Lawrence Livermore National Laboratory

Decision Trees for Analysis of Strategies to Deter Limited Nuclear Use in a Generic Scenario

Thomas Edmunds*, Jeffrey Garrett, and Craig Wuest

LLNL-TR-759558

September 2018

*edmunds@llnl.gov, (925) 423-8982

Disclaimer

Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344.

Executive Summary

We assert that analysis of a decision tree model can reveal key insights for deterring limited nuclear use in a conventional conflict. To illustrate the method, we develop a decision analysis model to explore the choices an attacker would make and how they can be influenced by his perceptions of the defender's capabilities and potential responses. Sensitivity analysis with the model can provide insights into strategies and actions the defender might take in order to influence the choices the attacker would make. This methodology is widely applicable and could provide insight into a variety of scenarios.

Basic and simplistic nuclear deterrence states that typically one nation would be deterred from attacking another if the defender possesses nuclear weapons. However, recent developments in the strategic landscape and a rise in tensions among some nuclear-armed states suggests some nations could be emboldened to take aggressive actions if they possess nuclear weapons, possibly even against a nuclear-armed adversary.

We illustrate the method with a multi-stage decision tree model of a hypothetical conflict between two generic nuclear-armed powers, Red and Blue. Red is the decision-maker in the tree, and its decisions and perceptions about how Blue might respond are represented by decision and chance nodes. Red objectives are represented by a multiattribute utility function that was developed with input from subject matter experts using standard elicitation techniques. The quantitative representation of Red objectives that determines decisions it would make under different circumstances allows us to do extensive sensitivity analyses with respect to many unknowns.

The general structure of the decision tree is:

1. Red initial decision to mobilize and observe Blue reaction
2. Blue may or may not mobilize (Red assessment of probabilities)
3. Red decision to attack or not
4. Blue counter-attack or capitulate (Red assessment of probabilities)
5. Red attack may stall or overrun Blue (chance event)
    a. If Red stalled, Red chooses to commit more conventional forces, use a tactical nuclear weapon, or be deterred
    b. If Red overruns, Blue may commit more conventional forces, use a tactical nuclear weapon, or capitulate (Red assessment of probabilities)
6. Tree ends with Red success, a choice to risk nuclear escalation, or be deterred

Examples of insights derived from exercising the model include:

- Red would be more likely to resort to the use of nuclear weapons if its assault is stalled.
- Red would be more likely to use nuclear weapons later in the conflict after Blue has demonstrated resolve by committing additional conventional forces.
- A Blue "no first use" policy or if Blue had no nuclear weapons at all would make Red more aggressive later in a conflict.

[This page intentionally left blank.]

Table of Contents

[This page intentionally left blank]

1 Introduction

During the cold war, strategic deterrence was achieved primarily through the principle of mutually assured destruction. The U.S. and Soviet Union had confidence that if one side were to launch an attack first, the other side would have enough surviving nuclear weapons to launch a devastating counterstrike. Neither side desired such an outcome.

Today, however, there is increasing concern over the likelihood of limited nuclear use in a conflict between nuclear powers. The 2018 U.S. Nuclear Posture Review (NPR) indicates the U.S. is concerned that "Russia may also rely on threats of limited nuclear first use, or actual first use, to coerce us, our allies, and partners into terminating a conflict on terms favorable to Russia." Further, "Moscow apparently believes that the United States is unwilling to respond to Russian employment of tactical nuclear weapons with strategic nuclear weapons." [OSD 2018]

Deterrence strategies have been extensively studied by military planners and academicians. Studies conducted by military planners generally include highly-detailed logistical and effectiveness models that are built and maintained by a large community of defense contractors. These models tend to be customized, propriety, and difficult for the community to access and use. At the other extreme, academic researchers tend to produce very high level, analytic models with simple solutions that can be communicated in a short journal article. In the following sections we describe modeling and analysis conducted at an intermediate level of resolution between these two extremes.

In addition, most nuclear deterrence modeling has been focused on cold-war era concepts, like large-scale strategic exchanges and mutually assured destruction. Less attention has been paid to modeling nuclear deterrence in the context of limited nuclear use in a regional conflict, a scenario that is much more likely today than a large-scale exchange. [Kent 2008, Oelrich 1988, Manzo 2018, Kuhn 2018]

In this report, we describe a decision tree model to assess the impacts of non-strategic weapons on regional deterrence in a conflict between two generic nuclear powers, Red and Blue. Red has undertaken a campaign of smaller scale conflicts to add resources and land or to achieve hegemony over neighboring countries. The conflict is over a Red incursion into the territory of an ally of Blue. The model includes a multiattribute utility function to represent Red preferences over possible outcomes, Red's uncertainty with respect to effectiveness of weapons systems, and Red's assessment of Blue's possible actions as the conflict unfolds. The models are used to find multi-dimensional surfaces in the space of Blue's possible decisions and weapon design parameters where Red switches strategies (e.g., commit more conventional troops, launch a tactical nuclear weapon at a military target, or withdraw from the conflict).

2 Political and technical/military background of Red-Blue conflict

The conflict we modeled is a Red incursion into an ally of Blue to annex land. In this hypothetical conflict, Red has tactical nuclear weapons that could be used against massed troops or logistical centers without incurring large numbers of civilian casualties. Blue also has tactical nuclear weapons and conventional superiority in many cases.

The motivation for the hypothetical conflict is general unrest in the Red populace due to economic stagnation and the perceived loss of prestige and influence in the global community. The government is attempting to unify the population by declaring a common enemy that is persecuting ethnic Red populations in nearby countries. In this scenario, Red claims that action must be taken to defend its allegedly persecuted people.

In general, because of its shared borders Red has easy access to Blue's ally. In contrast most of Blue's military forces that would be deployed to oppose Red forces are not within the ally's borders. Weeks would be required to transport the troops and equipment needed to drive Red forces out of captured territories. Finally, we note that in aggregate the Blue conventional forces are much larger than Red forces. If Red is to succeed, it must act quickly to occupy land and threaten devastating losses on Blue forces attempting to dislodge it.

Blue is assumed to have a portfolio of tactical nuclear options that covers a wide variety of capabilities. A key question of the analysis is which weapons would be best suited to deter Red aggression, and in particular Red's use of its nuclear weapons. One factor is the effectiveness of Red's integrated air defense system against Blue's aircraft and cruise missiles. In addition, different Blue weapons will produce different military losses. Both Red and Blue would be tempted to use nuclear weapons against chokepoints in the other's logistics support.

Finally, Red's situational awareness is not perfect. To account for this when planning, Red assigns probabilities to outcomes of individual engagements in the campaign and how these probabilities are influenced by the commitment of additional conventional forces or using nuclear weapons. Red also assigns probabilities to possible Blue actions at various points in the conflict.

3 Decision tree models for offensive/defensive scenarios

In this section we will provide a brief description of decision trees, a simple example of an application, and the type of insight that can be derived from decision tree analysis. Previous work has shown that decision trees are effective tools for analyzing offensive/defensive scenarios, such as an adaptive adversary attempting to deliver a weapon to target [Maurer 2009, Powers 2010]. Simple algorithms are used to find the best sequence of decisions the attacker can make given the available options at each point during the attack considering the uncertainties about offensive and defensive weapon system performance. In general, decision trees can have an arbitrary number of nodes representing decisions, uncertainties, and the values that attacker and defender assign to possible outcomes. Solution algorithms scale well so trees with thousands of branches can be solved in less than one minute with commercial software.

A simple example of an attacker's decision tree is shown in **Figure 1**. In this decision tree, an attacker is deciding which launch platform to use (air or sea), and which target to strike (1, 2 or 3). In the figure, square blue nodes correspond to attacker decisions, green circles are chance events, and red triangles are payoffs to the attacker. As indicated on the left-hand side of the figure, the first attacker decision is whether to launch the weapon from an air or sea platform. The probability of penetrating the defender's border depends upon which launch platform is chosen. In the example, the probability of success is 0.7 for an air launch and 0.5 for a sea launch, and these probabilities are controlled by the adversary's choice of defensive systems. If the border is penetrated successfully, one of three potential targets must be selected. As shown, the probability of success varies by target. The payoff nodes in this example are binary, but the software can accommodate an attacker multiattribute value function. The decisions and targets in the lower part of the tree corresponding to a sea launch are similar in this example, but the probabilities are different. Once the probabilities at the chance nodes and payoffs at the terminal nodes have been set, a solution algorithm identifies the best decision the attacker can make at each point in the tree.

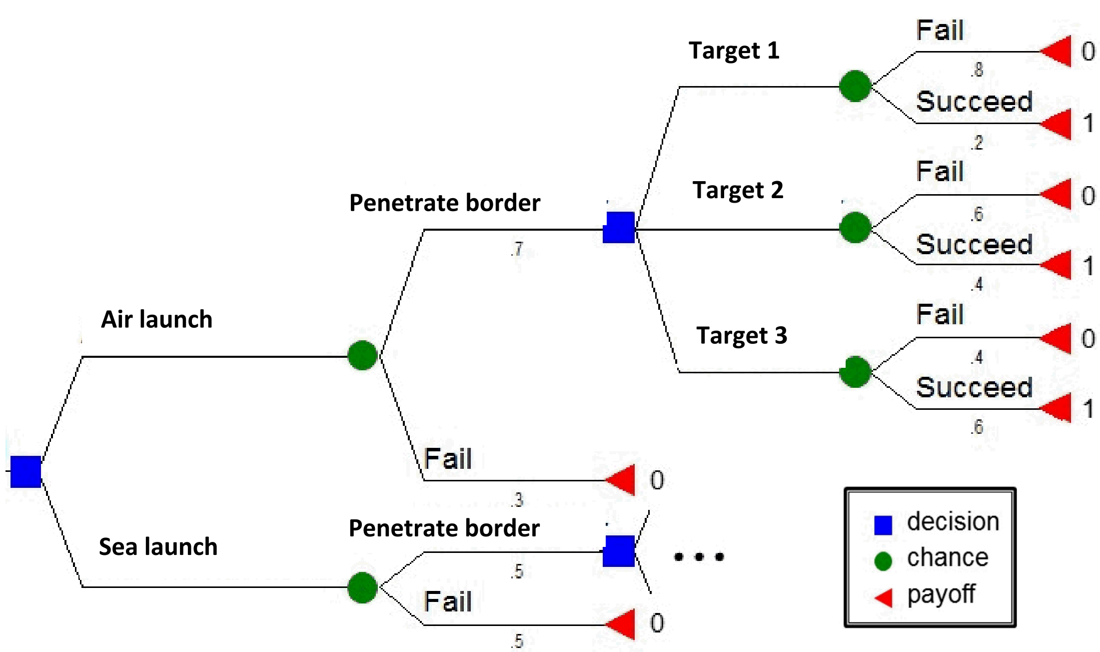

Figure 1 Decision tree for attack strategy analysis

Once the tree is complete and solvable, one can conduct sensitivity analyses. In this example, a parametric analysis with respect to the probabilities can reveal how much defense is necessary to cause the attacker to switch strategies. Results of a parametric analysis to determine the optimal target (i.e. Target 1, 2 or 3), where the tree was solved multiple times while varying the probabilities of penetrating the border from air and sea platforms, is shown in **Figure 2.** As indicated in **Figure 2**, if defenses against a sea launch are more effective than against an air launch, target 2 is the optimal choice. If defenses against both air and sea launch are highly effective, the attacker is driven toward target 1. If target 1 corresponds to a decision not to attack at all, the diagram shows what defensive performance is needed to achieve deterrence. In the following sections we describe in detail the model we built, how we populated it with data, and the analyses we performed on it.

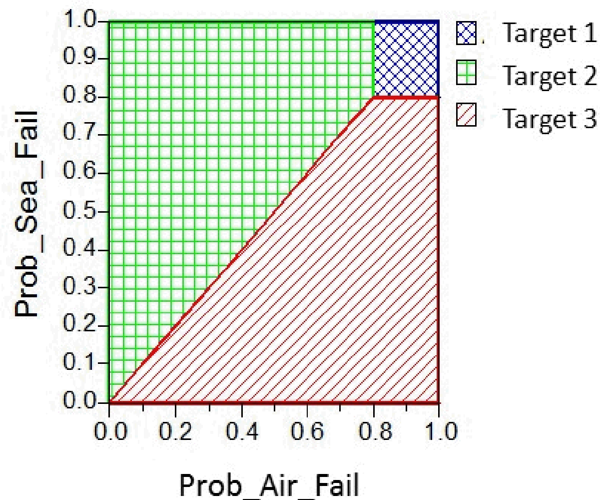

Figure 2 Optimal attacker strategies for different defense effectiveness levels

4 Application of decision tree to Red-Blue conflict

In this section, we describe a decision tree model of Red decisions, Red assumptions about the likelihood of various Blue responses, and the outcomes of chance events. *Red's overall objective is to limit the expected value of losses of credibility and military personnel and to avoid risk of escalation to a strategic nuclear exchange.* This model of Red behavior can provide insight on how Blue can configure its portfolio of tactical nuclear weapons, logistics, and other military capabilities in order to influence that behavior. Sensitivity analysis over the probabilities that Red assigns to Blue responses will help identify effective Blue strategies. Sensitivity analysis over quantitative models of Red objectives will help identify robust Blue strategies.

The decision tree presented here is a simplified model of a complex conflict with many stages. The general structure of the decision tree is:

1. Red initial decision to mobilize and observe Blue reaction
2. Blue may or may not mobilize (Red assessment of probabilities)
3. Red decision to attack or not
4. Blue counter-attack or capitulate (Red assessment of probabilities)
5. Red attack may stall or overrun Blue (chance event)
    a. If Red stalled, Red chooses to commit more conventional forces, use a tactical nuclear weapon, or be deterred
    b. If Red overruns, Blue may commit more conventional forces, use a tactical nuclear weapon, or capitulate (Red assessment of probabilities)
6. Tree ends with Red success, a choice to risk escalation, or be deterred

We made several simplifying assumptions for our intermediate resolution model. First, we limited each side to only one additional deployment of conventional forces and one use of a tactical nuclear weapon. This was done to manage both the size and complexity of the model and to reflect a conflict that would only last a few weeks before final resolution (mobilization and each additional deployment of conventional forces would take approximately a week). Allowing additional rounds of force deployments or nuclear weapon use would expand the model exponentially and greatly increase uncertainty in the results. Similarly, we did not allow either side to deploy additional conventional forces if it already chose to use a nuclear weapon, since we found this scenario to be unlikely. Finally, Blue's nuclear weapons were differentiated by the probability of their use, the probability that they would get to target, and binned in one of two yield classes (high or low).

The first stages of the decision tree are shown in **Figure 3** (a screen capture of the modeling software interface). The Red decision to mobilize or not is shown at the far left of the figure. In the model Red will make this decision by comparing the expected losses of mobilization (upper branch) with the status quo (lower branch) and choose the one with the minimum value. As shown in the upper right, Red estimates the probabilities of a full, partial, or no mobilization response from Blue at 0.1, 0.6, and 0.3, respectively. The data in the decision tree shown here were provided by an interview with Jacek Durkalec [Durkalec 2018], an expert in extended nuclear deterrence policy currently at the Center for Global Security Research at Lawrence Livermore National Laboratory. Dr. Durkalec's assessments are based upon some of his recent research documented in [Durkalec 2018a].

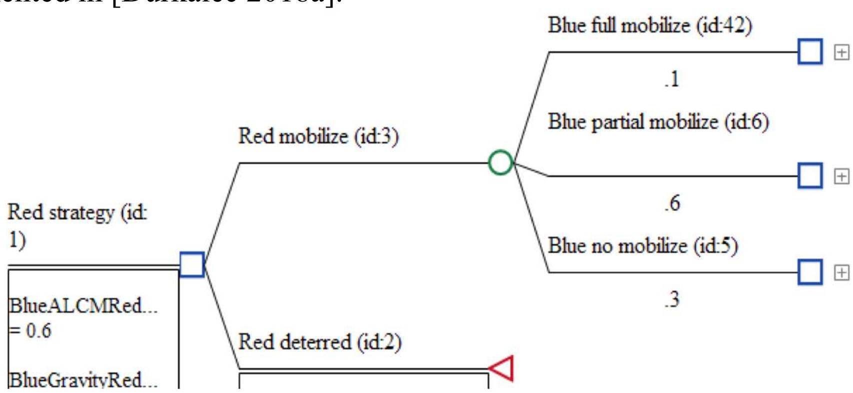

Figure 3 Red and Blue mobilization stages of decision tree

Red decisions and Blue responses assuming full Blue mobilization are shown in Figure 4. As indicated in the figure, after Red observes full Blue mobilization, it can choose to attack or be deterred, and realize a payoff at this point. An attack outcome is uncertain, with probability 0.3 that Red will overrun Blue forces and probability 0.7 that the attack will be stalled before reaching its objective. If the attack overruns Blue forces, Blue responses could be to use a tactical nuclear weapon (top branch), to commit more conventional forces, or to capitulate. Red assigns probabilities of 0.01, 0.6, and 0.39 to these actions, respectively. If the attack is stalled, Red makes a choice among these same three alternatives as shown at the bottom of the figure.

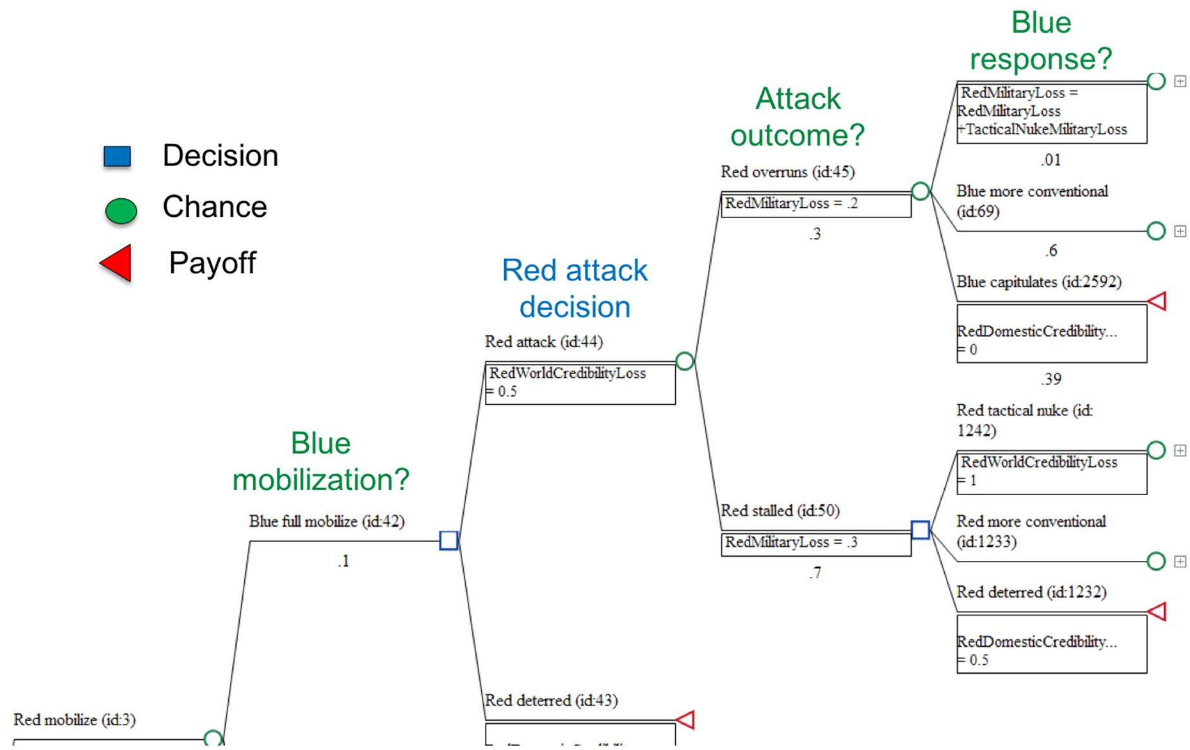

Figure 4 Red attack and Blue response stages of decision tree

As indicated in Figure 5, Red estimates a 0.01 probability that Blue would response with a nuclear weapon at this stage of the conflict. Red's assessment of Blue's possible choice among available tactical nuclear weapons is shown at the next stage in the figure. The highest probability is assigned to the use of a sea-launched cruise missile (SLCM), followed by an air launched cruise missile (ALCM), and a low-yield submarine launched ballistic missile (SLBM-LY). The probability of Blue use of an intercontinental ballistic missile (ICBM) is 0.005.

There is a much higher probability that Blue may instead choose to commit more conventional forces. As shown in the lower part of the figure, Red may still overrun Blue forces. At this point, Blue will either choose to launch a tactical nuclear weapon or capitulate. Note that at this later stage of the conflict Red assesses a higher probability that Blue will resort to the use of nuclear weapons.

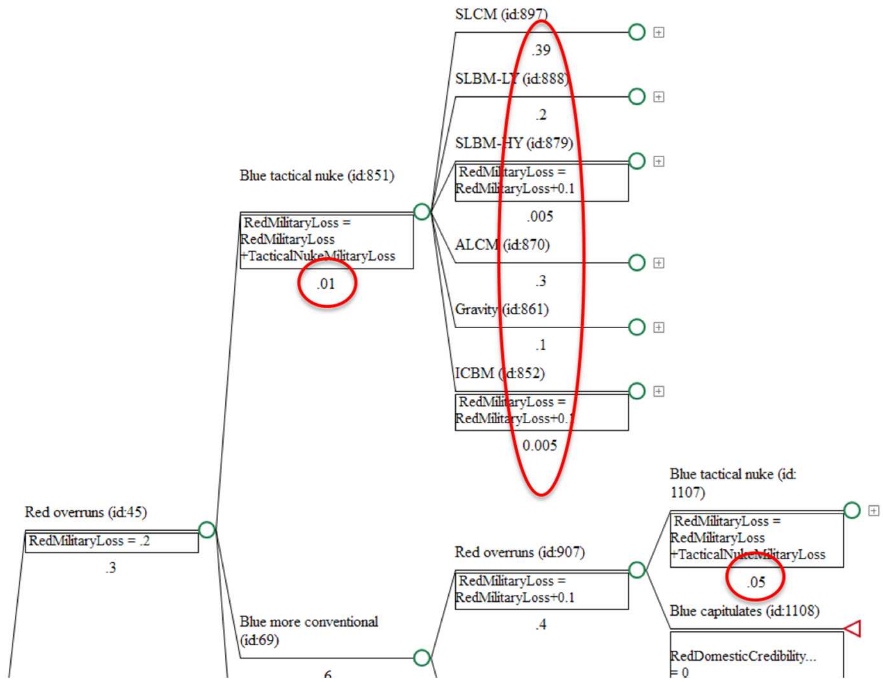

Figure 5 Red assessment of Blue's choice of nuclear weapon response

Figure 6 shows outcomes after Blue launches a nuclear weapon. As shown on the far right, if Red is stalled a second time after using a nuclear weapon, it may choose to risk escalation of the conflict by firing a second nuclear weapon.

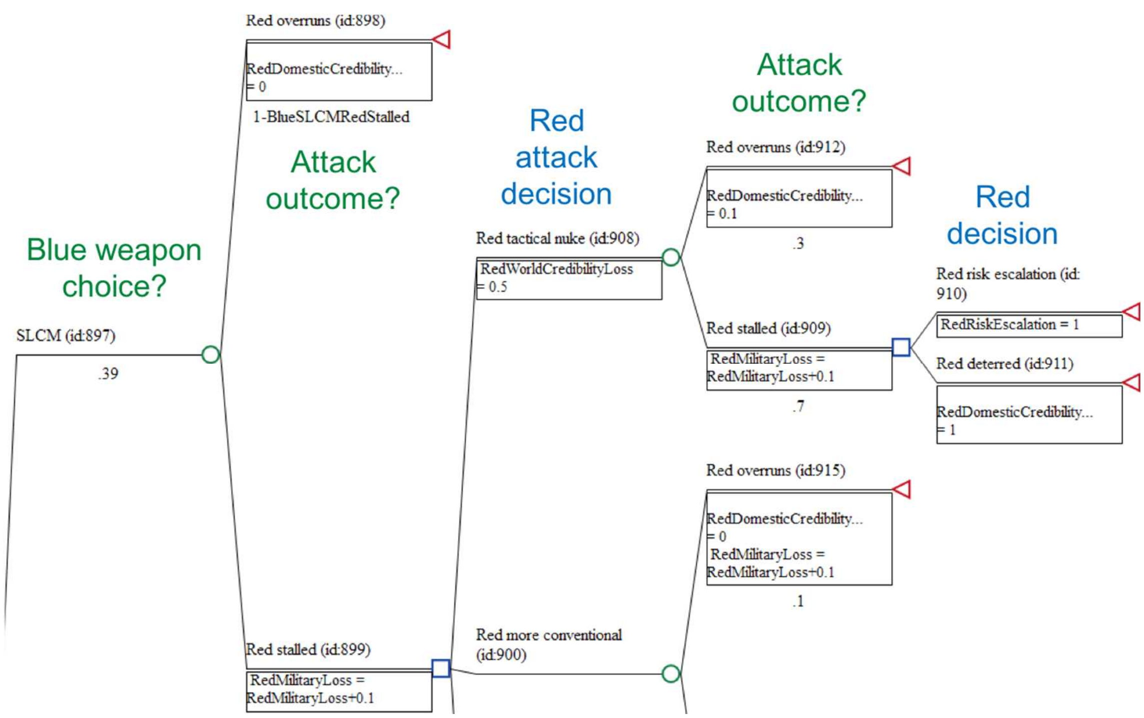

Figure 6 Red decision to escalate or be deterred

5 Metrics to quantify Red preferences

The probabilities in the decision tree model previously described are used to compute the expected value of choices throughout the tree in order to identify the optimal decisions at each point. The expected value computation uses numerical values for outcomes at each of the payoff nodes (red triangles).

A multiattribute utility function is used to represent Red's preferences for outcomes at each of the payoff nodes. This function consists of four attributes:

- Military casualties
- Domestic credibility loss
- World credibility loss
- Risk of nuclear escalation

A scale is constructed to measure outcomes in terms of each of these four components. The scale for military casualties is normalized for the range noted so that 100,000 casualties corresponds to a military loss of 1.0. The scale is assumed to be linear. The scale for risk of nuclear escalation is binary – Red either risks nuclear escalation by firing a second nuclear weapon or it is deterred – and is also normalized. Constructed scales with descriptive levels are used for domestic and world credibility losses.

The constructed scale for Red preferences with respect to domestic credibility loss as elicited from Dr. Durkalec is shown in Table 1. Each payoff node in the decision tree can be mapped into one of the levels shown in the table. In general, Red suffers the highest losses in domestic credibility when it attacks and is then deterred from further aggression after being repulsed by Blue forces. Its domestic credibility losses are zero if Blue capitulates after Red attacks or simply mobilizes for an attack. Losses are low (below 0.08) if Blue capitulates under any circumstances. Note that Red's domestic credibility loss is initially 0.3, which is the motivation for the aggression.

Table 1. Constructed scale for Red preferences for domestic credibility loss

Value (loss)	Description (mapped to decision tree)
1.0	Red capitulates after Blue reinforces conventionally and then uses a nuke
0.8	Red capitulates after Blue reinforces conventionally
0.8	Red capitulates after Blue uses a nuke, but hasn't reinforced conventionally
0.7	Red capitulates after Blue counterattacks once
0.3	Red chooses not to mobilize
0.08	Blue capitulates after Red uses a nuke, but hasn't reinforced conventionally
0.05	Blue capitulates after Red reinforces conventionally and then uses a nuke
0.05	Blue capitulates after Red reinforces conventionally
0	Blue capitulates after Red attacks once
0	Blue capitulates before direct conflict

The constructed scale for Red preferences with respect to world credibility loss is shown in Table 2. As shown in the table, Red suffers the highest loss in world credibility if it fires a nuclear weapon and Blue does not and world credibility loss is zero if Red does not mobilize.

Table 2. Constructed scale for Red preferences for world credibility loss

Value (loss)	Description (mapped to decision tree)
1.0	Red goes nuclear, Blue does not
0.9	Red goes nuclear first, Blue responds nuclear
0.7	Red reinforces conventionally, Blue does not
0.6	Both sides reinforce conventionally, neither goes nuke
0.6	Blue reinforces conventionally, Red does not
0.2	Blue goes nuclear, Red does not
0	Red does not mobilize (best Red outcome)

An overall figure of merit is needed to unambiguously determine which decision is best at a given node in the decision tree. This figure of merit is constructed by forming a weighted sum of the four attributes described above and multiplying by the probability that the outcome is achieved. The procedures for eliciting preferences and constructing these multiattribute utility functions are well established [Clemen 1996].

Tradeoff questions that compare two hypothetical alternatives that differ with respect to two of the attributes are used elicit preferences from decision makers. Responses to the tradeoff questions are used to compute the weights in the multiattribute utility function. In this process, the elicitor asks the decision maker which of the two alternatives is preferred. The elicitor then changes the attributes of one of the alternatives based upon the response in an effort to find two alternatives to which the decision maker is indifferent. This indifference point specifies one of the equations that can be used to find the weights. To find the four weights for our four attributes, three equations comparing three pairs of attributes are needed in addition to a normalization equation requiring the weights to sum to 1.0.

The results of the tradeoff elicitation comparing loss of Red domestic credibility and military losses are shown in Table 3. For this elicitation, the surrogate decision maker [Durkalec 2018] stated that he was indifferent between hypothetical Alternative A, which had a moderate loss in domestic credibility and zero military casualties, and Alternative B, which had a zero loss in domestic credibility and 12,500 Red military casualties. This indifference point establishes the following equations relating the weight of Red domestic credibility loss, W_{rdc}, to the weight of Red military losses, W_{rml}:

$$0.5W_{rdc} + 0W_{rml} = 0W_{rdc} + 12.5/100\ W_{rml}$$
$$W_{rdc} / W_{rml} = 0.25$$

where Red military losses have been normalized with respect to the maximum value of 100,000.

Table 3. Tradeoffs between Red domestic credibility and military losses

Alternative A

Value (loss)	Outcomes
0.5	Not Worst Red Domestic Credibility outcome (e.g., Some challenge to regime but change unlikely)
0	Red military loss (k)

Alternative B

Value (loss)	Outcomes
0	Best Red Domestic Credibility outcome (e.g., Overwhelming support for Red regime)
__12.5___	Red military loss (k)

The results of the tradeoff elicitation comparing loss of Red world credibility and military losses are shown in Table 4. For this elicitation, the surrogate decision maker was indifferent between hypothetical Alternative A, which had the largest loss in world credibility and zero military casualties, and Alternative B, which had a zero loss in world credibility and 20,000 Red military casualties. This indifference point establishes the following equations relating the weight of Red world credibility loss, W_{rwc}, to the weight of Red military losses, W_{rml}:

$$1.0W_{rwc} + 0W_{rml} = 0W_{rwc} + 20/100\ W_{rml}$$
$$W_{rwc} / W_{rml} = 0.2$$

where Red military losses have been normalized with respect to the maximum value of 100,000.

Table 4. Tradeoffs between Red world credibility and military losses

Alternative A

Value (loss)	Outcomes
1.0	Worst Red World Credibility outcome (e.g., Global pariah status)
0	Red military loss (k)

Alternative B

Value (loss)	Outcomes
0	Best Red World Credibility outcome (e.g., Red viewed overwhelmingly favorable by world)
__20___	Red military loss (k)

Finally, the results of the tradeoff elicitation comparing Red risking nuclear escalation and military losses are shown in Table 5. For this elicitation, the surrogate decision maker was indifferent between hypothetical Alternative A, which risked nuclear escalation and had zero military casualties, and Alternative B, which did not risk escalation and had 30,000 Red military casualties. This indifference point establishes the following equations relating the weight of Red risk of escalation, W_{rre}, to the weight of Red military losses, W_{rml}:

$$1.0W_{rre} + 0W_{rml} = 0W_{rre} + 30/100\ W_{rml}$$
$$W_{rre} / W_{rml} = 0.3$$

where Red military losses have been normalized with respect to the maximum value of 100,000.

Table 5. Tradeoffs between Red nuclear escalation risk and military losses

Alternative A

Value (loss)	Outcomes
1.0	Worst Red Risk Escalation outcome (e.g., Red launches second nuclear weapon risking all out nuclear war)
0	Red military loss (k)

Alternative B

Value (loss)	Outcomes
0	Best Red Risk Escalation outcome (e.g., Red deterred and therefore does not risk all out nuclear war)
__30___	Red military loss (k)

Invoking the weight normalization condition yields the following equations to compute the four weights:

$$W_{rdc} + W_{rwc} + W_{rre} + W_{rml} = 1.0$$

$$W_{rdc}/W_{rml} + W_{rwc}/W_{rml} + W_{rre}/W_{rml} + 1 = 1/\ W_{rml}$$

$$0.25 + 0.2 + 0.3 + 1 = 1/\ W_{rml}$$

$$W_{rml} = 0.571$$

$$W_{rdc} = 0.25 \times 0.571 = 0.143$$

$$W_{rwc} = 0.2 \times 0.571 = 0.114$$

$$W_{rre} = 0.3 \times 0.571 = 0.171$$

Hence, the multiattribute utility function at the payoff nodes in the decision tree is the following expression:

$$\text{Payoff} = 0.143\ \mathit{DomesticCredibilityLoss} + 0.114\ \mathit{WorldCredibilityLoss} + 0.171\ \mathit{RedRiskEscalation} + 0.571\ \mathit{MilitaryLoss}/100{,}000$$

6 Analysis and example results

The primary purpose of building the decision tree is to conduct sensitivity analyses to identify actions that Blue can take to reduce the risk of nuclear escalation. These Blue actions indirectly influence the risk through Red's perceptions of Blue's capabilities, intent, and likely actions.

There are eighty variables in the decision tree model representing probabilities, changes in values for the four metrics previously described, weights, and other factors. The complete list of variables is included in Appendix A. Sensitivity analyses can be performed on any one or combination of these variables. The following figures display results of three-way sensitivity analyses. The space explored can be viewed as a data cube in three dimensions. In the figures we display planes in the cube in which one of the variables is fixed. Subsequent planar views describe how results change throughout the data cube.

The sensitivity analysis reveals combinations of variable values at a decision node that cause Red to choose one strategy over another. These results provide insight into how Blue can change Red behavior to avoid the risk of escalation.

The first sensitivity analysis characterizes how the weights Red places on loss of domestic credibility, military casualties, and the risk of nuclear escalation affects its choice to commit more conventional forces, use a nuclear weapon, or be deterred from further aggression. In this case, Blue has partially mobilized, and Red's attack has stalled. For this bottom plane of the data cube, the weight on Red risking escalation is zero. Results are shown in Figure 7. As indicated in the figure, there is a range of weights Red might place on objectives that would lead it to using a tactical nuclear weapon.

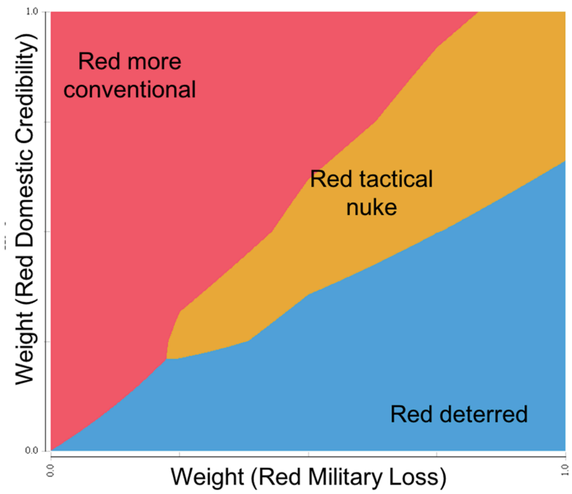

Figure 7. Red strategy choice for different attribute weights (escalation risk weight = 0)

The sensitivity analysis was repeated with the weight assigned to escalation risk increased to 0.25. At this higher level of importance, Red does not resort to the use of tactical nuclear weapons, as shown in Figure 8.

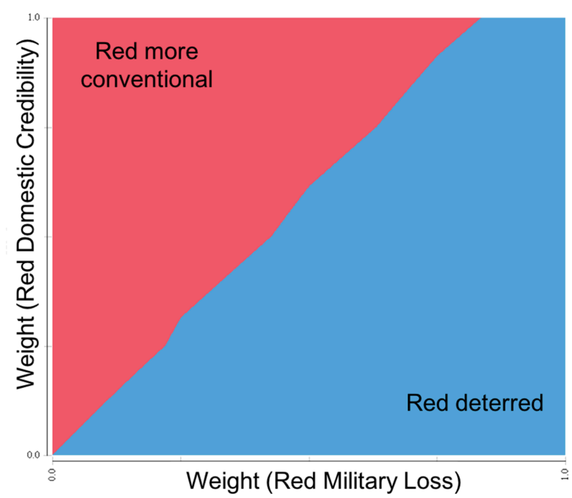

Figure 8. Red choice for different attribute weights (escalation risk weight = 0.25)

A Red setback later in the conflict may increase Red's propensity to resort to nuclear weapons. In the case shown in Figure 9, Red's intial attack has overrun Blue forces. However, Blue has demonstrated resolve by committing more convetional forces and the Red offensive has stalled. In this situation, Red is more likely to use a nuclear weapon to turn the tide (the orange region in Figure 9 is larger than the orange region in Figure 7.)

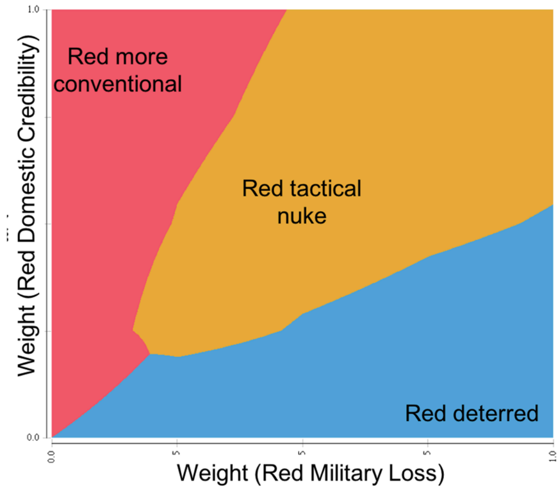

Figure 9. Red strategy choice after more Blue forces (escalation risk weight = 0)

Red's strategy choices when the weight on escalation risk is increased to 0.25 are shown in Figure 10. Note that there is no choice to use a nuclear weapon and the region where deterrence is chosen is much larger.

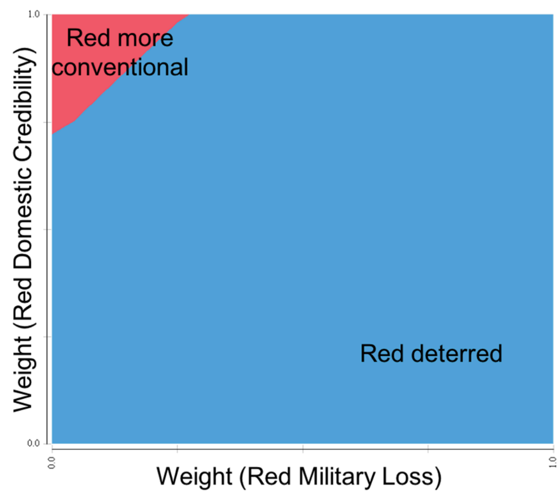

Figure 10. Red strategy choice after more Blue forces (escalation risk weight = 0.25)

By running multiple sensitivity analyses we can compare how different Blue options affects Red's decision space. In the following sections we will introduce our method for making this comparison, and the results of a few case studies.

7 Computing decision space volumes to compare strategies

In the previous section we showed how sensitivity analyses with respect to three weights on terms in the Red objective function[1] produce three-dimensional data cubes of Red's optimal decisions at a particular decision node in the model. In this section, we describe a method to summarize how different Blue strategies or weapon systems might affect Red's decisions by computing the volume of the data cube that is occupied by each of the Red strategy choices (e.g., more conventional, tactical nuclear weapon use, or be deterred).

We calculate the approximate volume of each strategy by counting the number of points in the discretized cube that corresponds to that strategy. A sample of a spreadsheet that implements this method is shown in Appendix B.

For example, we calculated the volume of Red's optimal choices at a particular point in the tree. At this point in the conflict, Blue has responded to Red's aggression with a partial mobilization, Red's initial attack has succeeded, and Blue has responded by deploying additional conventional forces which stop the Red offensive. Red then has to choose whether to deploy additional conventional forces, to use a nuclear weapon, or to be deterred and retreat to their borders. Figure 11 shows the volume of Red's choices. Note the decision space for using a nuclear weapon is so small it is nearly imperceptible – occupying approximately 0.2% of the cube in this example.

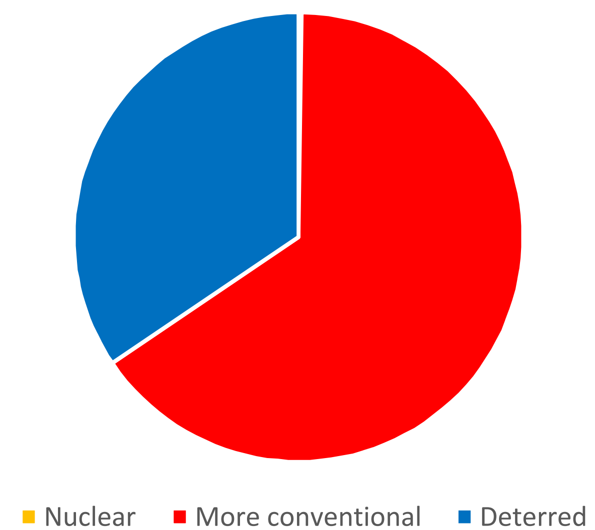

Figure 11. Volume of data cube for Red objective weights where each strategy is optimal

With this analysis model, we can make changes to probabilities in the decision tree to model a new Blue strategy and repeat the calculation of the volume associated with each Red decision. By comparing the decision volumes for two different Blue strategies, we can determine which Blue strategies induce desirable changes in Red's decisions. This analysis is discussed in detail in the next section.

[1] We consider these weights to be the most speculative parameters in the model. Accordingly, we present results aggregated over the entire range of possible values they may have.

8 Case studies using volume analysis

Using the volume analysis technique described in the previous section, we created several new models representing alternative strategies Blue could employ to investigate how they might affect Red's decisions. While some of these strategies may be considered extreme and are unlikely to be implemented in the real world, they do demonstrate the power of this analysis and how it could be applied to different strategies or entirely different scenarios.

For each new strategy we ran the sensitivity analysis at the decision node mentioned in the previous section, and at a node much earlier in the conflict where Blue has partially mobilized, and Red must decide whether to attack. To implement the new strategy, we created a copy of the tree where only the variables necessary for the new strategy were changed, e.g. in the case of a "no first use" policy the probability of Blue using a nuclear weapon first was changed to zero.

The sensitivity analysis at the point early in the conflict included Red's weights on military losses, domestic credibility, and risking nuclear escalation, as shown in Section 6. The sensitivity analysis at the point later in the conflict, however, included Red's weights on military losses, world credibility, and risking nuclear escalation. This was chosen because analysis showed that at this node domestic credibility had the least impact on Red's overall objective function.

Varying the sensitivity analysis over the weights of the objective function has the added benefit of accounting for the significant uncertainty in Red's decision-making. Determining what an adversary would value and how it would react in extreme circumstances is challenging, so by spanning all possible combinations of the three most influential weights, we average over this significant source of uncertainty.

The following strategies produced undesirable changes in Red's decision space:

- A "no first use" strategy, where Blue would definitely not use nuclear weapons first, made Red more aggressive. Specifically, later in the conflict the volume for Red to deploy more conventional forces increased by approximately 8% and the volume for Red to be deterred decreased by approximately 8%. There was no change in the volume for Red use of a nuclear weapon. There was no significant change in the decision space earlier in the conflict.

- If Blue had no nuclear weapons at all Red would be more aggressive. Specifically, later in the conflict the volume for Red to deploy more conventional forces increased by approximately 12% and the volume for Red to be deterred decreased by approximately 12%. There was no change in the volume for Red use of a nuclear weapon. Earlier in the conflict, the volume for Red to attack increased by a small amount, approximately 1.5%, and the volume for Red to be deterred decreased by a correspondingly small amount.

In the context of our model and assumptions, these changes are likely because Red's expected value from making aggressive choices is increased if the probability of Blue responding with a nuclear weapon is diminished or eliminated. This tracks with real world intuition, where a real-life adversary may be willing to be more aggressive if it need not fear nuclear reprisal.

The following strategy produced a desirable change in Red's decision space:

- Blue having only SLBMs (high and low yield) made Red less aggressive. Specifically, later in the conflict the volume for Red to deploy more conventional forces decreased by approximately 9% and the volume for Red to be deterred increased by approximately 9%. There was no change in the volume for Red use of a nuclear weapon. There was no significant change in the decision space earlier in the conflict.

In the context of our model and assumptions, this change is likely because SLBMs have a high likelihood of getting to target, and so Red's expected value decreases if it's guaranteed that Blue would choose to use an SLBM if it uses a nuclear weapon. In simplistic terms this also tracks with intuition. However, it is unclear if this would be the case in the real world, since our intermediate resolution model does not account for important additional weapon factors, such as self-deterrence, whether perceived or real.

The following strategies produced no changes in Red's decision space:

- Blue has no air-launched cruise missiles or gravity bombs. This can also be interpreted as meaning that Red believes its air defenses are so robust there is no chance Blue would choose to use air-launched cruise missiles or gravity bombs.
- Blue does not have the low-yield SLBM or the sea-launched cruise missile.

In the context of our model and assumptions, there likely was no change in these cases because Blue has other weapon options with similar or identical features, therefore there would be no change in Red's choices. This reveals a shortcoming in our intermediate resolution model, where Blue's nuclear weapon options are differentiated only by probability of use, probability of getting to target, and binned in one of two yield classes (high or low). More detailed factors, such as whether the weapon is visible and therefore useful for signaling or the locations from where it can be launched, were not included in our model.

9 Summary and conclusions

We have demonstrated how decision trees can provide a simplified approach to a difficult multi-stage game theory problem under uncertainty. The approach is to build the basic decision tree that is optimized from the perspective of one player, then use sensitivity analysis to gain insight into strategies and actions the other player might take in order to influence the actions of his competitor.

We illustrated the method with a simplified representation of a conflict between two generic nuclear powers, Red and Blue. Red decisions and perceptions about how Blue might respond are represented in the tree. Red objectives are represented by a multiattribute utility function that was calibrated with input from a subject matter expert using standard elicitation techniques. The quantitative representation of Red objectives that determines decisions it would make under different circumstances allows us to do extensive sensitivity analyses with respect to many of unknowns.

Some general insights derived from exercising the model, which are consistent with intuition, include:

- Red would be more likely to resort to the use of nuclear weapons if the assault is stalled
- Red would be more likely to use nuclear weapons later in the conflict after Blue has demonstrated resolve by committing additional conventional forces
- A "no first use" policy, or if Blue had no nuclear weapons at all, probably would make Red more aggressive later in a conflict.
- An arsenal of only SLBMs would make Red less aggressive later in a conflict, but it is unclear whether this would be reflected in the real world.
- Blue having no weapon systems that are highly vulnerable to air defenses, or some other select nuclear weapons, does not change Red's decision space. This may not reflect real world results, since our intermediate resolution model does not account for the more detailed aspects of these weapons that set them apart from other Blue nuclear weapons.

References

[Clemen 1996] Clemen, Robert, *Making Hard Decisions*, Duxbury Press, New York .

[Durkalec 2018] Durkalec, Jacek, Lawrence Livermore National Laboratory, *private communication*, 5/31/18 and 6/5/18.

[Durkalec 2018a] Durkalec, Jacek, The 2018 U.S. Nuclear Posture Review, Blue's Brussels Summit and Beyond, Center for Global Security Research, Lawrence Livermore National Laboratory (June 2018).

[Edmunds 2018] Edmunds, T., J. Garrett, C. Wuest, "Decision Tree Models of Deterrence Including Tactical Nuclear Weapons," LLNL-PROC-752489, Presented at Military Operations Research Society – 86th Symposium, Monterey, CA, June 21, 2018.

[Kent 2008] Kent, Glenn A., David Ochmanek, Michael Spirtas, and Bruce R. Pirnie, Thinking About America's Defense: An Analytical Memoir. Santa Monica, CA: RAND Corporation, 2008. https://www.rand.org/pubs/occasional_papers/OP223.html. Also available in print form.

[Kuhn 2018] Kuhn, Ulrich, Preventing Escalation in the Baltics – A Blue Playbook, *Carnegie Endowment for International Peace*, Washington, D.C.

[Maurer 2009] Maurer, Stephen M (Ed.), *WMD Terrorism: Science and Policy Choices*, MIT Press (2009) (Thomas Edmunds and Richard Wheeler contributed Chapter 7 Setting Priorities: Assessing Threats and Identifying Cost-Effective Responses to WMD Terrorism).

[Manzo 2018] Vince A. Manzo & John K. Warden (2018) After Nuclear First Use, What?, Survival, 60:3, 133-160, (June 2018)

[Oelrich 1988] Oelrich, I., Bracken, J., *A Comparison and Analysis of Strategic Defense Transition Stability Models*, Institute for Defense Analyses (1988)

[OSD 2018] Office of the Secretary of Defense, Nuclear Posture Review (February 2018)

[Powers 2010] Powers, Sarah, Daniel Faissol, Thomas Edmunds, and Richard Wheeler, A Model of Adaptive Adversary Decision Making: Application to International Nuclear Smuggling, LLNL-TR-462351 (October 2010)

[Zdanavicius 2015] Zdanavicius, Liudas and Matthew Czekaj, Eds., Red's Zapad 2013 Military Exercise – Lessons for Baltic Regional Security, *The Jamestown Foundation*, Washington, D.C. (December 2015).

Appendix A List of model variables

Eighty variables have been defined in the model for which sensitivity analyses can be conducted. A sensitivity analysis can be conducted on any combination of three or fewer variables, which provides a powerful tool for exploring the ways in which the tree's key factors affect the decision space. Variable names are intended to be descriptive. For example, the variable HYMilLoss refers to expected Red military losses from a Blue high yield nuclear weapon. The probability variables include references to the location in the decision tree. For example, Prob_ALCMAfterBlueSuccess refers to Red's estimate of the probability that Blue will launch an ALCM after an initial Blue success where Red is stalled and a subsequent attack by Red that overruns Blue forces (which provides incentive for a nuclear strike).

	Name
1	HYMilLoss
2	OverrunMilLoss
3	Prob_ALCM
4	Prob_ALCMAfterBlueSuccess
5	Prob_BlueALCMRedStalled
6	Prob_BlueALCMRedStalledAfterBlueSuccess
7	Prob_BlueCapitulate
8	Prob_BlueCapitulateAfterMoreConv
9	Prob_BlueCapitulateAfterRedNuke
10	Prob_BlueFullMobilize
11	Prob_BlueGravityRedStalled
12	Prob_BlueGravityRedStalledAfterBlueSuccess
13	Prob_BlueICBMRedStalled
14	Prob_BlueMoreConvAfterRedNuke
15	Prob_BlueMoreConventional
16	Prob_BlueNoMobilize
17	Prob_BlueNuke
18	Prob_BlueNukeAfterMoreConv
19	Prob_BlueNukeAfterRedNuke
20	Prob_BluePartialMobilize
21	Prob_BlueSLBMHYRedStalled
22	Prob_BlueSLBMLYRedStalled
23	Prob_BlueSLCMRedStalled
24	Prob_BlueSLCMRedStalledAfterBlueSuccess
25	Prob_Bomb
26	Prob_BombAfterBlueSuccess
27	Prob_ICBM
28	Prob_ICBMAfterBlueSuccess
29	Prob_RedMoreConvStalledAfterBlueMoreConvFull
30	Prob_RedMoreConvStalledAfterBlueMoreConvNo
31	Prob_RedMoreConvStalledAfterBlueMoreConvPartial
32	Prob_RedNukeStalledAfterBlueMoreConvFull
33	Prob_RedNukeStalledAfterBlueMoreConvNo
34	Prob_RedNukeStalledAfterBlueMoreConvPartial
35	Prob_RedOverrunBlueFullMobilize
36	Prob_RedOverrunBlueNoMobilize
37	Prob_RedOverrunBluePartialMobilize
38	Prob_RedStalledAfterBlueMoreConvFull
39	Prob_RedStalledAfterBlueMoreConvNo
40	Prob_RedStalledAfterBlueMoreConvPartial

Some of the variables refer to credibility losses by Red. For example, the variable RDC_BlueCapitulateRedMoreConv refers to the value assigned to Red domestic credibility loss when Blue capitulates after Red has committed more conventional forces.

	Name
41	Prob_RedStalledRedMoreConvAfterBlueSuccess
42	Prob_RedStalledRedNukeAfterBlueSuccess
43	Prob_SLBMHY
44	Prob_SLBMHYAfterBlueSuccess
45	Prob_SLBMLY
46	Prob_SLBMLYAfterBlueSuccess
47	Prob_SLCM
48	Prob_SLCMAfterBlueSuccess
49	RDC_BlueCapitulate
50	RDC_BlueCapitulateRedMoreConv
51	RDC_BlueCapitulateRedMoreConvAndNuke
52	RDC_BlueCapitulateRedNuke
53	RDC_RedDeterredAfterFirstAttackStalls
54	RDC_RedDeterredBlueMoreConv
55	RDC_RedDeterredBlueMoreConvAndNuke
56	RDC_RedDeterredBlueNuke
57	RDC_RedDoesntMobilize
58	RedDomesticCredibilityLoss
59	RedMaxMilLoss
60	RedMilitaryLoss
61	RedRiskEscalation
62	RedWorldCredibilityLoss
63	RRE_Deterred
64	RRE_Risk
65	RWC_BlueMoreConvRedNoMoreConv
66	RWC_BlueNukeRedNoNuke
67	RWC_RedAttack
68	RWC_RedDoesntMobilize
69	RWC_RedMoreConvBlueMoreConv
70	RWC_RedMoreConvBlueNoMoreConv
71	RWC_RedNukeBlueNoNuke
72	RWC_RedNukeBlueRespondsNuke
73	SFN
74	StalledMilLoss
75	SumW
76	TacticalNukeMilitaryLoss
77	Wrdc
78	Wrml
79	Wrre
80	Wrwc

Appendix B Volume analysis methodology

Figure 12 shows a screenshot of the spreadsheet used to calculate decision space volume. The first six columns (A-F) are the data as outputted from the sensitivity analysis. Since the sensitivity analysis outputs rows for each point we used a simple count in column G to partition each point. Then, in column H we determined which row for each point had the maximum expected value. Column I then matches the maximum expected value with a strategy in column D. Column J determines whether there are any ties. In this example there are none. Finally, columns K-N apply a weight based on whether or not there is a tie.

A	B	C	D	E	F	G	H	I	J	K	L	M	N
							3 decision node logic						
						Same	Max	Best	Num		Weighted	Weighted	Weighted
VARIABLE1	VARIABLE2	VARIABLE3	STRATEGYI	STRATEGY	VALUE	point?	value	strategy?	ties?	Weight	Strat 0	Strat 1	Strat 2
0	0	0	0	Red tactical	-0.30706136		-0.096269		1	1	0	0	0
0	0	0	1	Red more c	-0.09626922		-0.096269	1	1	1	0	1	0
0	0	0	2	Red deterre	-0.8	1	-0.096269		1	1	0	0	0
0	0	0.1	0	Red tactical	-0.398767796		-0.139556		1	1	0	0	0
0	0	0.1	1	Red more c	-0.139556002		-0.139556	1	1	1	0	1	0
0	0	0.1	2	Red deterre	-0.470781893	1	-0.139556		1	1	0	0	0

Figure 12. A sample of a spreadsheet used to calculate decision space volume

Figure 13 shows the aggregate results from this spreadsheet. This is the same result as described in section 7. The first row adds the sum of numbers in rows L-N to determine the number of points at which each strategy is optimal. The next row normalizes this number, and the results are displayed in a pie chart.

Nuclear	More conventional	Deterred			
3	869	459	sum		
0.0022539	0.6528926	0.3448535	divide by number of points in cube		

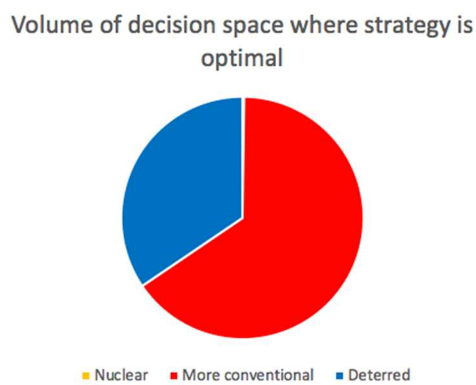

Figure 13. Volumes are calculated and results displayed