

Inherent Moral Hazards in Acquisition: Improving Contractor Cooperation in Government As The Integrator (GATI) Programs

*William E. Novak
Software Solutions Division
Software Engineering Institute
Carnegie Mellon University
4500 Fifth Avenue
Pittsburgh, PA 15213
wen@sei.cmu.edu
412-268-5519*

*Julie B. Cohen
Software Solutions Division
Software Engineering Institute
Carnegie Mellon University
4500 Fifth Avenue
Pittsburgh, PA 15213
jcohen@sei.cmu.edu
412-268-3950*

*Andrew P. Moore
CERT Division
Software Engineering Institute
Carnegie Mellon University
4500 Fifth Avenue
Pittsburgh, PA 15213
apm@sei.cmu.edu
412-268-5465*

*William A. Casey
CERT Division
Software Engineering Institute
Carnegie Mellon University
4500 Fifth Avenue
Pittsburgh, PA 15213
wcasey@sei.cmu.edu
412-268-3002*

*Bud Mishra
Courant Institute of Mathematical Sciences
New York University
251 Mercer Street
New York, N.Y. 10012
mishra@nyu.edu
212-998-1212*

Inherent Moral Hazards in Acquisition: Improving Contractor Cooperation in Government As The Integrator (GATI) Programs

Dr. William A. Casey

Dr. Casey has worked in the areas of threat analysis, code analysis, natural language processing, genomics, bio-informatics, and applied mathematics in academia, industry, and government. He received his PhD in applied mathematics from the Courant Institute at New York University. He also holds an MS in mathematics from Southern Illinois University Carbondale, a master's equivalency in computer science from the Courant Institute at New York University, and an MA in mathematics from the University of Missouri Columbia. Dr. Casey is a member of the American Mathematical Society (AMS) and Society for Industrial and Applied Mathematics (SIAM).

Julie B. Cohen

Julie B. Cohen is a member of the Client Technical Solutions Directorate at the Software Engineering Institute, where she leads diverse work with customers in satellite programs, medical scheduling, and sustainment. Prior to the SEI, Julie was a Program Manager at both Brashear and Marconi. In the Air Force, Julie worked in the F-16 program office, the Flight Training program office, the Air Force Operational Test and Evaluation Center, and the Air Force Research Laboratory. She received her B.S.E.E from Carnegie Mellon University, and an M.S.E.E from the Air Force Institute of Technology. She is a certified Program Manager Professional.

Dr. Bud Mishra

Dr. Mishra is a professor of computer science and mathematics at NYU's Courant Institute, NYU. He has degrees in Sciences from Utkal University, in Electronics and Communication Engineering from Indian Institute of Technology (IIT), Kharagpur, and MS and PhD degrees in Computer Science from Carnegie-Mellon University. Dr. Mishra is a fellow of National Academy of Inventors, an IIT Distinguished Alumnus, NYSTAR Distinguished Professor, AAAS Fellow (engineering: robotics, hardware verification and computational biology), IEEE fellow (robotics and automation) and a fellow of the ACM (computational biology and symbolic computation).

Andrew P. Moore

Andrew P. Moore is lead researcher in the CERT National Insider Threat Center at Carnegie Mellon University's Software Engineering Institute. He has thirty years of experience developing and applying mission-critical system analysis methods and tools. Andy has worked for the Naval Research Laboratory investigating high-assurance system development methods, has edited a journal Special Issue, and has authored a book, two book chapters, and a wide variety of technical journal and conference papers. Andy received a M.A. in Computer Science from Duke University, B.A. in Mathematics from the College of Wooster, and Graduate Certificate in System Dynamics from Worcester Polytechnic Institute.

William E. Novak

William E. Novak is a Principal Member of the Technical Staff at the Carnegie Mellon University Software Engineering Institute. He is a researcher, consultant, and instructor in the acquisition and development of software-intensive systems. Mr. Novak has over thirty years of experience with government acquisition, real-time embedded software and electronics product development, and business management. He has held positions with GE Corporate Research and Development, GE Aerospace, Texas Instruments, Tartan Laboratories, and GTE Automatic Electric Laboratories. Mr. Novak received his M.S. in Computer Engineering from Rensselaer Polytechnic Institute, and B.S. in Computer Science from the University of Illinois at Urbana-Champaign.

Inherent Moral Hazards in Acquisition: Improving Contractor Cooperation in Government As The Integrator (GATI) Programs

William E. Novak (wen@sei.cmu.edu)
Julie B. Cohen (jcohen@sei.cmu.edu)
Software Solutions Division
Software Engineering Institute
Carnegie Mellon University

Andrew P. Moore (apm@sei.cmu.edu)
William A. Casey (wcasey@sei.cmu.edu)
CERT Division
Software Engineering Institute
Carnegie Mellon University

Bud Mishra (mishra@nyu.edu)
Courant Institute of Mathematical Sciences
New York University

Abstract

In the Government As The Integrator (GATI) model of acquisition, prime contractors no longer hand-select the members of the acquisition team or consortium, as they often did in the Lead System Integrator (LSI) model. One drawback of GATI acquisitions, thus, is that independent contractors may have little incentive to cooperate by sharing data and supporting other contractors, potentially resulting in delays, overruns, and poor performance. These problems are considered in this work to be both breakdowns in cooperation and expressions of moral hazards. Since the need for cooperation among contractors is still critical to success, finding ways to motivate that cooperation to improve program performance and outcomes is key to effective GATI acquisition. In this research, potential incentive mechanisms were analyzed for their ability to promote cooperation by applying game theory framing and analysis to this GATI acquisition context, and using system dynamics and agent-based modeling to study the results for their ability to promote cooperation and improve program outcomes.

Keywords: game theory, buyer-supplier game, cooperation, system dynamics modeling, agent-based modeling, optimization, defense acquisition, Government As The Integrator, Government As Integrator

1. Introduction

Acquisition programs can suffer from recurring cost, schedule, and quality problems, failing to deliver on time or with full capability, and delaying or leaving the warfighter without needed capabilities (Buettner, 2014). Decision-making impacts acquisition outcomes due to their significant cost and performance implications. Many key decisions are initially made in defining the acquisition strategy (Ward, Elm & Kushner, 2006) (e.g., acquisition approach, solicitation type, monitoring activities, etc.), and often must be revisited as program circumstances change. Getting them wrong has consequences: "...Acquisition decision-making has not been well managed for these [enterprise IT] systems... and the result has not served today's leaders and soldiers well... The resulting operational impact is profound." (Dacus & Hagel, 2014) To wit, the DoD portfolio experienced \$469 billion in cost growth (over 48%) since programs established their first full estimates (USGAO, 2017).

Acquisition program managers often lack the data, evidence, and tools needed to make complex decisions associated with acquisition strategy trade-offs by identifying and evaluating acquisition solution options. They are also in need of better ways to encourage good performance on the part of key program stakeholders. They would be well-served by ways to improve both the quality of acquisition staff decision making, and ways to incentivize desirable acquisition program behaviors. One area that has increasingly come to the fore in terms of its importance to acquisition program outcomes is that of systems integration, which is the focus of the work described here.

The objective of the research described in this paper is to use modeling and simulation to describe and analyze problems with U.S. Government system acquisition, and to develop a means to improve decision-making in

the acquisition context. We identify different classes of incentive mechanisms—direct financial, future business, and team networking—that can be applied to influence contractor behavior by aligning their interests with those of the program. Specific incentive mechanisms can help to keep contractors focused on program work—but there is no single best incentive. Using a combination of types of mechanisms can achieve greater levels of influence on contractor behavior across a range of organizations with different business models, as will be discussed in detail.

This research uses system dynamics modeling to characterize the acquisition problem and the interrelationship of incentive mechanisms. It also uses game-theoretic agent-based modeling in a targeted way to analyze the ability of incentive mechanisms to mitigate system acquisition problems. Incentive mechanisms are analyzed for their ability to promote cooperation and thus improve program outcomes. Upon considering the mathematical models presented, we argue that there is an engineering science involving how best to apply a mix of incentive mechanisms, and in what context they may have the greatest effect. This and other areas of potential future work are discussed at the end of the paper.

1.1 Addressing Systems Integration

Largely as a result of the failures of a number of major acquisition programs using the Lead System Integrator (LSI) approach in the 1990s and 2000s, the U.S. Congress enacted the National Defense Authorization Act (NDAA) legislation of 2006, 2008, and 2009 that limited the use of LSI by barring the award of most new LSI contracts after FY 2010 to any contractor who had not already done LSI work previously, and prohibiting the use of an LSI for programs beyond the point of low-rate initial production (LRIP).

The poor track record of LSI programs (Young, 2010) and the NDAA legislation led acquisition program management offices (PMOs) to consider one of the most logical alternatives to LSI (DePillis, 2013), which was to have the government act as its own systems integrator, an approach known as Government as the Integrator (GATI). The GATI approach claims several key benefits over the use of an LSI, including government control of the design and architecture of the system and software, better visibility into program status and progress, and the development of higher technical expertise in the government acquisition workforce.

With these seemingly clear advantages, the question arises as to why GATI acquisition doesn't always perform as well as one might expect (Gansler et al., 2009). One key issue is the declining amount of technical expertise in the government acquisition workforce over the past 20 years, producing less than optimal outcomes when the government takes on technical responsibility. In addition, there is a difference in the way GATI teams are formed versus the way in which LSI teams, or "consortia," are formed. In an LSI acquisition the prime contractor deliberately and carefully chooses a team of select contractors that will bid on, and (if awarded) ultimately perform the work, taking into consideration their areas of expertise, past performance, competitive aspects, and other factors. In such a team there is an incentive for all members to contribute to the success of the team's work, and the prime helps to enforce this inclination.

In a GATI acquisition, however, no such team is chosen up front, and contracts with various companies are awarded individually and competitively to make up the government's contractor "team." One result of the GATI approach is that unnatural alliances may be inadvertently formed in the process of selecting "best of breed" contractors in each technical area, so that within the final team competing and contentious rivals may be expected to cooperate. Since there may be no pre-existing relationships among the contractors to encourage cooperation, and no prime contractor present to enforce it, mechanisms such as Associate Contractor Agreements (ACAs) may be used to mandate the sharing of information and the cooperation that will be required in a GATI context (AFFARS, 2013).

Unfortunately, but perhaps not surprisingly, contractors generally don't like using ACAs because they may be forced to support, work closely with, and exchange sensitive information with potential or actual competitors. The net effect in such situations can be superficial cooperation, with intense competition, mudslinging, and even backstabbing occurring behind the scenes. Since in acquisition programs the incentives driving the contractors' behaviors are not inherently aligned, rather than mandating cooperation, the most effective way to promote cooperation is to strongly incentivize contractors to behave that way towards one another to advance the program's goals.

1.2 The Principal Agent Problem and Moral Hazard

We can broadly characterize a defense acquisition following the GATI model as a government PMO that is coordinating the activities of a set of contractors who are each developing different portions, or *components*, of a system. We can refer to these contractors, or to the government-contractor entities responsible for developing these components (e.g., in cases where the major system components are developed by independent acquisition programs), as *component performers* (CPs). These component performers have specialized technical expertise and information that is beyond that of the PMO, which is precisely why they have been engaged to do the work. In the world of microeconomics and game theory this type of arrangement is referred to as the “Principal-Agent” problem, in which expert “agents” (i.e., contractors) apply their expertise to perform work for the “principal” (i.e., the PMO).

The disparity of expertise (or more generally information) between principals and agents is referred to as *information asymmetry*, and creates an imbalance of power between the two parties. When the agents’ knowledge exceeds that of the principal, the informational advantage can yield what economists refer to as “utility gains.” As an example consider the schedule to complete a task for which the agent has a high level of experience, while the principal has a low level of experience: the principal finds it inherently difficult to assess the proposed schedule without additional information. The agent, having more knowledge, may propose a schedule that benefits themselves over the principal, choosing an overly optimistic schedule to improve their chances of getting the contract, or a pessimistic one to provide scheduling flexibility should other lucrative side opportunities arise that the agent will want to pursue. Such situations are referred to as a “problem” because while the ideal goal is to have the agent act transparently as the principal needs, the agents’ interests are not necessarily aligned with those of the principal, and honesty and high quality work may not be the agents’ best strategy.

A concept that is embedded within the “Principal-Agent” problem is that of *moral hazard*. Moral hazard refers to situations in which the agent decides how much risk they are willing to accept, but the principal bears the cost if things go poorly—and so the agent may decide to take on more risk than they would have if they bore the full cost of that risk themselves. In the acquisition context a moral hazard may happen when the contractor and the government have different objectives for the engagement, and the government may not be able to tell whether the contractor is acting in the government’s interests (e.g., by providing honest estimates of cost, schedule, quality), or is acting in its own self-interest and forcing the government to accept the risk associated with inaccurate information.

The presence of moral hazard in acquisition can have many different manifestations and consequences. One that is of specific interest to GATI acquisitions deals directly with cooperation. Consider the situation where

Illustrative Moral Hazard in Acquisition

An example of this type of problem occurred on an actual acquisition program, but the names and technologies have been changed in this description. A next-generation cruise missile will be available from its program early next year. However, the program developing the airborne launcher for the missile, which was originally scheduled to be ready at the same time as the missile, has encountered technical problems, and due to the resulting delays the launcher won’t be available until the following year at best. Nevertheless, the cruise missile program plans to proceed with the originally scheduled production and deployment of a large quantity of the new cruise missiles.

To make matters worse, the sensitive electronics and advanced sensors in the missile are known to degrade over time, impacting the missiles’ performance, and ultimately rendering the missiles unusable after five years. The sponsoring military service doesn’t understand the decision to proceed with production of the missiles now, but they suspect it may be so that the missile program office will be able to say they met their schedule, even though it will waste a significant portion of the expensive missiles’ expected useful life.

This is a moral hazard in that component performers are making decisions either purposefully disregarding or passively ignorant of the needs of the overarching program. The risks to the program associated with those decisions are born primarily by the program, rather than by the performers making those decisions.

CPs may decide (without the government’s knowledge) to promote the interests and success of their particular component at the possible expense of

- others achieving their own component objectives, thereby adding risks to other team members and indirectly to the program’s goals, or
- the global program achieving its objectives, thereby adding risks to program goals.

This constitutes a moral hazard in that these CPs are making decisions that are either deliberately disregarding, or passively ignorant of, the needs of the overarching program. They are thus incurring risks to the program that will ultimately be transferred to and borne by the PMO, rather than by the CPs making those decisions. There could be many different rationales behind such a decision, ranging from optimizing the component’s own Earned Value Management (EVM) performance to achieve an incentive fee (despite adversely impacting the performance of other components), to temporarily diverting needed contractor staff to work on a more lucrative “side” opportunity external to the program (while in the process delaying the global program schedule). We will discuss the implications of such side opportunities in greater detail throughout the paper.

Such a situation is made even more acute in a system-of-systems acquisition context where each CP could be a separate acquisition program in its own right. If each program were a separate financial entity, there would be that much more incentive to pursue self-interest, even to the detriment of the overall system.

2. A Dynamic Characterization of the Problem

Among the critical dilemmas of GATI acquisition is that within a CP with weakly aligned utilities, local objectives can take priority over global objectives either in coordination with the government, or in coordination with other CPs. Figure 1 illustrates a “Growth and Underinvestment” dynamic (Senge, 1990) that can occur during acquisition, yielding increased propensities for CPs to pursue local objectives at the expense of global (i.e., GATI) goals. Since this growth and underinvestment behavior unfolds dynamically based on feedback to the acquisition decision makers, we use a conceptual system dynamics model to represent it.¹

As shown in the upper left portion of the figure, the GATI perspective promotes the desire for government integration and a demand for the government integration function. Government integration fulfills the needs of the CPs in the early stages as seen in the self-balancing feedback loop B1 (dark blue). However, as use of government integration increases naturally as part of an acquisition involving multiple CPs, the CP’s satisfaction with the integration decreases, driving the need to grow the government integration capability in the self-balancing B2 (purple) feedback loop. The inherent limitations on the speed of this growth relative to the increased need represents the problematic pattern of “Growth and Underinvestment” mentioned previously. The resulting decline in satisfaction with government integration promotes the CP to try to take on more of the integration function, leading to a design that is optimized more for that CP’s local objectives. This behavior is self-reinforcing (see the green feedback loop R1) in that the CP comes to prefer their own local decisions over those of the government integrator. The negative implications of this for the acquisition program achieving program goals is illustrated by the (red) influence arrow, which represents an undermining of the government acquisition program’s ability to achieve its global goals.

The stocks and flows in the middle-right portion of the figure illustrate staff moving from development to integration functions, development to support functions, or development to other project work. Moving staff to support functions is beneficial to the acquisition program in that it improves the productivity of the other performers. In addition, as shown in the (orange) feedback loop R2, it is self-reinforcing in that as a CP supports other performers, those other performers are more likely to support the original CP’s own needs. This mutual

¹ The notation used in this figure is summarized in the appendix.

figure shows that individual performance primarily depends on the CP's schedule and quality performance. The right half of the figure shows that the collaboration among CPs depends on the trust and accountability in the CP relationships. Collaboration can also be greatly enhanced if the CPs do not have to worry about losing key intellectual property (IP) on which their future business depends.

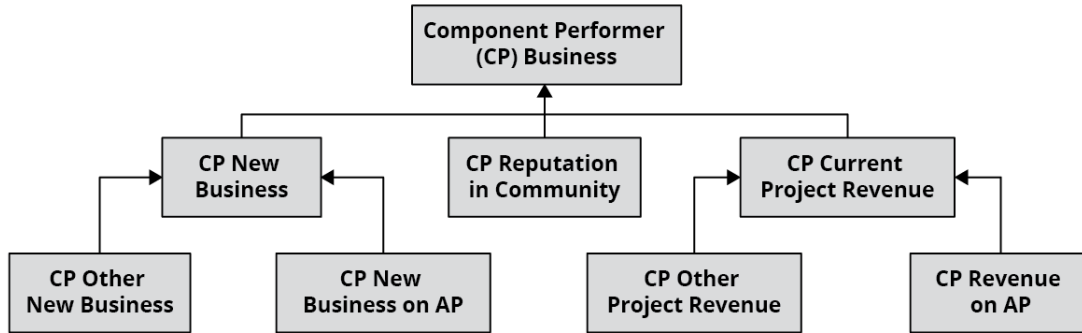


Figure 2: Drivers of Component Performer (CP) Behavior

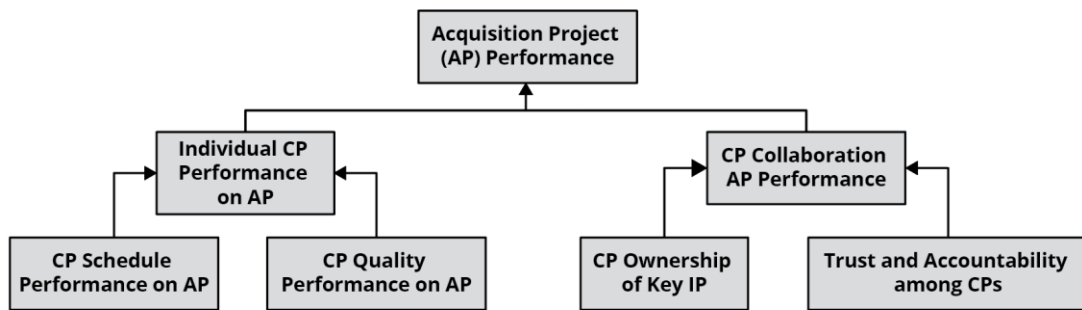


Figure 3: Drivers of Acquisition Program Management Office Behavior

The goal of this section is to show how incentive-based mechanisms can help align the objectives among CPs and between CPs and the PMO. To do that, we need to describe the types of incentive mechanisms we believe are important. Incentives that influence whether CPs take a local or a global perspective in their decision making have three primary types depending on whether the incentive focuses on the organization's *business*, on the *money*, or on the acquisition program *team* members in the organization. As we discuss each of these categories below, Figure 4 ties example incentives back to their influence on achieving CP and PMO objectives as stated previously.

Business: *Future business* incentives encourage desired behaviors by increasing the potential for CPs to earn future business. While this ultimately may be financially rewarding for the CP, it is primarily about increasing the performer's competitive edge so as to make money in the future, rather than about paying them money directly for good performance now. Figure 4 shows two future business incentives: *reputation tracking* (Padovan et al., 2001) and *intellectual property/non-disclosure agreements (IPA/NDA)*.

- Reputation tracking mechanisms track the performer's performance on the current program and make the results available for consideration in future acquisition program awards. The Contractor Performance Assessment Reporting System (CPARS) is one example of this in current use (CPARS, 2017) (USDODIG 2017). Good performance on the current tasking can also be incentivized by allowing options for additional years of contract performance when the CP's performance satisfies certain criteria.

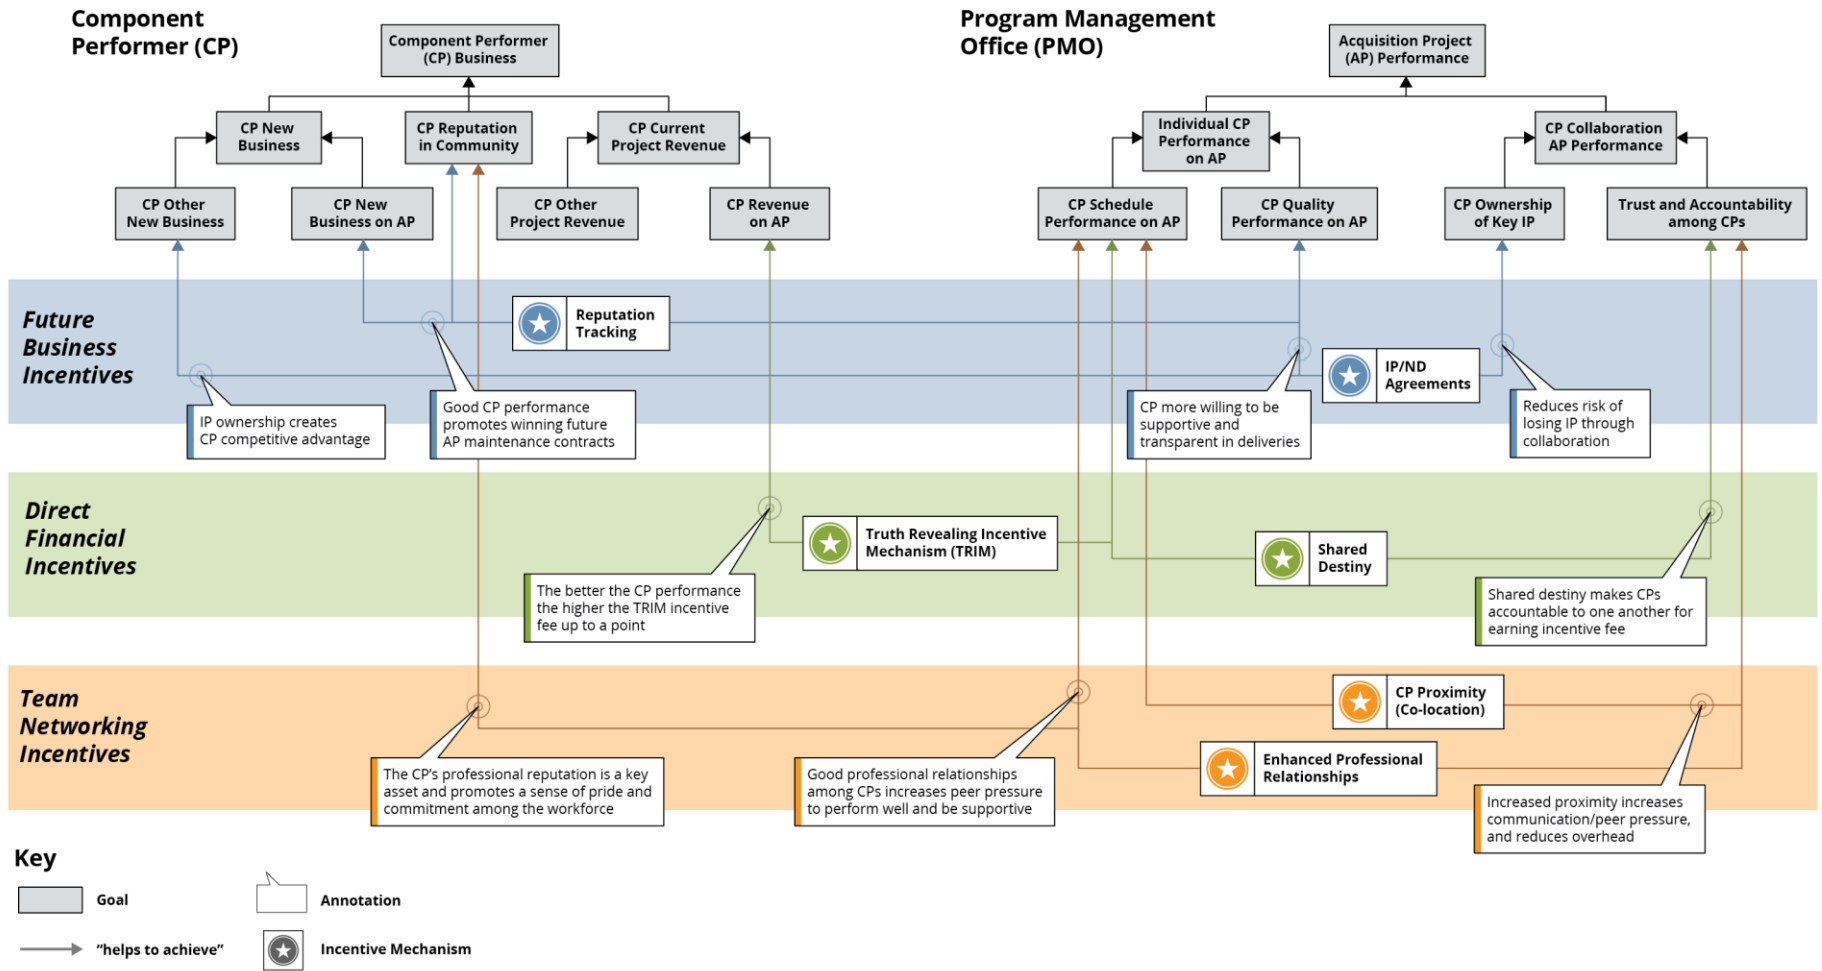


Figure 4: Mechanism-Based Incentive Alignment

IP agreements allow CPs to have certain exclusive rights to IP developed as part of the program activities, enabling them to use this IP in future developments or negotiations. NDAs ensure that collaborators will keep confidential communications that, if divulged, could hamper the CPs business or reputation on current and future ventures. IP agreements and NDAs can help ensure transparent communication and sharing among CPs that is needed for program success.

Money: *Direct financial* incentives encourage desired CP behaviors by providing financial rewards for those behaviors (Rendon, 2012). Figure 4 shows two such direct financial incentives: the *Truth Revealing Incentive Mechanism* (TRIM) (Coughlan et al., 2009) and *Shared Destiny* (USGAO, 2006).

- TRIM promotes the accuracy of performer cost and schedule estimates by making sure that accurate estimates optimize the cost-plus incentive fee awards provided to the CP (Coughlan et al., 2009). TRIM consists of a sliding incentive fee based on the promised schedule date, and the fee decreases more rapidly as the schedule moves further out. The smaller the schedule variation, the larger the fee, thus incentivizing on-time delivery. The mechanism encourages the CP to “reveal the truth” regarding their ability to meet the schedule, since by doing so they will earn the largest fee. Proper use of TRIM can thus ensure both accurate estimation that will benefit the PMO, and maximize the incentive fee, to the benefit of the CPs
- Shared Destiny promotes performer collaboration to achieve overall program objectives by financially rewarding performers according to the success of the program. All performers only receive as much fee as the lowest-performing team receives, so all teams are incentivized to improve the capability of the lowest-performing team, who without that help might reduce everyone’s fee (USGAO, 2006). This motivates both collaborative effort and individual CP accountability for achieving overall program success. CPs must balance the needs of achieving their own local objectives with the support, communication, and sharing needed to achieve global objectives. Trust among CPs is an essential ingredient of this success, which is promoted by the team networking incentives.

Team: *Team networking* incentives are social mechanisms to promote positive attitudes among individuals associated with different CPs. Previous research shows that social incentives can be even more influential in shaping behavior than market (financial or business) incentives (Pentland 2015). Figure 4 shows two team networking incentives: *CP proximity* and *enhanced professional relationships*.

- Professional relationships can be enhanced through team building or other social functions that permit different teams across the CPs to get to know and respect each other both professionally and personally.
- Team networking can be improved by co-locating team members of different CPs. While this can be done virtually using networking and collaboration tools available, physical co-location, if possible, is likely to improve team networking, interaction, and familiarity to the greatest extent (Pentland 2015, Pentland 2012).

4. A Game-Theoretic Basis for Acquisition

Game theory has been characterized as “the study of mathematical models of conflict and cooperation between intelligent rational decision makers” (Myerson, 1991). While early discussion of two-person games goes back at least as far as the 1700s, game theory did not exist as a unique field until John von Neumann and Oskar Morgenstern’s descriptions of the foundations in their book in 1944, *Theory of Games and Economic Behavior* (Leonard, 2010). The 1950s saw the field expand into the logical side of decision science by many scholars in the areas of economics, political science, psychology, computer science, and biology. Cooperation has long been studied in applications of game theory with fundamental insights in repeated games occurring in the 1980s (Axelrod, 1984). Leveraging these insights for business gained significant traction in the 1990s by extending cooperation into competition (so-called “co-opetition”) (Axelrod, 1997) (Nalebuff and Brandenburger, 1997) and also by generalizing the business ecosystem (Moore, 1993). Interest and publications in these areas have continued through the 2000s into the recent literature (Peltoniemi and Vuori, 2004) (Axelrod, 2006) (Dix-

it and Nalebuff, 2008) (Brandenburger and Nalebuff, 2011) (Daidj, 2017). Another aspect of game theory, evolutionary game theory, as described in *Evolution and the Theory of Games*, (Smith, 1982) introduces the idea of the Evolutionarily Stable Strategy (ESS) that can be identified by analyzing evolutionary games using simulation.

Our application of game theory at the most abstract level builds on the earlier work on combining competitive and cooperative business strategies as laid out in Nalebuff and Brandenburger (1997). As in that paper, we ask the question of how we could change the current acquisition “game” to the benefit of acquisition programs more broadly. Figure 5 extends and adapts the value net for the business relationships specified in that paper to the acquisition context. As seen on the left side of that figure, CPs and the PMO are the key strategic players in the game, loosely governed by the acquisition contract. However, each CP is desired to collaborate with other CPs and achieve overall program objectives, usually through vehicles such as an Associate Contractor Agreement (ACA) and possibly additional Memoranda of Understanding (MOU).

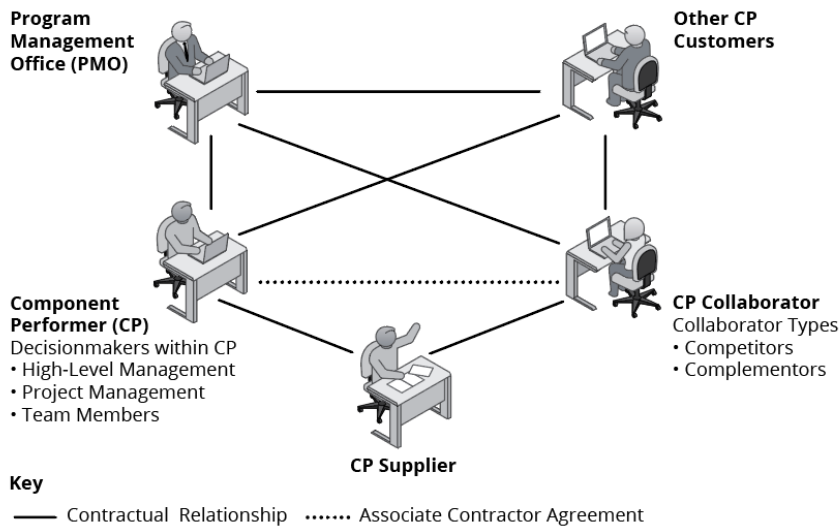


Figure 5: Acquisition Game Players

Both CPs and their collaborators have contractual relationships with the PMO, but will likely have other customers (and suppliers) as well. The figure is *not* meant to imply that all CPs necessarily have the same customers (or suppliers), but rather that all CPs have other customers besides the acquisition PMO. These other customers, both current and potentially in the future, are in some sense in competition with the PMO in terms of adequacy of staffing the acquisition program, since those other customers can draw resources (capabilities and people) away from the program by engaging with the CPs. A focus of our modeling and simulation efforts is on incentive mechanisms that try to avoid this, and can help make CPs concentrate on the PMO’s acquisition program.

A key challenge is the extreme diversity of the players: different stakeholders within a CP’s organization are likely to be influenced by the incentive types in different ways. High-level management within the CP organization is likely to be most concerned with future business incentives, with direct financial incentives a close second. Project management personnel are likely to be most influenced by the direct financial incentives, since their performance evaluation will presumably depend most on turning a profit. At the lowest level in the organizational hierarchy, project team members of one performer may be most influenced by the professional relationships among project team members of other performers. For example, if management approves of sharing data, but there’s little trust of other CPs, a CP may send as little as possible, causing the other CP to continually request more data. Similarly, if team members really trust their colleagues in other CPs, they may share information with them without explicit approval from management (McAllister, 1995). This same dynamic may be true about a CP sharing data with the PMO. The level of influence wielded by these different stakeholders may vary by organization, but all of them likely need to be considered at some level to determine overall organization behavior.

Finally, the quality of the relationship between the CPs on an acquisition program is key to its success, and characterizing it is a central component of our modeling and analysis approach. Keeping the relationship healthy is more challenging if the collaborators are also competitors. A collaborator CP is a *competitor* if they pursue the same business as the other CP. A collaborator CP is a *complementor* if they go after business in a different and complementary domain as the other CP. As characterized in Nalebuff and Brandenburger (1997), a CP's service/product is worth *more* when combined with a complementor's service/product (because of the increased customer value of the combination), whereas it is worth *less* when combined with a competitor's service/product (because of the competition for the same customer space).

While acquisition can be viewed as a game, the very large complexity associated with acquisition programs suggests that a game-theoretical view of acquisition may be constructed from a set of simple, but essential games. Taking this approach, more complex programs are considered as a set of causal connections among essential games. Section 4.1 describes a particular game relevant to acquisition. Section 4.2 provides an overview of the specification of the CP agent as it is affected by incentive mechanisms discussed previously. Simulation results from a detailed system dynamics model of the agent describe some interactions of incentive mechanisms. Later sections describe the vision for possible future extensions of this work.

4.1 An Acquisition Game

Since the contractors within an acquisition program are voluntary and work toward shared objectives with common or reciprocal interest, a primary consideration is their coordination and cooperation to improve efficiencies. For this purpose, we consider a class of games known as “coordination” games and note its ability to characterize a large class of social cooperation problems, including those arising within acquisition programs, and specifically GATI acquisitions. In the acquisition context, the CPs may be modeled as rational (utility optimizing) agents. Coordination relies on the actions of others, and naturally presents uncertainty and risk. One example of this is the uncertainty regarding the level of cooperation. Due to uncertainty and risk, CPs may choose a course of action to maximize their individual rewards over that which achieves the best outcome for the acquisition program. To illustrate this dilemma we use a version of the familiar Stag Hunt game (explained below) to describe how the perception of risks may lead a rational agent (i.e., the CP) to potentially select non-cooperative actions, thus burdening the acquisition program with suboptimal outcomes.

This situation is essential to acquisition strategies, and represents the aforementioned “moral hazard,” as often the government (not the CPs) bears the cost of those risks when program outcomes are poor. Thus, it is important to better understand the nature and degree of the risks, and how those risks are transferred. They are examined here in the context of a specific GATI program which

Acquisition as a Stag Hunt

The problem of agents pursuing self-interest despite clear direction from the principal to cooperate in order to achieve more global objectives is well represented in game theory terms by the Stag Hunt game. Stag Hunt is a type of “coordination” game that characterizes a wide range of social cooperation problems. In the game two players go out hunting together, and each player can individually choose to hunt either a stag or a hare, but they may not be aware of what the other player has chosen to do. A player can catch a hare alone. However, a stag is worth more than twice as much as a hare, but (due to its larger size) both players must choose to hunt the stag if they are to be successful. Thus, the Stag Hunt game has two outcomes (i.e., Nash equilibria) that are most likely: either both players hunt a stag, or each hunts a hare.

The Stag Hunt game clearly describes a situation containing a “moral hazard.” It also directly parallels an acquisition scenario where the component performers make decisions more focused on maximizing their individual gain (i.e., pursuing hares in Stag Hunt parlance), and less on achieving the cooperative best outcome for the overall system (i.e., hunting a stag). Even if a hare is worth only half of a stag, hunting a hare represents work that can be done successfully alone, vs. working with a team of others where there's risk and uncertainty around what the collective outcome will be.

However, it should be noted that if the Stag Hunt game is played repeatedly, there may be reputational impacts if one player/contractor repeatedly abandons their partner(s) to pursue another opportunity (i.e., hunting a hare), rather than cooperating to achieve the agreed-upon goal of the program (i.e., hunting the stag). Such “defections” may be met by others being unwilling to engage with them in the future.

seeks to obtain the highest satisfaction when all CPs avoid moral hazard and follow through with schedule commitments. As such, the GATI program manager may offer an incentive so that each CP also gains the greatest satisfaction in that case. However, note that the program manager can do little to manage the CP's perception of risk, such as the uncertainty in other CPs' actions.

In summary, we describe a model of cooperation by Stag Hunt games which involves the essential nature of cooperation and risk. To consider coordination problems more directly on a GATI task we return to the problem of maintaining a CP's follow-through to schedule commitments, even when they (and their peers, i.e., other CPs supporting the GATI program) are tempted with potentially lucrative side-offers from other customers. Since all CPs on a GATI task are likely to have similar attractive side offers and are aware of this possibility, the action of following through on task schedule commitments (and declining lucrative side-offers from other customers) will be risky as it can yield losses if other CPs fail at their schedule commitments. The greatest reward achievable for the GATI program is obtained only when *all* CPs follow through on commitments.

The simplest Stag Hunt game involves two players who select how to allocate their effort to either hunt stag (S) or hare (H). If both players hunt hare, the reward is for each h . If both hunt stag, then the reward is $s/2 > h$. However, a stag can only be captured when both players work together, so if one hunts stag while the other hunts hare, the rewards are 0 for the player hunting stag, and h for the one hunting hare (see the payoff matrix in Table 1). The most satisfactory outcome for both players is therefore to cooperate in hunting stag. Unfortunately, it entails risk (as the partner may commit to hunting stag, but not follow through, deciding instead to hunt hare). Its Nash equilibria can be analyzed by computing the expected reward and variance conditioned on selected effort. Letting P_S and P_H represent the probability that a partner will hunt stag and hare (note that $P_H = 1 - P_S$), then by selecting H the expected reward is h independent of the other's selection (i.e., having variance zero). By selecting S the expected reward is $s/2 * P_S$ with variance $P_S P_H$. Uncertain of what a partner player may select, the value P_S (with $0 \leq P_S \leq 1$), may represent a subjective belief of their propensity to play S, and its ability to diminish the expected return can be seen by noting that the expected reward for choosing H is greater than that of selecting S when $P_S < 2(h/s)$. A risk-averse player may also trade off reward to decrease risk, thereby selecting H for higher values of P_S in order to control the downside risk.

With S representing a CP's choice to follow through on commitments, and H otherwise (perhaps delaying the GATI schedule by accepting a lucrative side project), the single-shot Stag Hunt game combines the elements of risk, reward, and cooperation found in the aforementioned GATI scenario. However, the single-shot game lacks the temporal dimension where repeated outcomes may result in dynamic behaviors such as the emergence of reciprocity. We therefore contend that the repeated game form of the Stag Hunt game may be more relevant to capture the dynamic behaviors on a GATI program, and thus suggest that properly monitoring the behavior and reputational metrics of the CPs may form important tools for GATI PM. Additionally, other incentive mechanisms may be considered, e.g., the TRIM and Shared Destiny incentive rewards.

An analysis of how players consider acting in a cooperative game may be extended to consider the following: (1) how various incentive mechanisms alter the dynamics of the game, and (2) how changing mechanism parameters affect the expected behaviors within the game. This type of analysis may be applied to optimize the mechanisms governing how a GATI PMO may best consider various incentives or scheduling actions to yield the best program outcome. To reiterate, in modeling the moral hazard to GATI programs we only need to consider who assumes the cost arising from risk when outcomes are poor, and also by what means the moral hazard may enter the GATI program. While above we model the *defection* (or the selection H, i.e., to not follow through on a schedule commitment) as facilitated by a lucrative side deal, we note that it could likewise arise from various other causes. For example, the CP may choose to spend more time to improve the quality, or may decrease the quality to free up time to pursue a lucrative side offer, yet still abide by the schedule commitment (albeit with a lower quality and higher profit). Still, the subject of obtaining a truthful and accurate schedule estimate is presumed as asymmetric information yielding leverage for the CP over the government, as the CP retains the expertise needed to accurately estimate cost and schedule. The estimate of schedule, being asymmetric information similar to other cases listed before, is known to the CP but not to the government, and thus facilitates how moral hazard may enter the scenario. In future work we wish to consider the entry of moral hazard generally, and summarize how risks may be transferred within these or specific GATI scenarios.

Stag Hunt (Player 1, Player 2)		Player 2	
		Stag (S)	Hare (H)
Player 1	Stag (S)	$(s/2, s/2)$	$(0, h)$
	Hare (H)	$(h, 0)$	(h, h)

Table 1: Stag Hunt Game Normal Form

Recognizing that the normal form of the game (above) is ideal, but not necessarily realistic, and that we cannot fully know the CPs' true preferences, we then used a statistical approach to approximate a range of CP preferences regarding the specific incentive mechanisms that the PMO chooses to employ. The image shown in Figure 6 is a single frame from a movie developed by the project to illustrate the outcomes of the Stag Hunt cooperation game. The left-hand side of the figure shows the utilities of Player 1 and Player 2 for various outcomes of a Stag Hunt game with incentives, which in the movie version change depending on how the PMO "weights" or controls parameters of four different incentive mechanisms including TRIM and Shared Destiny. The right-hand side of the figure provides a statistical view considering the range of contractor preferences of the program outcomes (in a color scale), which in the movie is further explored by varying parameters of the different incentive mechanisms. The outcomes for the PMO are shown with better outcomes (i.e., where both players cooperate, or hunt a stag) in lighter colors, and worse outcomes (i.e., both players defect, and hunt hares) in darker colors. The X and Y axes both represent increasing CP satisfaction, so locations in the upper right represent the greatest overall satisfaction among the CPs. A dark area in the lower left is where the CPs are generally incentivized to move, which indicates that poor (uncooperative) outcomes for the program as a whole are most likely. The size of the light or dark areas indicates the variance or risk (uncertainty) of that outcome.




Figure 6: Image from Stag Hunt Analysis Movie

The movie illustrates a key turning point in the outcomes of the Stag Hunt game. In Stag Hunt terms, this is the point at which the value of a hare becomes larger than the value of half of a stag (which is what each hunter would receive for his or her efforts in capturing a stag). This threshold determines where cooperation would break down, as it becomes no longer worthwhile to work together to hunt stag. In acquisition terms, an equivalent scenario would be the simultaneous use of incentive mechanisms such as TRIM (promoting self-interest) and Shared Destiny (promoting cooperation). An overly heavy emphasis on the use of TRIM will undermine the effectiveness of the cooperative effects of the Shared Destiny mechanism. The full movie can be viewed at the SEI website at <https://www.sei.cmu.edu/go/moralhazards>.

4.2 A Systemic View of the Acquisition Game

We now move to a more detailed descriptive model of the dynamic behavior of a CP agent, and its interaction with both the PMO and other CPs. This systemic view of an acquisition program has a number of key dynam-

ics as depicted in Figure 8. We detail these dynamics below, starting with the basics of getting work done by CPx in dynamics D1 and D2:

- **D1: CPx Getting Work Done on AP** – D1 is the primary self-balancing loop (shown in dark blue) that drives CPx to do work. CPx allocates FTE to the AP and does work at a level of schedule performance as deemed acceptable by the acquisition PMO.
- **D2: CPx Getting Work Done on Other Projects** – Of course, CPs are likely to have other development work on which they need to progress. D2 involves a self-balancing loop (also shown in dark blue) very similar to D1 to indicate progress on this other work. Allocating personnel among projects is a key function of CPx management. As other project work grows, personnel may need to be taken from the AP to staff that work.

The dynamics for each of the incentive mechanism categories are specified in dynamics D3 through D5:

- **D3: CPx Benefit from AP Direct Financial Incentives** – Dynamics associated with Direct Financial incentives are shown (in green) in the figure and include both the TRIM and Shared Destiny mechanisms. The PMO agent controls TRIM through the incentive fee calculation which is designed to discourage CPx from taking other business opportunities that would take resources away from the AP to a level that would prevent CPx from meeting its schedule estimates for its AP tasking. Likewise, the PMO controls the Shared Destiny incentive fee calculation, but here it motivates CPx to work with its AP collaborators sharing information and providing assistance, so that all CPs on the AP can perform in an effective and efficient manner.
- **D4: CPx Benefit from AP Future Business Incentives** – Dynamics associated with Future Business incentives are shown (in light blue) in the figure and include both Reputation Tracking and IP/ND Agreement mechanisms. Both of these mechanisms motivate CPx decision makers by enabling and supporting the development of new business ventures, which can draw personnel and capabilities away from the AP. Therefore, the use of these mechanisms needs to be balanced with other incentives that promote CPx focusing on achieving AP objectives.
- **D5: CPx Benefit from AP Team Networking Incentives** – Dynamics associated with Team Networking incentives are shown (in orange) in the figure and include mechanisms based on CP proximity and enhancing professional relationships. CP Proximity mechanisms increase interaction frequency and the potential for observing collaborator behaviors, which have been shown to reinforce behavior and increase trust (Pentland, 2015). Combined with mechanisms for Enhancing Professional Relationships, good foundations for trust can be enabled which spur the CP interaction necessary to achieve program goals. Good collaboration can be a self-reinforcing dynamic, in the positive direction. However, poor relations between CPs can cause this feedback to go in the negative direction, resulting in a downward spiral that reflects dynamics seen in actual programs.

No matter what incentives are adopted, over time the players will likely find ways to game the resulting system to their own advantage, and to the disadvantage of the overall program. The ways to game a particular system are likely numerous and certainly hard to predict; one possible way is described by dynamic D6. Modeling and simulation can be used to help analyze these schemes by applying concepts from evolutionary game theory.

- **D6: Gaming the System – Shortcutting Collaboration Support to Enhance own Business** – The dynamic associated with one potential way of gaming the system of incentive mechanisms introduced so far is shown (in light purple) in the figure. It is based on the notion that CP reputation is relative to the reputation of other CPs, and a reputation tracking system may have the unintended effect of having deficiencies in one CP's performance being blamed on other CPs as a way for the guilty CP to keep their own reputation intact. Even worse, CPs may disrupt their own support to other CPs in subtle ways so as to bolster their claims. Another more obvious potential motivation for shortcutting a CP's support of its collaborators is that this frees up resources to allow a CP to take on other opportunities.

Although the constraints of this paper will not allow a description of the details of the simulation model, we can discuss a few simulation results that show the conceptual value of modeling and simulation in

evaluating combinations of incentive mechanisms in the acquisition context. We use a measure called Composite Program Performance, defined as the product of three component measures: two measuring the CP’s schedule performance and productivity, and one measuring the extent to which the overall program requirements are satisfied. Each of the component measures are normalized to be between 0 and 1, as is the Composite Program Performance. While a more general model may weight the component measures differently, we’ve assumed equal weighting in the following analysis.

Figure 7 shows the complementary aspects of team networking and Shared Destiny incentives. Both types of incentives are intended to improve collaboration among CPs, but they do so through different means. Team networking incentives stimulate the social connections between personnel in different CPs, e.g., through proximity, team building, etc. Shared Destiny stimulates collaboration through direct financial incentives. These incentives are complementary—they target different decision makers within a CP. Team networking is more likely to affect team members, whereas Shared Destiny will have a bigger impact on management. Since we do not yet know the individual impact of these incentives on CP behavior, we can simulate them over a range of possible impacts, from low to high as shown in the figure. We note that there are diminishing returns with respect to applying more powerful team networking or Shared Destiny incentives, especially as the cost of those incentives grow. As we shall see next, incentive mechanisms are not always complementary in nature—stronger incentives can, in some instances, actually cause worse outcomes in terms of program performance.

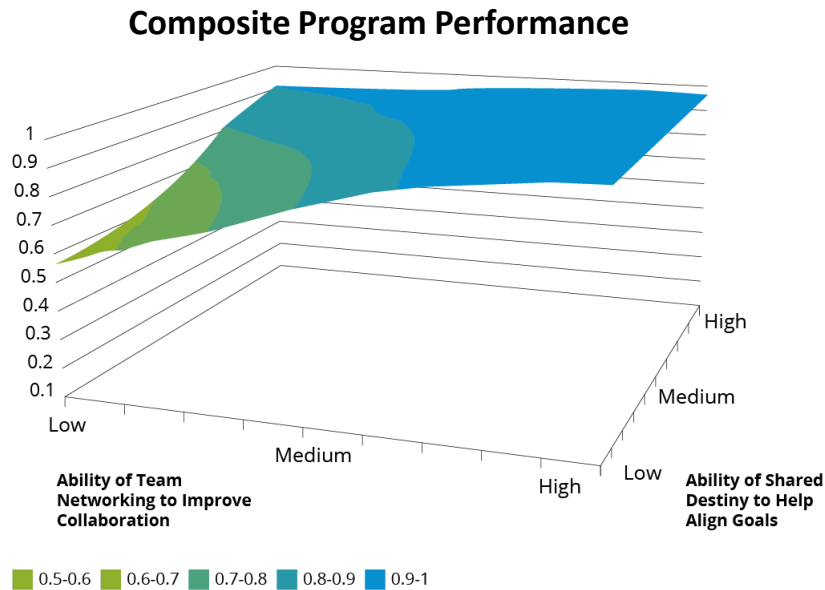


Figure 7: Complementary Effects of Team Networking and Shared Destiny

To some extent, TRIM and Shared Destiny work at different ends of the acquisition problem. TRIM motivates the CP to generate accurate estimates of the schedule, and to keep their development activities focused on meeting that schedule. Shared Destiny motivates meeting the needs of the larger multi-CP acquisition program, providing support for a CP’s collaborators by potentially pulling resources away from their own (self-interested) development activities. These incentives can also influence the extent to which a CP accepts integration decisions made by the PMO, especially if those decisions do not favor the CP’s own development efforts and local objectives.

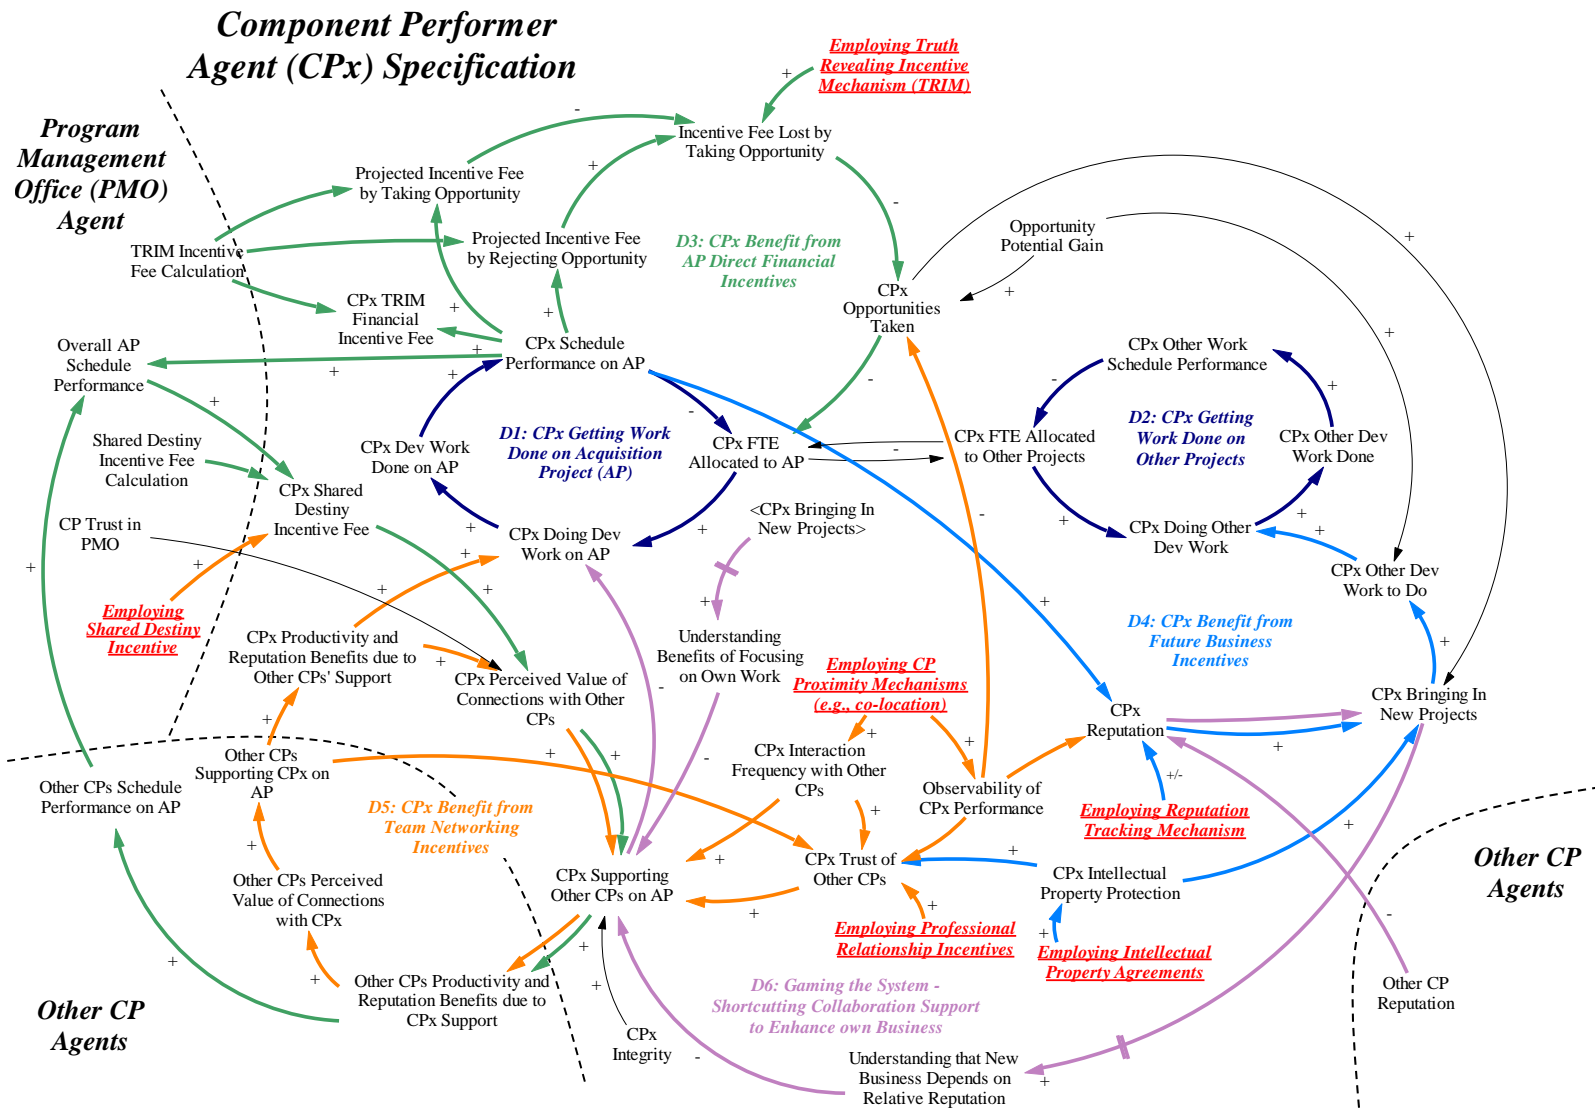


Figure 8: Specification of Component Performer Agent (CPx)

Figure 9 shows Composite Program Performance as the level of collaboration support and the level of acceptance of PMO integration decisions (each varies from low to high). The resulting shape is a gently-sloping ridge with a peak that is not at the endpoint. The bottom line is that improperly balanced TRIM and Shared Destiny incentives can result in too much support being provided to CP collaborators (drawing needed resources away from development) or blind acceptance of PMO integration decisions (when greater CP involvement in integration could produce a better overall system).

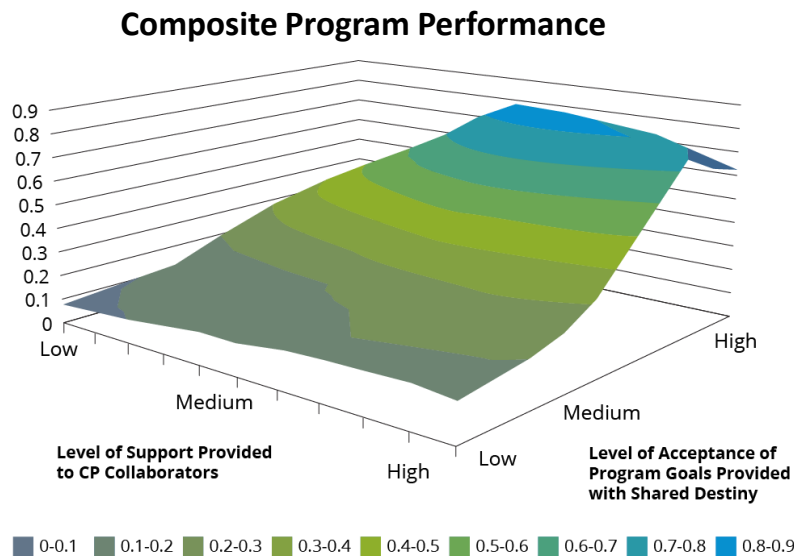


Figure 9: Counterbalancing Effects of TRIM and Shared Destiny

5. Future Vision of Acquisition Wargaming

An incentive is one weapon in an ongoing war where weapons must evolve. You must plan for the obsolescence of an incentive mechanism as the war escalates. Evolutionary Game Theory defines a framework of strategies and analytics in which Darwinian competition is modeled to allow analysis of the dynamic behavior in a population of acquisition programs in which this escalating conflict occurs among PMOs and their CPs.

It is widely recognized that acquisition programs suffer from recurring cost, schedule, and quality problems, wasting taxpayer dollars, often failing to deliver on-time or with full capability, and delaying or leaving the warfighter without needed capabilities. One key factor behind these issues is that acquisition program managers lack the data, evidence, and tools they need to make complex decisions associated with multi-dimensional acquisition strategy trade-offs, so that they can identify and evaluate acquisition solution options. One area of proposed acquisition research would be to create a virtual acquisition modeling laboratory to improve the quality of acquisition staff decision-making through the ability to analyze the consequences of acquisition decisions, and by designing and testing incentive mechanisms to improve acquisition outcomes over conventional approaches.

Acquisition PMOs historically manage the CP through oversight, gathering costly progress and status information without addressing the fundamental underlying fact that the CP inherently has more complete information about the effort, and as a result may not be acting in the program's interests. This is an instance of a cooperation problem which exists in many forms within acquisition programs, where decisions can be more focused on maximizing individual returns than on overall outcomes.

Hybrid modeling techniques combining agent-based and system dynamics modeling can model a game theory formalism to characterize the program's behavior. The laboratory could treat acquisition decision-making as an

optimization problem searching through the decision space, enabling better decisions with simulated “what if?” questions (e.g., “If A occurs, then in 3 months it will be N% more likely to see outcome X,” or “What is the impact of reducing staff by N%?”) related to program decision scenarios, and identifying potential tipping points where key shifts in cost, schedule, and performance occur.

The acquisition program model, after being parameterized to represent a program’s historical performance, would be validated using retrodiction model runs with no incentive mechanism, comparing it to historical performance data, and refining the model to accurately reflect the historical data.

Many incentive mechanisms for mitigating cooperation problems exist in the literature, and could be applied to acquisition contexts. The laboratory would include models of such incentive mechanisms that address the cooperation problem from the relevant acquisition scenario. Explicitly identified assumptions of these incentive mechanisms would be matched with the acquisition context characteristics.

The use of modular models would allow for incremental model development, creating components characterizing many different aspects of modeled acquisition programs that could be plugged into a model architecture to instantiate a model of a single acquisition program with a specific set of characteristics. This would let program-specific acquisition models be built more quickly from existing components that are developed over time, answering PMO questions sooner than older, more monolithic models could.

The objective would be to improve the quality of PMO decision-making by

1. enabling staff to make evidence-based decisions, rather than relying on past solutions and intuition
2. anticipating likely consequences of different decisions, using validated computational analysis
3. mitigating traditionally intractable cooperation problems among program entities/stakeholders that undermine productivity and hinder progress, by using verified mechanisms to align incentives

This technology provides a type of wargaming capability for acquisition programs, allowing them to study “what if” scenarios related to governance and program execution, to anticipate problematic behaviors in program execution for resolution, and generalize and extend acquisition modeling and simulation tools to other program contexts. These tools and methods could provide DoD with a virtual acquisition laboratory that could be used to simulate the likely response of programs to proposed changes in rules and governance, and thus test the efficacy of proposed policy approaches before their implementation.

This paper has analyzed the interactions of various mechanisms, revealing in the process that the properties of various mechanisms can reveal both compatibilities and counter-intuitive conflicts in their effects when they are combined. We have started to consider how studying the individual properties of mechanisms may be used to predict how well different mechanisms may work together. Believing that this could be done in a rigorous way, we hypothesize that a “mechanism calculus” could be developed as a future research topic, to further assist PMO decision-making in the choice of multiple incentive mechanisms, avoiding counter-productive combinations, and maximizing the intended beneficial effects.

Another area of future work is the characterization of CP preferences for cost, schedule, and quality trade-offs as what is known as a *Pareto surface*. A Pareto surface is a mathematical formalization of the trade-offs agents prefer to make. It represents the result of a mathematical optimization analysis of a problem in which more than one objective must be simultaneously satisfied, and trade-offs must be made among multiple conflicting objectives. CPs know their own Pareto surface (i.e., what trade-offs they would be willing to make, such as money received vs. effort required), but the government PMO does not. Eliciting a CP’s Pareto surface reveals ways (e.g., incentives) to influence their behavior.

6. Summary and Conclusion

The work described in this paper is part of an effort to equip acquisition program managers with more powerful tools to understand and control the behavior of large, complex development efforts. While this work is incomplete, there are some key insights that have been gained from this study of acquisition program analysis through simulation, and the application of a wider range of incentive mechanisms to this domain. Some of the insights from this work can be summarized as follows:

- Incentive mechanisms to promote trust can create a positive self-reinforcing dynamic—but poor relationships can cause this same feedback to be negative, resulting in a downward spiral that reflects the counter-productive dynamics seen in many actual acquisition programs.
- Specific incentive mechanisms can be used to achieve specific objectives (such as helping to keep CPs focused on program work)—but there is no single perfect incentive. They should be used in combination to maximize their positive effect on program performance.
- Different organizational roles are also influenced by different incentive mechanism types (i.e., future business incentives appeal to executives, direct financial incentives appeal to project management, and team networking incentives appeal to engineers and developers). The most effective incentives for a given organization will depend on that organization’s values and business priorities.
- Using a combination of different types of incentive mechanisms can achieve greater levels of influence on CP behavior across a range of organizations that employ different business models—especially when information on those preferences is incomplete.
- Using incentives in combination is complex and often non-intuitive. Certain types of incentive mechanisms can undermine the effectiveness of others, such as the way mechanisms promoting cooperation (e.g., Shared Destiny) can undermine those that improve cost and schedule performance (e.g., TRIM).
- Acquisition modeling and simulation could not only help predict the likely results of specific program decisions, but could also analyze the expected responses of acquisition programs to proposed policy and regulation changes, and evaluate their effectiveness.

Perhaps the most significant area of future work would be to provide a virtual acquisition laboratory by constructing a computational model, informed by causal modeling of the relationships, using an extensible, component-based acquisition program model architecture, and validated against historical program behavior data. This model could be used together with appropriate existing incentive mechanisms that could address the cooperation issues involved with the program under study. This capability could be a “wargaming” capability that models acquisition behaviors to help leaders avoid problems through better-informed decision-making. Such a capability could provide acquisition program leadership with a new level of insight that lets them look into the near future of their program’s performance to anticipate issues, make *evidence-based* decisions, and thus avoid serious problems.

Acknowledgments

The authors would like to thank the U.S. Department of Defense for providing funding for this work, and the Software Engineering Institute’s Director’s Office and Software Solutions Division for their help in encouraging and conducting this research.

Copyright 2018 Carnegie Mellon University. All Rights Reserved. This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. [DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution. DM18-0442

References

- Air Force Federal Acquisition Regulations System. (2013). *Informational guidance: Associate Contractor Agreements*.
- Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.
- Axelrod, R. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton University Press.
- Axelrod, R. (2006). *The evolution of cooperation* (Revised ed.). New York: Perseus Books Group.

- Brandenburger, A. M., & Nalebuff, B. J. (2011). *Co-opetition*. Crown Business.
- Buettner, D. (2014). *Acquisition games and software defects: Dynamics of the software developer's dilemma*. Lambert Academic Publishing.
- Coughlan, Peter J., & Gates, William. (2009). *Innovations in defense acquisition: asymmetric information and incentive contract design*. No. NPS-CM-09-127. Naval Postgraduate School, Monterey CA.
- Coughlan, Peter, Gates, William, & Lamping, Jennifer. (2008). *Innovations in defense acquisition auctions: Lessons learned and alternative mechanism designs*. No. NPS-AM-08-013. Naval Postgraduate School, Monterey CA.
- CPARS. (2017). *Guidance for the Contractor Performance Assessment Reporting System (CPARS)*. CPARS Program Office. <https://www.cpars.gov/pdfs/CPARS-Guidance.pdf>
- Dacus, C., & Hagel, S. (2014). A conceptual framework for defense acquisition makers: giving the schedule its due. *Defense ARJ*, 21(1), 486-504.
- Daidj, N. (2017). *Cooperation, coopetition and innovation*. John Wiley & Sons.
- DePillis, Lydia. Government did a poor job on HealthCare.gov. A private firm might've been worse. *Wonkblog*. October 24, 2013.
- Dixit, A. K., & Nalebuff, B. (2008). *The art of strategy: a game theorist's guide to success in business & life*. WW Norton & Company.
- Gansler, J., & Lucyshyn, W., & Spiers, A. (2009). The role of lead system integrator. *6th Annual Acquisition Research Symposium of the Naval Postgraduate School: Volume II: Defense acquisition in transition*.
- Leonard, R. (2010). *Von Neumann, Morgenstern, and the creation of game theory*. New York: Cambridge University Press.
- McAllister, D. (1995). Affect- and cognition-based trust as foundations for interpersonal cooperation in organizations. *Academy of Management Journal*, 38(1), 24-59.
- Moore, J. F. (1993). Predators and prey: a new ecology of competition. *Harvard Business Review*, 71(3), 75-86.
- Myerson, R. (2007). *Game theory*. Cambridge, Mass.: Harvard Univ. Press.
- Nalebuff, B., & Brandenburger, A. (1997). Co-opetition: Competitive and cooperative business strategies for the digital economy. *Strategy & Leadership*, 25(6), 28-33.
- Padovan, B., Sackmann, S., Eymann, T., & Pippow, I. (2001). A prototype for an agent-based secure electronic marketplace including reputation tracking mechanisms. In *Proceedings of the 34th Annual Hawaii international conference on system sciences*: IEEE.
- Peltoniemi, M., & Vuori, E. (2004). Business ecosystem as the new approach to complex adaptive business environments. In *Proceedings of eBusiness research forum* (Vol. 2, pp. 267-281).
- Pentland A. (2015). *Social physics: How social networks can make us smarter*. Penguin.
- Pentland A. (2012). The new science of building great teams. *Harvard Business Review*, 90(4), 60-9.
- Rendon, R. G., Huynh, T. V., & Osmundson, J. S. (2012). Contracting processes and structures for systems-of-systems acquisition. *Systems Engineering*, 15(4), 471-482.
- Senge, Peter M. (1990). *The fifth discipline: The art and practice of the learning organization*. New York: Doubleday/Currency,
- Smith, John M. (1982). *Evolution and the Theory of Games*. Cambridge University Press.
- United States DoD Inspector General (2017). *Summary of audits on assessing contractor performance: Additional guidance and system enhancements needed*.
- United States Government Accountability Office. (2006). *Defense acquisitions: Business case and business arrangements key for future combat system's success (GAO-06-478T)*.
- United States Government Accountability Office. (2017). *Defense acquisitions: Assessments of selected weapon programs (GAO-16-329SP)*.
- Ward, Mary Catherine, Elm, Joseph P.; & Kushner, Susan. (2006). *Techniques for developing an acquisition strategy by profiling software risks (CMU/SEI-2006-TR-002)*. Software Engineering Institute, Carnegie Mellon University.
- Young, S. (2010). Lead systems integrator role for government, presentation. *NDIA Systems Engineering Conference*.

Appendix: System Dynamics Modeling Notation

Figure 10 summarizes the notation used in our system dynamics model. The primary elements are variables of interest, stocks (which represent collection points of resources), and flows (which represent the transition of resources between stocks). Signed arrows represent causal relationships, where the sign indicates how the variable at the arrow's source influences the variable at the arrow's target. A positive (+) influence indicates that the values of the variables move in the same direction, whereas a negative (-) influence indicates that they move in opposite directions. A connected group of variables, stocks, and flows can create a path that is referred to as a feedback loop. There are two types of feedback loops: balancing and reinforcing. The type of feedback loop is determined by counting the number of negative influences along the path of the loop. An odd number of negative influences indicates a balancing loop; an even (or zero) number of negative influences indicates a reinforcing loop.

Significant feedback loops identified within the model described here are indicated by a loop symbol and a loop name in italics. Balancing loops—indicated with the label B followed by an identifying number in the loop symbol—describe aspects of the system that oppose change, seeking to drive variables to some equilibrium goal state. Balancing loops often represent actions that an organization takes to manage, or mitigate a problem. Reinforcing loops—indicated with a label R followed by a number in the loop symbol—describe system aspects that tend to drive variable values consistently either upward or downward. Reinforcing loops often represent the escalation of problems, but may include problem mitigation behaviors.

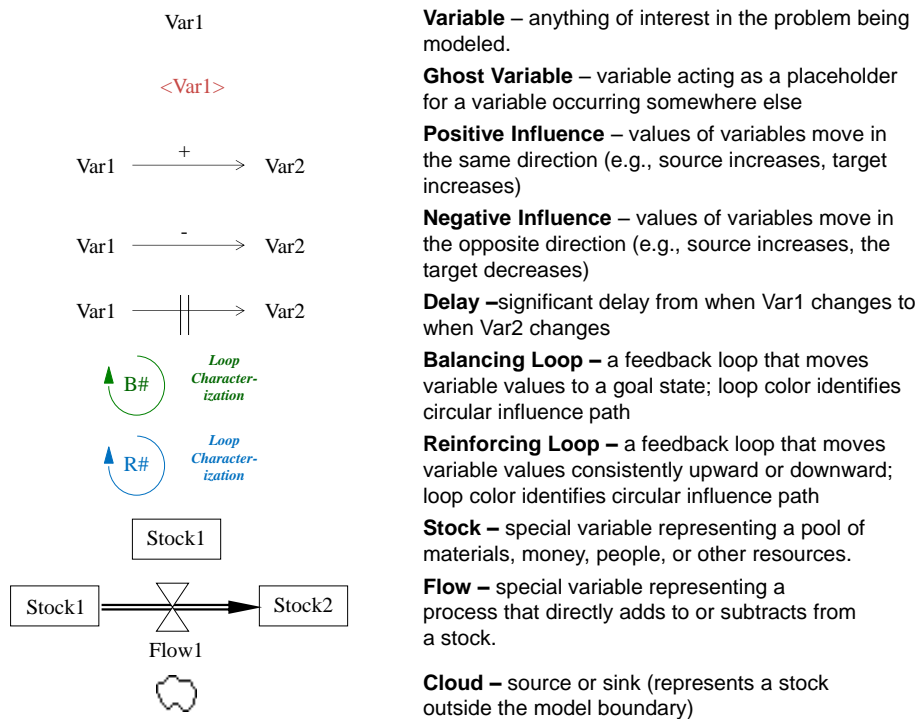


Figure 10: System Dynamics Notation