center (MCC-M) in the HSG, and interviewed Russian and American controllers and program officials in Moscow and Baikonur. Landis is a former ISS flight controller and a former Operations Lead at the HSG. Matushin directs the SRP (Flight Director) Office at MCC-M. Unless otherwise specified, interviews cited in the paper were conducted by Bianco or by Bianco and Landis in Houston, Moscow, Washington, Baikonur, and over the phone during 2012–2014. However, none of our analysis or conclusions reflect the official or unofficial position of any office or individual inside NASA or the Russian space agencies.
2. The USOS also includes components and modules constructed and operated by the ESA, the Canadian Space Agency (CSA), and the JAXA. In practice, NASA controls the USOS, and the Russians treat the other international partners as subcontractors to NASA. Hence, our focus here is on bilateral interactions between NASA and the various Russian space agencies and corporations.
3. The key to this result is the assumption that a hierarchy either exists or can be imposed that produces a leader who is well-informed about follower actions and has an appropriate reputation for rewarding cooperation and punishing defection, and controls the distribution of benefits. These expectations make sense in the context of an established firm or organization, where there is a known structure of leaders and followers, and the scope of cooperation is defined in advance.
4. Note that if $p = q = 1$, the C2 strategy can be eliminated as dominated, and the game is identical to Figure 1.
5. The Memorandum of Understanding is available at ftp://ftp.hq.nasa.gov/pub/pao/reports/1998/nasa_russian.html (accessed 27 January 2015).
6. Hierarchies are also common in other high-tech programs: management of the Large Hadron Collider (LHC) and the ITER International Fusion Energy Organization are in the hands of multi-nation councils that use weighted voting schemes to make authoritative decisions about the programs, with some exceptions for specific hardware (e.g. large detectors that search for particles created at the LHC) that are managed as separate programs controlled by a single nation.
7. This section of the paper draws from, but does not cite, some 60 interviews with NASA and Russian personnel. The reason is that managers in NASA's Mission Operations Directorate and Flight Crew Operations Directorate refused to approve the use of quotes from any of their subordinates unless NASA had a veto on any publications emerging from the project.
8. In practice, this decision created a new venue for conflict, as NASA occasionally asks Russian specialists for help, only to find that their counterparts are much less inclined to investigate suspicious noises or intermittent malfunctions.
9. Consider the disagreements about flight controller shifts. The Russian practice is backed by in-house studies that suggest reducing handovers increases controller performance, while the NASA practice reflects studies that indicate long shifts degrade controller performance. The practices are also shaped by experience: particularly during the Mir program, Russian controllers were only in contact with their station for short periods of time, giving them long periods of downtime for rest, while NASA controllers had essentially full-time contact with astronauts on the Shuttle. The Russian practice was also shaped by societal realities: few controllers had cars and public transportation did not operate continuously, making evening and nighttime shift changes problematic.
10. Few American controllers are proficient in Russian. While NASA has a rigorous certification program for its interpreters, there is widespread agreement among the people we interviewed that even highly trained interpreters often have difficulty with the details of technical material as well as slang and colloquialisms.
11. At present, the HSG is active for two shifts per day – the Moscow day shift and Houston day shift. During the remaining eight hours, if coordination is required, the shift flight director in Houston uses an

interpreter to communicate with his or her Russian counterpart, with a RIO (Russian Interface Officer, a NASA controller who speaks Russian) added in case of predictable real-time operations such as a Soyuz docking or Russian spacewalk. The Russians have a similar group, the Moscow Support Group, located in Mission Control Center – Houston. One important function of Moscow Support Group (MSG) controllers in Houston is to participate in NASA simulations, where they control operations of the simulated RS. One Russian manager argued that this practice provided crucial operational knowledge to controllers on both sides – specifically, when Russian staff at the MSG return to Moscow, they have a much better understanding about how NASA controllers will react to different situations, both the specific responses and the rationale behind these actions.

References

Axelrod R (2006) *The Evolution of Cooperation*. 2nd ed. New York, NY: Basic Books.

Bianco WT and Bates RH (1990) Cooperation by design: Leadership, structure, and collective dilemmas. *The American Political Science Review* 84: 133–147.

Brannen MY and Salk JE (2000) Partnering across borders: Negotiating organizational culture in a German-Japanese joint venture. *Human Relations* 53: 451–487.

Brehm J and Gates S (1999) *Working, Shirking, and Sabotage: Bureaucratic Response to a Democratic Public*. Ann Arbor, MI: University of Michigan Press.

Burke WW (2010) *Organization Change: Theory and Practice*. New York, NY: SAGE.

Dempsey DP (2009) NASA Proposal on Hatch Configuration for Soyuz Safe Haven. Mission Operations Directorate Presentation, 23 April 2009. Available at: NASASpaceflight.com (accessed 11 March 2016).

Kotter JP (1996) *Leading Change*. Cambridge, MA: Harvard Business Review Press.

Matushin M (2010) Interaction between mission control centers during International Space Station flight operation: The example of ISS debris avoidance operations. In: *the 2010 American Institute of Aeronautics and Astronautics Annual meetings*.

Miller G (2004) The political evolution of principal-agent models. *Annual Review of Political Science* 8: 203–225.

Oberg J (2002) *Star-Crossed Orbits: Inside the U.S.–Russian Space Alliance*. New York: McGraw-Hill.

Ostrom E (2005) *Understanding Institutional Diversity*. Princeton, NJ: Princeton University Press.

Schein E (2010) *Organizational Culture and Leadership*. Chicago, IL: Jossey-Bass.