Environmental Modelling and Software 122 (2019) 104524

Contents lists available at ScienceDirect

Environmental Modelling and Software

journal homepage: http://www.elsevier.com/locate/envsoft

Conceptualizing and implementing an agent-based model of information flow and decision making during hurricane threats

Joshua Watts[a], Rebecca E. Morss[b,*], C. Michael Barton[a], Julie L. Demuth[b]

[a] *Center for Social Dynamics and Complexity, Arizona State University, Tempe, AZ, USA*
[b] *Mesoscale and Microscale Meteorology Laboratory, National Center for Atmospheric Research, Boulder, CO, USA*

A R T I C L E I N F O

Keywords:
Hurricanes
Evacuation decisions
Agent-based modeling
Natural hazards
Information networks
Risk communication

A B S T R A C T

This article introduces an agent-based modeling laboratory for investigating how evolving hazard information, propagated through forecaster, media, public official, and peer information networks, affects patterns of public protective-action decisions during hurricane threats. The model, called CHIME ABM, provides a platform for integrating atmospheric science, social science, and computer and information science knowledge and data to explore the complex socio-ecological dynamics of modern hazard information and decision systems from a new perspective. First, the model's interdisciplinary conceptualization and implementation is described. Results are then presented from experiments demonstrating the model's behaviors and comparing patterns of evacuation decisions when key agent parameters and the geographical population distribution, forecast skill, and storm are varied. The article illustrates how this type of theoretically and empirically informed digital laboratory can be used to develop new insights into the interactions among environmental hazards, information flow, protective decisions, and societal outcomes.

1. Introduction

As hurricanes such as Harvey, Irma, Maria, Florence, and Michael demonstrated during the 2017 and 2018 Atlantic hurricane seasons, improving hazard risk communication and decision making is critical for scientists and society (NASEM, 2017a, b; NWS, 2019). Most research on societal information flow and decision making for environmental risks is observational, using data from surveys, interviews, and other empirical methods. In the context of hurricanes (tropical cyclones), such research has developed a broad base of empirical understanding about how at-risk members of the public access and use risk information and make protective decisions (see, e.g., reviews in Baker, 1991; Dash and Gladwin, 2007; Lazo et al., 2015; Huang et al., 2016a). However, because this work typically focuses on specific populations or situations, findings can vary widely across studies (e.g., Huang et al., 2016a), and it is difficult to identify broader spatial and temporal patterns.

Many of these existing studies also utilize data at the individual or household levels, which limits their ability to elucidate how the socially interactive processes underlying hazard risk communication and response scale up to influence outcomes (e.g., Drabek, 1999; Dash and Gladwin, 2007; Taylor et al., 2009). Moreover, in today's world, forecast and warning information and its communication evolve rapidly as a hazard approaches, in conjunction with interacting individuals' information behaviors, risk perceptions, and decisions (Lee et al., 2009; Morss and Hayden, 2010; Sherman-Morris et al., 2011; Morss et al., 2017; Bica et al., 2019). Collecting the empirical data needed to analyze and understand these intersecting social, spatial, and temporal dynamics is challenging (Meyer et al., 2014; Morss et al., 2017; Demuth et al., 2018).

Computational modeling provides an opportunity to investigate hazard information flow and decisions using a fundamentally different approach—one that complements and builds on empirical studies. As discussed in Morss et al. (2017), weather prediction, communication, and decision making are interconnected components of a dynamic coupled natural-human system. When hazardous weather threatens, multiple types of actors interact to create, communicate, interpret, and use forecasts, warnings, and other risk information (Gladwin et al., 2007; Demuth et al., 2012; Morss et al., 2015; Bostrom et al., 2016; Bica et al., 2019). Within this system, the risk information available and associated uncertainties change across space and through time. Further, in the modern, hyper-connected environment, information can be widely exchanged almost instantaneously; little is known about how

* Corresponding author.
E-mail address: morss@ucar.edu (R.E. Morss).

https://doi.org/10.1016/j.envsoft.2019.104524
Received 28 September 2018; Received in revised form 4 June 2019; Accepted 24 September 2019
Available online 25 September 2019

new hazard information propagates and influences protective decisions. Agent-based modeling provides a useful toolkit for exploring the dynamics of this type of system.

This article introduces a new agent-based modeling platform for investigating relationships among environmental hazards, hazard information, information flow, and patterns in people's protective decisions. The Communicating Hazard Information in the Modern Environment (CHIME) ABM was developed through collaboration among physical scientists, social scientists, computer and information scientists, and agent-based modelers, as part of a larger, multi-method project studying the dynamic, interconnected processes that characterize the modern hazard prediction, communication, and decision-making system (Morss et al., 2017). CHIME ABM was created to provide a platform for integrating knowledge and data across disciplines to build understanding about the dynamical system of interest, connecting concurrent geophysical predictive modeling and empirical social and information science research (e.g., Anderson et al., 2016; Bica et al., 2019; Demuth et al., 2018; Fossell et al., 2017; Kogan et al., 2015; Kogan and Palen, 2018; Morss et al., 2017; Wilenski, 1999). In doing so, we aimed to develop an interdisciplinary digital laboratory for running controlled experiments with different configurations of hazard information and social interactions, under various, dynamic scenarios (see, e.g., Waldrop, 2017; Verburg et al., 2016; Rovere et al., 2016; Magliocca and Ellis, 2016).

The modeling platform was designed to represent key aspects of the real-world system of interest given our research goals, while remaining sufficiently simple to allow meaningful exploration and knowledge building (Sun et al., 2016; Buchmann et al., 2016; Allison et al., 2018). It is currently implemented for a hurricane threatening the US coastline over a five-day period, but it was designed to represent general features of weather hazard forecasting, warning, and response in order to enable a variety of future extensions. A core component of CHIME ABM is a model of hazard information flow and protective decisions with heterogeneous agents who interact via peer and media networks. The agent types represent five types of key actors in hazard forecast and warning systems: forecasters, public officials, media broadcasters, other media aggregators and communicators, and citizens with varying individual characteristics. The agents interact with each other and with a virtual world with ocean and land areas, a hurricane that moves across the landscape, and forecast and other risk information that evolves as the storm approaches.

Although agent-based modeling has previously been used in hazards research, the research discussed here differs from this previous work in several ways. One body of related work uses agent-based modeling to study evacuation planning for hurricanes (e.g., Chen et al., 2006; Zhang et al., 2009; Yin et al., 2014; Ukkusuri et al., 2017) or other hazards (e.g., Dawson et al., 2011; Wang et al., 2016; Bernardini et al., 2017). Such work focuses primarily on how collective evacuation behaviors interact with built infrastructure, in order to explore issues such as traffic demand, evacuation routing, and strategies for improving evacuation effectiveness. The emphasis of the agent-based modeling in these studies is therefore not on who decides to take protective action, when, and why, but on how people and vehicles move and interact after they decide to evacuate. Although evacuation logistics are important, they are not examined here; in CHIME ABM V1 (version 1), the agents do not physically move. Instead, we focus on how inter-agent dynamics and human-information-environment interactions influence patterns in evacuation *decisions*.

Other studies have used agent-based models to examine how hazard information diffusion or protective decision making is influenced by a population's characteristics, such as social networks or inter-individual heterogeneity (e.g., Widener et al., 2013; Rand et al., 2015; Dixon et al., 2017). Again, the research presented in this article has a complementary but different emphasis. Although our work also focuses primarily on members of the public, we design and utilize a model with multiple types of interacting agents representing key roles in the warning communication and response system. In addition, we add a new type of dynamical interaction as a central theme by modeling different types of evolving hazard information — especially forecast information and associated uncertainty — interacting with the dynamical human behaviors simulated by an agent-based model.

Another previous application of agent-based modeling for hazards investigates hazard risk management decisions on longer (multi-year) time scales (e.g., Haer et al., 2016; Reilly et al., 2017; Tonn and Guikema, 2018). The modeling presented here expands on this type of work (and much of the agent-based modeling work on human-environment interactions; e.g., Parker et al., 2003; Boone et al., 2011; An, 2012; Rounsevell et al., 2012; Filatova et al., 2013; Barton et al., 2016; Groeneveld et al., 2017; Schulze et al., 2017) by examining interactions on much shorter time scales, where different types of information and decisions are important. The research presented in this article focuses on the time scale of a single hurricane threat, although CHIME ABM could be adapted to investigate social learning and other longer-term issues.

Building on and extending previous related work, here we describe the conceptualization and implementation of a new agent-based modeling laboratory for investigating how interactions among geospatially distributed, heterogeneous agents influence and are influenced by hazard information flow, decisions, and outcomes. We then present results from a set of experiments illustrating how large-scale patterns of hazard risk management decisions can emerge from the decisions of many simplified, heterogeneous agents as they interact with each other and with their evolving physical and informational environment. For the initial model development and research shown here, we used a case study approach, performing simulations for the state of Florida, US, and two hurricanes that previously made landfall in the state. The model was designed to be flexible, however, and thus the modeling framework can readily be modified to study other regions, agent configurations, hazards, and hazard information.

Section 2 provides an overview of CHIME ABM V1 and its key components. Section 3 describes the experimental design and implementation and the analysis of the experimental output. Section 4 presents results, beginning with the spatial and temporal patterns in evacuation decisions exhibited by the model. It then investigates the sensitivity of the model's behavior to changes in three key parameter sets (the Citizen agents' weighting of different types of information, the timing of public official agents' evacuation orders, and the geographic distribution of the Citizen-agent population) and compares results for forecasts with two different levels of skill for two different hurricanes. Section 5 summarizes key results and discusses implications of this work for hurricane risk communication and future research.

2. CHIME ABM: conceptual model and implementation

CHIME ABM was developed and all reported experiments were run in the NetLogo 5.3 modeling environment (Wilenski, 1999; see Fig. 1). This section describes the major components of the model and key elements of their design. For further details, the commented code, an ODD specification (a formal, detailed model description), and supporting input files are available for download at the CoMSES model library (https://www.comses.net/codebases/5504/releases/1.4.0/).[1]

CHIME ABM V1 includes a spatially explicit modeled world representing a geographical area of interest (described in section 2.1), a dynamic hazard (i.e., a hurricane in this implementation) that moves through that world (section 2.2), evolving forecast information about that hazard (section 2.3), and a multi-agent model in which different types of hazard-related information are accessed and interpreted by

[1] Version 1.4.0 is a NetLogo 6.0.2 version of the code, which has been peer reviewed through the CoMSES Network. This version of the code also includes supporting files for model simulations in the Texas region of the US Gulf Coast (tested for Hurricane Harvey in 2017).

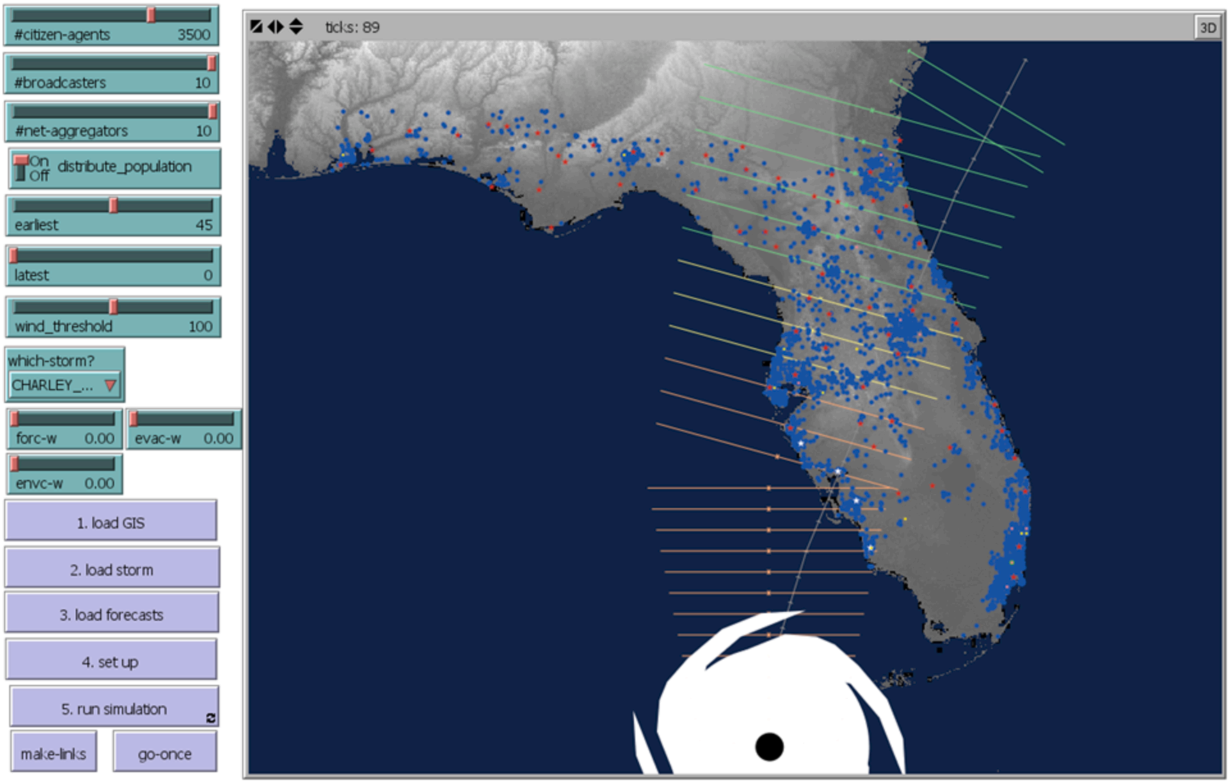

Fig. 1. CHIME ABM V1 NetLogo interface, including various user-adjustable parameters (left) and depiction of the modeled world (right). The image is a screenshot taken from a single time step during a simulation of Hurricane Charley approaching Florida, US, with historical forecasts and 3500 Citizen agents distributed according to Census data. The depiction of the modeled world includes land elevation (greyscale) and ocean (dark blue), the storm (large white tropical cyclone symbol) and actual storm track (hatched line), the most recent forecast of the storm track with uncertainty bars demarcating the "cone of uncertainty" (sequence of colored lines with asterisk at center), and agents (small colored symbols over land: blue circles = citizens, red and white stars = public officials, green circle = forecaster, yellow circles = broadcasters, pink circles = information aggregators; see section 2.4).

different agents, circulated among them, and used to assess risk and make protective decisions (section 2.4). These components are conceptually and numerically interconnected as shown in Fig. 2. The simulations shown here use a one-hour time step and run for 120 time steps, representing 5 days.

The time step, geography, and representations of the hazard, forecasts, and agents in the current implementation of the model were designed to simulate decision-making in response to a tropical cyclone threat. However, the model was designed to represent fundamental features of hazard information and decision systems and thus to be adaptable to other geophysical hazards.

To design and implement the model, we integrated expertise in agent-based modeling with knowledge and data from research and operational meteorology, emergency management, risk communication, information science, social vulnerabilities, and protective decision making. As in any modeling effort of this type, many aspects of CHIME ABM V1 are abstracted or simplified from the real world, and some real-world features and processes are not represented. Decisions about what to include in the model and how to represent it were based on our research questions, cross-disciplinary discussions among our research team, interactions with the larger research project discussed in Morss et al. (2017), and review of relevant literature. Elements of the model version presented here could readily be modified or expanded to address additional research questions, and several features of the model were designed to facilitate such future work.

2.1. Modeled world

The modeled world is a cellular representation of a geographic area of interest, which can be real or imagined, and includes land and ocean surface. In the simulations shown here, the world is the state of Florida (US) and surrounding oceans, derived from a GIS digital elevation model (DEM) in ESRI ASCII raster format with a spatial resolution of 0.5 km. We used a realistic rather than abstract geographical domain in order to provide a starting point for more complex future experiments. We selected Florida for initial model development and testing because it is an area of the mainland US that is highly susceptible to hurricanes, although it had not experienced a major hurricane landfall in several years when we began our study.

To create the modeled world used here, the DEM is imported into NetLogo and interpolated to model environment cells with a spatial resolution of 5.91 km (3.67 mi). Elevation is not used in the current version of the model algorithms but is included for potential use in future experiments, e.g., for estimating a cell's level of flood risk. Data for Florida counties, county seats, and population density (derived from US Census data for the year 2000) are also imported into NetLogo and applied to the model cells.

In CHIME ABM V1, none of the cellular landscape characteristics is dynamic, nor are features of the built environment such as buildings or roads represented. The modeled world does include hazard zones, representing geographical areas at different levels of risk, that abstractly simulate hurricane evacuation zones. In the simulations reported in this article, a single evacuation zone is defined as all locations within 1.5 cells (approximately 9 km) from the coast. As with other components of the model, these aspects of the modeled world can be revised in future experiments.

2.2. Modeled storm

The CHIME ABM V1 modeled world also includes a simulated

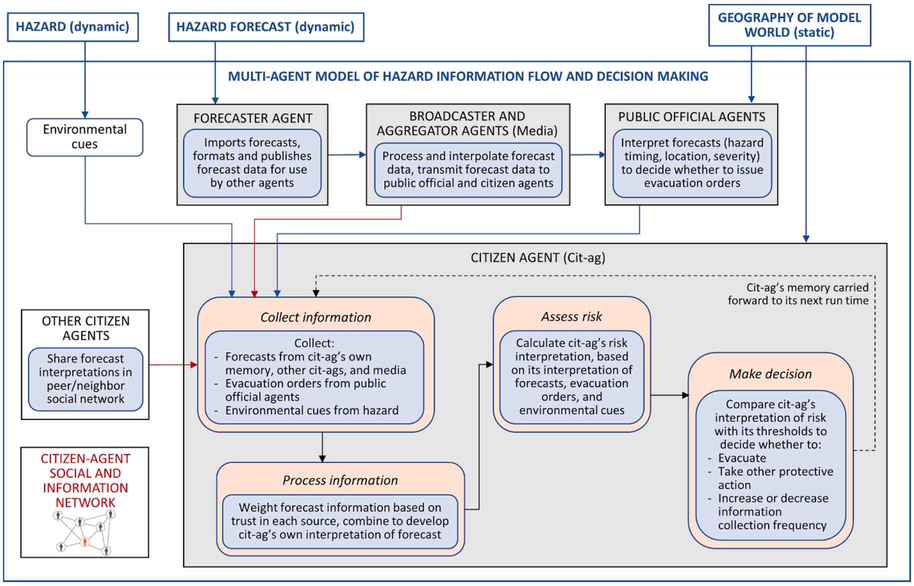

Fig. 2. Overview of the major components of CHIME ABM V1 and their interactions. The black arrows depict information flow within a Citizen-agent's decision module; the solid black arrows represent information flow within the module at a single time step and the dashed black arrow represents information retained by a Citizen agent for use at a subsequent time step. The red arrows represent information flow through a Citizen-agent's social and information network, and the blue arrows represent other types of information flow within the modeling laboratory.

tropical cyclone that approaches and then (potentially) affects the model domain. The storm can be real or synthetic; here we simulate historical storms using the US National Hurricane Center's (NHC's) Tropical Cyclone Best Track data (https://www.nhc.noaa.gov/data/?#hurdat), which include a sequence of locations of the storm's eye and other characteristics (Table 1). Future versions of the model could include more complex representations of the storm and associated hazards and impacts (e.g., areas experiencing storm surge or inland flooding, transportation disruptions, or power outages), which could then also influence information flow and decision making.

In this article, we perform simulations for Hurricanes Charley and Wilma, which made landfall in Florida in 2004 and 2005, respectively. These were selected to represent different types of storms (e.g., Charley had a small wind field for a tropical cyclone and Wilma a large wind field), with different tracks and forecast errors.

2.3. Forecast information

Along with an evolving storm, CHIME ABM simulates forecasts of the future state of the storm, which again can be real or synthetic. In CHIME ABM V1, the information in these forecasts (Table 1) is based on the format of the forecasts provided in NHC's Tropical Cyclone Forecast/Advisory products (Fig. 3). Each forecast typically provides predictive information that is valid at 12 or 24-h intervals in the future,[2] with new forecasts typically issued every 6 hours. This means that the most current forecast information available in the model evolves with time, although older forecast information may still be circulating in the information network of the multi-agent model. As discussed in the introduction and Morss et al. (2017), this is an important dynamical element of real-world modern weather forecast and warning systems.

The experiments reported here use two types of forecast information: *historical* and *ideal*. The historical forecasts were simulated using data from the NHC Tropical Cyclone Forecast/Advisory products that were available in real time as the storm being studied approached (e.g., Fig. 3). The ideal forecasts simulate perfect forecast information available at all lead times and were generated from the model's representation of the evolution of the actual storm.

Along with the forecasts, the model also includes information about forecast uncertainty. In V1, this is represented by the NHC "cone of uncertainty", which estimates uncertainty in the track forecast at different forecast lead times based on average errors in recent track forecasts (see http://www.nhc.noaa.gov/aboutcone.shtml). On average, tropical cyclone track forecast errors decrease as a storm approaches, and so the track forecast uncertainty increases with lead time. A location is defined as in the cone of uncertainty if its distance from the eye of the storm is less than the track uncertainty for that forecast lead time (Fig. 1). Here, we use historical track uncertainty data from 2005, obtained from the NHC. Note that for the experiments in this article, the same cone of uncertainty is used for both historical and ideal forecast configurations. Future versions of the model could include more complex representations of hurricane forecast information, including more detailed spatial representations of the forecast of the storm, forecasts of storm-related hazards and impacts (e.g., areas at risk from storm surge or inland flooding or power outages), and associated uncertainties.

In the multi-agent model, agents extract several variables from the forecast information for use in the risk assessments discussed in section 2.4. They search the forecast information to identify their anticipated *closest distance to storm track* (smallest distance between their location and the forecasted trajectory of the storm's eye, not considering the cone of uncertainty). They identify the anticipated *time of storm arrival* (the

Table 1
Representations of the evolving storm and evolving forecast information in CHIME ABM V1 (see sections 2.2, 2.3). Not all of the data shown are used in the current version, but they were retained for potential use in future versions. Consistent with the wind speeds in the NHC data, winds are discussed here in the unit knots (nautical miles per hour, equivalent to approximately 1.15 mph or 1.85 km/h).

Storm Characteristics (typically updated every 6 hours)	Forecast Information (updated forecasts typically issued every 6 hours)
Geographic coordinates (representing location of hurricane eye)	Forecasted geographic coordinates (location of eye), typically available at 12-24-h intervals into the future
Maximum sustained winds (representing hurricane intensity)	Forecasted maximum sustained winds, at same forecast times as geographic coordinates
Radius of 64, 50, and 34 knot wind speeds in each of 4 quadrants	Forecasted radius of 64 and 34 knot winds in each of 4 quadrants, at same forecast times as geographic coordinates
Minimum central barometric pressure	

[2] The difference between the forecast valid time and the issuance time is referred to as "*lead time*"; for example, in the NHC forecast product shown in Fig. 3, forecasts are provided at 9, 21, 33, 45, and 69 hour lead times (as well as 93 and 117 hour lead times, not shown in the figure).

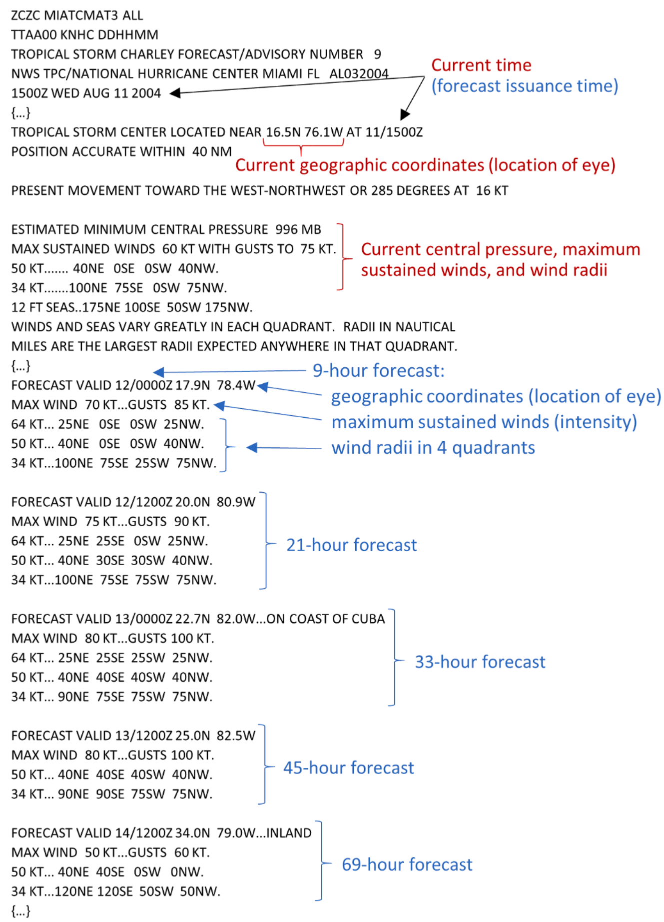

ZCZC MIATCMAT3 ALL
TTAA00 KNHC DDHHMM
TROPICAL STORM CHARLEY FORECAST/ADVISORY NUMBER 9
NWS TPC/NATIONAL HURRICANE CENTER MIAMI FL AL032004
1500Z WED AUG 11 2004
{...}
TROPICAL STORM CENTER LOCATED NEAR 16.5N 76.1W AT 11/1500Z
POSITION ACCURATE WITHIN 40 NM

PRESENT MOVEMENT TOWARD THE WEST-NORTHWEST OR 285 DEGREES AT 16 KT

ESTIMATED MINIMUM CENTRAL PRESSURE 996 MB
MAX SUSTAINED WINDS 60 KT WITH GUSTS TO 75 KT.
50 KT....... 40NE 0SE 0SW 40NW.
34 KT.......100NE 75SE 0SW 75NW.
12 FT SEAS..175NE 100SE 50SW 175NW.
WINDS AND SEAS VARY GREATLY IN EACH QUADRANT. RADII IN NAUTICAL
MILES ARE THE LARGEST RADII EXPECTED ANYWHERE IN THAT QUADRANT.
{...}
FORECAST VALID 12/0000Z 17.9N 78.4W
MAX WIND 70 KT...GUSTS 85 KT.
64 KT... 25NE 0SE 0SW 25NW.
50 KT... 40NE 0SE 0SW 40NW.
34 KT...100NE 75SE 25SW 75NW.

FORECAST VALID 12/1200Z 20.0N 80.9W
MAX WIND 75 KT...GUSTS 90 KT.
64 KT... 25NE 25SE 0SW 25NW.
50 KT... 40NE 30SE 30SW 40NW.
34 KT...100NE 75SE 75SW 75NW.

FORECAST VALID 13/0000Z 22.7N 82.0W...ON COAST OF CUBA
MAX WIND 80 KT...GUSTS 100 KT.
64 KT... 25NE 25SE 25SW 25NW.
50 KT... 40NE 40SE 40SW 40NW.
34 KT... 90NE 75SE 75SW 75NW.

FORECAST VALID 13/1200Z 25.0N 82.5W
MAX WIND 80 KT...GUSTS 100 KT.
50 KT... 40NE 40SE 40SW 40NW.
34 KT... 90NE 90SE 75SW 75NW.

FORECAST VALID 14/1200Z 34.0N 79.0W...INLAND
MAX WIND 50 KT...GUSTS 60 KT.
50 KT... 40NE 40SE 0SW 0NW.
34 KT...120NE 120SE 50SW 50NW.
{...}

Fig. 3. Excerpt from an example Tropical Cyclone Forecast/Advisory product issued by the NHC as Charley approached the US in 2004. This text is part of a single NHC forecast product issued at 15:00 UTC on August 11, 2014. It is depicted in the original NHC product format, with { … } indicating text not shown here. The full product includes forecasts out to 12:00 UTC on August 16, 2014 (117 hours). The information annotated in red represents the current state of the storm; similar information is available in the NHC's Tropical Cyclone Best Track data and is used in the model's representation of the storm (left-hand column of Table 1). The information annotated in blue is extracted into the model's representation of the historical forecast information (right-hand column of Table 1) for the forecast issued at 15:00 UTC on August 11, 2014. Similar forecast products issued by the NHC at earlier and later times are used to extract the full set of historical forecasts for the 5-day time period simulated for the Hurricane Charley experiments in this article.

forecast time corresponding to the anticipated closest distance to storm track), and they use this to calculate the anticipated *time until storm arrival* (time until the storm's eye is anticipated to arrive at their location). They also extract the forecasted maximum sustained winds at the anticipated time of storm arrival, which is referred to as the anticipated *storm intensity at arrival*.

2.4. Multi-agent model of hazard information flow and decision making

The multi-agent model of hazard information flow and decision making was designed to represent key elements of modern US weather forecast, warning, and response systems, specifically for hurricanes (e.g., Gladwin et al., 2007; Demuth et al., 2012; Bostrom et al., 2016), with features that are sufficiently general that the model could readily be adapted for other types of hazardous weather (e.g., Parker and Fordham, 1996; Brotzge and Donner, 2013; Morss et al., 2015). It includes five types (or *breeds*, in NetLogo) of agents: weather forecasters who initiate forecast information; media broadcasters who adapt and communicate information; other media aggregators and communicators; public officials who provide protective action recommendations; and citizens (members of the public) who collect and share information, assess risks, and make protective decisions. Given the goals of our research, the Citizen agents and their decision-making processes are more complex than the other agents. The other agent breeds are purposefully simple in CHIME ABM V1, but their roles can be expanded in the future.

An overview of each of the agent breeds is provided in Table 2, and each is described further below. More detailed descriptions of the agents' rules and algorithms can be found in the supporting documentation for the model.

2.4.1. Forecaster agent

Forecaster agents function as weather forecasters (modeled after, e.g., the US National Weather Service's NHC) who provide forecast information about the hazard. In CHIME ABM V1, there is one Forecaster, and its only job is to publish new forecast information (including uncertainty estimates) as it is updated, for use by other agents (Tables 2 and 3).

In the simulations shown here, the Forecaster runs every time step. Although the Forecaster has no need for a physical presence in the current implementation, it is placed at a random location in the model's populated geographic domain (represented by a small green circle, e.g., in Fig. 1). Forecaster agents were given a location to provide a starting point for potential future versions of the model in which different forecasters provide forecast information for different regions, as in the real world.

Table 2
Overview of the five agent breeds in CHIME ABM V1 and their implementation in the experiments discussed in this article.

Breed	Count	Behaviors	Geographic distribution	Temporal scheduling
Forecaster	1	Publishes forecast and forecast uncertainty information into model world	Random	Active at each time step (every hour)
Broadcaster	10	Republishes forecast information from Forecaster, reorganized and temporally interpolated for use by other agents	Random	Each Broadcaster is active at each time step
Aggregator	10	Same as Broadcasters	Random	Each Aggregator has a 1 in 3 (random) chance of being active at each time step
Public Official	67	Interprets forecasts and assesses risk to decide if and when to issue evacuation orders	One in each of the 67 Florida counties	Each Official is active at each time step
Citizen (Cit-ag)	1000	Collects forecasts from various sources (media, other Cit-ags, own memory) and other information (evacuation orders, environmental cues); processes forecasts; shares forecasts; assesses risk; and makes decisions about evacuation and other actions	Random or realistic	Each Cit-ag is active every 1–32 time steps; scheduling is set randomly at initialization and is changed if the Cit-ag decides to increase or decrease information collection frequency

Table 3
Key variables for the Forecaster, Broadcaster, Aggregator, and Public Official agent breeds and their implementation in the experiments discussed in this article. In V1, the static variables listed in this table are global (apply to all agents of that breed) and set at initialization.

Breed	Variable (s)	Type	Definition	Notes
Forecaster	*forecast*	Dynamic	Imported forecast information from data file	Historical (from NHC) or ideal
Broadcaster	*broadcast*	Dynamic	Agent's representation of the forecast	Reorganized and temporally interpolated from *forecast*
Aggregator	*info*	Dynamic	Agent's representation of the forecast	Reorganized and temporally interpolated from *forecast*
Public Official	*wind-threshold*	Static	Anticipated storm intensity (maximum sustained wind speed) over which an Official will issue an evacuation order	93-116 knots for simulations shown here
	earliest	Static	Earliest lead time at which an Official can issue an evacuation order (based on the anticipated time until storm arrival at coastal locations in its county)	0–54 hours for simulations shown here
	latest	Static	Latest lead time at which an Official can issue an evacuation order (based on the anticipated time until storm arrival at coastal locations in its county)	0 hours for all simulations shown here
	orders	Dynamic	1 if that Official has issued an evacuation order; 0 if not	Boolean

2.4.2. Broadcaster agents

Broadcaster agents simulate the role of traditional media, such as television, in communicating hazard information to a broad audience (Demuth et al., 2012); in CHIME ABM V1, their primary role is to convey forecast information to the Public Official and Citizen agents. Without changing the content of the information, Broadcasters reorganize the forecasts published by the Forecaster into a format that can be more readily used by the Public Official and Citizen agents (Tables 2 and 3). They also linearly interpolate the forecasts from the temporal resolution available in the NHC-provided information (e.g., 12- or 24-h time steps as shown in Fig. 3) to 1-h time steps, which facilitates interpretation by Public Official and Citizen agents.

In the simulations shown here, there are 10 Broadcasters, and each Broadcaster runs at each time step. Although Broadcasters have no need for a physical presence in the current implementation, they are placed at random locations in the populated domain (represented by small yellow circles). Broadcasters were given locations for potential use in future versions, e.g., to represent major media markets. In order to simplify possible influences on evacuation patterns for the experiments reported here, the Broadcasters were constrained not to modify the forecast content nor to potentially introduce delays in providing updated forecast information to other agents; this can also be modified in future experiments.

2.4.3. Information Aggregator agents

Aggregator agents are intended to simulate the roles of "new media" information actors who access, process, and redistribute hazard information, e.g., on the internet or mobile devices. These agents were included in the model to provide the ability to represent the ways in which many people currently obtain and combine information from multiple types of sources – a key feature of the modern information environment that we aim to explore (Dow and Cutter, 2000; Gladwin et al., 2007; Morss et al., 2017).

In the simulations shown here, Aggregators function identically to Broadcasters, with two exceptions: they do not run at every time step (Table 2), and they do not provide information to Public Official agents. The random activation was designed to simulate internet-based sources who may aggregate and communicate information intermittently, compared to traditional media actors who tend to communicate on a more regular schedule. In future versions of CHIME ABM, Aggregators can be revised to play more complex information roles similar to those of real-world internet-based information sources, and experiments can be run to investigate the influence of these different types of sources creating and conveying information in different ways.

The simulations shown here have 10 Aggregators, placed at random locations in the populated model domain and depicted as small pink circles. As with the Broadcasters, Aggregators were given a physical location to allow agents to have location-based preferences for Aggregator information sources in future experiments, but this location does not influence the model's behavior in the current version.

2.4.4. Public Official agents

Public Official agents (also referred to as Officials) simulate government personnel who help protect the public and inform people about protective actions (Demuth et al., 2012). In V1, Officials' only role is to decide whether to issue evacuation orders, which are conveyed to Citizen agents for use in their risk assessment (section 2.4.5).

For the simulations shown here, one Official is located in each county, at the county seat, and only Officials located in coastal counties (those exposed directly to the ocean) can issue evacuation orders. This is modeled after real-world hurricane evacuation orders, which are typically issued for areas at risk of inundation from storm surge, i.e., areas near the coast. As a simulation evolves, the modeled Officials decide whether and when to issue evacuation orders based on their assessment of the risk that the hurricane poses to coastal locations in their county, using updated forecast information obtained from Broadcasters.

At each time step, Officials assess risk by obtaining and processing forecast information and comparing it to three global parameters whose values are set at initialization: *earliest*, *latest*, and *wind-threshold* (Table 3). At each coastal cell, three criteria are evaluated: 1) is the anticipated closest distance to storm track within the cone of uncertainty (in other words, less than the track forecast uncertainty corresponding to that lead time)? 2) is the anticipated time until storm arrival within the lead time window defined by *earliest* and *latest*? and 3) is the anticipated storm intensity at arrival greater than *wind-threshold*? If all three criteria are met at any of its coastal cells, an Official will decide to issue an evacuation order (set *orders* from 0 to 1), which becomes active at that time step and remains active for the remainder of the simulation.

The track cone of uncertainty is used here as a proxy for areas that warrant evacuation orders because at lead times of more than a day or two, location-specific storm surge predictions have low skill (Fossell et al., 2017). Thus, the storm track and associated uncertainty is a reasonable first-order approximation of coastal areas at risk of significant impacts. Future versions of the model could include more complex estimates of storm surge risk at different locations based, e.g., on more complex representations of coastal geography and topography and/or translation of the atmospheric hurricane forecasts into forecasts of surge inundation.

In all of the simulations shown here, *latest* is set to 0 hours, which means that Officials can issue evacuation orders up until the storm's eye arrives at its county's coastline. The values of *wind-threshold* used in these simulations are equivalent to high Category 2 – low Category 4 winds on the Saffir-Simpson Hurricane Winds Scale (https://www.nhc.

noaa.gov/aboutsshws.php?).

Officials are depicted in Fig. 1 as small red stars, which turn white if that Official issues an evacuation order. As with the Broadcasters and Aggregators, we chose to make the Officials' behaviors relatively simple here to simplify interpreting the results of these initial experiments; the Officials' algorithms could be made more complex or their roles expanded in the future.

2.4.5. Citizen agents

Citizen agents (Cit-ags) model members of the public (or public households) who dynamically collect and process information about a potential hazard, assess the risk posed by the hazard, and decide whether the risk is sufficient to warrant changing their behaviors (Fig. 2). Cit-ags' information sources include other agents (connected through their social and information network, described in section 2.4.6) and the model's physical environment. The Cit-ags also serve as disseminators by relaying information to other Cit-ags.

The design of the Cit-ag algorithms was adapted from conceptual models of protective decision-making for hazards, such as the Protective Action Decision Model (PADM; Lindell and Perry, 2004, 2012), as well as findings from empirical research by members of our project team and other scholars on information flow and protective decision making for hurricanes (e.g., Baker, 1991; Dow and Cutter, 2000; Gladwin et al., 2001; Gladwin et al., 2007; Dash and Gladwin, 2007; Morss and Hayden, 2010; Lazo et al., 2015; Huang et al., 2016a, Morss et al., 2016a, Demuth et al., 2016; Cuite et al., 2017; Bostrom et al., 2018; Demuth et al., 2018; and references therein) and for other weather-related hazards (e.g., Mileti and Sorensen, 1990; Sorensen, 2000; Brotzge and Donner, 2013; Ruin et al., 2014; Lazrus et al., 2016; Morss et al., 2016b). One challenge of this project was to translate parsimonious theoretical models, such as that provided by PADM, and the more detailed, but often incomplete, information available from empirical analyses into simple yet sufficiently specific instructions for agents. To do so, we synthesized the relevant literature to identify key behavioral features to implement given our research goals, informed by the cross-disciplinary expertise within our research team. Existing research indicates that, in general, people evacuate when they believe that an approaching hurricane poses a risk to their own or their family's safety, and that different people can both perceive risk differently and have different evacuation barriers or constraints (e.g., Baker, 1991; Gladwin et al., 2001; Dash and Gladwin, 2007; Lazo et al., 2015). Thus, we formulated the Cit-ag module in terms of combining and translating information obtained from multiple sources into a risk assessment, which is then compared with decision thresholds that vary across the Cit-ag population.

At model initialization, CHIME ABM V1 has two options for geographically distributing Cit-ags: *random* and *realistic*. The realistic population simulations distribute Cit-ags according to the real-world population density (based on Census data, here from the year 2000). Along with location, each Cit-ag is assigned multiple variables that influence its hazard information collection, information processing, risk assessment, and decisions (Table 4). These variables were designed to capture the real-world heterogeneity among members of the public in factors such as interest in and access to hazard information, social and information connectedness, trust in information sources, risk perceptions, and interest in and capacity for evacuating and taking other protective actions.

Values for these variables are assigned individually to each Cit-ag at initialization and do not change during a simulation, except for the variable that controls scheduling when the Cit-ag is active (*feedback1*). This parameter can change if, during an active time step, a Cit-ag's risk assessment is sufficiently high or low that it decides to increase or decrease its information collection frequency. Cit-ags check for information on average every 12 hours towards the beginning of a simulation, until they assess that they may be at risk, at which point they begin to run their algorithms more frequently but never more often than the hourly time-step of the model. If they assess that they are at low risk,

Table 4
Key variables for Cit-ags and their implementation in the experiments discussed in this article. All variables are local to each Cit-ag, in other words, vary across the Cit-ag population.

Variable	Type	Definition	Notes
network-list	Static	List of other Cit-ags in Cit-ag's social and information network	See section 2.4.6
broadcaster-list	Static	List of Broadcasters in Cit-ag's social and information network	Random subset of Broadcasters (section 2.4.6)
aggregator-list	Static	List of Aggregators in Cit-ag's social and information network	Random subset of Aggregators (section 2.4.6)
trust-score	Static	Cit-ag's trust in each source of information (Broadcasters, Aggregators, other Cit-ags) in its social and information network	For each source, random between 0 and 1
self-trust	Static	Cit-ag's trust in its own previous interpretation of the forecast (*memory*)	Random between 0.6 and 1
trust-authority?	Static	Cit-ag's trust in information from Officials, used to weight evacuation orders	Random between 0 and 1
evac_zone	Static	Cit-ag's representation of whether it lives in an evacuation zone	Evacuation zone is defined as locations 1.5 grid cells or less from the coast; each Cit-ag has a 1 in 5 chance (random) of incorrectly determining its zone
risk-life	Static	Cit-ag's threshold for risk to life; if *risk-estimate* is above this threshold, the agent will decide to evacuate	Random-normal with mean 14, std. dev. 2
risk-property	Static	Cit-ag's threshold for risk to concerns other than life (e.g., property); if *risk-estimate* is above this threshold, the agent will decide to take other (non-evacuation) protective actions	Random-normal with mean 0.7**risk-life*, std. dev. 0.5
info-up	Static	Cit-ag's threshold for increasing frequency of information collection; if *risk-estimate* is above this threshold, the agent will decide to collect information more frequently (decrease *feedback1*)	Random-normal with mean 0.4**risk-life*, std. dev. 0.5
info-down	Static	Cit-ag's risk threshold for decreasing frequency of information collection; if *risk-estimate* is below this threshold, the agent will decide to collect information less frequently (increase *feedback1*)	Random-normal with mean 0.1**risk-life*, std. dev. 0.5
feedback1	Dynamic	Number of time steps before Cit-ag's next active time	Integer; set at initialization to random-normal with mean 12, std. dev. 2; changes if Cit-ag decides to increase or decrease information collection frequency; min 1, max 32
env-cues	Dynamic	1 if Cit-ag's location is currently experiencing gale-force (34-knot) or greater winds, based on	Boolean

(*continued on next page*)

Table 4 (*continued*)

Variable	Type	Definition	Notes
		the current characteristics of the storm; 0 if not	
interp	Dynamic	Cit-ag's own interpretation of the forecast	Output of Cig-ag's process information algorithm (section 2.4.5)
memory	Dynamic	Cit-ag's previous interpretation of the forecast (from the last time step it was active)	Stored previous value of *interp*
C	Dynamic	time until storm arrival at which the Cit-ag's risk function based on forecast information (eq. (1), Fig. 5) peaks	Random with mean 36, std. dev. 3; reset every run of the Cit-ag's algorithms
σ	Dynamic	Width of the Cit-ag's risk function based on forecast information (eq. (1), Fig. 5), representing how sensitive the Cit-ag's risk assessment is to the difference between anticipated time of storm arrival and *C*	Random with mean 24, std. dev. 12; reset every run of the Cit-ag's algorithms
risk-estimate	Dynamic	Cit-ag's current assessment of risk, based on forecast information and, if present, evacuation orders and environmental cues	See Table 5 and section 2.4.5

Table 5
Overview of the Cit-ag risk assessment algorithm.

Variable	Definition
Cit-ag's assessment of risk based on forecast information	Calculated using eq. (1), with *time* = Cit-ag's anticipated time of storm arrival
Error in Cit-ag's risk assessment	Random with mean 0, std. dev. 0.5; set every run of the Cit-ag's algorithms
Cit-ag's assessment of risk based on evacuation orders, if issued	If Official nearest to Cit-ag has issued evacuation orders (*orders* = 1): set to 6**trust_authority?* if Cit-ag believes it is in an evacuation zone (*evac_zone* = 1), 2.4**trust_authority?* if not; 0 otherwise
Cit-ag's assessment of risk based on environmental cues, if present	If the Cit-ag's location is exposed to > 34 knot (tropical-storm force) winds based on the current storm conditions (*env-cues* = 1): set to 3; 0 otherwise
risk-estimate, Cit-ag's final assessment of risk	Sum of the above 4 variables; passed to Cit-ag's decision-making algorithm

they increase the interval between active time steps. This dynamic scheduling for individual Cit-ags was designed to represent the ways that real people may become more (or less) attuned to their physical and informational environment as a hazard threat evolves (Mileti and Sorensen, 1990; Lee et al., 2009; Sherman-Morris et al., 2011; Lindell and Perry, 2012; Morss et al., 2017; Demuth et al., 2018).

At each time step when a Cit-ag is active, it executes four algorithms: 1) collect hazard-related information, 2) sort and process that information, 3) use the information to assess risk, 4) decide whether to change its information collection frequency, evacuate, or take other (e.g., property-protective) action (Fig. 2). An overview of each of these algorithms is below; additional details can be found in the supporting documentation.

A Cit-ag collects information by querying: a) the Official agent in its network about whether an evacuation order has been issued, b) its environment about whether its location is currently experiencing winds of 34 knots or greater (the minimum threshold for a tropical storm), c) Broadcasters and Aggregators in its media network for their forecast information, d) other Cit-ags in its peer network for their forecast information,[3] and e) its memory for its own most recent forecast interpretation. The first four (a-d) represent several of the major external sources of information that members of the public use in hurricane risk assessments and decision making: evacuation orders, physical environmental cues, and forecasts from professional and social sources (e.g., Dow and Cutter 1998; Dash and Gladwin, 2007; Morss and Hayden, 2010; Petrolia and Bhattacharjee, 2010; Lindell and Perry, 2012; Demuth et al., 2018). The last (e) represents people's tendencies to use new information (such as a newly accessed forecast) to update previous interpretations.

The Cit-ag then selects a random subset of the collected forecast information to process and combines that information, weighted by its trust in each of the information sources, to generate its own updated forecast interpretation. This forecast interpretation is saved into the Cit-ag's memory for use the next time it seeks information, and it is also now available to other linked Cit-ags in the social and information network.

Next, the Cit-ag assesses its risk (Table 5). It evaluates its risk based on forecast information using the equation:

$$\text{risk function} = H^{*}e^{-\frac{(time-C)^2}{2^{*}\sigma^2}}, \quad (1)$$

where *time* is time until storm arrival, *H* is the peak height of the curve, and *C* and σ are random variables defined in Table 4. An example risk function is shown in Fig. 4. *H* is a function of the Cit-ag's representation of whether it lives in an evacuation zone, its interpretation of the forecast, and the track forecast uncertainty (see full model description for details). Functionally, *H* increases when the Cit-ag thinks that it is in an evacuation zone, that the storm intensity at arrival will be higher, and that it is in the cone of uncertainty (or, if not in the cone of uncertainty, that its closest distance to storm track is greater relative to the track forecast uncertainty at time until arrival). This formulation was developed based on previous and concurrent work indicating that these are important factors influencing how people interpret hurricane forecasts to evaluate risk (e.g., Dow and Cutter 1998, 2000; Gladwin et al., 2001; Zhang et al., 2007; Dash and Gladwin, 2007; Morss and Hayden, 2010; Petrolia and Bhattacharjee, 2010; Huang et al., 2016a; Morss et al., 2016a; Bostrom et al., 2018).

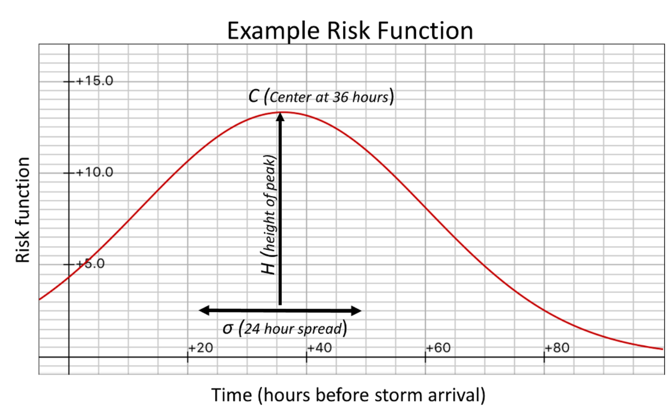

Fig. 4. Example time-sensitive risk function calculated by each Cit-ag based on its current forecast interpretation, using eq. (1). The example shown uses a typical value for *H* and mean values for *C* and σ (Table 4).

[3] Cit-ags make information available to other Cit-ags at all time steps, not only when they are active.

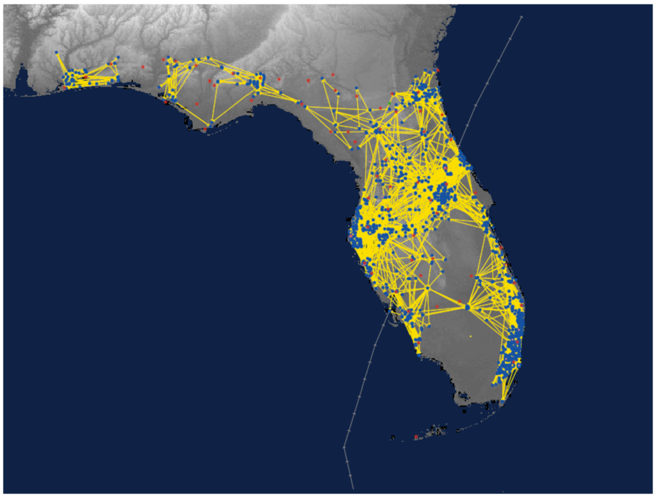

Fig. 5. Example Citizen-agent peer social and information network (section 2.4.6), for a run with 1250 Cit-ags (blue circles) distributed geographically according to Census data. The connections between Cit-ags are depicted with yellow lines.

The value of eq. (1) at the Cit-ag's anticipated time until storm arrival then becomes the Cit-ag's risk assessment based on forecast information. As shown in Fig. 4, a Cit-ag assesses higher risk when H is greater and when the anticipated time until arrival is closer to C, modulated by $\sigma\sigma$. This formulation and the values of C and σ were selected to abstractly represent heterogeneous public preferences for evacuation timing, with most members of the public evacuating between 12 and 72 hours in advance of landfall, depending on the situation (Lindell et al., 2005; Gudishala and Wilmot, 2010; Czajkowski, 2011; Wu et al., 2012; Huang et al., 2016b).[4]

The Cit-ag's final risk assessment, *risk-estimate*, is calculated by adding the risk assessment based on forecast information, a small amount of random error, and factors based on evacuation orders and environmental cues (if present), as summarized in Table 5. Evacuation orders are given greater weight if the Cit-ag has greater trust in Officials and if it believes it is in an evacuation zone (e.g., Mileti and Sorensen, 1990; Gladwin et al., 2001; Cuite et al., 2017; Thompson et al., 2017). In the current formulation, the presence of environmental cues adds a constant value to the risk assessment.

Finally, the Cit-ag decides whether and how to modify its behaviors by comparing its final risk assessment with the risk thresholds in Table 4. Specifically, *risk-estimate* values greater than *risk-life*, *risk-property*, and *info-up* trigger the Cit-ag to decide to evacuate, take other protective action,[5] and increase its information collection frequency (decrease *feedback1*), respectively. If *risk-estimate* is less than *info-down*, the Cit-ag decreases its information collection frequency (increases *feedback1*). If the Cit-ag decides to evacuate, it will no longer run its risk assessment and decision algorithms at subsequent time steps; however, it will continue to collect and interpret forecast information according to its schedule and to make its evolving forecast interpretation available to other Cit-ags in its peer network.

At initialization, each Cit-ag is visually depicted by a small blue circle (Fig. 1). If a Cit-ag decides to take a non-evacuation protective action, its color changes to green, and then to orange if it decides to evacuate.

One important simplification of the Cit-ags' algorithms in CHIME ABM V1 compared to the real world is that Cit-ags do not share or remember information other than their interpretation of the forecast. Another is that they do not consider social cues such as observations of others' protective behaviors. Although these processes are known to influence people's hazard-related risk assessments and behaviors (e.g., Mileti and Sorensen, 1990; Dash and Gladwin, 2007; Taylor et al., 2009; Lindell and Perry, 2012; Demuth et al., 2018), we chose not to include them in V1 because doing so would add free parameters and dynamics, complicating interpretation of results from the experiments. Such features could be added in future model versions to explore additional system dynamics.

2.4.6. Citizen-agent social and information network

Information sharing among agents in CHIME ABM is implemented through a social and information network that connects Cit-ags with each other and with other agent breeds. To support interpretation of other aspects of the system's dynamics, the network structure used here was designed to capture some aspects of real social and information networks, while also being relatively simple. The network changes from simulation to simulation, but in the current model formulation, it is static during a simulation. In addition, all links between Cit-ags are non-directional, such that information can flow both ways. This network formulation can be extended in complexity in the future or dynamic elements added.

In the experiments conducted here, each Cit-ag selects a random-sized, randomly selected subset of Broadcasters and Aggregators to include in its network. Each Cit-ag is also connected with the geographically closest Official.

The peer network builds connections between Cit-ags using two algorithms. First, a standard preferential attachment routine is run that creates a relatively small number of highly connected nodes in the network, in other words, a scale-free social network. This type of network is typically sparser than real social networks, and real social networks often exhibit transitivity. Thus, a second algorithm is then run which makes additional connections among Cit-ags to complete triads.

An example peer social and information network created using these algorithms is depicted in Fig. 5. Additional detail about the network-building algorithms can be found in the supporting model documentation.

3. Experimental methodology and data analysis

3.1. Initializing the model and running simulations

When CHIME ABM V1 is initialized, the inputs needed to run a simulation are loaded, including map layers, storm information, and forecasts. Next, the population of agents is distributed in the modeled world and the agents' variables are initialized. Running a simulation starts the clock and triggers the storm and forecasts to evolve and every agent to run breed-specific instructions according to its schedule. As the model runs, key variables are stored, and the depiction on the model interface is updated. At the end of each simulation, relevant data are output for further analysis. To run large numbers of simulations in parallel, we used the NetLogo Behaviorspace tool.

3.2. Experimental design

The experiments shown in this article begin with a set of model parameters that produced quasi-realistic behaviors of interest, and then systematically modify those parameters to explore key model behaviors and sensitivities. Because CHIME ABM includes a number of stochastic elements, it can exhibit significant run-to-run variability. Thus, we ran multiple repetitions for each model configuration and aggregated results across these simulations.

Table 6 provides an overview of the different sets of experiments

[4] Note that CHIME ABM V1 does not simulate the fact that in the real world, people tend to evacuate during daylight hours. This diurnal cycle in evacuation timing would be important to add if the model were used to examine issues such as travel demand and evacuation routing.

[5] This is modeled as a proxy for non-evacuation protective actions such as boarding windows, protecting other property, or gathering supplies, which people typically engage in at lower risk thresholds than evacuation.

Table 6
Parameter settings for the experiments shown in section 4. The parameters that were modified in each set of experiments are indicated by italics.

Parameter	Settings for simulations with "standard" parameter set (section 4.1)	Settings for experiments varying Cit-ags' information weighting (section 4.2)	Settings for experiments varying potential timing of evacuation orders (section 4.3)	Settings for experiments varying geographic distribution of Cit-ag population (section 4.4)	Settings for experiments varying the storm and forecasts (section 4.5)
Storm	Charley	Charley	Charley	Charley	*Charley, Wilma*
Forecasts	Ideal	Ideal	Ideal	Ideal	*Ideal, Historical*
Cit-ag population distribution	Random	Random	Random	*Random, Realistic*	Realistic
Cit-ags' weight of forecast information	1	*0, 0.25, 0.5, 0.75, 1, 1.25, 1.5, 1.75, 2*	1	1	1
Cit-ags' weight of evacuation orders	1	*0, 0.25, 0.5, 0.75, 1, 1.5, 2, 3, 4*	1	1	1
Cit-ags' weight of environmental cues	1	*0, 0.25, 0.5, 0.75, 1, 2, 3, 4.5, 6*	1	1	1
earliest possible issuance of evacuation orders (hours of lead time)	54	54	*54, 48, 42, 36, 30, 24, 18, 12, 6, 0*	54	54
wind-threshold over which Officials will issue evacuation orders (knots)	116	116	116	116	*116 (Charley Ideal), 95 (Charley Historical), 108.75 (Wilma Ideal), 93 (Wilma Historical)*
Repetitions per configuration	100	100	100	1000	1000

reported in this article. The first set of results, shown in section 4.1, is from 100 simulations for Hurricane Charley with random geographical distribution of Cit-ags and ideal forecasts. This configuration was selected as a starting point for the experiments to minimize the complicating effects of non-uniform population distribution, forecast errors, and evolving forecast information when interpreting the results. We ran these initial experiments with Charley because, unlike Wilma, the full diameter of Charley's >34 knot winds remained over land as the storm crossed Florida, and so the model's land domain encompasses more of the highest-impact zones discussed in section 3.3.

The next two sets of experiments, shown in sections 4.2 and 4.3, investigated the sensitivity of the model's behavior to key parameters in the Cit-ag and Official agent algorithms, respectively. For the experiments varying the Cit-ags' information weightings, we ran 100 simulations for each combination of the information weightings shown in the second Settings column in Table 6 (729 configurations), for a total of 72,900 simulations.[6] For the experiments varying the timing of Officials' evacuation orders, we shifted the earliest lead time at which Officials could issue evacuation orders between 54 hours and 0 hours prior to anticipated storm arrival, as shown in Table 6, and ran 100 simulations for each of the 10 configurations.

We then examined the effects of changing the geographical distribution of Cit-ags from random to realistic, holding all other settings constant (Table 6). For this set of experiments, with only two model configurations, we ran 1000 simulations per configuration (section 4.4).

Finally, we explored the evacuation patterns produced by the model when the ideal forecasts are changed to historical forecasts, and when the storm is changed from Charley to Wilma (section 4.5). As shown in Table 6, all other settings were kept constant as in previous experiments, except that *wind-threshold* was adjusted so that each of the configurations had, on average, similar numbers of Officials issuing evacuation orders. Evacuation orders are an important driver of evacuation decisions in the model, and so this modification removes some of the variability between otherwise parallel scenarios while still allowing the Officials to respond to forecast information.[7] For each of the four configurations, we ran 1000 simulations.

3.3. Data analysis

To facilitate comparing evacuation patterns quantitatively across simulations, CHIME ABM V1 tracks Cit-ag decisions within multiple *impact zones*, designed as first-order approximations of areas likely to experience different levels of impacts from the storm. Here, we use six impact zones, defined by whether a location: a) is coastal (in an evacuation zone, 1.5 grid cells or less from the ocean) or inland, and b) experiences maximum storm winds during the simulation that are greater than 64 knots (hurricane-force), between 34 and 64 knots (tropical-storm-force), or less than 34 knots. The storm tracks and six impact zones for the Hurricane Charley and Wilma simulations are shown in Fig. 6.

The primary model output data analyzed here are the percent of Cit-ags in each impact zone that decided to evacuate in each simulation. These were analyzed across the simulations run for each model configuration by examining statistics such as the mean and inter-simulation variability. We also examined Officials' evacuation order decisions, more detailed spatial and temporal patterns in Cit-ags' decisions, and other aspects of the model's behavior and outcomes. Most of the results are presented in summary tables or figures to provide a compact summary of broad patterns, together with graphics depicting more detailed aspects of the model's behavior to support more in-depth interpretations discussed in the text.

4. Results

4.1. Spatial and temporal patterns of Cit-ag evacuation decisions

First, we examine results from simulations with the model configuration shown in the first Settings column in Table 6. These results provide a first-order assessment that the agents in CHIME ABM are behaving in a structurally valid way, based on the processes included in the model. They also illustrate several key aspects of the model's behavior, which provides a starting point for interpreting subsequent

[6] The values of the weightings for the different types of information are relative scales and do not have any independent meaning. The maximum weightings used in the sensitivity tests were selected based on the increase in weighting of each type of information that was required to significantly affect evacuation rates.

[7] Note that because the storm tracks and forecasts are different in the four configurations, different sets of Officials may issue evacuation orders.

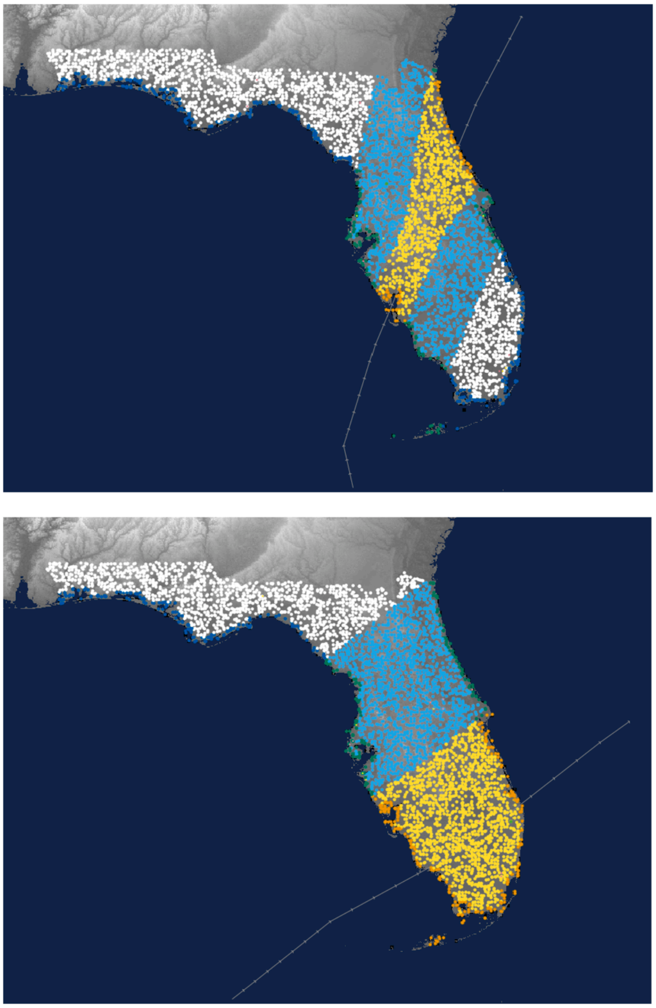

Fig. 6. Map of the storm track (hatched line) and six impact zones (represented by different-colored randomly-distributed Cit-ags) for Hurricanes Charley (upper) and Wilma (lower). The six impact zones are: coastal and >64 knot winds (orange), coastal and 34–64 knot winds (green), coastal and <34 knot winds (dark blue), inland and >64 knot winds (yellow), inland and 34–64 knot winds (light blue), inland and <34 knot winds (white).

results.

The top row of Table 7 summarizes Cit-ags' evacuation decisions in the six impact zones, averaged across the 100 simulations. To illustrate spatial patterns in the model's output in greater detail, Fig. 7 presents a map of Cit-ag evacuation decisions for a single (randomly selected) completed simulation. Note that random placement of the Cit-ags combined with Charley's small size results in few Cit-ags within the coastal >64 knot impact zone.

These results show several patterns that are similar to real hurricane evacuation behaviors. First, Cit-ags' evacuation rates are higher near the coasts (in evacuation zones) than inland. They are also higher in areas closer to the storm's track. This pattern arises because, with the ideal forecasts in this simulation, the Officials in coastal counties that will experience strong winds issue evacuation orders shortly after their evacuation order window opens, and Cit-ags in these areas receive forecasts that the storm will track near their region throughout the simulation. Cit-ags in coastal areas are much more likely to believe they are in an evacuation zone, which increases their sensitivity to both evacuation orders and forecast information. Thus, between the evacuation orders, the forecast information, and the environmental cues that they experience as the storm approaches, many of the Cit-ags in coastal >34 knot impact zones decide to evacuate.

As expected, Cit-ag evacuation rates are greater in the >34 knot zones than in the lower-impact <34 knot zones. Counterintuitively, however, modeled evacuation rates are greater in the 34–64 knot zones than in the higher-impact >64 knot zones. A more in-depth investigation indicates that this occurs because the statistics presented here average across the western coast of Florida, where the storm makes landfall and a higher percentage Cit-ags evacuate, and the eastern coast of Florida, where a lower percentage of Cit-ags evacuate. As shown in Fig. 6, the >64 knot wind area expands as the storm crosses Florida; evacuation statistics in the >64 knot impact zones are therefore more heavily weighted towards the lower evacuation rates in eastern Florida. At the same time, the model's land domain fully encompasses the 34–64 knot and <34 knot wind zones in the western part of Florida, but not in the eastern part of Florida; evacuation statistics in the 34–64 knot impact zones are therefore more heavily weighted towards the higher evacuation rates in western Florida. Together, this decreases evacuation rates in the >64 knot zones compared to those in the 34–64 knot zones.

As explained by Baker (1991), "evacuation rates vary from place to place in the same hurricane and from storm to storm in the same place" (p. 291), which complicates comparing the model results with real-world evacuation rates. Despite these limitations, this general pattern – evacuation rates that are highest in the highest-risk areas, along the coast near the storm's track, and that decrease as one moves inland and away from the storm – is broadly similar to that found in real-world hurricane evacuations (e.g., Baker, 1991; Lindell et al., 2005; Morrow and Gladwin, 2005; Huang et al., 2012, 2016b). One major difference is that the evacuation rate in the model decreases much more rapidly as one moves inland than it typically does in the real world, due in part to the simplified formulation of the influence of coastal proximity on risk assessments and decisions; this could be modified in future versions.

A second pattern illustrated by Table 7 and Fig. 7 is the inter-agent variability in Cit-ag evacuation decisions: although many Cit-ags in the highest-impact zones decide to evacuate, some do not, and a small percentage of Cit-ags who are located in low-risk areas decide to evacuate. This is consistent with real-world hurricane evacuations, and more generally with the heterogeneity exhibited by real-world U.S. individuals and households in hurricane evacuation decisions (e.g., Hasan et al., 2011; Dixon et al., 2017). In the model, this variability arises from the individual Cit-ags' different values for the randomly generated variables in Table 4. For example, a Cit-ag in a high-risk area may decide not to evacuate because it has a very high value of the risk threshold *risk-life*, or because it has very low values of the *trust-authority?* parameter or erroneously thinks it is not in an evacuation zone. In these simulations, none of the agents misinterpret the forecasts and Cit-ags do not consider evacuation orders from distant Officials. Thus, Cit-ags who evacuate from very low-risk areas (such as the one represented by the white dot in northern Florida in Fig. 7) do so primarily because they have low values of the risk threshold for evacuation, although preferences for evacuating early (when the cone of uncertainty covers a larger area) and other factors can play a role.

Variability in Cit-ag behaviors can also result from more complex interactions among components of the model's dynamics. For example, a Cit-ag in a high-risk area may not evacuate because the times at which it collects information (and thus receives information indicating that it is at high risk) do not coincide with the timing of peaks (C) in its risk function curve. This is more likely if a Cit-ag collects information infrequently (*feedback1* is large and/or *info-up* and *info-down* are high) and has a narrow risk function curve (σ is small) when it runs the risk assessment algorithm.

To illustrate temporal patterns in the model's behavior, Fig. 8 depicts the timing of Cit-ag evacuation decisions in each of the six impact zones, averaged across the 100 simulations. These results show two peaks in evacuation timing: one at 30–54 hours before anticipated storm arrival, and a smaller bump at 0–6 hours before arrival. A small percentage of Cit-ags decide to evacuate prior to the issuance of evacuation orders

Table 7
Percentage of Cit-ags deciding to evacuate in each of the six impact zones and overall, for the model configuration in the first Settings column in Table 6 and a subset of the model configurations in the second Settings column. For each model configuration and impact zone, the mean and standard deviation are shown, calculated across the 100 simulations for that model configuration. Fc = weighting of forecast information, Ord = weighting of evacuation orders, ECue = weighting of environmental cues.

Weighting of 3 information types used in Cit-ags' risk assessment algorithm		Mean (std dev) % of Cit-ags deciding to evacuate						
		Coastal >64 knot zone	Coastal 34–64 knot zone	Coastal <34 knot zone	Inland >64 knot zone	Inland 34–64 knot zone	Inland <34 knot zone	All zones
"Standard" information weighting	Fc = 1; Ord = 1; ECue = 1	32.4 (13.2)	44.3 (9.5)	20.4 (5.6)	1.2 (0.8)	1.4 (0.6)	1.0 (0.6)	3.9 (0.6)
Forecasts only	Fc = 1; Ord = 0; ECue = 0	16.4 (10.1)	20.4 (7.6)	11.8 (4.9)	1.1 (0.7)	1.1 (0.6)	0.6 (0.5)	2.2 (0.4)
Orders only	Fc = 0; Ord = 1; ECue = 0	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
Env. cues only	Fc = 0; Ord = 0; ECue = 1	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
Forecasts + Orders	Fc = 1; Ord = 1; ECue = 0	31.4 (12.9)	45.7 (10.6)	21.1 (5.9)	1.1 (0.8)	1.5 (0.7)	1.1 (0.5)	4.0 (0.6)
Forecasts + Env. Cues	Fc = 1; Ord = 0; ECue = 1	18.1 (10.0)	21.5 (7.8)	11.6 (4.9)	1.1 (0.9)	1.1 (0.5)	0.5 (0.5)	2.2 (0.4)
Orders + Env. Cues	Fc = 0; Ord = 1; ECue = 1	0.0 (0.0)	0.04 (0.4)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
Forecasts only – maximum weighting	Fc = 2; Ord = 0; ECue = 0	90.2 (9.0)	91.7 (5.2)	86.2 (4.2)	48.6 (4.1)	49.0 (2.6)	37.7 (2.6)	48.2 (1.5)
Orders only – maximum weighting	Fc = 0; Ord = 4; ECue = 0	15.3 (9.7)	21.8 (8.2)	9.5 (3.5)	0.3 (0.5)	0.4 (0.5)	0.4 (0.5)	1.7 (0.4)
Env. cues only – maximum weighting	Fc = 0; Ord = 0; ECue = 6	22.0 (9.6)	8.3 (4.9)	0.0 (0.0)	21.0 (3.4)	8.8 (1.5)	0.0 (0.0)	7.1 (0.8)

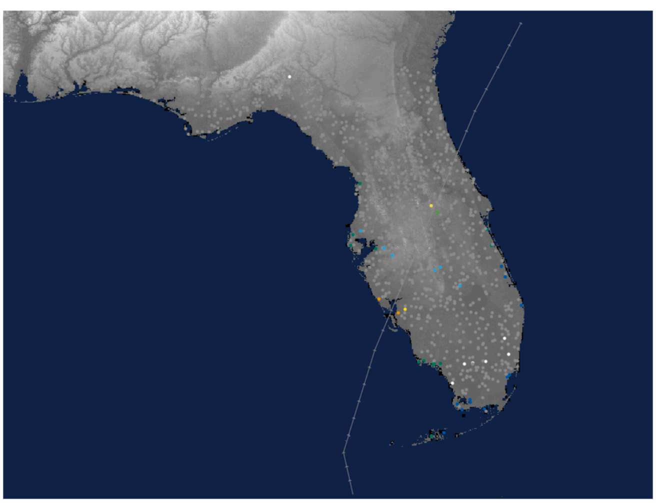

Fig. 7. Spatial pattern of Cit-ag evacuation decisions for a single CHIME ABM V1 simulation with random geographical distribution of Cit-ags, ideal forecasts for Hurricane Charley, and other parameter settings in the first Settings column in Table 6. Cit-ags deciding to evacuate in this simulation are color-coded according to their location in one of the six impact zones depicted in Fig. 6. Cit-ags deciding not to evacuate are colored gray.

(prior to approximately 54 hours), based on their interpretations of forecast information. This then leads into the first peak in evacuation decisions, which results from Cit-ags' use of evacuation orders combined with forecast information. The second peak in evacuation decisions occurs as the storm approaches close enough to provide Cit-ags with environmental cues, along with evacuation orders and forecasts.

The first peak in evacuation timing is much larger in coastal impact zones than inland, because Cit-ags in coastal zones are more likely to both receive evacuation orders and believe they live in an evacuation zone (the latter of which makes them more sensitive to both forecasts and evacuation orders). The second peak occurs only in the >64 knot and 34–64 knot impact zones, because only Cit-ags in those impact zones receive environmental cues. The variability in when Cit-ags decide to evacuate arises for reasons similar to those discussed above. These include differences in Cit-ags' scheduling and the timing of their risk function peak, as well more complex interactions such as a higher *risk-life* threshold that leads some Cit-ags to need to accumulate more information indicating that they are at high risk before they decide to evacuate.

4.2. Varying Citizen-agents' weighting of different types of information

Building on the results in section 4.1, next we investigate the effects of modifying Cit-ags' weightings of the three main types of hazard information they use to assess risk in CHIME ABM V1: forecast information, evacuation orders, and environmental cues (Table 6). These results further elucidate key aspects of the modeled system's dynamics, and they provide additional insight into the roles of different types of information in Cig-ags' risk assessments and decisions. Table 7 summarizes Cit-ags' evacuation decisions in the six impact zones for a subset of the perturbed information weightings. To examine aspects of these results in greater depth, Fig. 9 depicts evacuation rates and peak evacuation timing for Cit-ags in the highest-impact zone, across all 729 model configurations, and Fig. 10 compares more detailed evacuation timing results for the highest-impact zone for several of the information weightings.

The results in Table 7 and the overall pattern of symbol sizes in Fig. 9 indicates that in the model's current formulation, the Cit-ags' risk assessments and decisions are more sensitive to forecast information than they are to evacuation orders and environmental cues. For example, when Cit-ags use only forecast information at its standard weighting ("Forecasts only" row in Table 7), approximately 16% of Cit-ags in coastal impact zones evacuate. In contrast, when Cit-ags use only evacuation orders and environmental cues at their standard weightings, alone or together, very few Cit-ags evacuate. In the absence of forecast information, Table 7 and Fig. 9 indicate that increasing the weighting of evacuation orders and environmental cues by a factor of 4 or more is needed to motivate a substantial percentage of high-impact Cit-ags to evacuate.

Aspects of these model behaviors can be modified by changing the information weightings. However, Cit-ags' evacuation decisions are also more sensitive to forecast information due to features of the different

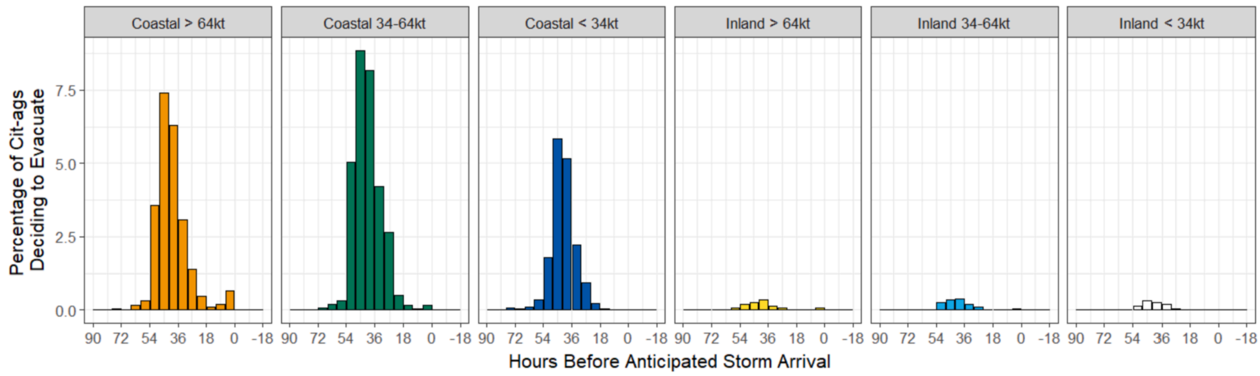

Fig. 8. Histograms of percentage of Cit-ags deciding to evacuate in sequential 6-h bins as a hurricane approaches and arrives in CHIME ABM V1, for each of the six impact zones. Results are for the model configuration in the first Settings column in Table 6, averaged over 100 simulations. The x-axis is the Cit-ags' anticipated time until the storm's arrival at their location (equivalent to actual time until storm arrival at their location, in these runs with ideal forecasts).

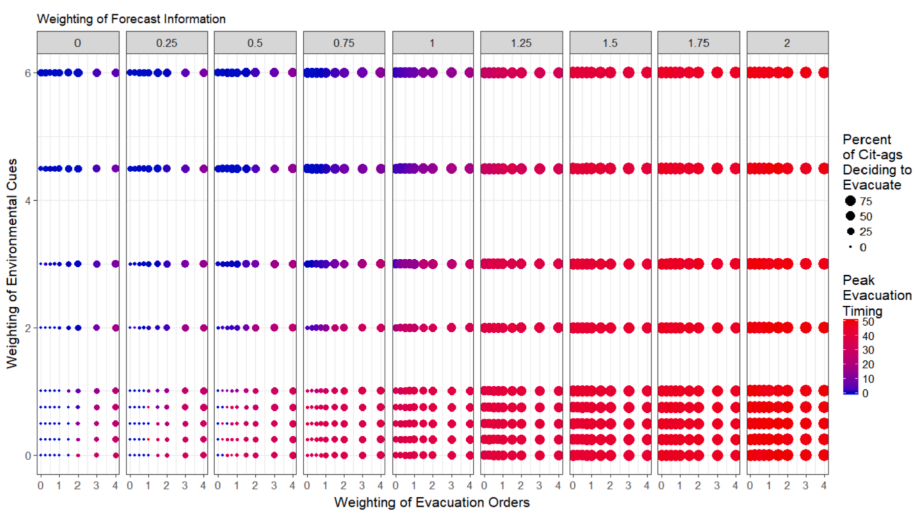

Fig. 9. Summary of how evacuation decisions for Cit-ags in the highest-impact zone (coastal >64 knots) vary with the information weightings in the Cit-ags' risk assessment algorithm (second Settings column in Table 6). The scale on the top represents the weight of forecast information (0–2), the scale on the left represents the weight of environmental cues (0–6), and the scale on the bottom represents the weight of evacuation orders (0–4). For each set of information weights, the size of the circle depicts the percentage of Cit-ags in the highest-impact zone deciding to evacuate, and the color of the circle depicts the peak timing of Cit-ag evacuation decisions, both averaged over 100 simulations for each model configuration.

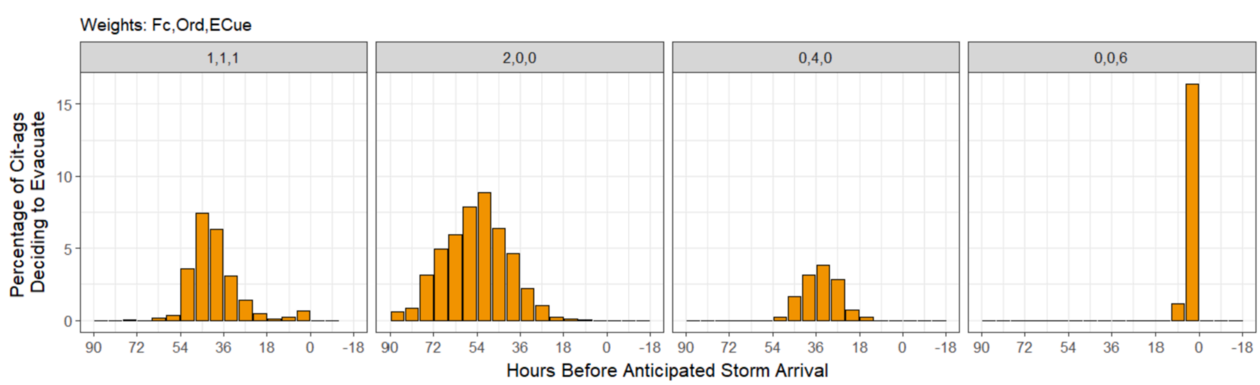

Fig. 10. Histograms of percentage of Cit-ags in the coastal >64 knot impact zone deciding to evacuate in sequential 6-h bins, for 4 of the model configurations with different Cit-ag information weightings (second Settings column in Table 6). The panels depict results for, from left to right, the standard information weightings (same as the left-most panel in Fig. 8) and forecast information only, evacuation orders only, and environmental cues only, each at the maximum weight tested. Results are averaged over 100 simulations with random geographical distribution of Cit-ags and ideal forecasts for Hurricane Charley.

types of information and the model's formulation. First, forecast information is available to all Cit-ags throughout the simulation, while evacuation orders and environmental cues only become available to a subset of Cit-ags (those whose closest Official issues evacuation orders or who experience >34 knot winds, respectively), later in the simulation. Second, forecast information is more influential because Cit-ags can modify their information behaviors in response to the hazard risk. More specifically, when Cit-ags receive forecast information that begins to signal risk, they may decide to collect information and assess risk more frequently. They are then more likely to obtain information from evacuation orders and/or environmental cues soon after it becomes available, at times that are closer to the peak of their risk function. Without forecast information playing this role, many Cit-ags wait 12 or more hours between active times, which may lead to significant delays before they obtain and process evacuation orders and environmental cues.

In this way, the forecast information in the model helps Cit-ags become more attuned to the risk and access risk information more frequently by the time evacuation orders are issued and environmental cues are felt, increasing their likelihood of evacuating. Similar behaviors occur in the real world, in which risk information available at earlier stages of a hazard threat primes people to obtain and understand subsequent information (e.g., Mileti and Sorensen, 1990; Dash and Gladwin, 2007; Morss et al., 2017; Demuth et al., 2018). However, relatively little is known about how these types of behaviors influence system-level patterns. This illustrates how modeling laboratories such as CHIME ABM V1 can provide a toolkit for studying such interactions in a simplified context, to build understanding that can be used to interpret real-world information and decision dynamics (Morss et al., 2017).

The results in Table 7 also illustrate how the influence of different types of information in CHIME ABM V1 varies spatially, based on how the spatial attributes of the information intersect with the model algorithms. As indicated by the "Forecasts only" and "Orders only – maximum weighting" rows, both forecasts and evacuation orders affect Cit-ag evacuations in all six impact zones. However, their influence is much larger in coastal zones, where Cit-ags are more likely to receive evacuation orders and believe they are in an evacuation zone. Within coastal areas, forecasts and evacuation orders have less influence farther from the storm's track (in the <34 knot zone[8]). As indicated by the "Env. cues only – maximum weighting" row, on the other hand, environmental cues only affect evacuations in the 34–64 knot and >64 knot zones. They have no effect in the <34 knot zones (where Cit-ags do not receive any environmental cues), and they are most influential in the >64 knot zones (where winds are typically greater than 34 knots for a longer period of time than in the 34–64 knot zones). Unlike forecasts and evacuation orders, environmental cues have similar effects on evacuations in coastal and inland zones.

The effects of varying the Cit-ag information weightings on the timing of Cit-ag evacuations is depicted by the colors of the symbols in Fig. 9 and the timing histograms in Fig. 10. When Cig-ags weight forecast information highly (right side of Fig. 9, second plot from left in Fig. 10), evacuations peak between approximately 78–30 hours before storm arrival. When evacuation orders are Cit-ags' primary source of risk information and weighted highly (third plot from left in Fig. 10), the evacuation peak shifts to 24–48 hours before arrival. This occurs because Cit-ags have no information available to signal risk until after Officials' evacuation order window opens and because Cit-ags are not primed by the forecasts to quickly obtain and assess this risk information. When environmental cues are Cit-ags' primary information source (top left of Fig. 9, right plot in Fig. 10), evacuations peak only a few hours before the storm's arrival, because this is when environmental cues manifest and again because Cit-ags have not been primed by earlier information. These results further illustrate the important role of forecasts and evacuation orders for motivating timely protective behaviors in the modeled system, as in the real world.

4.3. Varying potential timing of evacuation orders issued by Public Officials

Now we investigate the effects of modifying a key component of the Public Official agent algorithms: the timing of their evacuation orders (Table 6). These experiments build on the results examined in sections 4.1 and 4.2, related to the model's behavior and the roles of different types of information in the system. They also begin to explore interactions between the evolving forecast uncertainty and the dynamics within the multi-agent model.

The results are summarized in Table 8, which presents Cit-ag evacuation rates in the six impact zones for the full set of experiments. To examine aspects of these results in greater detail, Fig. 11 depicts the timing of Officials' evacuation orders (black dots) and Cit-ag evacuation decisions (histograms) for a subset of the experiments. The results for the Officials' evacuation orders in Fig. 11 show that, as expected given the ideal forecasts and the model's formulation, shifting the opening of Officials' evacuation order window later (closer to storm arrival) leads to later issuance of evacuation orders. It also typically leads to fewer Officials issuing evacuation orders. This occurs primarily because as the time until storm arrival decreases, the track forecast uncertainty decreases, which means that fewer coastal counties intersect with the forecast cone of uncertainty; fewer Officials therefore decide that evacuation orders are needed.

The results in Table 8 show that shifting the potential timing of evacuation orders later also leads to fewer Cit-ags deciding to evacuate. This effect is most prominent in the coastal zones, where (as discussed in section 4.2) evacuation orders have the largest influence on evacuations. Fig. 11 indicates that, as one would expect, shifting the timing of evacuation orders changes the timing of Cit-ag evacuations. Together, Table 8 and Fig. 11 show that shifting evacuation orders later decreases earlier evacuation decisions and shifts some (but not all) of those evacuations later (closer to storm arrival).

Building on section 4.2, this set of experiments further elucidates the influence of different types of information on evacuations. For example, for the coastal >64 knot zone results with *earliest* = 24 hours in Fig. 11, 3 peaks in Cit-ag evacuation timing are evident, at 30–54 hours, 18–24 hours, and 0–6 hours prior to storm arrival. The first peak is due to forecast information, and the latter 2 peaks correspond to the times at which evacuation orders and environmental cues become available to signal risk.

In addition, these results depict how the model simulates the trade-offs between Officials issuing evacuation orders earlier (when the track uncertainty is greater) versus waiting until closer to the storm's arrival (when the forecast uncertainty is reduced). The former leads to more Officials issuing evacuation orders, which increases the percentage of Cit-ags evacuating from both high-impact areas and areas that end up not experiencing significant impacts from the storm. The latter leads to more geographically targeted evacuation orders but a lower percentage of Cit-ags evacuating from high-impact areas; it also leads to later evacuations, giving Cit-ags less time to complete the evacuation process. This illustrates how this type of model can be used to explore, in a simplified context, the effects of different communication and decision strategies by professionals during hazardous weather threats.

4.4. Varying geographical distribution of Citizen-agent populations

Next, we relax one of the idealizations in the experiments in sections 4.1-4.3 and investigate the model's behavior when the random geographical distribution of Cit-ags is changed to a realistic geographical distribution (Table 6). We present these results to explore the effects of using a more realistic (less idealized) model setup, and to help interpret the results from subsequent experiments with realistic population distributions.

Fig. 12 summarizes Cit-ag evacuation decisions for the two model configurations. To examine the spatial patterns in greater detail, Fig. 13 depicts a map of Cit-ag evacuation decisions for a single completed simulation with a realistic Cit-ag geographical distribution. Overall, the simulations with a realistic population distribution exhibit some patterns similar to those discussed above for a random population distribution. For example, Cit-ag evacuation rates remain higher in coastal than in inland impact zones. However, on average, a much larger percentage of Cit-ags decide to evacuate in the realistic population distribution simulations. There is also a shift in the spatial pattern of

[8] As discussed in section 4.1, evacuation rates are higher in the 34–64 knot zones than in the >64 knot zones because of the asymmetric, evolving nature of the storm (which leads to an asymmetric distribution of these zones across the populated model domain; Fig. 6) combined with the limited model domain and the averaging of evacuation rates across the west and east portions of Florida.

Table 8

Percentage of Cit-ags deciding to evacuate in each of the six impact zones and overall, for the model configurations in the third Settings column in Table 6. For each model configuration and impact zone, the mean and standard deviation are shown, calculated across the 100 simulations for that model configuration.

Value of *earliest* possible issuance of evacuation orders (hours prior to anticipated storm arrival)	Mean (std dev) % of Cit-ags deciding to evacuate						
	Coastal >64 knot zone	Coastal 34–64 knot zone	Coastal <34 knot zone	Inland >64 knot zone	Inland 34–64 knot zone	Inland <34 knot zone	All zones
54 (standard)	31.3 (13.6)	44.9 (10.4)	20.3 (5.5)	1.4 (0.7)	1.4 (0.7)	1.0 (0.6)	3.8 (0.6)
48	32.2 (13.9)	45.3 (10.7)	20.5 (5.9)	1.4 (0.9)	1.5 (0.7)	0.9 (0.6)	3.9 (0.6)
42	30.2 (14.8)	32.5 (8.9)	20.2 (5.5)	1.4 (1.0)	1.4 (0.7)	0.9 (0.5)	3.4 (0.7)
36	34.7 (12.2)	31.4 (9.0)	13.5 (4.8)	1.3 (0.8)	1.4 (0.7)	0.6 (0.5)	3.0 (0.5)
30	29.2 (13.4)	31.7 (9.2)	14.3 (4.8)	1.3 (0.8)	1.2 (0.6)	0.7 (0.5)	3.0 (0.6)
24	28.6 (12.0)	29.8 (8.8)	12.6 (4.7)	1.3 (0.7)	1.3 (0.6)	0.5 (0.5)	2.8 (0.5)
18	25.6 (11.0)	25.6 (8.6)	11.5 (4.2)	1.2 (0.8)	1.2 (0.7)	0.6 (0.5)	2.5 (0.5)
12	22.9 (11.4)	23.2 (9.0)	12.1 (3.9)	1.3 (0.8)	1.1 (0.6)	0.6 (0.5)	2.4 (0.5)
6	18.0 (10.3)	21.1 (7.2)	11.7 (4.1)	1.1 (0.8)	1.1 (0.6)	0.5 (0.5)	2.2 (0.4)
0	16.7 (10.1)	21.9 (7.4)	10.7 (3.9)	1.2 (0.9)	1.1 (0.6)	0.5 (0.5)	2.2 (0.4)

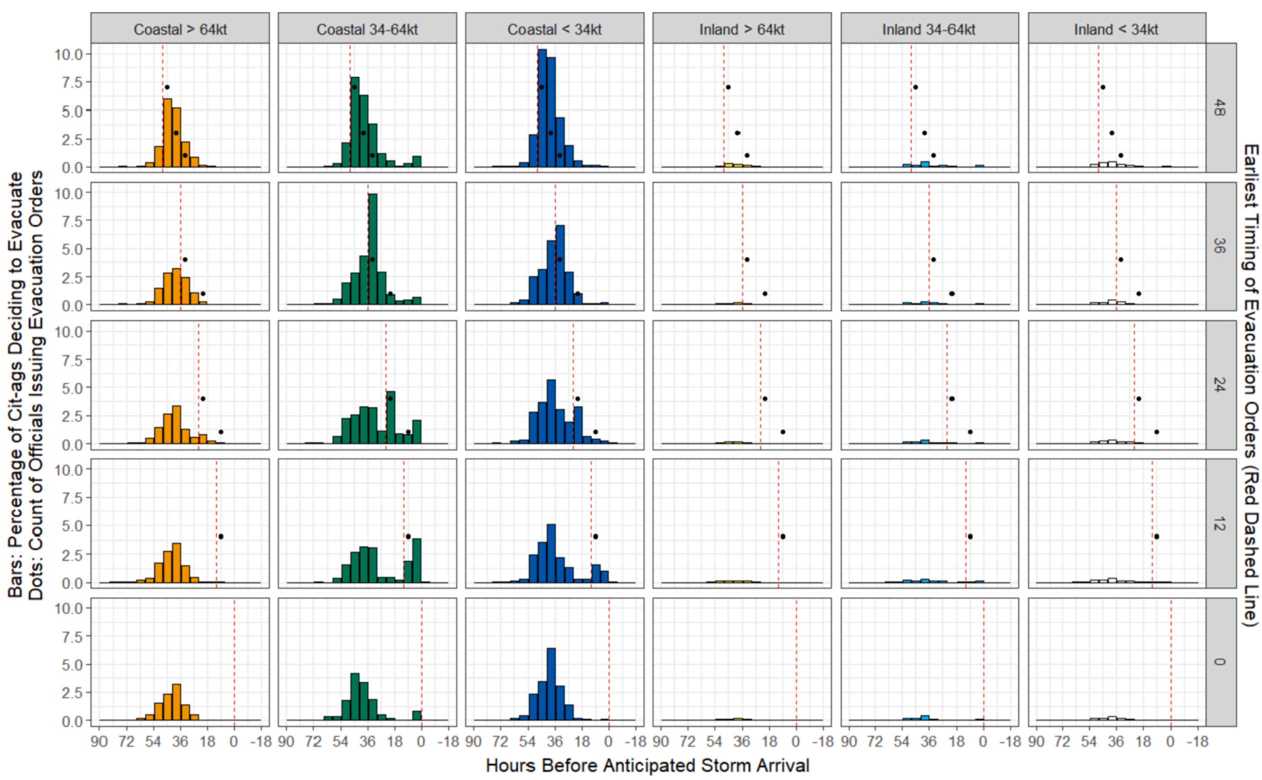

Fig. 11. Histograms of percentage of Cit-ags deciding to evacuate in sequential 6-h bins as in Fig. 8, for different values of the parameter *earliest*. The dashed red lines depict time = *earliest* (the earliest lead time at which Officials can issue evacuation orders). The black dots represent the number of Officials who issued evacuation orders in each 6-h bin, averaged over the 100 simulations for each value of *earliest*; for 6-h bins with no black dot, no Officials issued evacuation orders. Results are shown for a subset of the model configurations in the third Settings column in Table 6, with *earliest* set, from top to bottom, to 48, 36, 24, 12, or 0 hours before anticipated landfall.

evacuations, including a decreased evacuation rate in the coastal >64 knot zone and an increased evacuation rate in the coastal <34 knot zone.

More in-depth investigation reveals that this counter-intuitive pattern arises from how the geographical distribution of Florida's population intersects with the forecasts, the forecast uncertainty, and the storm's eventual track and impact zones. For example, in the realistic-population-distribution simulations, a large number of Cit-ags are located in the Miami area, in southeast Florida (Fig. 13). This region is in the forecast cone of uncertainty for the idealized Charley forecasts for much of the simulation period, and so evacuation orders are issued in this area and many coastal Cit-ags decide to evacuate (blue circles in southeast Florida in Fig. 13). Based on the storm's track and size,

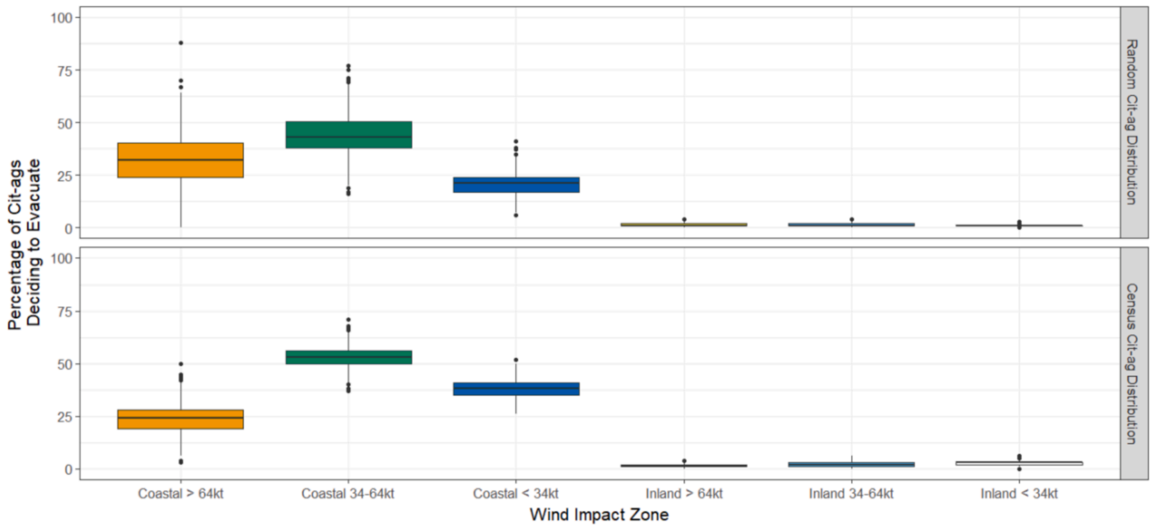

Fig. 12. Percentage of Cit-ags deciding to evacuate in each of the six impact zones, for the model configurations in the fourth Settings column in Table 6: ideal forecasts for Hurricane Charley and random (upper) or realistic (lower) geographical distribution of the Cit-ag population. The box plots show the mean percentage evacuating and the distribution across 1000 simulations for each model configuration. The mean percentage of Cit-ags evacuating across the domain is 3.9% for random and 13.7% for realistic population distribution.

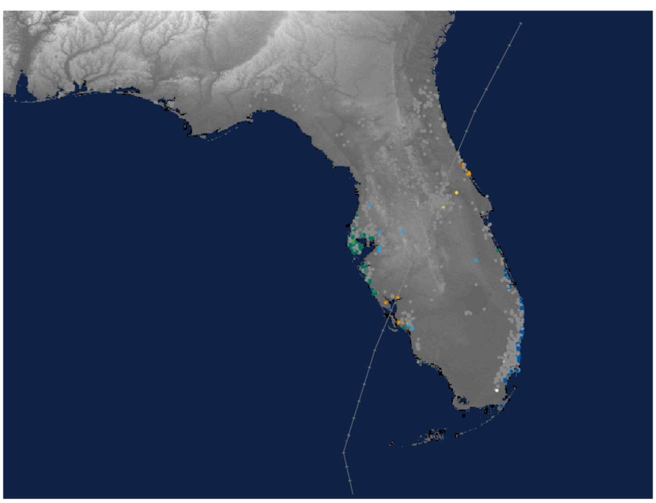

Fig. 13. As in Fig. 7, for a single simulation with realistic geographical distribution of Cit-ags and ideal forecasts for Hurricane Charley.

however, southeast Florida ends up not experiencing >34 knot winds. Thus, in the realistic-population-distribution simulations, a significant number of Cit-ags located in evacuation zones in the Miami area make decisions to evacuate that, in retrospect, were unnecessary. In Fig. 12, this appears as a larger evacuation rate in the coastal <34 knot impact zone.

As these results illustrate, the uneven distribution of population in the real world complicates analyzing spatial patterns of evacuation decisions using aggregated metrics. More generally, the intersection between specific hazard tracks or forecasts, complex coastal geography, and clustered populations can lead to decision patterns that are difficult to understand and attribute, even in this simplified model world. In the real world, which has many additional complexities, understanding patterns in protective decision making is even more challenging. This further indicates the potential value of this type of modeling laboratory, where different components can be simplified or modified systematically to run a suite of experiments.

4.5. Varying the storm and forecast skill: historical and ideal forecasts for Hurricanes Charley and Wilma

Finally, we explore the effects of using historical rather than idealized forecasts and of modifying the storm. All experiments shown in this section use a realistic rather than random geographical distribution of Cit-ags (Table 6). Using realistic forecast information adds a further complexity to the model simulations by adding a new dynamical component — evolving, imperfect forecast information — that is present in real weather forecast information and decision systems. These experiments also begin to investigate scenarios of interest to stakeholders such as meteorologists or emergency managers, by exploring how differences in storm characteristics and forecast information can propagate through the multi-agent system and translate into different patterns in public evacuations.

The top panel of Fig. 14 shows results for Hurricane Charley (the storm used in the experiments in section 4.1-4.4) with historical forecasts. Comparing these results with those for parallel simulations with ideal forecasts (lower panel of Fig. 12), we see that the mean evacuation rate is lower in the historical forecast simulations, especially in the coastal <34 knot impact zone. In other words, as one might expect given the less consistent risk information, imperfect forecasts lead to fewer Cit-ags deciding to evacuate.

The middle and lower panels in Fig. 14 show results for simulations with Hurricane Wilma for ideal and historical forecasts. Like Charley, Wilma made landfall in southwestern Florida, but farther south and with a more west-to-east track (Fig. 6). Wilma was also a much larger storm than Charley, and so its >64 knot and 34–64 knot winds cover a much larger portion of the model domain.

Comparing the Wilma and Charley results for ideal forecasts (middle panel of Fig. 14 and lower panel of Fig. 12), the overall pattern of evacuation rates is similar, except that the Wilma evacuation rates are much lower in the coastal <34 knot zone. In other words, a much lower percentage of Cit-ags in the coastal <34 knot zone decide to evacuate unnecessarily in the ideal-forecast simulations for Wilma than for Charley. This occurs because, unlike Charley, Wilma is a large enough storm that most of the areas that are within the cone of uncertainty several days before landfall – including the Miami area – end up experiencing >34 knot winds. Thus, few Cit-ags in the <34 knot impact zones for Wilma receive information indicating that they are at high risk and decide to evacuate. This illustrates how and why, given similar forecast

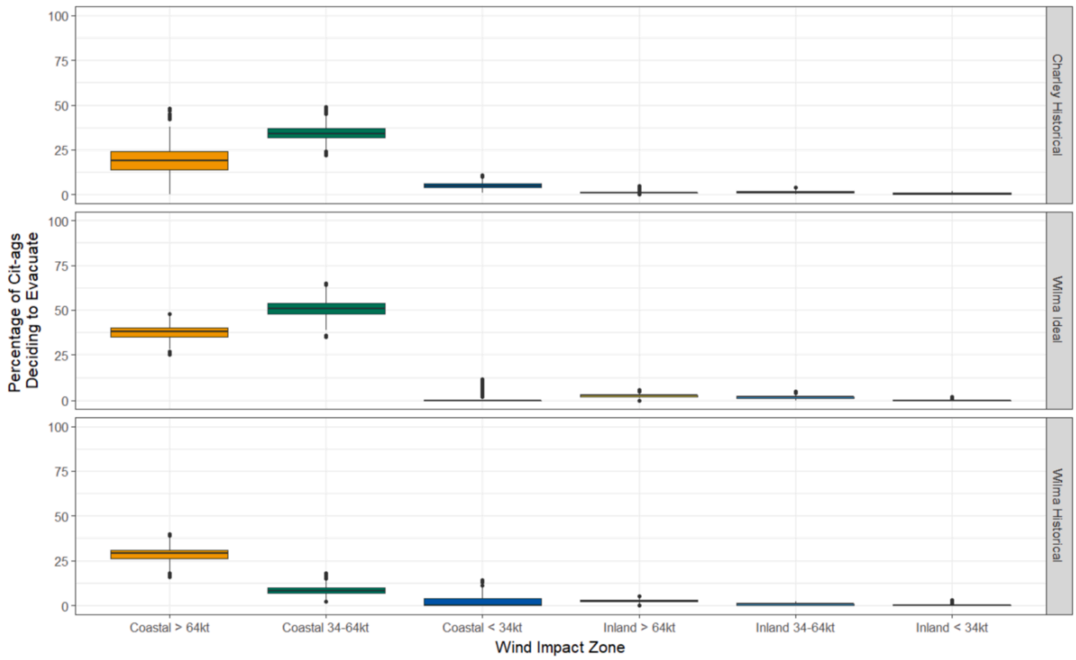

Fig. 14. As in Fig. 12, for the model configurations in the rightmost Settings column in Table 6: historical forecasts for Hurricane Charley (upper), ideal forecasts for Hurricane Wilma (middle), and historical forecasts for Hurricane Wilma (lower), all with a realistic geographical distribution of Cit-ags. The mean percentage of Cit-ags evacuating across the domain is 5.8% for Charley with historical forecasts, 12.8% for Wilma with ideal forecasts, and 6.1% for Wilma with historical forecasts.

track uncertainty and decision algorithms, more people are likely to make evacuation decisions that turn out to be unnecessary for a smaller hurricane: fewer people end up experiencing strong winds and other impacts.

Similar to the Charley results discussed above, running Wilma simulations with historical forecasts produces lower evacuation rates than simulations with ideal forecasts (middle and lower panels of Fig. 14). In fact, the Wilma simulations with historical forecasts produce an overall Cit-ag spatial evacuation pattern similar to that which hurricane forecast and evacuation professionals might wish to see, with higher evacuation rates in higher-impact zones. As the other results presented show, however, this pattern is not an inherent outcome of the model's dynamics. Instead, it is produced by how the evolving, imperfect, and uncertain forecasts and the storm's track and winds overlay onto the unevenly distributed population. This illustrates the challenges of understanding the dynamics that lead to different outcomes in realistic hurricane situations. By enabling systematically perturbed experiments in more simplified contexts, this type of modeling laboratory can help build new understanding about the interactions among evolving environmental hazards, hazard information, information flow, and protective decisions.

Given the goals of the work presented here, we did not try to adjust the model to match real-world evacuation rates, and robust empirical data on spatially distributed evacuation rates for these two storms is not publicly available. However, the available empirical data indicates that the model is, to first order, generating reasonable evacuation rates in coastal high-impact zones. For example, a survey conducted by Smith and McCarty (2009) after Charley made landfall found that 36% of the sample in Charlotte County, Florida (the coastal county where Charley made landfall with >64 knot winds) reported evacuating. Another post-storm study of Charley, by Baker (2005), found that 22–53% of the sample in areas similar to the model's coastal >34 knot zones and 12–33% of the sample in areas similar to the model's inland >34 knot zones reported evacuating. For Wilma, a post-storm survey conducted by Solis et al. (2010) in 3 southeastern Florida counties, all of which experienced >64 knot winds, found that 32% of the sample evacuated. Comparision with Fig. 14 indicates that these evacuation rates are similar to those generated by the model in our coastal 34–64 knot and >64 knot zones, for historical forecasts for the 2 storms. As noted in section 4.1, however, the evacuation rates in inland >34 knot zones produced by the model in its current configuration are lower than those in the real world.

5. Summary and discussion

This article conceptualizes and implements an agent-based model for studying the modern hazard information and decision system, in the context of hurricanes approaching the US coastline. The model includes multiple types of agents who interact with each other and with their physical and informational environments to access, interpret, and decide how to respond to evolving hazard information in a theoretically and empirically informed way. The resulting digital laboratory provides opportunities to study this complex dynamic system from a new perspective, complementing recent related work using other methods (e.g., Lee et al., 2009; Gudishala and Wilmot, 2010; Meyer et al., 2013, 2014; Ruin et al., 2014; Morss et al., 2015, 2017; Lazrus et al., 2016; Bostrom et al., 2018; Demuth et al., 2018).

We use the modeling laboratory to ask: How are the spatial and temporal patterns of protective decisions during hazardous weather threats affected when heterogeneous agents with semi-realistic decision rules access, share, and interpret evolving forecasts and other hazard information? Specifically, we perform experiments investigating how the model's behavior and outcomes change when key agent parameters and the geographical population distribution, hurricane evolution, and forecast skill are varied. The results provide insight into how and why evacuation patterns can arise when interacting agents exchange and respond to evolving, uncertain information from different environmental and social sources. They also illustrate how interactions among evolving information, uncertainty, and decisions can produce complex, emergent dynamics. For example, as agents respond to information indicating decreasing uncertainty about the potential threat, feedback loops can lead to rapidly increasing risk assessments as a storm approaches.

As the experiments further show, including factors that add complexity and realism to the model — such as the coastal geography of a region such as Florida, the impacts of an asymmetric evolving storm, non-uniform geographical population distributions, and evolving imperfect forecasts — can complicate interpreting the model output. This demonstrates the value of this type of modeling laboratory for building in-depth understanding about hazard information and decision dynamics, by enabling systematic manipulation of factors that cannot easily be controlled in the real world. It also underscores the potential of this type of interdisciplinary modeling for addressing questions of interest to forecasters, emergency managers, and other stakeholders as well as researchers, by allowing experiments in a wide range of scenarios.

The research described here advances scientific capabilities and knowledge in several ways. First, by adapting existing theoretical models and empirical understanding of hazard information flow and decision making for use in computational agent-based models, the work provides a new approach for exploring how evolving hazard information and decisions interact to create broader patterns of interest. Second, the modeling framework discussed here provides a unique toolkit for exploring the effects of different hazard forecast information (including timing and uncertainty), risk-related decision-making, and information network topologies on patterns in social decisions — experiments that are impossible to perform in the real world. Further, this research demonstrates how agent-based modeling can be used to study systems in which coupling with evolving environmental and social information contributes to the system dynamics along with coupling between the natural and human system. In these ways, this study aims to extend model-based hazards research toward work with theoretically and empirically informed agent-based modeling in complex, dynamic information contexts.

Based on the research reported here, we propose several areas for future related work. First, CHIME ABM can be used to address additional research questions related to hurricane forecasting, information communication, and evacuations. In particular, the forecast information currently represented in the model is much simpler than that typically available in the real world today. Thus, one possible extension for future experiments is incorporating additional or more complex forecast information and/or representations of forecast uncertainty. This might include coupling the agent-based components of the model with more complex hazard and forecast information inputs, e.g., from numerical modeling of hurricanes and their impacts, to develop a more complete coupled physical-social modeling laboratory. Another possible extension is to simulate different interpretations of forecast information content (e.g., by having agents modify their representations of the forecast) to explore how more complex aspects of information flow and interpretation interact to influence decisions. The model could also be extended to study sequences of hurricanes during one or multiple years, in order to explore the dynamics of how hurricane-related experiences influence attitudes and behaviors in subsequent storms or other longer-term aspects of hazard risks and resilience.

Another area for future research is revising the model structure and its components, including the agent algorithms, hazard information, social and information networks, and evaluation of impacts, to address idealizations in the current version. In doing so, it is important to consider the potential trade-offs of adding different forms of realism and complexity, in the context of the research goals. We designed CHIME ABM to be capable of using quasi-realistic geography, populations, storms, and forecasts, imported using real data. Nevertheless, as the results in section 4 show, abstractions and simplifications can enhance the interpretability of the model output and the value of the model in elucidating key dynamics of the system of interest.

The model can also be adapted to study other hazards for which the evolution of different types of information and its exchange among multiple types of actors play important roles. For example, floods, wildfires, and volcanic eruptions take place at temporal and spatial scales similar to those of hurricanes; the CHIME modeling environment could readily be adapted to investigate decision-making in these contexts. Tsunamis, tornados, and flash floods typically unfold more quickly and affect smaller regions; CHIME ABM would require more extensive modification to explore information flow and decision-making for such hazards.

An additional potential area for future work is comparing quasi-realistic model simulations with observed social data from specific historical hurricanes to refine the model structure and parameters. In its current form, CHIME ABM is well suited for exploring aspects of the system's behavior by comparing results across sets of simulations. Qualitative comparisons with prior related research and empirical evacuation data suggest that the model produces reasonable evacuation patterns, but it does not attempt to realistically represent the multitude of factors that influence real-world individual and household evacuation decisions in specific situations. More in-depth quantitative comparisons with real-world data may therefore help improve the model's capabilities to address practical questions of interest. However, given the limited availability of the types of comprehensive empirical data required to perform such comparisons, new data sets may need to be collected or compiled.

Real-world decisions are influenced by many complicated factors that must be simplified in any modeling approach. Since little is known about the emergent dynamics of the type of system being studied here, our aim was to begin developing fundamental understanding that could form building blocks to be expanded on in future related modeling work. Given these goals, CHIME ABM is purposefully abstracted from the real world in multiple ways, and so it has many limitations if evaluated from the perspective of simulating actual hurricane decisions and outcomes. The model was developed, however, using theory, research findings, expertise, and data from several relevant disciplines. The modeling effort is also intersecting with ongoing research on hurricane hazard predictability, information flow, and decision making being conducted as part of a larger multi-method research project (Morss et al., 2017). Thus, the model is both informed by and feeds back into empirical research. Interpreted in conjunction with other work investigating real-world hazard information flow and decision making, we propose that modeling research such as that conducted here has significant potential to develop new understanding, identify strengths and weaknesses in hazard forecasting and risk communication, and recommend areas for improvement.

Software/data availability

CHIME ABM was implemented in the freeware agent-based modeling platform NetLogo version 5.3.1 and later updated to NetLogo version 6.0. The model code, supporting documentation, and input files are archived on the ComSES.net library (OpenABM.com) at the URL: https://www.comses.net/codebases/5504/releases/1.4.0/.

Declarations of interest

None.

Acknowledgements

The authors acknowledge Sean Bergin, Wendy Cegielski, Nicholas Gauthier, and Grant Snitker for their assistance with the development and testing of the agent-based model. We also thank our collaborators on the larger project, especially Heather Lazrus, Olga Wilhelmi, Christopher Davis, Kathryn Fossell, David Ahijevych, Chris Snyder, Leysia Palen, and Kenneth Anderson, for their contributions to this research. The work reported here was supported by National Science Foundation award AGS 1331490. The National Center for Atmospheric Research is sponsored by the National Science Foundation.

References

Sorensen, J.H., 2000. Hazard warning systems: review of 20 years of progress. Nat. Hazards Rev. 1, 119–125.

Thompson, R.R., Garfin, D.R., Silver, R.C., 2017. Evacuation from natural disasters: a systematic review of the literature. Risk Anal. 37, 812–839. https://doi.org/10.1111/risa.12654.

Allison, A.E.F., Dickson, M.E., Fisher, K.T., Thrush, S.F., 2018. Dilemmas of modelling and decision-making in environmental research. Environ. Model. Softw 99, 147–155. https://doi.org/10.1016/j.envsoft.2017.09.015.

An, L., 2012. Modeling human decisions in coupled human and natural systems: review of agent-based models. Ecol. Model. 229, 25–36. https://doi.org/10.1016/j.ecolmodel.2011.07.010.

Anderson, J., Kogan, M., Bica, M., Palen, L., Anderson, K., Stowe, K., Morss, R., Demuth, J., Lazrus, H., Wilhelmi, O., Henderson, J., 2016. Far far away in Far Rockaway: responses to risks and impacts during Hurricane Sandy through first-person social media narratives. In: Proceedings of the 13th International Conference on Information Systems for Crisis Response and Management. Rio de Janeiro, Brazil.

Baker, E.J., 1991. Hurricane evacuation behavior. Int. J. Mass Emergencies Disasters 9, 287–310.

Baker, E.J., 2005. Hurricane Charley Behavioral Analysis. Prepared for the U.S. Army Corps of Engineers.

Barton, C.M., Ullah, I.I.T., Bergin, S.M., Sarjoughian, H.S., Mayer, G.R., Bernabeu-Auban, J.E., Heimsath, A.M., Acevedo, M.F., Riel-Salvatore, J.G., Arrowsmith, J.R., 2016. Experimental socioecology: integrative science for Anthropocene landscape dynamics. Anthropocene 13, 34–45. https://doi.org/10.1016/j.ancene.2015.12.004.

Bernardini, G., Postacchini, M., Quagliarini, E., Brocchini, M., Cianca, C., D'Orazio, M., 2017. A preliminary combined simulation tool for the risk assessment of pedestrians flood-induced evacuation. Environ. Model. Softw 96, 14–29. https://doi.org/10.1016/j.envsoft.2017.06.007.

Bica, M., Demuth, J.L., Dykes, J.E., Palen, L., 2019. Communicating hurricane risks: multi-method examination of risk imagery diffusion. In: Proceedings of the ACM 2018 Conference on Human Factors in Computing Systems. Association for Computing Machinery, Glasgow, Scotland, UK.

Boone, R.B., Galvin, K.A., BurnSilver, S.B., Thornton, P.K., Ojima, D.S., Jawson, J.R., 2011. Using coupled simulation models to link pastoral decision making and ecosystem services. Ecol. Soc. 16 (2), 6 [online].

Bostrom, A., Morss, R.E., Lazo, J.K., Demuth, J.L., Lazrus, H., Hudson, R., 2016. A mental models study of hurricane forecast and warning production, communication, and decision making. Weather Clim. Soc. 8, 111–129. https://doi.org/10.1175/WCAS-D-15-0033.1.

Bostrom, A., Morss, R.E., Lazo, J.K., Demuth, J.L., Lazrus, H., 2018. Eyeing the storm: how residents of coastal Florida see hurricane forecasts and warnings. Int. J. Disaster Risk Reduct. 30A, 105–119. https://doi.org/10.1016/j.ijdrr.2018.02.027.

Brotzge, J., Donner, W., 2013. The tornado warning process: a review of current research, challenges, and opportunities. Bull. Am. Meteorol. Soc. 94, 1715–1733. https://doi.org/10.1175/BAMS-D-12-00147.1.

Buchmann, C.M., Grossmann, K., Schwarz, N., 2016. How agent heterogeneity, model structure and input data determine the performance of an empirical ABM – a real-world case study on residential mobility. Environ. Model. Softw 75, 77–93. https://doi.org/10.1016/j.envsoft.2015.10.005.

Chen, X., Meaker, J.W., Zhan, F.B., 2006. Agent-based modeling and analysis of hurricane evacuation procedures for the Florida Keys. Nat. Hazards 38, 321–338. https://doi.org/10.1007/s11069-005-0263-0.

Cuite, C., Shwom, R., Hallman, W., Morss, R.E., Demuth, J.L., 2017. Improving coastal storm evacuation messages. Weather Clim. Soc. 9, 155–170. https://doi.org/10.1175/WCAS-D-16-0076.1.

Czajkowski, J., 2011. Is it time to go yet? Understanding household hurricane evacuation decisions from a dynamic perspective. Nat. Hazards Rev. 12, 72–84. https://doi.org/10.1061/(ASCE)NH.1527-6996.0000037.

Dash, N., Gladwin, H., 2007. Evacuation decision making and behavioral responses: individual and household. Nat. Hazards Rev. 8, 69–77. https://doi.org/10.1061/(ASCE)1527-6988(2007)8:3(69).

Dawson, R.J., Peppe, R., Wang, M., 2011. An agent based model for risk-based flood incident management. Nat. Hazards 59, 167–189.

Demuth, J.L., Morss, R.E., Hearn Morrow, B., Lazo, J.K., 2012. Creation and communication of hurricane risk information. Bull. Am. Meteorol. Soc. 73, 1133–1145. https://doi.org/10.1175/BAMS-D-11-00150.1.

Demuth, J.L., Morss, R.E., Lazo, J.K., Trumbo, C., 2016. The effects of past hurricane experiences on evacuation intentions through risk perception and efficacy beliefs: a mediation analysis. Weather Clim. Soc. 8, 327–344. https://doi.org/10.1175/WCAS-D-15-0074.1.

Demuth, J.L., Morss, R.E., Palen, L., Anderson, K., Anderson, J., Kogan, M., Stowe, K., Bica, M., Lazrus, H., Wilhelmi, O., Henderson, J., 2018. "sometimes da #beachlife ain't always da wave": understanding people's evolving hurricane risk communication, risk assessments, and responses using Twitter narratives. Weather Clim. Soc. 10, 537–560. https://doi.org/10.1175/WCAS-D-17-0126.1.

Dixon, D.S., Mozumder, P., Vásquez, W.F., Gladwin, H., 2017. Heterogeneity within and across households in hurricane evacuation response. Netw. Spat. Econ. 17, 645–680.

Dow, K., Cutter, S.L., 2000. Public orders and personal opinions: household strategies for hurricane risk assessment. Environ. Hazards 2, 143–155. https://doi.org/10.1016/S1464-2867(01)00014-6.

Drabek, T.E., 1999. Understanding disaster warning responses. Soc. Sci. J. 36, 515–523. https://doi.org/10.1016/S0362-3319(99)00021-X.

Filatova, T., Verburg, P.H., Parker, D.C., Stannard, C.A., 2013. Spatial agent-based models for socio-ecological systems: challenges and prospects. Environ. Model. Softw 45, 1–7. https://doi.org/10.1016/j.envsoft.2013.03.017.

Fossell, K.R., Ahijevych, D., Morss, R.E., Snyder, C., Davis, C., 2017. The practical predictability of storm tide from tropical cyclones in the Gulf of Mexico. Mon. Weather Rev. 145, 5103–5121. https://doi.org/10.1175/MWR-D-17-0051.1.

Gladwin, C.H., Gladwin, H., Peacock, W.G., 2001. Modeling hurricane evacuation decisions with ethnographic methods. Int. J. Mass Emergencies Disasters 19, 117–143.

Gladwin, H., Lazo, J., Morrow, B.H., Peacock, W.G., Willoughby, H.E., 2007. Social science research needs for the hurricane forecast and warning system. Nat. Hazards Rev. 8, 87–95.

Groeneveld, J., Müller, B., Buchmann, C.M., Dressler, G., Guo, C., Hase, N., Hoffmann, F., John, F., Klassert, C., Lauf, T., Liebelt, V., Nolzen, H., Pannicke, N., Schulze, J., Weise, H., Schwarz, N., 2017. Theoretical foundations of human decision-making in agent-based land use models - a review. Environ. Model. Softw 87, 39–48. https://doi.org/10.1016/j.envsoft.2016.10.008.

Gudishala, R., Wilmot, C., 2010. Development of a time-dependent, audio-visual, stated-choice method of data collection of hurricane evacuation behavior. J. Transp. Saf. Secur. 2, 171–183. https://doi.org/10.1080/19439962.2010.488315.

Haer, T., Botzen, W.J.W., Aerts, J.C.J.H., 2016. The effectiveness of flood risk communication strategies and the influence of social networks—insights from an agent-based model. Environ. Sci. Policy 60, 44–52. https://doi.org/10.1016/j.envsci.2016.03.006.

Huang, S.-K., Lindell, M.K., Prater, C.S., Wu, H.-C., Siebeneck, L.K., 2012. Household evacuation decision making in response to Hurricane Ike. Nat. Hazards Rev. 13, 283–296. https://doi.org/10.1061/(ASCE)NH.1527-6996.0000074.

Hasan, S., Ukkusuri, S., Gladwin, H., Murray-Tuite, P., 2011. Behavioral model to understand household-level hurricane evacuation decision making. J. Transp. Eng. 137, 341–348.

Huang, S.-K., Lindell, M.K., Prater, C.S., 2016. Who leaves and who stays? A review and statistical meta-analysis of hurricane evacuation studies. Environ. Behav. 48, 991–1029. https://doi.org/10.1177/0013916515578485.

Huang, S.-K., Lindell, M.K., Prater, C.S., 2016. Multistage model of hurricane evacuation decision: empirical study of Hurricanes Katrina and Rita. Nat. Hazards Rev. 18, 283–296. https://doi.org/10.1061/(ASCE)NH.1527-6996.0000237.

Kogan, M., Palen, L., 2018. Conversations in the eye of the storm: at-scale features of conversational structure in a high-tempo, high-stakes microblogging environment. In: Proceedings of the ACM 2018 Conference on Human Factors in Computing Systems. Canada, Association for Computing Machinery, Montreal.

Kogan, M., Palen, L., Anderson, K., 2015. Tweet local, retweet global: retweeting by the geographically-vulnerable during Hurricane Sandy. In: Proceedings of the 18th ACM Conference on Computer-Supported Cooperative Work and Social Computing. Association for Computing Machinery, Vancouver, Canada.

Lazo, J.K., Bostrom, A., Morss, R.E., Demuth, J.L., Lazrus, H., 2015. Factors affecting hurricane evacuation intentions. Risk Anal. 35, 1837–1857. https://doi.org/10.1111/risa.12407.

Lazrus, H., Morss, R.E., Demuth, J.L., Bostrom, A., Lazo, J.K., 2016. "Know what to do if you encounter a flash flood": mental models analysis for improving flash flood risk communication and public decision making. Risk Anal. 36, 411–427. https://doi.org/10.1111/risa.12480.

Lee, K.L., Meyer, R.J., Bradlow, E.T., 2009. Analyzing risk response dynamics on the web: the case of Hurricane Katrina. Risk Anal. 29, 1779–1792. https://doi.org/10.1111/j.1539-6924.2009.01304.x.

Lindell, M.K., Perry, R.W., 2004. Communicating Environmental Risk in Multiethnic Communities. Sage, Thousand Oaks, CA.

Lindell, M.K., Perry, R.W., 2012. The protective action decision model: theoretical modifications and additional evidence. Risk Anal. 32, 616–632. https://doi.org/10.1111/j.1539-6924.2011.01647.x.

Lindell, M.K., Lu, J.C., Prater, C.S., 2005. Household decision making and evacuation in response to Hurricane Lili. Nat. Hazards Rev. 6, 171–179.

Magliocca, N.R., Ellis, E.C., 2016. Evolving human landscapes: a virtual laboratory approach. J. Land Use Sci. 11, 642–671. https://doi.org/10.1080/1747423X.2016.1241314.

Meyer, R., Broad, K., Orlove, B., Petrovic, N., 2013. Dynamic simulation as an approach to understanding hurricane risk response: insights from Stormview lab. Risk Anal. 33, 1532–1552. https://doi.org/10.1111/j.1539-6924.2012.01935.x.

Meyer, R.J., Baker, J., Broad, K., Czajkowski, J., Orlove, B., 2014. The dynamics of hurricane risk perception: real-time evidence from the 2012 Atlantic hurricane season. Bull. Am. Meteorol. Soc. 95, 1389–1404. https://doi.org/10.1175/BAMS-D-12-00218.1.

Mileti, D.S., Sorensen, J.H., 1990. Communication of emergency public warnings: a social science perspective and state-of-the-art assessment. Oak Ridge National Laboratory Manuscript #ORNL-6609, [Available online at: http://emc.ed.ornl.gov/publications/PDF/CommunicationFinal.pdf.

Morrow, B., Gladwin, H., 2005. Hurricane Ivan Behavioral Analysis. Prepared for the Federal Emergency Management Agency and U.S. Army Corps of Engineers.

Morss, R.E., Hayden, M.H., 2010. Storm surge and "certain death": interviews with Texas coastal residents following Hurricane Ike. Weather Clim. Soc. 2, 174–189. https://doi.org/10.1175/2010WCAS1041.1.

Morss, R.E., Demuth, J.L., Bostrom, A., Lazo, J.K., Lazrus, H., 2015. Flash flood risks and warning decisions in Boulder, Colorado: a mental models study of forecasters, public officials, and media broadcasters. Risk Anal. 35, 2009–2028. https://doi.org/10.1111/risa.12403.

Morss, R.E., Demuth, J.L., Lazo, J.K., Dickinson, K., Lazrus, H., Morrow, B.H., 2016. Understanding public hurricane evacuation decisions and responses to forecast and warning messages. Weather Forecast. 31, 395–417. https://doi.org/10.1175/WAF-D-15-0066.1.

Morss, R.E., Mulder, K.J., Lazo, J.K., Demuth, J.L., 2016. How do people perceive, understand, and anticipate responding to flash flood risks and warnings? Results from a public survey in Boulder, Colorado, USA. J. Hydrol. 541, 649–664. https://doi.org/10.1016/j.jhydrol.2015.11.047.

Morss, R.E., Demuth, J.L., Lazrus, H., Palen, L., Barton, C.M., Davis, C.A., Snyder, C., Wilhelmi, O.V., Anderson, K.M., Ahijevych, D.A., Anderson, J., Bica, M., Fossell, K. R., Henderson, J., Kogan, M., Stowe, K., Watts, J., 2017. Hazardous weather prediction and communication in the modern information environment. Bull. Am. Meteorol. Soc. 98, 2653–2674. https://doi.org/10.1175/BAMS-D-16-0058.1.

NASEM (National Academy of Sciences, Engineering, and Medicine), 2017. Integrating Social and Behavioral Sciences within the Weather Enterprise. National Academies Press, Washington, DC.

NASEM, 2017. Emergency Alert and Warning Systems: Current Knowledge and Future Research Directions. National Academies Press, Washington, DC.

NWS (National Weather Service), 2019. NWS Weather-Ready Nation Strategic Plan (2019 - 2022). https://www.weather.gov/media/wrn/NWS_Weather-Ready-Nation_Strategic_Plan_2019-2022.pdf.

Parker, D., Fordham, M., 1996. An evaluation of flood forecasting, warning and response systems in the European Union. Water Resour. Manag. 10, 279–302.

Parker, D.C., Manson, S.M., Janssen, M.A., Hoffmann, M.J., Deadman, P.J., 2003. Multi-agent systems for the simulation of land-use and land-cover change: a review. Ann. Assoc. Am. Geogr. 93, 314–337. https://doi.org/10.1111/1467-8306.9302004.

Petrolia, D.R., Bhattacharjee, S., 2010. Why don't coastal residents choose to evacuate for hurricanes? Coast. Manag. 38, 97–112. https://doi.org/10.1080/08920751003605365.

Rand, W., Herrmann, J., Schein, B., Vodopivec, N., 2015. An agent-based model of urgent diffusion in social media. J. Artif. Soc. Soc. Simul. 18 (2), 1.

Reilly, A.C., Tonn, G.L., Zhai, C., Guikema, S.D., 2017. Hurricanes and power system reliability–The effects of individual decisions and system-level hardening. Proc. IEEE 105, 1429–1442. In: https://doi.org/10.1109/JPROC.2017.2689720.

Rounsevell, M.D.A., Robinson, D.T., Murray-Rust, D., 2012. From actors to agents in socio-ecological systems models. Philos. Trans. R. Soc. Biol. Sci. 367, 259–269. https://doi.org/10.1098/rstb.2011.0187.

Rovere, S.L., North, M.J., Podesta, G.P., Bert, F.E., 2016. Practical points for the software development of an agent-based model of a coupled human-natural system. IEEE Access 4, 4282–4298. https://doi.org/10.1109/ACCESS.2016.2592418.

Ruin, I., Lutoff, C., Boudevillain, B., Creutin, J.D., Anquetin, S., Rojo, M.B., Boissier, L., Bonnifait, L., Borga, M., Justin, L.C., Créton-Cazanave, L., Delrieu, G., Douvinet, J., Gaume, E., Gruntfest, E., Naulin, J.P., Payrastre, O., Vannier, O., 2014. Social and hydrological responses to extreme precipitations: an interdisciplinary strategy for post-flood investigation. Weather Clim. Soc. 6, 135–153. https://doi.org/10.1175/WCAS-D-13-00009.1.

Schulze, J., Müller, B., Groeneveld, J., Grimm, V., 2017. Agent-based modelling of social-ecological systems: achievements, challenges, and a way forward. J. Artif. Soc. Simul. 20 (2), 8. https://doi.org/10.18564/jasss.3423.

Sherman-Morris, K., Senkbeil, J., Carver, R., 2011. Who's Googling what? What internet searches reveal about hurricane information seeking. Bull. Am. Meteorol. Soc. 92, 975–985. https://doi.org/10.1175/2011BAMS3053.1.

Smith, S.K., McCarty, C., 2009. Fleeing the storm(s): An examination of evacuation behavior during Florida's 2004 hurricane season. Demography 46, 127–145.

Solis, D., Thomas, M., Letson, D., 2010. An empirical evaluation of the determinants of household hurricane evacuation choice. J. Dev. Agric. Econ. 2, 188–196.

Sun, Z., Lorscheid, I., Millington, J.D., Lauf, S., Magliocca, N.R., Groeneveld, J., Balbi, S., Nolzen, H., Müller, B., Schulze, J., Buchmann, C.M., 2016. Simple or complicated agent-based models? A complicated issue. Environ. Model. Softw 86, 56–67. https://doi.org/10.1016/j.envsoft.2016.09.006.

Taylor, K., Priest, S., Fussell Sisco, H., Banning, S., Campbell, K., 2009. Reading Hurricane Katrina: information sources and decision-making in response to a natural disaster. Soc. Epistemol. 23, 361–380. https://doi.org/10.1080/02691720903374034.

Tonn, G.L., Guikema, S.D., 2018. An agent-based model of evolving community flood risk. Risk Anal. 38, 1258–1278. https://doi.org/10.1111/risa.12939.

Ukkusuri, S.V., Hasan, S., Luong, B., Kien, D., Zhan, X., Murray-Tuite, P., Yin, W., 2017. A-rescue: an agent based regional evacuation simulator coupled with user enriched behavior. Netw. Spat. Econ. 17, 197–223.

Verburg, P.H., Dearing, J.A., Dyke, J.G., van der Leeuw, S., Seitzinger, S., Steffen, W., Syvitski, J., 2016. Methods and approaches to modelling the Anthropocene. Glob. Environ. Chang. 39, 328–340. https://doi.org/10.1016/j.gloenvcha.2015.08.007.

Waldrop, M.M., 2017. News Feature: special agents offer modeling upgrade. Proc. Natl. Acad. Sci. 114, 7176–7179. https://doi.org/10.1073/pnas.1710350114.

Wang, H., Mostafizia, A., Cramer, L.A., Cox, D., Park, H., 2016. An agent-based model of a multimodal near-field tsunami evacuation: decision-making and life safety. Transp. Res. C Emerg. Technol. 64, 86–100. https://doi.org/10.1016/j.trc.2015.11.010.

Widener, M.J., Horner, M.W., Metcalf, S.S., 2013. Simulating the effects of social networks on a population's hurricane evacuation participation. J. Geogr. Syst. 15, 193–209.

Wilenski, U., 1999. NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. http://ccl.northwestern.edu/netlogo/.

Wu, H.-C., Lindell, M., Prater, C., 2012. Logistics of hurricane evacuation in hurricanes Katrina and Rita. Transp. Res. F Traffic Psychol. Behav. 15, 445–461. https://doi.org/10.1016/j.trf.2012.03.005.

Yin, W., Murray-Tuite, P., Ukkusuri, S.V., Gladwin, H., 2014. An agent-based modeling system for travel demand simulation for hurricane evacuation. Transp. Res. C Emerg. Technol. 42, 44–59. https://doi.org/10.1016/j.trc.2014.02.015.

Zhang, F., Morss, R.E., Sippel, J.A., Beckman, T.K., Clements, N.C., Hampshire, N.L., Harvey, J.N., Hernandez, J.M., Morgan, Z.C., Mosier, R.M., Wang, S., Winkley, S.D., 2007. An in-person survey investigating public perceptions of and response to Hurricane Rita forecasts along the Texas Coast. Weather Forecast. 22, 1177–1190.

Zhang, B., Chan, W., Ukkusuri, S., 2009. Agent-based modeling for household level hurricane evacuation. In: Proceedings of the IEEE 2009 Winter Simulation Conference. Institute of Electrical and Electronics Engineers, Austin, TX, pp. 2778–2784.