Laboratory for Unconventional Conflict Analysis & Simulation

User Manual for the Course of Action Simulator

Technical Paper 2014-2v2

Introduction

The Course of Action and Simulator (CAS) is an analytic tool that models the courses of action that two states take given certain assumptions about their goals, resources, capabilities, and interpretations of these elements. CAS is an agent-based model and simulation (AMBS) in that it uses a simple set of rules and goals to model the interaction of independent actors. CAS serves a dual educational and analytical purpose. It can (and has) been used as a training tool to help students of security studies unpack the interplay of goals, resources, capabilities, and interpretation, and it can also help introduce them to advanced complexity-based modeling for strategic analyses. AMBS provides a way to use simple sets of rules to represent how states make decisions, and models like CAS can provide a systematic tool for representing a complex, large set of assumptions about these rules and the various ways they can affect outcomes. Once the assumptions are initially specified and validated, the model can be deployed fairly rapidly to analyze new scenarios. The model's second purpose is to provide a tool to supplement strategic analyses of bi-lateral interactions. The simulator's utility to this second purpose will grow over time as the model benefits from further critique, specification, and validation through applications to new scenarios.

This *User Guide* seeks to provide educators and practitioners with the information necessary to basically understand how the model works and how to apply it to analyze a scenario. Because this user guide aims to be accessible for non-experts, the lab has also published a companion technical report that details the mathematical foundation for the algorithms undergirding the model design. This documentation intends to help make the

The CAS model was developed by Dr. Tony Rivera, Harrison Richard, and Luke Maier. Also, students from the Joint Special Operations Masters of Arts provided crucial feedback on the model during a training seminar on agent-based modelling held at Fort Bragg, NC, in May 2014.

model transparent to empower users to evaluate its validity and spark discussion on how to further enhance AMBS modelling for strategic studies. Also, an overarching challenge of designing models like this assuring the model design accurately represents relevant theory—in this case, theories about international relations and strategic influence. Models yield little useful information if they are under- or mis-specified, so another purpose of this documentation is to provide users with enough information to evaluate and potentially provide feedback on how well CAS represents theory from strategic studies and international relations.

The next two sections provide a walkthrough of the model's basic assumptions and mechanics and how to run the model and interpret its results. An example of how to apply this model to deterrence will be provided by a forthcoming chapter in an edited volume on constructivist methods in international relations (Rivera and Maier 2014, forthcoming University of Michigan Press). The final section discusses opportunities to expand and improve this model, including integrating artificial intelligence planning through machine learning.

How CAS Works

Agent-based models like CAS have three basic elements: inputs, mechanics (also known as "processing), and outputs (Holland 1995). These roughly correspond to the initial state of the system, the set of rules that guide the actors' decisions when changing the system's state, and next state of the system that results from the decisions made. CAS draws heavily from concepts in artificial intelligence planning, in which "an agent is given a set of possible states and is asked for a sequence of actions that will lead it to a desirable world state" (Vidal 2007). But unlike a true artificial intelligence planning simulator, the actors in CAS only look one step ahead, as opposed to how the world could be given various combinations of steps over longer periods of time.

In CAS, states make series of moves, and we assume the mechanisms by which they select moves remain stable over time. CAS represents this by using the same move-selection algorithm each time interval and iteration. Of course, the inputs to each iteration can vary, so the outputs depend on the inputs, despite an identical decision-making process. This essentially parallels the idea that states follow rules of decision-making that remain fairly stable despite short-term events and conditions. Because their rules of action remain constant, the two states cannot adapt their decision-processes over time the way agents can in complex adaptive systems (Holland 1995). Making the agents adaptive is a point of future development for the Lab's modeling suite.

CAS begins with the assumption that states have goals regarding certain "objects", which can be thought of as material or ideological issues that provide focal points of interaction among states. The states have their own goals for these objects. A goal is represented as a desired deployments of elements of national power toward an object, which correlates with a desirable perception of that object. A state's goal might be, for example, to pursue economic gains through an object. Another goal could be to take actions regarding a particular object to gain prestige in the international system or legitimacy at home. The model enables these goals to be harmonious or conflictual, which allows the analyst to decide for which objects gains are zero sum or cooperative. States can base these

perceptions on their deployment of power resources *relative* to the deployment of the other actor; as Clausewitz describes, an "object" cannot be conceptualized "in itself, but so in relation to both belligerent states" (Clausewitz 2006, 26).

Drawing from Lebow and Onuf, the model focuses on the effects associated with the state's perceptions of fear, honor, and interest. It allows each state to have a different system for perceiving how each type of resource affects its sense of fear, honor, and interest, considered discretely. Also, "one and the same political object may produce totally different effects upon different people, or even upon the same people at different times" (Clausewitz, 2006, 26). These three drivers each can have different levels of salience for each state's decision-making process. Some states will act more based on their sense of fear, others more on their sense of interest, or others will act with relatively equal focus on all three drives. And these preference sets will change over time. The model represents the level at which a state is driven by fear, for example, on a scale from 0-100, with 100 representing the maximum influence (not complete influence, however) and 0 no influence. These scales are relatively weighted, such differences between the weights represent the degree of drive.

In addition, the model assumes that the state have various elements of national power that can be deployed to achieve their goals for these political objects. Drawing from Morgenthau's DIME theory and other concepts (e.g., see Art 1996; Art and Waltz 2009) of national power such as COIN, CAS considers the states have available varying initial amounts of eight types of resources: diplomatic, informational, military, economic, natural, legitimacy, human, and technological. Each type of resource has a particular degree of salience for the state's perception of fear, honor, and interest. For example, some states associate the possession of military resources with their sense of prestige, whereas another state might associated military resources more closely with fear. Using a scale of 0-100, like with the drivers, the analyst specifies the relative extent to which each actor associates the resource types with fear, honor, and interest. So if the state equally associates, for instance, its technological resources with interest and fear, then actions that change their amount of natural resource capital deployed toward an object will result in an equal change in their perception of interest and fear regarding the object. Also, the model represents the amount of resources deployed toward an object as a number that ranges from 0 to as high as the analyst desires. Analysts are encouraged to use empirical data to inform these estimates, but empirical data often cannot entirely capture these largely unquantifiable concepts of power. Of course, given the imprecision of the analyst's estimations, is it useful to run the model multiple times using a variety of assumptions to see how sensitive the results are to particular assumptions, but considerations such as this are described in the next section.

Running the Model

At the beginning of the first iteration, the analyst sets the initial inputs, and the model's algorithms processes these moves into the first set of outputs. The second iteration uses the same set of mathematical constructs to process the first set of outputs into the second set, and so on until the model produces the final, twenty fourth sets of outputs, which marks the end of the simulation. In essence, in simulations of changes over time, the outputs from the first time (t) interval provide inputs for the next time interval ($t+1$), which in turn feed the following time interval and so on (until $t+n$). The simulation makes the analyst's initial

inputs for the first time interval (i.e., the start of the simulation) cascade through later intervals until the end of the simulation.

The target state can take various actions to change the resources it has deployed toward each object. Similarly, the sending state can take various actions that change the resources the other state (i.e., the target state) has deployed. Each state makes a discrete move for each object for each time interval. These power resource deployments come in the form of "moves." Over time, many moves compose a "course of action." The process by which they select which move is described next.

States have a range of actions (interchangeably referred to here as "moves") that affect the objects regarding fear, honor, and interest. States select the "best" of 20 possible discursive and active moves: Do Nothing, Threaten, Insult, Defend, Acknowledge, Bluff, Promise, Commitment, Apology, Blow, Strike, Attack, War, Arm, Buy, Sell, Sanction, Support, Fund, Concede, and Wait. In the modelling sense, each move can be considered an "operator" that has a specified effect on the target state's amount of each type of resources (Vidal 2007). Recall that each state has a goal for how much of each resource type in each object the target has deployed. The states compare how close each move gets to that goal, and the move that gets the closest is generally selected as the best move. However, the resource effects of the move for an object also affect the state's perceptions of fear, honor, and interest vis-à-vis that object. CAS assumes that two states could cause asymmetric effect through the same move. The set of effects is unique to each actor, and we have somewhat arbitrarily set the effects of each move to give the analyst a starting point for refining and validating these key elements of the model. We caution that in order for the model to run to completion (i.e., for it not to return #N/A as moves after only a few time intervals), the effect of these moves should be fairly small compared to the total resources initially deployed. (If the distance between the ideal resource deployment levels and deployment level at any time interval—including the initial deployment—is small, then the moves that are selected could only represent "noise"; this occurs for fairly complex statistical reasons that result from the model's use of averaging, standardized scoring, and three-dimensional distances. In plain terms, this means the analyst should be aware the model's reliability depends on how exactly and minutely the effects of the moves were specified in comparison to the starting deployments.)

Furthermore, also recall that the agent can be driven to different degrees by fear, honor, and interest. The strength of the driver correlates to the priority that it has for move selection. For example, highly honor-driven states will prioritize moves that achieve their goals for honor, even if that comes at the expense of failure to meet fear and interest goals. So, the move rankings are adjusted based on how well each advances the given state's priorities. Each state follows this same process, but each is operating under the perception frameworks specified by the analyst. A more mathematically rigorous discuss of this move-selection algorithm is provided in the accompanying technical supplement.

Once the best move is determined for each side, the model applies the effects of the moves on the target's resource deployment. The result could be that the target state's move results in a general increase in its deployment, but the sending state counters with a move that

results in a general decrease in the target's deployment. For objects for which the states have harmonious goals, the states will probably take similar courses of action.

Courses of action emerge as patterns of types of moves. If a state is taking many militaristic moves, for example, its general line of action could be considered militarily aggressive. It could also take what could be considered more submissive moves, for example. A key takeaway from the model is how the states alter each other's courses of action, which is an example of how the model produces information that can aid analysis of the ways states interact given certain goals, perceptive frameworks, and abilities.

Interpreting the Results

Moves
Think of moves as deployments of certain resources that affect fear, honor, or interest. For example, highly realist moves deploy hard power resources, whereas constructivist moves deploy honor-based resources.

- Fear-related resources: Military, Economic, Human, Technology
- Honor-related resources: Diplomatic, Informational, Legitimacy
- Interest-related resources: Diplomatic, Economic, Natural, Human, and Technology

Move strength

The "strongest" moves make significant progress along their priority axis, but they also have the largest tradeoffs (i.e., they have losses along the non-prioritized axes.). "Weaker" moves make some progress along their priority axis, but they have little or no trade-offs on other axes. In this way, the more strongly an actor identifies with one of the axes, the stronger their moves will be. An actor with a highly complex identity that prioritizes all of the axes will take fairly weak moves that make some progress on all of their goals.

Move groupings and explanations: (each group's moves are listed from weakest to strongest)

The moves can be categorized into three groups; each group prioritizes a different axis. Realist moves prioritize reducing threats to fear, liberal moves primarily reduce threats to interest, and constructivist moves primarily reduce threats to honor.

Each move has an effect on the fear, honor, and interest threat levels, and it also has an effect on Actor A's resources. These resource effects can be considered representative of the moves' "costs" and "payoffs" for each resource.

- **Constructivist Moves** (i.e., moves that gain honor):
  - **Apology**
    - FHI effect: Very weakly reduces threats to honor and fear; no effect on interest.
    - Resource effect: Weakly reduces Diplomatic, Economic, and Legitimacy resources.
  - **Concede** (give into the other actor's stance on an issue)
    - FHI effect: Weakly reduces honor threats; very minimally increases threats to interest (due to the consequences of concessions).

        - Resource effect: Weakly reduces Diplomatic, Economic, and Legitimacy resources. Slightly increases information.
    - **Promise** (agree to do something, such as uphold a treaty or obligation)
        - FHI effect: Weakly decreases threats to Honor; increases threats to fear and interest (due to greater vulnerability to entanglement in ally's conflicts)
        - Resource effect: Weakly reduces Human; increases Technology; very weakly increases Military
    - **Acknowledge** (acknowledge a program exists)
        - FHI effect: Moderately decreases threats to honor; slightly increases threats to fear and interest (due to the loss of secrecy and strategic ambiguity)
        - Resource effect: Increases Diplomatic, Informational, Legitimacy resources
    - **Bluff** (fake out the other actor)
        - FHI effect: Decreases threats to Honor (by saving face at home, while proving cleverness abroad);
        - Resource effect: Decreases Diplomatic and Legitimacy; slightly decreases Economic resources; increases Informational; slightly increases Military and Technology resources
    - **Insult** (levy insult against the other actor, either to signal or galvanize domestic support)
        - FHI effect: Significantly reduces threats to honor (by galvanizing the homefront and rallying them against other powers); moderately increases threats to fear and interest (due to rhetorical escalation)
        - Resource effect: Slightly decreases Military and Technology resources; increases Diplomatic, Informational, Economic, and Legitimacy resources
    - **Threaten** (threat a course of action in response to certain behaviors by the other actors)
        - FHI effect: Strongly decreases threats to honor; moderately increases threats to interest and fear (due to possible retaliatory sanctions or action)
        - Resource effect: Strongly increases Diplomatic, Informational, and Legitimacy resources; moderately reduces Military and Technological resources
- **Realist Moves** (i.e., moves that reduce fear):
    - **Arm** (build up military and technology resources)
        - FHI effect: Very slightly reduce threats to fear and honor; no effect on threats to interest
        - Resource effect: Slightly Increases Military; slightly increases Diplomatic, Economic, and Legitimacy resources
    - **Defend** (retrench or against an small or large attack)
        - FHI effect: Slightly reduces threats to fear; very slightly reduces threats to honor and interest

        - Resource effect: Slightly decreases diplomacy
    - **Blow** (undermine
        - FHI effect: Moderately decrease threats to fear; moderately increase threats to honor and interest
        - Resource effect: Moderately increase military resources; Decrease Diplomatic, Legitimacy, and Economic resources; slightly decrease Natural and Human resources.
    - **Strike** (conduct a targeted tactical mission, such as a one-off bombing of a nuclear facility)
        - FHI effect: Decrease threats to fear; increase threats to honor and interest
        - Resource effect: Significantly increase military resources; significantly decrease Diplomatic, Economic, and Legitimacy; decrease Information and Human resources
    - **Attack** (conduct a strategic, medium but limited operation that is larger than a strike)
        - FHI effect: Significantly decrease threats to fear; increase threats to honor and interest
        - Resource effect: Significantly increase Military resources; moderately increase technological resources; substantially decrease Diplomatic, Legitimacy, Information resources; moderately decrease Economic and Human resources
    - **War** (wage a long-term, full-spectrum military engagement)
        - FHI effect: Significantly reduce threats to fear; significantly increase threats to honor and interest (due to cost and potential for humiliation)
        - Resource effect: Significantly increase Military and Technological; decrease Diplomatic, Information, Legitimacy; slightly decrease Economic, Human and Natural resources
- **Liberal Moves** (i.e., moves that support interest):
    - **Support** (send limited economic aid)
        - FHI effect: Slightly decrease threats to interest
        - Resource effect: Slightly increase Military, Natural, and Technological resources
    - **Commit** (make a trade or general economic agreement)
        - FHI effect: Slightly decrease threats to honor; very slightly decrease threats to fear.
        - Resource effect: Slightly increase Economic, Diplomatic, Legitimacy, and Natural
    - **Buy** (encourage the domestic economy to buy goods)
        - FHI effect: Decrease threats to interest; Slightly increase threats to fear and honor
        - Resource effect: Moderately increase Legitimacy, Technology, Natural, and Economic resources; Decrease Diplomatic and Military resources
    - **Fund** (send significant economic aid to another actor to gain influence)

        - FHI effect: Significantly decrease threats to interest; increase threats to honor and fear
        - Resource effect: Moderately Increase Diplomatic, Informational, Natural, Legitimacy, and Technological resources
    - **Sell** (sell goods to the other actor)
        - FHI effect: Substantially decrease threats to interest; increase threats to honor and fear
        - Resource effect: Increase Economic and Informational resources; slightly increase Diplomatic, Natural, Legitimacy, and Technological resources
    - **Sanction** (sanction the other actor)
        - FHI effect: Substantially decrease threats to interest; increase threats to fear and honor
        - Resource effect: Increase Economic resources; Slightly decrease Diplomatic, Informational, Military, Legitimacy, and Technology resources
- **Do nothing**
    - **Wait** (sit tight)
        - FHI effect: None
        - Resource effect: Very slightly increase Informational, military, and Economic resources (due to their modest natural build up over time)

Metric Graphs

After the model is run, it presents a series of graphs that display information about how several metrics changed over the course of the simulation. The first set of graphs provide information on the target actor's change in resource capital per move. This can be thought of as the rate of change in possession of elements of power. The large graph shows this information aggregated for all of the objects, and the smaller graphs show the changes in total resource capital in each object for each move. Upward trends indicate the target state is gaining resources. Steady gains are indicated by consistent slope, whereas spikes and highly uneven slopes indicate volatility.

The second set of graphs shows the change in the total capital in each type of resource over time. This information empowers the analyst to observe how the courses of action that the states adopted resulted in changes in the target's resources over the course of the simulation. The sender state might have goals that can be met by changing the amount of particular types of resources, such as military resources. Furthermore, these graphs provide information

The final set of graphs depicts the salience of the objects to the target over time. Salience is defined by the proportion of the target's total resources deployed to a particular object. High salience means, for example, that the object occupies a large proportion of the target's resources.