Computational and Mathematical Organization Theory
https://doi.org/10.1007/s10588-021-09352-x

S.I. : GROUND TRUTH: IN SILICO SOCIAL SCIENCE (GTIS3)

Groups, governance, and greed: the ACCESS world model

Scott Rager[1] · Alice Leung[1] · Shannon Pinegar[2] · Jennifer Mangels[3] · Marshall Scott Poole[4] · Noshir Contractor[5]

Abstract
ACCESS—the Agent-based Causal simulator with Cognitive, Environmental, and Social System factors—is an agent-based simulation of an alternate world that is designed to test social science methodologies' abilities to explain, predict, and prescribe policies for complex social systems. The ACCESS world model includes behaviors based on behavioral and cognitive sciences within and across individuals, groups, and the society to create a multi-level model that exhibits emergent phenomena. In this paper, we detail the logic underlying our conceptualization of the entities (individuals, groups, and the world) and their interactions. We also provide details on how we used the ACCESS model to challenge and score social scientist teams' abilities to explain, predict, and prescribe in the artificial world as part of the DARPA Ground Truth program.

Keywords Agent-based modeling and simulation · complex social systems · emergent behavior

1 Introduction

The ACCESS simulation was designed to create a testbed for evaluating the accuracy and robustness of various social science methodologies, including those that explain the causal model of a social system, predict a future state of that system, or prescribe interventions to steer the system toward a desired future state. This work was done as one of four "simulation teams" performing in Technical Area 1 (TA1) on the DARPA Ground Truth program. Our role was to develop simulations that exhibited complex, socially plausible behaviors to provide data to test social science methodologies. Performers in Technical Area 2 (TA2) on the program, the "research teams," were teams of world-class social and data scientists that applied state of the

✉ Scott Rager
scott.rager@raytheon.com

Extended author information available on the last page of the article

Published online: 02 December 2021

art techniques attempting to uncover the causal model and answer the predict and prescribe "Challenge Questions" that we posed. We also worked closely with a Testing and Evaluation (T&E) team throughout the course of the program.

This paper is focused on explaining the details of the world model being simulated in ACCESS. We first attempt to give some context for the design in Sect. 2, explaining the high-level goals for our world model. Next, we provide necessary details about the ACCESS simulation framework in Sect. 3 that give the reader an understanding of the larger mechanics of the simulation. Then, the majority of the paper focuses on explaining the details of all of the actions, decisions, and updates that make up the ACCESS world in Sect. 4. We believe the details provided in this section should be sufficient to interpret any analysis and findings of ACCESS data and provide context for analysis of the research teams' performance in the Ground Truth program, which is one of the primary goals of this paper. Finally, we explain our approach to providing realistic accessibility to data in Sect. 5 and our approach to providing explain, predict, and prescribe tests for the research teams in Sect. 6.

2 ACCESS design goals

We had a number of high-level goals that we set out to achieve when designing ACCESS world models. First, we wanted the fundamental elements of the world model (i.e., decisions, behaviors, and changes in beliefs and feelings) to exhibit a level of complexity comparable to what is observed in the real world. Therefore, as inspiration for actions in ACCESS, we selected a subset of theories from the social and behavioral sciences literature that have been shown to influence individual and group decision making . These specific theories and their translation into ACCESS actions are detailed in Sect. 4. Anchoring our model in the social and behavioral science literature increased the likelihood that the complexity of the decisions, interactions, and observable data from the simulation would exhibit realistic levels of complexity, while also reducing the potential for counter-intuitive processes or outcomes that would have posed unrealistic challenges for the research teams.

On the other hand, although the main ideas associated with these theories were represented in the model architecture, it was not essential for us to capture every nuance of these theories, nor was it even desirable to do so. Rather than designing the model to simulate a specific aspect of real-world behavior veridically, the goal of the project was to provide a plausible testbed for social science and data science analysis methodologies. If we used or modeled our world directly after an existing set of findings or data, then we would run the risk that our model would have been too transparent, and the social scientists on the research team side of the program, who would likely to be familiar with these theories, would easily discover the original model or data set. Thus, it is probably most accurate to describe our use of these theories for inspiration rather than for replication in designing the detailed world model below.

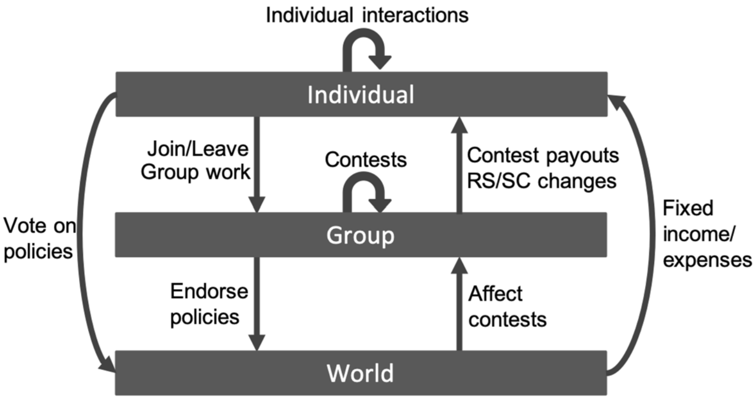

Fig. 1 ACCESS is a multi-level model including interactions within and between the individual, group, and society levels. The specific actions shown here that link these different levels will be described in more detail in Sect. 4

Another high-level goal of ACCESS was to design a social model that represented interaction within and between each of the individual, group, and society levels as shown in Fig. 1. In addition to adding realism to the model, these interactions require that methods used to explain, predict and prescribe the work taken into account multi-level dynamics. These connections also increase the opportunity for the simulation to exhibit emergent behavior.

3 ACCESS framework overview

While explaining the comprehensive details of the ACCESS simulation architecture and codebase are out of scope for this paper, a short introduction to its underlying framework can guide the reader to how the world model operates in general. As is typical in agent-based modeling and simulation, we use the term “agent” to refer to an autonomous actor that has some internal state and that can perform actions and interact with other agents. In ACCESS, we refer to agents and individuals interchangeably, but also consider groups, locations, and the world itself as having an internal state and the ability to perform actions/interactions. We refer to these types of instances collectively as *entities*. In ACCESS, entities could have three types of variables: traits, attributes and actions. Figure 2 provides a list of all entities, including those entities’ traits and attributes, and their most important actions, that our world models implemented in the three Ground Truth Challenges.

Across the course of the Ground Truth program, we progressed through a total of 8 major (i.e., version 1.0, 2.0, etc) and 2 minor versions (i.e., versions 5.1) of the world model. Figure 2a shows the variables that were included in the models implemented for Challenge 1 (version 5) and 2 (version 5.1), and Fig. 2b lists the variables included in the model implemented for Challenge 3 (version 8.0). In both figures, static traits are italicized and dynamic attributes are not.

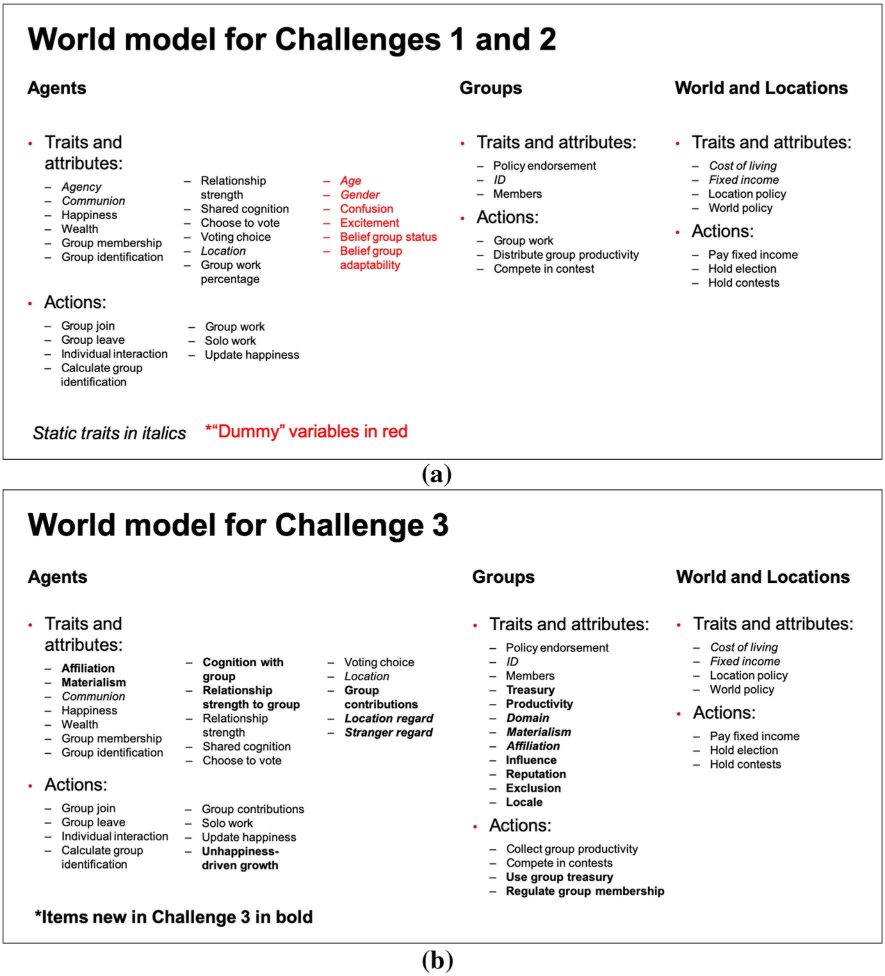

Fig. 2 **a** the final world model for Challenges 1 and 2; **b** final world model for Challenge 3

Any specific examples we use throughout this paper use the traits, attributes, and actions from the Challenge 3 version of the model (v8.0). While this model retains many of the basic elements from earlier versions, there are some notable differences. First, the world model for Challenges 1 and 2 included "dummy variables" that had no causal relationships with any other variable. When it appeared that dummy variables might be causing research teams undue difficulty in uncovering our model's causal structure, we opted to eliminate this type of variable. Second, the version of the model used in Challenge 3 greatly extended the number of group and group-related traits and attributes, as a reflection of the more detailed group structures that emerged when we increased the number of agents in our world from 100 (Challenges 1 & 2) to 10,000 (Challenge 3). Finally, given the central role that Happiness and Wealth play in our model, we developed "motivational" goals for agents and groups that would result in prioritizing Affiliation (relationships) or Materialism (wealth) and a number of new traits, attributes and actions related to these goals.

First, we will describe the concepts of entities, traits, attributes, and actions more generally, along with the general mechanics of the simulation framework. After this introduction, we will take the reader through the causal relationships that define two of the most central agent-based dynamic traits in the model-happiness and wealth-particularly as they relate to interactions with others (i.e., groups). We will then discuss voting and policies, a core aspect of world-level action and how these also circle back to influence agent happiness and wealth.

3.1 Entities

An *entity* contains a set of immutable internal variables called *traits* that are set randomly on the initialization of the simulation run, but then do not change throughout the simulation. At the level of individuals, many variables can be thought of as fixed personality traits or qualities. For example, the degree of "communion" associated with the agent is a fixed trait associated with qualities that establish and maintain social relationships and is loosely based on the agency-communion dichotomy initially outlined by Bakan (1966) (for more recent review see Paulhus and Trapnell 2008). The location where the agent "lives" can also be thought of as a static quality because while agents may do work with different groups, they always return "home" to the same location throughout the simulation. Traits can also exist at the group or societal level, where they can be thought of as distinguishing characteristics that do not change within the timeframe of the simulation, such as global prosperity and the cost of living in their particular home location. Although traits do not vary across the simulation, their initial values are sampled from a uniform distribution that varies as a function of the specific trait (e.g., Communion is sampled from 0 and 1).

Entities can also contain mutable internal variables called *attributes*, which can be thought of as feelings, resources, or other qualities. These values are also set randomly at initialization time, but then can change as the simulation progresses. For example, at the level of individuals, Happiness and Wealth are two resources that can change across time as a function of actions in which the agent performs or is involved. They may also experience changes in ?behavior guiding values? such as their need for Affiliation or Materialism. These behavior-guiding values (Affiliation and Materialism) can also guide behavior at the group level. Additionally, groups may increase or decrease factors such as Influence or Reputation. At the world level there may be changes in Policies. As with traits, the distribution used to initialize attributes varies depending on the specific variable (e.g. initial Wealth values are sampled from a distribution with mean 100 and standard deviation 25), but unlike traits, attributes change over the course of the simulation, although for most attributes these changes are bounded. For example, attributes like Happiness, Materialism, and Affiliation are bounded between 0 and 1. The exceptions are Wealth and GroupTreasury, which are unbounded attributes, and Group Domain, which can take an integer value up to the total number of domains in the simulation.

3.2 Actions

Any kind of decision, behavior, interactions, etc. in the ACCESS world is called an *action*. This set of actions includes all activity within and across entities, such as:

- Update of a single entities' attributes, e.g. agent action of updating their Happiness based on the current state of their other traits and attributes and anything that recently happened
- Decisions/actions of a single entity that only affect that entity, e.g. agent deciding how much time to spend working on their own project (Solo Work)
- Decisions/actions of a single entity that interact with a different entity, e.g. an agent interacting with another agent (individual interaction) or an agent joining or leaving a group
- Decisions/action of a group that affect the group, e.g., competing in contests or decisions by the group to remove a member
- Scheduled events/actions, e.g. holding Elections or Group Contests

Each action is performed by all entities of a specified type, e.g. agents, groups, locations, or the world. Furthermore, all entities perform the same logic when executing an action. The differences in outcomes for different entities performing the same actions is due to the values of the entities' internal traits and attributes being unique. Most actions also include a number of configuration parameters, which affect the behaviors in three different ways. First, some parameters are weights that determine the strength of the causal relationship of one variable to another. A second type of parameter is a threshold that decides a point at which an action will take place or an input has an effect in an action, e.g. agents only compare their Wealth to others if their Happiness is below a threshold. The last type of configuration parameter is one that determines the frequency with which actions will take place, e.g. how often Group Contests and Elections are held. These configuration parameters are used to control the world from a macro-level, ensuring a desired balance between how often observed behaviors occur as well as keeping the volatility of key observable variables within a realistic range.

3.3 Simulation mechanics

The ACCESS simulation is a time-based event-driven system, built on the Repast Symphony simulation framework (North et al. 2013), in which actions are scheduled to execute on a certain time step in the future. Most actions are scheduled to execute every time step, but some are less frequent, e.g. contests and elections. Each time step in the ACCESS simulation is a day in the ACCESS world model, and a day is the smallest level of granularity with which an observer or researcher can collect data. Within a time step, the actions being executed have a preordained order in which they are executed. This ordering of actions is decided by us, the model designers, and is not visible to the researchers/observers of the world data. Each action is designed to be executed in parallel by all entities that are performing the action. Therefore, the actions or updated status of a single entity performing an action will not affect any other entities' results from executing the same action in the same time period. In the case that a variable is an input and an output of the same action, all entities use the state of the system as it exactly is before any of the entities perform that action. While these details do not affect the macro-level actions and behaviors, they can be relevant when understanding the details of fine-grained analyses.

The overall causal model of the system is derived from the comprehensive view of causal relationships in the actions. Each action has a set of inputs, i.e. the variable(s) whose values are read and used in the action logic, and a set of outputs, i.e. the variable(s) whose values are set or changed during or as a result of the action. We derive causality in ACCESS by drawing edges from all input variables of an action to all output variables of an action. We then merge these causal edges from all the actions into a single global causal diagram. A full presentation of the final causal model for World Model v8.0 can be found in Fig. 3.

3.4 Scale and complexity

As noted above, the variables and relationships in the ACCESS model evolved over the course of the program in accordance with the program's goals. The main motivation for these changes was to control and increase the complexity of the simulated world over the three challenges. Accordingly, the ACCESS model increased in both scale, i.e. the number of entities in a simulation run, and in complexity of the model, measured by several metrics designed by the testing and evaluation team. Included in these complexity metrics were the numbers of nodes and edges in the causal diagram and the cyclomatic complexity of the diagram (McCabe 1976). Table 1 provides values for both scale and complexity of the ACCESS world model over the three challenges in the Ground Truth program, as extracted by the ACCESS team from the simulation code.

4 Details of actions in the ACCESS world model

4.1 General information about ACCESS world

The construction of our world was driven by consideration of some of the most powerful forces motivating individuals in a democratic, capalist society: the desire for wealth, the desire for socialization and a sense of belonging to a group, the desire for choice and control, and the general pursuit of happiness. As mentioned previously, the specific ways in which these motivations were operationalized in the model drew from well-established theories in the social sciences, but there were also important differences. In some cases, these differences were necessary to limit both the complexity of the world and the transparency of the theories. In others, adaptations were necessary to fit with our particular world structure.

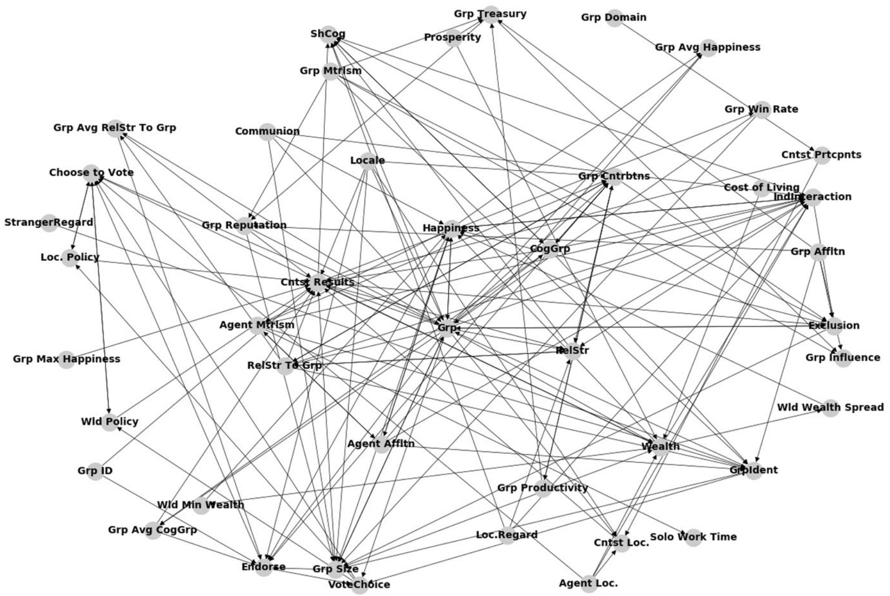

Fig. 3 Causal model of World Model v8.0 showing directional relationships between variables. Names of variables correspond to those listed in Fig. 2b, although some names are abbreviated here

Table 1 The scale of the ACCESS world used in the Ground Truth program increased across the three Challenges

	Challenge 1	Challenge 2	Challenge 3
Number of agents	100	200	10,000
Number of groups	10	10	153
Number of locations	4	4	50
Nodes in causal model graph	34	34	45
Nodes in causal model graph	80	80	144
Cyclomatic complexity of causal model graph	4	4	50

In this first section, we list a set of details that were provided to research teams at the start of the first challenge (unlike the causal relationships between variables which had to be discovered through the challenges themselves). This list contained some of the important ways in which the ACCESS world differed from our own real world and specifically, noted constraints on some key variables. These basic world principles were held consistent across the world models used in all three Challenges.

– People do not seem to ever be born or die, and seem to stay physically healthy all the time.

- People do not pair up into couples or cluster in family households, but they do have individual, dyadic relationships and varying degrees of participation in groups.
- Some type of contest with a resulting wealth transfer is occasionally observed between groups.
- The ACCESS world has neighborhoods, or *locations*. Agents are placed randomly in different locations at initialization and do not have the freedom to relocate, i.e. there is no immigration/emigration between locations or in/out of the world.
- Each location has its own cost of living, which can be set by configuration parameters or sampled from a uniform distribution within a defined range. Agents pay the cost of living for their location every day.
- There are economic disparities. No one is ever denied necessities due to poverty, but this can cause some individuals to go into debt.
- There does not seem to be the concept of a work week, or weekends. Individuals seem to do work every day, though not always the same work.
- There are no seasons, and the weather is monotonous.

There were some additional standardized aspects of the larger World Model v8.0 used in Challenge 3, where we scaled up to 10,000 agents. In this version of the model, there were 50 locations, each with approximately 200 individuals. Groups had a minimum of 2 members; if the number dropped below this threshold, it was emptied and repopulated. Agents could belong to both "small" location-based groups with a maximum of about 50 members each (similar to a neighborhood association), and "large" non-location based groups that could host up to about 1000 members, crossing multiple locations (similar to a national organization). In Challenge 3, there were 150 location-based groups, 3 per location, and 3 non-location based groups. We also introduced the concept of "domain," which defined the groups that could participate in contests with each other. All groups in a single domain were either all location-based or all not location-based, to ensure that small "neighborhood" groups would not be competing with large "national" groups.

4.2 Simulation set-up

At the initiation of the model, static attributes and baseline levels of dynamic attributes were set at the individual, group, and society/world level using a specified random seed and, where appropriate, configurable parameters. For example, the static attributes of individuals (i.e., Communion, Agency, Location Regard and Stranger Regard) were set at a value between 0 and 1, with the aim of a normal distribution across the population, using the specified random seed. The behavior-guiding attributes of Affiliation and Materialism were also set for individuals to produce a normal distribution across the population, although these could change during the simulation.

Agents were then randomly assigned to "live" in a particular location for the duration of the simulation and provided with a Wealth of 10(±5) and a Happiness level of 0.5(±.1). They also "joined" some number of groups based on both random factors and fit between their behavior-guiding values and those of the group (i.e., Affiliation and Materialism). Specifically, an individual starts with an X probability (e.g., 0.15; configurable parameter) of joining a group. Each group then conducts a random roll with the specified random seed to determine if the individual will be considered for membership in the group. If the individual will be considered, the group then evaluates whether to accept (or reject) the individual based on whether it exceeds (or not) the group's "standards" (i.e., static attributes) for both Materialism and Affiliation.

After joining groups, the model checked to make sure that each group had at least 2 members; if the number dropped below this threshold, the group was emptied and repopulated. It is also checked that no one was in too many groups; if the number of groups exceeded the GroupMembershipLimit, the agent would quit the group with which it had the lowest Group Identification. For each group the individual may join, the simulation provided a randomly set level of Relationship Strength to Group and Cognition to Group that guided future actions the individual could take with the group.

4.3 Overview of model actions

As the simulation ran, agents periodically engaged in the following actions: updating their Happiness, initiating one-on-one interactions, deciding whether to join or quit groups, deciding whether and how much time to spend on Solo and Group work, and deciding whether and how to vote in elections. We also introduced a "metacognitive" action in which agents could reflect on how well their current state was meeting their goals and if this match was consistently poor, change their goals to align better with their current state. At the group level, there were periodic Group Contests with other groups. Groups also used individual contributions to group work to improve their Group Treasury which they could then use to improve the group's reputation and influence. Groups also periodically endorsed policies that had bearing on how their members voted during Elections. Finally, at the Society/World level, the primary actions involved deactivating old policies and activating the policies at the location and global level that had won the most recent elections. In the following sections we describe these actions in more detail.

4.4 Updating happiness

Happiness is an important attribute in the ACCESS world and occupies a central node in the model (see Fig. 3) and is always bounded between 0 and 1. In addition to isolated events that can influence Happiness, such as Individual Interactions (see below) and being kicked out of a group that loses a contest (see Group Contests-Social Factors), agents regularly update their happiness once a day based on actions whose parameters were informed by multiple social and behavioral theories, as well as some intuitive mechanisms added by the authors.

The first influential theory was the Hedonic Treadmill Theory (Brickman et al. 1978; Diener et al. 2009), which generally espouses that individuals' experience rises and drops in their happiness as they experience positive or negative external events, but that these effects wear off over time, returning the individuals to an inherent baseline happiness value. The second theory incorporated in the UpdateHappiness action is the Social Comparison Theory (Festinger 1954; Wood 1989). In this context, agents with current Happiness values under a given threshold compare their Wealth value to other agents' Wealth values (see Acquiring and Losing Wealth Section for details on factors influencing agent Wealth). In addition to these theories, the authors allowed recent changes in Wealth and Socialization (i.e., number of groups they belong to, see Group Interactions section) to impact Happiness. The UpdateHappiness actions sets Happiness values for an agent as follows: ($F_{RB}, F_{CR}, F_{SC}, F_S$ are the Return to Baseline, Change in Resources, Social Comparison, and Socialization (Number of Groups) factors, respectively, and $W_{HT}, H_B, W_{CR}, T_{SC}, WSC, W_{SF}$, and T_{SF} are all configuration parameters.)

$$
\begin{aligned}
F_{RB} =& W_{HT} \times H_B + (1 - W_{HT} \times h(t-1) \\
F_{CR} =& W_{CR} \times (w(t) - w(t-1)/w(t-1)) \\
F_{SC} =& \begin{cases} W_{SC} \times \frac{(w(t) - w_{Layer1Avg}(t))}{w_{GlobalMax}} & h(t-1) < T_{SC} \\ 0 & o.w. \end{cases} \\
F_S =& \begin{cases} a \times W_{SF} & agent\ in \geq T_{SC}\ groups \\ 0 & o.w. \end{cases} \\
h(t) =& max(0, min(F_{RB} + F_{CR} + F_{SC} + F_S, 1))
\end{aligned}
$$

Figure 4 shows sample data from an agent in an ACCESS simulation run that illustrates several of the factors in the update happiness action. Significant changes in wealth around ticks 110, 125, and 175 in Figure 4b show corresponding changes in happiness in Fig. 4a caused by the Change in Resources factor. After these spikes, the Return to Baseline factor's effect is evident by the happiness value drifting back toward a baseline. Finally, the Socialization factor's effects can be seen by noting that the agent's happiness drops around tick 150 when the agent's number of groups drops to zero and then recovers when the number of groups exceeds one again around tick 195.

4.5 Social interactions

Social interactions were a critical component of the world model and could take place at both a dyadic level (one-on-one) and group level. Moreover, these were interrelated. For example, Relationship Strength to group and Cognition with Group values for all members to a group are the average of the members' dyadic

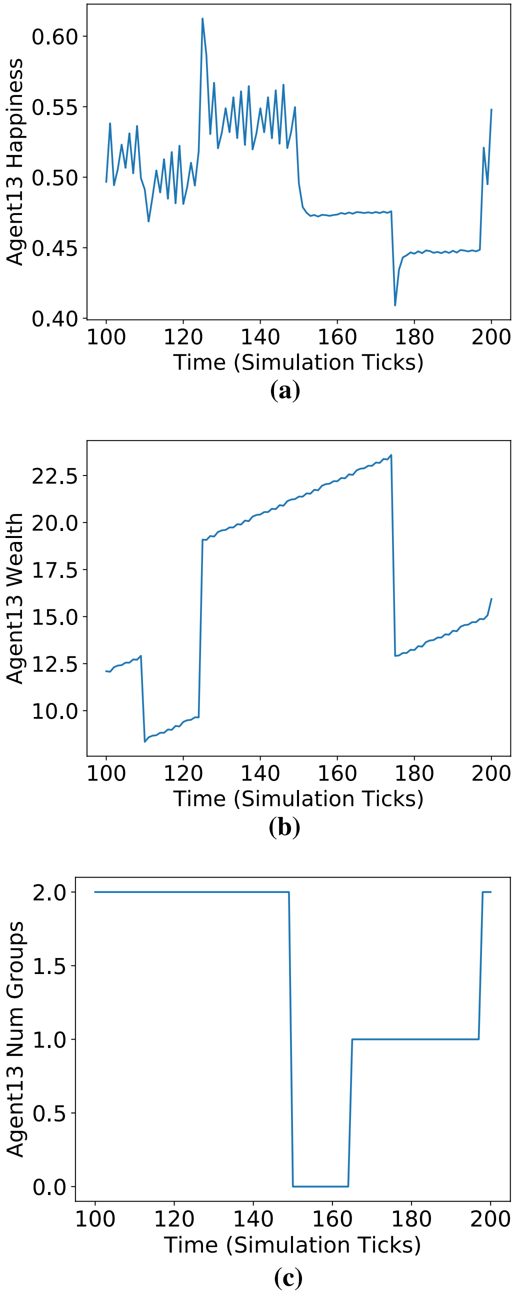

Fig. 4 Example of the factors in the update happiness action. Happiness spikes when the agent experiences large changes in wealth and drops when the agent's number of groups is less than 2, but otherwise tends to follow a baseline value

Relationship Strength and Shared Cognition values, respectively, calculated at the end of each day for use in the next day's actions. At the group level, the number of groups to which an agent belonged was an important measure of Socialization (i.e., number of groups to which an individual belonged; see above), with factors such as Group Identification determining whether that agent would stay with that group or leave. Although there were active means by which these values changed over the course of the simulation, a selection of which we describe in detail below, we also introduced a passive decay function to reflect a time-based weakening of social attributes that was based on lack of interaction. Specifically, the following attributes decayed by a certain value, specified by a decay percentage configuration parameter for each, at the end of each day: Relationship Strength, Shared Cognition, Group Reputation, and Group Influence. Along with representing real-world decay of relationship strength, it also ensured that these values did not asymptote at a ceiling level.

With the increase in scale of the model to 10,000 in Challenge 3 (v 8.0), the number of potential interactions multiplied correspondingly. To keep the model manageable at both a conceptual and computational level, it was necessary for agents in the ACCESS world to categorize their relationships to other agents into one of three social circles, based on the concept of Dunbar Layers, each with different levels of interaction possible. We describe these Social Circles first, then discuss some actions associated with individual interactions, followed by actions associated with group interactions, including the factors that influence whether individuals choose to join or leave groups. Finally, we conclude this section with a discussion of group contests and the specific social factors that influence whether groups will win or lose in these contests, as well as the consequence of wins and loses on social ties between group members.

Social Circles: Robin Dunbar noted that humans have a limit on the numbers of individuals that they know and with whom they can maintain close relationships, which is based on the limitations of human memory (Dunbar and Dunbar 1998). This notion is intuitive, as it is not psychologically realistic for individuals to track feelings toward all other individuals in the world. From the perspective of the ACCESS model, limiting the number of dyadic relationship variables that must be stored helped increase the scalability of the simulation by reducing the memory requirements of these variables from $O(n^2)$ to $O(n)$, where n is the number of agents.

In ACCESS we constructed three mutually exclusive layers. In other words, agent X can be in agent Y's Layer 1 or Layer 2 or Layer 3, but not more than 1 of these layers. The following explanation is summarized in Fig. 5.

Layer 1 comprises an agent's colleagues or regular friends. These are individuals an agent has one-on-one interactions with, or with whom they've participated in (location-based) group work. Generally, this category is no more than 50 people, but for some agents, it may be 0. A given agent or target have both dyadic Relationship Strength toward, and Shared Cognition with other agents in their Layer 1 (i.e., Cognition with Group). This layer is the pool of agents with whom a target prefers to initiate one-on-one interactions. An agent can move up into this layer from Layer 2 based on performing co-group work (with a location-based group) or experiencing one-on-one interactions. An agent can also move down out of this layer, into Layer 2 (if still a co-group member or co-located) or into Layer 3 (otherwise), if the target's dyadic Relationship Strength toward the agent drops below a threshold value, given by a configuration parameter. When this happens, the target stops tracking dyadic

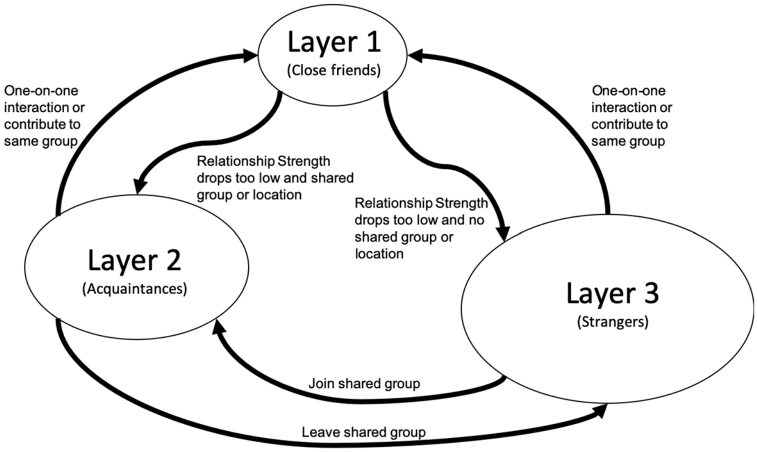

Fig. 5 Agents' relationships to other agents places them into one of three layers, but certain actions or criteria trigger moving an agent to a different layer

Relationship Strength and Shared Cognition with that agent. Layer 1 status is asymmetric because dyadic Relationship Strength is asymmetric, although Shared Cognition is symmetric. It is possible for an agent to be in a target's Layer 1, even if the two are no longer co-members of a group. Consequently, these non-grouped agents are tracked by a list.

A target's Layer 2 includes agents in the same social milieu, or acquaintances with whom the target may have occasionally shared membership in a group. Generally the population of this category is no larger than the low hundreds (depending on how many agents are in a given location and average group size/memberships). While Layer 2 represents the agents a target shares a group or location with, the target have not had significant enough contact with them to have a distinct, individual relationship. Nonetheless, the target may feel generically friendly toward them based on their general feeling about the group or location; targets have group-based dyadic Relationship Strength and group-based shared cognition (i.e., Cognition with Group) if they co-belong to any groups and location-based dyadic Relationship Strength and Shared Cognition with any agents in their shared location. Note that it is possible for a given individual to co-belong to multiple groups with another agent, or co-belong to both a group and a location with that agent. Layer 2 status is asymmetric because dyadic Relationship Strength is asymmetric.

The union of Layers 1 and 2 creates the pool of agents with whom a target may initiate one-on-one interactions, even though targets will prefer to initiate interactions with Layer 1. Indeed, by virtue of initiating an individual interaction, the initiator automatically moves into the recipient's Layer 1, and the recipient is moved into the initiator's Layer 1. If the other agent was previously in Layer 2, moving into Layer 1 establishes an individual dyadic Relationship Strength and Shared Cognition between the agents, which if they were co-group members, is based on the Relationship Strength and Shared Cognition to Group (using the group with lowest ID if they are co-members in multiple groups). If they do not belong to any of the same groups, it is based on Location Regard.

An agent moves up into Layer 2 from Layer 3, or down out of Layer 2 into Layer 3, based on changes in group membership. Notably, since location is fixed, the co-location of agents is also fixed, and therefore, co-located agents are always in each others' Layer 2. An agent who is not a member of any group is still a resident of a location, so an agent's Layer 2 is never empty. These agents would use the static attribute of Location Regard as a stand-in for both dyadic Relationship Strength and Shared Cognition in actions where those attributes are used.

All other agents that are not in Layer 1 or 2 are in an agent's Layer 3 by default. Layer 3 agents are essentially strangers. Agents have only generic "Stranger Regard" toward their Layer 3. That is, some agents may be friendlier toward strangers than others as defined by a static, but configurable parameter between zero and the minimum thresholds for Relationship Strength and Shared Cognition. However, this level of friendliness is applied uniformly across people in their Layer 3 and does not distinguish between them. Stranger regard serves as a default value for dyadic Relationship Strength and Shared Cognition with people in the world in general, when needed.

Individual Interactions: Two agents can have a one-on-one interaction on any given day. The details of what happens in these interactions are left abstract, e.g., what they say to each other, the topic and emotional tone of the interaction. We simply define how the interactions are initiated, i.e. the "who" and the "when" of the interaction, and the effects of the interaction. The high level mechanics of how interactions are initiated is that every day each agent first decides if it would like to initiate an interaction. If the agent decides to initiate an interaction, then it decides with whom it will interact and the interaction occurs. The chosen recipients of the interaction are always willing participants, i.e., agents do not refuse interactions initiated by another agent. In this way, an agent can initiate at most one interaction per day, but can be the recipient of any number of interactions initiated by other agents.

The decision of whether an agent initiates an interaction in a given day is driven by a schedule that depends on the agent's Communion value. Specifically, each agent initiates interactions every x days, where $x = 5 + ceiling(-5 \times Communion)$. In other words, agents with Communion between 0.0 – 0.2 will initiate actions every 5 days, agents with Communion in 0.21 – 0.4 initiate interactions every 4 days, and so on, ending with agents who have Communion between 0.81 – 1.0 initiating an interaction every day. Then, if an individual agent is going to interact, it would pick their interaction partner according to the following rules:

- If there are individuals in Layer 1 who haven't been interacted with in the last 5 days, pick from that set with the following criteria:

  - Pick the recipient that maximizes the following score, which is the sum of the agent's Materialism (m) multiplied by the prospective recipient's normalized Wealth (w) and the agent's Affiliation (a) multiplied by the Relationship

Strength (rs) (the subscripts i and r represent value of the initiator and recipient, respectively, and the subscript ir denotes the dyadic relationship value from i to r):

$m_i \times (w_r/max(w_r, w_i)) + a_i \times rs_{ir}$

This choice reflects the initiator's desire to interact with other individuals who are perceived as successful and with whom it has a good relationship, with the balance between the two being decided by the initiator's Materialism and Affiliation values. In other words, an initiating agent who was stronger on Materialism than Affiliation would select a recipient who had higher normalized Wealth, whereas an agent who was high on Affiliation would place greater weight on the Relationship Strength to the recipient.

– If there are no individuals in Layer 1 who haven't been interacted with in the last 5 days, choose an agent from Layer 2. Pick a group or location based on the following:

  – For each group the individual is a member of, average the initiator's Relationship Strength to group and the initiator's Cognition with group
  – For the location, take the initiator's Location Regard
  – Pick the highest scoring group or location based on those amounts.
  – Then, within the group or location, pick the individual based on which individual has the highest value of the following score, which is the same as above except that it replaces Relationship Strength with Happiness (h) since agents do not keep dyadic relationship values for other agents outside of their Layer 1:

    $m_i \times (w_r/max(w_r, w_i)) + a_i \times h_r$

After the individual interaction, if either individual has Communion greater 0.5, that individual's Happiness value increases 0.05 after an interaction. Both initiator and recipient can gain Happiness in this manner. In this way, individuals who "value" Communion receive Happiness from interactions. Additionally, if the individuals are in each others' Layer 1 at the time of the interaction, they each increase their relationship strength value to the other individual. The initiator increases their relationship strength to the recipient by 10% of the difference between the current value and 1, and the recipient increases their relationship strength value to the initiator by 5% of the difference between the current value and 1.

Group Membership and Interactions: Group membership and interactions are another important aspect of agent "life" in ACCESS. As described previously, most groups are based in a single location, but a small fraction are world-wide. Only agents that live in the location of a Location-based group can join that group, but any agent can join any or all of the world-wide groups.

Agents consider leaving and joining groups every D days:

$$D = ceil(D_{min-join} - agency \times D_{max-join})$$

where $D_{min-join}$ and $D_{max-join}$ are configuration parameters setting the minimum and maximum number of days that agents will wait between considering leaving or joining a group. On the days that agents consider leaving and joining groups, they first find the group that they are currently in with the lowest Group Identification value. If that Group Identification value is less than a configuration parameter threshold, they leave that group. Agents also leave the group with the lowest Group Identification value if the total number of groups they belong to is greater than the pre-determined GroupMembership limit, even if all those groups are over the configuration parameter threshold.

Group Identification values are based on the Theory of Optimal Distinctiveness (Brewer 1991), which posits that agents desire to have a balance of Inclusion (I) and Distinctiveness (D) between social groups. Group Identification (gi) is calculated as follows:

$$gi = (D^2 + I^2)/2$$

where Distinctiveness (D) is how much the agent feels this group is different from strangers. This factor is calculated using Stranger Regard (sr), Relationship Strength to group (rsg), and Cognition with group (cg) as follows:

$$D = \frac{\sqrt{(sr - rsg)^2 + (sr - cwg)^2}}{\sqrt{2}}$$

and Inclusion (I) is how similar the individual is to the group. Inclusion is a function of the agent's Affiliation value (a_{agent}), the group's Affiliation value (a_{group}), the agent's Materialism value (m_{agent}), and the group's Materialism value (m_{group}):

$$I = \frac{\sqrt{(a_{agent} - a_{group})^2 + (m_{agent} - m_{group})^2}}{\sqrt{2}}$$

Then, regardless of whether the agent left a group or not, it also considers joining a new group. In this process, the agent first chooses a prospective group by ruling out location-based groups of other locations, groups it had quit within the last day, and groups it was already in. Then from the remaining groups, the agent chooses the one with which its Groud Identication value would be highest and decides if it will join that group.

The decision process that an agent uses when determining if it will join a prospective group is inspired by the Theory of Planned Behavior (Fishbein et al. 1975; Feldman and Lynch 1988; Ajzen and Fishbein 2007) in which a behavior concerning a controllable action follows from an intention, which itself is a combination of an individual's perceived behavioral control (PBC) subjective norm (SN) and attitude (A). PBC also directly influences behavior, as there are cases where regardless of one's intention, one is unable to control a behavior (e.g. trying to vote when your car breaks down). For the purposes of our model, however, our interpretation of PBC relates less to external barriers and instead is a function of an agent's internal mental state. Specifically, PBC is an average of agent Happiness and Communion values, which essentially determines whether the individual thinks it is emotionally capable of joining a new group. SN is set to the prospective group's Group Reputation value, which represents how other agents view a decision to join that group. Finally, we represent attitude (A) by the average of the candidate group's win rate in Group Contests (see following section) over the last month, which represents the extent to which the agent thought joining the group would serve to improve their finances (if high in Materialism), and the average of the individual's Cognition with Group and Relationship Strength to Group, which represents whether the agent strongly thinks the group will improve their relationships (if high in Affiliation). Intention (I) is the weighted sum of PBC, SN, and A, where weights for each factor are configuration parameters. The agent will join the group, i.e., the final behavior (B), if I is greater than a threshold (given by a configuration parameter). Otherwise, the agent will not join the group and waits until the next interval to consider joining any groups again.

Group Contests - Social Factors: Groups in the ACCESS world interact with each other by competing periodically in Group Contests, which are held by default every five days (but can be changed via a configuration parameter) between pairs of groups. In this section we focus on the relationship of contests to social interaction factors, both in terms of factors that influence winning and how winning (or losing) in turn affects these intra-group factors. However, these contests also have wagers riding on them, and thus, can influence wealth of individuals in the group, which we will discuss later (see "Group Contests-Wagers" in Sect 4.6).

The actual tasks or nature of the competition of the group contests are left abstract, but we notionally determine that teamwork, cohesiveness, and common understanding are important for group tasks. Therefore, the winner of a group contest in ACCESS is determined by a group-based performance score. The baseline group performance score is calculated by taking the sum of the average Relationship Strength to Group and average Cognition with Group values of all members in each group, and thus, is a measure of group cohesiveness. The group with the higher score is deemed the winner, unless an active local and/or global (World) Policy is in effect that influences these scores (see Sect. 4.7 below).

With regard to Socialization and Social Interaction factors, winning (or losing) a contest will increase (or decrease) Group Reputation by a configurable percentage, which in turn influences the likelihood an agent will join a group. The outcome of a group contest also influences the dyadic attributes of agents on the winning and losing teams. All agents on the losing team decrease their Relationship Strength values to all other agents on the winning team by 10% of the difference between their current relationship strength and 0. They also decrease their Cognition with Group value by 10% of the differences between their current Cognition with Group value and 0. On the winning side, agents on the winning team increase their Relationship Strength values towards other members of the winning team by 10% of the difference between their current Relationship Strength and 1, and they increase their Cognition with Group value to the winning group by 10% between their current Cognition with Group and 1.

Loss of a contest by a group also triggers a group membership regulation action. Specifically, if the group values Materialism over Affiliation, it kicks out the member who has contributed the least to group work over the past 30 days, whereas if it values Affiliation more, it kicks out the member who has interacted 1:1 with the current members the fewest times over the last 30 days. Any agent that was kicked out of a group receives a one-time decrease in Happiness determined by a configuration parameter.

Rules also determine which groups will compete against each other. Each group belongs to one of a number of domains. Groups only compete against other groups in the same domain, although there could be multiple domains in a given location. Within each domain, a round-robin schedule is followed each time it is that domain's turn to hold a contest. Determination of the two groups that will compete against each other in a particular contest and where that contest would take place are then specified as follows:

- The domain that will hold the next contest rotates such that if domain d holds a contest on day t then domain $d + 1$ will hold the next contest on day $t + 5$.
- The contest is then held at the location in the world in which the greatest number of the participants, i.e. the union of the members of the competing groups, live.
- After all domains have held a single contest, the schedule returns to the beginning of the list and repeats through all the domains again.

4.6 Acquiring and losing wealth

Wealth is another critical factor in the ACCESS model and highly related to Happiness through the Social Comparison action (comparison of an agent's wealth to others). Although the Global Prosperity function of the ACCESS world automatically provides each agent with a fixed income every day,[1] there are also two primary ways in which individual agents in the ACCESS model can accumulate wealth through more "active" means: a) Solo Work and b) being a member of a winning group in a Group Contest. It should be noted that while agents can also contribute work to a group, this work is unpaid (i.e., "volunteer"). Nonetheless, the Group Contributions of an individual to a group work indirectly influences an individual's wealth via the likelihood that the group will win a Group Contest and thus, confer wealth to the individual. The baseline conditions for winning a contest are the average Relationship Strength and Cognition of the Group for group members, which increase for an individual agent when that individual chooses to volunteer their time toward Group Contributions. Group Contributions also directly increase the Group's Treasury.

There are also means by which individuals lose Wea-lth. They are charged daily "rent" to live in their world location, with certain locations being more costly than others (although agents cannot move). Additionally, if their group loses a group contest they must pay the winning group. If the loss of Wealth outpaces the gains, the

[1] The original plan for this feature was to add complexity to the world by modifying it to be a payout received only by low wealth agents and have the amount change based on a separate set of policies that all agents would vote on. The model was deemed to be sufficiently complex before this feature was implemented, though, so while the fixed income remained, it represents a rather innocuous action in the world.

Wealth value can become negative. In ACCESS, however, this does not result in loss of housing or access to basic necessities (see Sect. 4.1), but it can negatively impact Happiness especially for an individual who weighs Materialism as being more important than Affiliation.

Solo Work: Each day, every agent decides how much of its free time it would like to spend performing Solo Work. Here, we normalize the concept of free time to 1 unit of free time per day and limit the Solo Work time to a maximum of 0.5, or half of the agent's free time in a day. Agents receive compensation for their Solo Work, so agents decide how much to work based on their Materialism: $SoloWorkTime = 0.5 \times materialism^2$. Then, as a result of this Solo Work, it gain Wealth according to the time spent working and a configuration parameter, Global Prosperity, which is the same for all agents: $Wealth \mathrel{+}= SoloWorkTime \times prosperity$. In other words, agents all earn the same wage, but can earn different amounts of money each day by working more.

Group Contributions: Agents can choose to spend some fraction of the other 0.5 of their free time contributing to a group, which can be thought of as volunteering for that group. Each day, the agent first chooses which group it might contribute to according to the following process:

- If the agent is in any location-based groups, the agent chooses from these groups. Otherwise, if the agent belongs to any world-wide groups, it chooses from these groups.
- If the agent's Materialism is greater than its Affiliation, then the agent will want to contribute to the group most likely to win contests in the future and boost the agent's Wealth, so it picks the group it belongs to that had the highest productivity in the previous day. Otherwise, if the agent's Affiliation is greater than its Materialism, then the agent opts to help the group that it belongs to that most needs help, and thus choose the group it belongs to that had the lowest productivity in the previous day.

This approach to choosing which group and agent will contribute to is partially designed to help keep a small number of groups from becoming overly dominant since high affiliation agents will altruistically opt to contribute to the groups that have the most need. Once the group is chosen, the agent determines the amount of time to contribute using the formula:

$$GroupContributionsTime = m_{agent} \times (0.25) \times (rsg) + a_{agent} \times (0.25) \times (cg)$$

At the end of each day, all agents who contributed to the same group increase their dyadic Relationship Strength and Shared Cognition Values by 1% of the difference between their current values and 1. Agents also increase their Cognition with Group values toward the groups that they contributed to by 10% of the difference between their current value and 1. Further, according to a parameter, such as every 10 days, the model updated an agent's Relationship Strength to Group and Cognition with Group. Relationship Strength to Group was set to the average of that individual's average dyadic Relationship Strength with group members in Layer 1 and current Relationship to Group; Cognition with Group for an individual toward a group was set to the average of that individual's average Shared Cognition with Group for members in Layer 1 and current Cognition with Group.

Group Contests-Wagers: Every group contest has in effect a wager riding on the outcome. After a contest, the losing team transfers a configurable amount of money to the winning team (default value = 500). Each losing agent pays an equal share of the contest exchange amount, and the winnings are evenly distributed among all members of the winning team. It is possible for one or more agents to be on both the losing and winning teams. In that case, those agents would both pay their share and receive their share of the transfer.

Group Treasury: As noted above, each group has its own Group Treasury, which is the Wealth held by the group itself. Wealth is a function of Group Productivity, the sum of individual time contributed to the group by all individuals in the group, thus generally giving larger groups an advantage. It is also multiplied by two configuration parameters, the Global Prosperity factor (which also affect individuals' fixed income payments) and a Group Productivity Factor, *gpf* (which is only used in this calculation).

The following equation specifies how a group's earned money contributes to their Treasury, i.e. the group's bank account:

$$\begin{aligned} &\mathit{GroupTreasury} + \\ &\quad = \mathit{GroupProductivity} \times \mathit{gpf} \times \mathit{prosperity} \end{aligned}$$

When the Group Treasury value exceeds a threshold, a group can choose to use some of its money in one of two ways. If the group's Materialism is greater than or equal to its Affiliation, then the group will try to recruit new members by spending money on advertising to boost the group's Reputation. If the group's Affiliation is greater than its Materialism, it will instead try to increase its Influence over current members by spending money on boosting the group's Influence. The threshold at which groups spend funds and the exact costs and gains to reputation and influence are all given by configuration parameters.

4.7 World policies and elections

Policies: Each location and the world itself can have one of four policies active or no policy active on any given day. Policies are activated through elections that occur every 30 days. These policies can have direct effects on group contests (see "Group Contests-Social Factors" in Sect. 4.5 above) and influence a group's performance scores during an active contest. Only a maximum of one policy can affect a single contest. If policies are active at both the location and at the world, then only the location policy is used. If the location has no active policy, but the world has an active policy, then the world policy is used.

The four policies, their notional goals, and their specific impact on group performance scores are as follows:

- DropLowDyadicScoresPolicy: The idea of this policy is to advantage groups that are more diverse or inclusive, i.e. groups that are not only made up of small cliques of individuals with strong dyadic relationships. When this policy is active, groups may drop up to the 10 lowest dyadic scores when calculating group cohesiveness and Shared Cognition though at least one score must remain.
- GroupOnePolicy: This policy represents the possibility of an inherently advantaged group. When this policy is active, any time that Group 1 is in the contest, it receives a 10% bonus to their result in a contest. If this contest policy is enforced between two groups that are not Group 1, it has no effect.
- LeastMembersPolicy: This policy is designed to give an advantage to smaller groups. When this policy is active, the group with the least members receives a 10% bonus to their result in a contest.
- MaxHappinessPolicy: The idea of this policy is that having happy group members improve group performance. When this policy is active, the group with the individual member that has the highest happiness value receives a 10% bonus to their performance result in a contest.

If a policy affects the group contest, then the winner is the group with the highest group performance score after the above adjustments are made.

Elections: Every 30 days, there is an election for these policies in the ACCESS World. Agents make two decisions related to voting, *whether* they choose to vote and, if they do choose to vote, *how* they choose to vote on each policy.

The agent's intention to vote or not is based on the Theory of Planned Behavior (Fishbein et al. 1975; Feldman and Lynch 1988; Ajzen and Fishbein 2007) and thus, is a function of Intention (I) and whether it meets the threshold to initiate Behavior (B).

As with the decision to join a group, three parameters influence Intention (I): Perceived Behavior Control (PBC), Subjective Norms (SN) and Attitude (A). However, here PBC represents whether the person thinks voting will be easy or hard, which is calculated using the agent's money as a proxy for whether it can take time off work to go vote: If $wealth \leq 0$, $PCB = 0$, otherwise $PCB = 1$. The SN factor notionally captures whether a person will vote on a policy against their social circle's overall belief in that policy. For example, if an agent plans to vote yes on Policy X, because their favorite group endorses it, but the agent also belongs to 3 other groups that are against Policy X, then the agent's larger social circle is not enthusiastic about the agent going to vote on the policy. We use this factor as a proxy for an agent's friends having an opinion about whether it is important for everyone in the group to vote. Finally, the A factor represents whether the person thinks their vote will have an impact. We define A as a value in the set $\{0, 0.5, 1\}$, based on the current situation. If a person would vote yes on a policy that is currently not activated at both the local and global levels, then $A = 1$. If the policy is active at the local, but not global level, or vice versa, then $A = 0.5$. If it is already activated at both levels, then $A = 0$. The Intention (I) is calculated as the weighted sum of PCB, SN, and A (where weights are configuration parameters) for each policy. The final decision of whether to vote, the Behavior (B), depends on the maximum Intention value. If that value is greater than 0.5, then the agent votes. Otherwise, the agent does not vote.

When an agent decides it will vote, it decides how to vote based on group endorsements of policies. Specifically, the agent determines its favorite group, i.e., the group that it has the highest Group Identification value toward among all the groups to which it belongs. Then, if that group endorses a given policy, the agent will vote “Yes” on that policy. If the group does not endorse a given policy, the agent will vote “No” on that policy. If the agent does not belong to any groups, it will vote “No” for all policies.

With regard to choices the group makes on policies, each group could choose to endorse none, some, or all policies, and change their endorsements at each election. The decision process that groups use to determine if they will an endorse a policy is different for each policy:

- DropLowDyadicScoresPolicy: Only groups who have both a low average Relationship Strength to group and an average Cognition with Group based on a threshold given by a configuration parameter, endorse this policy.
- GroupOnePolicy: Only Group 1 endorses this policy.
- LeastMembersPolicy: Only groups that are smaller than average (comparing location-based groups to other location-based groups and world-wide groups to other world-wide groups) endorse this policy.
- MaxHappinessPolicy: Two groups endorsed this policy: groups that have a member with Happiness above a threshold, given by a configuration parameter, and location-based groups that have average happiness above a threshold number, also a configuration parameter.

Finally, after all the votes are cast, they are counted in each location and across the entire world to determine which, if any, policy should be activated at the local and global levels, respectively. The policy that received the most “Yes” votes in the location or world is the only candidate for being activated, provided that the percentage of “Yes” votes among all votes cast in that location/world is at least 50%. This majority vote policy is activated in that location/world until the next election. Ties are decided randomly among the policies that received the highest number of votes. If no policy receives 50% or more “Yes” votes, then no policy is active in that location/world until the next election.

4.8 Metacognitive-based actions

Version 8.0 of ACCESS introduces a type of metacognition for agents, where they monitor if their behavior-guiding values of Materialism and/or Affiliation are being met and if not, engage control processes to better align their current state with their goals. Specifically, if an agent’s Happiness value is below a certain threshold for a certain number of days as given by a configuration parameter then the agent performs some self reflection and decides to adjust their individual Materialism and Affiliation values. This adjustment happens as follows: If the Materialism and Affiliation values are less than 0.6 apart, swap the two values and move each value 10% of the available scale further away from each other. For example, if

$materialism = 0.4$ and $affiliation = 0.7$, set $materialism = (0.7 + 0.1 \times (1.0 - 0.7))$ and $affiliation = (0.4 - 0.1 \times 0.4)$. If Materialism and Affiliation are greater than 0.6 apart then set the currently higher one to 0.55 and the lower one to 0.45.

As a result, the agent will swap its Materialism and Affiliation values and increase the spread between them, unless the individual's values have already become polarized. In that case, the individual will "reset" their polarized attributes to 0.45 and 0.55. This update of Materialism and Affiliation values generally amounts to a rather large "personality change" since these two attributes are central in many other individual decisions, e.g. choosing with whom to interact, which group to contribute to, how much one identifies with a group, etc. This update will swap which of the two attributes is higher for the individual.

5 Accessibility to social scientists

5.1 Extracting and providing realistic data sets

The ACCESS simulation outputs all relevant information about states and actions in the simulation at every time tick to a single comprehensive log. From this log, all discoverable data can be extracted. The ACCESS team built a number of library tools to fulfill many actions that can be combined to create realistic datasets mimicking ones that social scientists can access in the real world, like census data, surveys, etc. To build these data sets, we filter and sample the complete log to extract whatever data is relevant to answering each question on the prospective data request. This process can include fulfilling many different specifications from the data requestor, such as sampling agents at random or according to specified demographic criteria.

The ACCESS team, relying on our experience performing social and behavioral science, also added the capability of a second level of data realism in the form of realistic response rates and data obfuscation. We added this realism to the provided data sets by taking actions like specifying a realistic percentage of the population that could be accessed for each method, applying realistic response rates to requests, providing some answers on a realistic scale, and adding response bias to certain data points.

Specifically, in Challenge 1, we included a number of "dummy" variables, listed in Figure 2a, that were discoverable by the research teams but had no causal impact in the world. We also included a set of data obfuscations that we designed with the help of our social science team members' expertise to emulate the real-world difficulties of collecting data from actual people. For example, we instituted mechanisms like the following, transforming data from the simulation output to research request responses:

- Survey fatigue: agents were less likely to respond to surveys over time
- Self-reporting bias: agents would lie or exaggerate certain values like under-reporting their age and over-reporting their Wealth

– Loss of granularity: some values could only be reported as where they fall in a set of ranges instead of specific values, e.g. Wealth of $1,234 reported as "$1000-$2000"

We also defined restrictions on the number and size of research requests that were meant to reflect the real-world limitations of budgets and reach of studies.

During the Ground Truth program, we started by providing access to data that included all of the above limitations, biases, obfuscations, etc., in an attempt to create realism in the accessibility to data by the research teams. Over the course of the program, however, the complexity of the data and difficulty of the challenges posed to the research teams caused us to remove these limitations and transformations.

5.2 Executing interventions

The ACCESS simulation framework was built with the ability to execute interventions, i.e. exogenously change state or force actions or behaviors to happen, in order to find answers to hypothetical questions and counterfactuals, conduct experiments in the world, and create and test predict and prescribe scenarios. Examples of specific state variables and behaviors that could be manipulated for those purposes included:

– Agents' Wealth values
– Agents' Locations
– One-on-one Interactions (force or prevent agent interactions)
– Group Membership (force agents to join or leave groups)
– Group Contests (which groups participate and which group wins)
– Group Policy Endorsements
– Governance Event Results (which policies are activated at locations and globally)

In the ACCESS simulation code, interventions are specified by an interventions file that is provided as an input to the simulation. Each intervention in the file specifies what intervention to run, what time it should be executed, and what entities (if any) are involved. In the absence of any interventions, the simulation will execute the exact same steps in successive runs when the seed to the random number generator is kept the same. Therefore, introducing an intervention allows users to witness the results of one or more interventions with everything else remaining equal. This capability is crucial to running realistic experiments, which is another novel form of accessibility that we were able to offer to the research teams. In fact, the ability to isolate interventions in this way is more controlled than real-life experiments where individual variability cannot truly be controlled, particularly in complex social environments such as those we were evaluating. This intervention mechanism also provides the ability to enact hypothetical situations and compare the outcomes to a baseline timeline where the hypothetical action is absent, thus testing outcomes of prescribed actions. This feature of ACCESS enabled us to create comprehensive predict and prescribe challenge questions that we could then objectively score in terms of accuracy.

6 Explain, predict, and prescribe challenges

The Ground Truth program contained three types of challenges: explain, predict, and prescribe. We describe our approach to executing these challenges from the ACCESS point of view. For the challenge questions, we provide only example questions. Please contact the authors for more details on exact questions used in the Ground Truth program as well as the optimal answers generated by the ACCESS simulation.

6.1 Explain test

For the Explain Test, the research teams were tasked with providing a reconstruction of the ACCESS simulation's causal model. To facilitate this test, we executed the ACCESS simulation, created a comprehensive Initial Data Package that provided an illustrative view of the ACCESS world and introduced the research teams to most of the discoverable data variables. We then provided additional data in response to requests throughout the test. As noted in Sect. 3, ACCESS's design provides a basis for constructing a ground truth causal diagram by extracting the causal relationships in each action and merging them across all actions. This ground truth causal diagram was used to score the research teams' submissions. Here, we also performed a manual inspection of submitted nodes and edges to match with the actual nodes and edges in our ground truth diagram to overcome any issues of differences in naming or semantics. This matching was done manually using descriptions provided by research teams and our own knowledge of the ACCESS world models, subjectively interpreting the submissions.

6.2 Predict test

In order to test the research teams' ability to predict outcomes in the ACCESS world, we created and provided a number of prediction challenge questions spanning the individual, group, and society levels. These questions also covered both cases in which the simulated world continued normally and when some hypothetical intervention was interjected to test research teams' abilities more comprehensively. Table 2 provides an example set of challenge questions issued by the ACCESS team for the Predict Test.

The ACCESS simulation was built to produce the exact same output when run twice with the same configuration parameters and random seed. Therefore, to provide the correct answers to the Predict Challenge Questions, we simply ran the ACCESS simulation once for each question, intervening with a hypothetical event when the questions that included one. Then we extracted the answer to the question from the output log.

Table 2 Predict questions in the ACCESS challenges covered steady-state and hypothetical predictions across the individual, group, and society levels of the world to provide a comprehensive test of prediction methods and tools

	Steady-state predictions (future events based on current state)	Hypothetical predictions (contingent on specific events/interventions)
Individual behavior	How much wealth will agent 96 have on Day 1100?	What would Agent 18's happiness be on Day 1100 if it left groups 1 and 3 on Day 1001?
Group characteristics	How many members will Group 6 have on Day 1200?	If Group 2 did not accept any new members for 200 days, what would the average relationship strength between members be on Day 1200?
Societal outcomes	Will the wealth be more equally or less equally distributed after 1000 days?	If Policy 1 is activated in all locations and globally in the election on Day 1040, will the average happiness of the population of Location 1 be higher or lower than that of the population of Location 2 at Day 1140? By how much?

6.3 Prescribe test

The ACCESS team ran the Prescribe Tests in the same manner as the Predict Tests, explained above, except providing a different set of questions that would require the research teams to specify optimal prescriptions across a variety of levels and about a number of different entities and attributes. Table 3 provides examples of challenge questions issued by the ACCESS team for a Predict Test.

Answers to the Prescribe Test challenge questions were generated similarly to those generated to answer the Predict test challenge questions. In this case, though, for each question, we ran a trial of the simulation for each possible prescription option. We then extracted the value of the quantity of interest in the question across all the possible options and found which prescription performed the best. Options were constrained by discrete choices, which simplified the calculation to the ratio of choosing the correct choice over all possible options.

7 Conclusion

The ACCESS World model was designed modularly with agents, groups, and the world itself all acting as autonomous entities that perform their own actions. The logic of these actions was inspired in many cases by established theories in the social and behavioral sciences literature to provide a level of complexity that plausibly approximates human decision-making and behavior in the real world. The entire set of actions across the different entities created interaction within and between individual, group, and society levels, creating a complex world that exhibits emergent behaviors. The combination of logging internal entities' states and intervening into the simulation provided comprehensive accessibility into the world allowing us to mimic all forms of social scientists' data gathering approaches, ranging from surveys and polls to controlled lab experiments. The combination of the simulation's complex model and accessibility allowed us to design and run the described explain, predict, and prescribe challenges. While the results of these challenges provide a number of insights into the capabilities of modern social and data science techniques, there is still much to learn from continued experimentation with ACCESS and the other Ground Truth simulated worlds.

Table 3 Prescribe questions in the ACCESS challenges covered steady-state and hypothetical predictions across the individual, group, and society levels of the world to provide a comprehensive test of prescription methods and tools

	Steady-state prescriptions (future events based on current state)	Hypothetical predictions (contingent on specific events/interventions)
Individual behavior	Which location should Agent 83 live in from Days 1001 to 1200 to maximize their group identification among all the groups they are members of on Day 1200?	If Group 8 recruited all individuals who were not in any other group to join Group 8 at Day 1001, which policy should Group 8 endorse to maximize the size of all groups on day 1100?
Group characteristics	What group should Agent 14 join to be happiest on Day 1200?	If Group 1 added all of the members of Group 2 on Day 1001, which policy should they endorse in the next election to maximize their reputation at Day 1200?
Societal outcomes	If the richest 10% of individuals transferred wealth to the poorest 10%, how much should they transfer to maximize global average happiness?	If the MaxHappiness policy was activated over the next 200 days, how often should contests be held to minimize global wealth inequality?

Acknowledgements This document does not contain technology or technical data controlled under either the U.S. International Traffic in Arms Regulations or the U.S. Export Administration Regulations. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001118C0022. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the DARPA.

References

Ajzen I, Fishbein M et al (2007) Theory of reasoned action-theory of planned behavior. Univ South Florida 67–98:1988

Bakan D (1966) The duality of human existence: an essay on psychology and religion. Rand McNally

Brewer MB (1991) The social self: on being the same and different at the same time. Personal Soc Psychol Bull 17(5):475–482

Brickman Philip, Coates Dan, Janoff-Bulman Ronnie (1978) Lottery winners and accident victims: is happiness relative? J Personal Soc Psychol 36(8):917

Diener ED, Lucas RE, Scollon CN (2009) Beyond the hedonic treadmill: Revising the adaptation theory of well-being. In: The science of well-being, pp 103–118. Springer

Dunbar R, Dunbar RM (1998) Grooming, gossip, and the evolution of language. Harvard University Press, Cambridge

Feldman JM, Lynch JG (1988) Self-generated validity and other effects of measurement on belief, attitude, intention, and behavior. J Appl Psychol 73(3):421

Festinger L (1954) A theory of social comparison processes. Hum Relat 7(2):117–140

Fishbein M, Ajzen I et al (1975) Intention and behavior. An Introd Theory Res 50:179–221

McCabe TJ (1976) A complexity measure. IEEE Trans softw Eng 4:308–320

North MJ, Collier NT, Ozik J, Tatara ER, Macal CM, Bragen M, Sydelko P (2013) Complex adaptive systems modeling with Repast Simphony. Complex Adapt Syst Model 1(1):3

Paulhus DL, Trapnell PD (2008) Self-presentation of personality: an agency-communion framework

Wood Joanne V (1989) Theory and research concerning social comparisons of personal attributes. Psychol Bull 106(2):231

Scott T. Rager is a Research Scientist in the Network and Communications Technologies business unit of Raytheon BBN Technologies. He earned his Ph.D. in Computer Science and Engineering from Penn State University. Before that he received a B.A. in Physics from the Slippery Rock University of Pennsylvania and a B.S. in Computer Engineering from Penn State both in 2009 after completing a cooperative dual-degree undergraduate program. He has published in a variety of conferences and journals including a best paper award from IEEE GLOBECOM, and he has received funding from DARPA and the Army Research Laboratory while at BBN. He is a member of both IEEE and ACM. His research interests include stochastic network optimization, distributed protocol design and analysis, and network modeling and performance analysis, including not only physical networks, but also social and information networks.

Alice Leung is a Lead Scientist in the Intelligent Software and Systems business unit of Raytheon BBN Technologies. She holds a Ph.D. in Polymer Science and Technology and a B.S. in Chemical Engineering from the Massachusetts Institute of Technology. Her research focuses on the use of computer games and simulations to shape, study, and model human behavior. Her research projects have spanned a variety of human behaviors including the effects of culture and personality on team performance, information sharing in distributed teams, indicators of betrayal and insider threat, mitigating cognitive biases, and improving counterfactual thinking.

Shannon Pinegar is a visiting assistant professor in the Department of Psychology at Miami University. She completed a Postdoctoral Fellowship in the Dynamic Learning Lab at Baruch CUNY, which included research with the Army Research Lab and Defense Advanced Research Projects Agency. She earned her PhD in Experimental Psychology from Ohio University with an emphasis on social cognitive psychology and quantitative methods. Her areas of research include the topics of metacognition and intuition, false confessions in interrogation settings, and narrative identification.

Jennifer Mangels is a Full Professor in the Psychology Department at Baruch College, a senior college of the City University of New York (CUNY), and is Doctoral Psychology Faculty at the Graduate Center of CUNY. She received her Ph.D. in Psychology (specialization in Cognitive Neuroscience) at the University of California (Berkeley), after which she completed a post-doctoral fellowship at the Rotman Research Institute (Baycrest Centre for Geriatric Care and the University of Toronto, Canada). Her work broadly investigates the influence of executive control functions on declarative memory and has received near continuous federal funding for the past 20 years from sources such as NIH, NSF, IES, ARL and DARPA, with publications in high impact journals such as *Science*, the *Journal of Cognitive Neuroscience*, and *Social, Cognitive and Affective Neuroscience*. Recent work has been highly interdisciplinary and incorporates additional social, affective and computational components.

Marshall Scott Poole is David L. Swanson Professor of Communication, Director of the Institute for Computing in the Humanities, Arts, and Social Sciences (I-CHASS), Senior Research Scientist at the National Center for Supercomputing Application at the University of Illinois Urbana-Champaign. Scott received his Ph.D in 1980 from the University of Wisconsin-Madison. Scott has taught at the University of Illinois, the University of Minnesota, and Texas A&M University. His research interests include group and organizational communication, information systems, collaboration technologies, organizational innovation, and theory construction. Scott is the author of more than 150 articles, book chapters, and proceedings publications which have appeared in publications in the Communication, Management, Organization Studies, and Information Systems disciplines. Scott has co-authored or edited twelve books including Communication and Group Decision-Making, Organizational Change and Innovation Processes: Theory and Methods for Research, and Working Through Conflict. Scott has been named a Fellow of the International Communication Association, a Distinguished Scholar of the National Communication Association, and is recipient of the Steven A. Chaffee Career Productivity Award from the International Communication and the Joseph E. McGrath Award for Lifetime Achievement in the Study of Groups from the Interdisciplinary Network for Group Research (INGRoup).

Noshir Contractor is the Jane S. & William J. White Professor of Behavioral Sciences in the McCormick School of Engineering, the School of Communication and the Kellogg School of Management at Northwestern University. He investigates how social and knowledge networks form and perform. He is the President-Elect-Select of the International Communication Association. He is also a Distinguished Scholar of the National Communication Association and a Fellow of the American Association for the Advancement of Science and the Association for Computing Machinery. He received the Lifetime Service Award from the Organizational Communication & Information Systems Division of the Academy of Management. He received the Distinguished Alumnus Award from the Indian Institute of Technology, Madras where he received a Bachelor's degree in Electrical Engineering. He has a Ph.D. from the Annenberg School of Communication at the University of Southern California.

Authors and Affiliations

Scott Rager[1] · Alice Leung[1] · Shannon Pinegar[2] · Jennifer Mangels[3] · Marshall Scott Poole[4] · Noshir Contractor[5]

Alice Leung
alice.leung@raytheon.com

Jennifer Mangels
jennifer.mangels@baruch.cuny.edu

Marshall Scott Poole
mspoole@illinois.edu

Noshir Contractor
nosh@northwestern.edu

[1] Raytheon BBN Technologies, 10 Moulton St., Cambridge, MA 02138, USA

[2] Miami University, Oxford, OH 45056, USA

[3] Baruch College and The Graduate Center, City University of New York, New York, NY 10010, USA

[4] University of Illinois, Urbana, IL 61801, USA

[5] Northwestern University, Evanston, IL 60208, USA