# Exploring the Effect of the Diffusion of Geo-Targeted Emergency Alerts

## The Application of Agent-Based Modeling to Understanding the Spread of Messages from the Wireless Emergency Alerts System

July 2015

Homeland Security

# Exploring the Effect of the Diffusion of Geo-Targeted Emergency Alerts

## The Application of Agent-Based Modeling to Understanding the Spread of Messages from the Wireless Emergency Alerts System

Homeland Security

Science and Technology

*July 2015*

*Task Leads*

Andrew M. Parker

Brian A. Jackson

*Task Staff*

Angel Martinez

Ricardo Sanchez

Shoshana R. Shelton

Jan Osburg

Henry H. Willis

*Director, Homeland Security and Defense Center*

*National Defense Research Institute*

Jack Riley

*Vice President and Director, National Security Research Division*

*Director, National Defense Research Institute*

# Acknowledgments

The National Defense Research Institute (NDRI), a division of the RAND Corporation, performed this analysis for the Department of Homeland Security Science and Technology Directorate. The NDRI team offers its gratitude to the emergency responders whose dedication and commitment ensure the safety of our families, our communities and our nation. This report is a tribute to their service.

In addition, we would like to extend our appreciation for constructive peer reviews provided by Brooke Liu of University of Maryland and by Raffaele Vardavas of RAND.

# Table of Contents

# Executive Summary

Emergency alerts and warnings are an element of emergency plans designed to provide information that members of the public can use to protect themselves. Whatever the event, the general goal of emergency alerting is similar: successfully transmit information to the potentially affected population, have that message spread or diffuse to the people who need it, and do so far enough before the incident occurs that they can act accordingly (Sorensen, 2000,).

Changes in technology and the way people consume media raised concerns that legacy alert systems were no longer sufficient to achieve the goals of alerting (U.S. Government Accountability Office, 2007). Mobile devices are the entry point for much of current media consumption and they have moved people away from traditional channels that could be used to transmit emergency alerts. Because these mobile devices are always at least somewhat location aware, they can provide a way to transmit alert information relevant to the position where a person is at the time the emergency is occurring, rather than to a home address or other less precise location. To take advantage of this potential, the Wireless Emergency Alerts (WEA) system was implemented in 2012 to deliver short (up to 90 text characters) emergency messages to individual mobile devices in a designated warning area (Federal Communications Commission, 2014). WEA is designed to be a new addition to the toolbox for emergency alerting in the United States — not to replace other channels for transmitting information to the public (Federal Emergency Management Agency, 2015).

Technology could, in principle, make it possible to geo-target alert messages at a very specific level — e.g., delivering them to mobile subscribers within a single cell tower's footprint, or even narrower by taking advantage of technologies resident on the mobile device itself (Committee on Geotargeted Disaster Alerts and Warnings, 2013). Other factors can pull in a different direction, however, promoting the diffusion of messages beyond the area where they are delivered initially. In a connected age, individuals can be broadcasters of information, sending such messages to friends (e.g., through a phone call or email) or, more broadly, via the Internet and social media (Vieweg et al., 2010; Gao et al., 2014; Sutton et al., 2014). Technology has also made it possible for individual cities or entities within a community (e.g., schools, colleges or universities) to have their own alerting mechanisms. While such mechanisms provide alternative ways for individuals to receive the information they need on channels they pay attention to, they also provide ways for a geo-targeted alert that was designed for one area to "leak out" and diffuse more generally.

This project took on the fundamental question of how important diffusion behavior was for understanding the value of geo-targeting WEA messages. The study used agent-based modeling (ABM) to examine the diffusion of alerts within an area populated by individual recipients who move from place to place in the course of their daily activities. The use of ABM made it possible to simulate increases and decreases in the magnitude of different behaviors and processes to see the effects of diffusion on geo-targeted alerts. Such modeling also makes it possible to explore dynamic behaviors that are sufficiently complex that they are difficult to solve mathematically — while still maintaining simplicity in the design of the model itself so that the behaviors can be understood.

The research, and therefore this report, sought to answer three main questions:

- How does the interaction of geo-targeting and message-diffusion behavior affect the overall effectiveness of alerting?
- How do the effects of both geo-targeting and message diffusion vary for emergencies (and therefore alerting requirements) of different geographic sizes?
- How does the interaction of geo-targeting and message diffusion affect the ability to use sophisticated alerting strategies, such as staging in time or delivering different emergency instructions to populations in different areas of risk?

# Building a Model of Emergency Alerting and Diffusion

A deep literature about how people react to warnings provided the foundation for our efforts to build a model of individual human behavior relating to emergency alerting and diffusion. An ABM is a form of computer simulation where a population is represented by individual "agents" (in our model, representing people in at-risk areas) who interact and make individual choices about their behavior based on a set of variables or rules. In our model, the agents could move around geographically, since understanding the effect of individuals moving in and out of areas where geo-targeted alerts were sent was the central aim of our work.

For understanding alerting and diffusion behavior, we needed to understand and represent in the model a chain of events:

(1) receipt and understanding of an alert message, whether from WEA or another alerting channel;
(2) the decision by the agent to share or not share that alert with others, the central mechanism for alert diffusion in the model; and
(3) the decision by the agent whether or not to act on the alert and take whatever protective action it recommended.

For modeling, the requirements were to both represent each of these steps or behaviors at an appropriate degree of simplification — since models, by definition, must significantly simplify reality — and make estimates of the various parameters and variables needed to define the behavior of the agents at each step.

Agents lived and moved in a realistic geographic landscape, including a network of roads, modeled loosely after Dubuque, Iowa. Our model contained 10,000 agents; they moved about this region in the course of routine activities. Because one of the central behaviors we were interested in regarding alert diffusion was alert transmission, a key characteristic of our agents was how and how much they were connected to one another through social networks. In terms of the agents' social networks, we considered three different categories of connections where communications (i.e., forwarded alerts) might have different levels of influence on the recipients' behavior: a family social network (estimated from literature as four to five connections, on average), a close social network (six to 15 connections), and a broad social network (150 to 300 connections; see Hamill and Gilbert, 2010; Roberts et al., 2009; McCarty et al., 2001). In the course of our analysis, we also explored implications for expanding or contracting the portion of these social networks that forwarded messages might influence.

Alerts were transmitted to individual citizens through multiple alerting channels, themselves stylized representations of general alerting effects. The WEA channel delivered an alert to everyone with a compatible device in one geo-targeted area. The media delivered an alert indiscriminately to the entire region. Technical forwarding included tools that delivered alerts to one targeted group or group within an area. We included two technical forwarders: one was a subscription-based service that delivered alerts to subscribers scattered across the entire city (e.g., an emergency-management Twitter feed). The other was geographically concentrated (e.g., parents close to a single school). The final channel was individual forwarding. A citizen within the model, upon receiving an alert, had the potential to forward that alert to those in the citizen's personal social network. This forwarding could include phone, text, email, social media (e.g., Facebook, Twitter) or other mechanisms.

Within the model, citizen agents moved about the road network within a simulated area. Once an agent received an alert, he or she decided (1) whether or not to comply with the message directive, and (2) whether or not to forward the message to those in his or her personal social network. These events and agent decisions were affected by several agent characteristics, represented as parameters in the model. Phone compatibility and agents choosing to opt out of the system limited WEA's reach, specifically. Specific situations might vary how likely it is that an individual received an alert or the extent to which that receipt was delayed. The probability of understanding an alert could be fundamentally different for non-English speakers or those with sensory disabilities. People might differ in their motivation to comply with alert guidance, which was reflected in the likelihood of compliance and any delay in compliance. Similarly, people might differ in their likelihood to forward an alert that they received. Finally, these parameters were varied across individuals, as well as by the alerting channel, the nature of the alert event and other contextual factors. Broadly speaking, these parameters correspond to the sequence of steps that Mileti and colleagues proposed for how individuals respond to an alert or warning (Mileti and Sorensen, 1990; Mileti and Peek, 2000).

# Defining the Emergencies and WEA Performance in the Study

Beyond the structure of the model itself, which defined the geography and agent behavior, two other elements were required for the study: the emergency scenarios in which warning and agent behavior were tested and the measures against which outcomes (and changes in those outcomes as a result of alert diffusion) were assessed.

We based our selection of emergency scenarios on data describing the WEA messages that had been issued to date, provided by the Integrated Public Alert and Warning System (IPAWS) at the Department of Homeland Security. For our simulations, we used four scenarios, which was a large enough set to cover a number of emergency characteristics while remaining practical given the scope of the effort. The scenarios were:

1. *Flash flood* — a frequent scenario that has been the focus of large numbers of messages from WEA, covering a small area and with a short timeline, where individuals at risk would be directed to evacuate the affected area;
2. *Tornado* — also a frequent focus of past messages from WEA, but covering a larger area and with a potentially very short alert-to-emergency timeline, where individuals at risk would be directed to shelter in place until the emergency ended;
3. *Hazmat plume* — a small-sized event, but one that could involve multiple hazard areas where different protective actions (evacuation versus sheltering in place) would be most appropriate; and
4. *Major flood* — a rarer but larger event, potentially affecting a very large population and therefore requiring a significant multistage evacuation over a longer emergency timeline.

These four scenarios provided a set of emergencies with distinct characteristics that would make it possible to test different elements of alerting strategy and performance, as well as cases where the ability to geo-target might facilitate more-complex responses (e.g., in a phased evacuation or to deliver different protective-action directions to agents in distinct hazard areas). The scenarios also were explicitly selected to focus on cases that could be modeled below the county level, which was the previous level of targeting resolution for WEA and many legacy emergency alerting systems and where precise geo-targeting was most relevant and interesting.

Generally considering WEA performance, and more specifically the role of geo-targeting, we took into consideration four basic concepts. First, the proportion of agents receiving the WEA inside versus outside the emergency zone (*specificity*) reflects the precise ability to deliver an alert to very specific groups. Second, and in contrast to specificity, *coverage* is reflected in the number of agents receiving alerts, inside and outside the emergency zone. This reflects the fact that, within IPAWS, WEA does not exist in isolation, so receipt of an alert outside the zone is not simply leaking into a population that has no other means of being alerted. Third, and what could be considered the bottom line, is *alert response*, or the proportion of agents within the emergency zone following directions (in our simulations, either to evacuate or shelter in place). Finally, *unnecessary alert response* is the proportion of agents outside the emergency zone following directions to evacuate or shelter.

In our analyses, we focused primarily on the final two concepts, operationalizing them in two ways. (The first two concepts acted as validations of our model and are presented in Appendix C.) The first way was used primarily for evacuation-based scenarios, where we traced the number of agents in the emergency zone over the course of the simulation. The second way that we measured response was through assessment of agent goals. Goals to evacuate or shelter, once adopted, remained with the agent for the rest of the simulation. Hence, we assessed how many agents had adopted goals consistent with the scenario (either evacuation or shelter) between the time when alerts were first transmitted and the projected time of the emergency.

# How Do Geo-Targeting and Alert Diffusion Affect Alerting Effectiveness?

In principle, WEA should reach those who need to be alerted, but not those who do not need to be alerted. Geo-targeting has the potential to serve this goal by reducing over-alerting, and diffusion of a geo-targeted alert has the potential to water down such gains. Many prominent arguments for geo-targeting emphasize the benefits of limiting over-alerting, including reduction in both the annoyance caused by unnecessary alerts and subsequent opt-out behavior. There is also concern that over-alerting, especially with short WEA messages that convey limited information, could cause unnecessary compliance, such as evacuation or sheltering in place, by those outside an emergency zone (Department of Homeland Security, 2013a).

To address this, we first used our model to examine how geo-targeting and alert diffusion affect alert effectiveness. We started with a set of simulation results using realistic, first-approximation parameters, drawing from the literature where possible. Our initial focus was on the flash flood scenario, which was both very small geographically and very quick-moving. We found limited effectiveness of the alerts, especially within a small, quick event. Increasing the size of the alerting zone to extend beyond the emergency zone and the forwarding of alerts by people (two key diffusion mechanisms) had very little effect on this result. The location and context (e.g., timing) of an event, it turned out, were critical to understanding who are the ideal targets of an alert.

We then provided a contrasting set of flash flood cases, where we reduced a key set of parameters — delay of alert receipt and delay of alert compliance. Effectively, this increased the churn of forwarding by allowing for more iterations. With greatly reduced delays, we saw marked increases in message effectiveness within the emergency zone.

To better understand the generality of these results, we then contrasted the flash flood case with a larger and longer event — a major flood. We found that, with the increase in the number of agents in the emergency zone (and the number alerted via WEA) and the more protracted event, WEA forwarding had a distinct impact on evacuation behavior, even with the longer delay parameters.

An examination of the parameter structure of the model highlights how forwarding can act as an alerting magnifier, amplifying the number of alerts received and hence increasing compliance over what would otherwise be much lower baseline levels. This explanation, based on the logic of the model, was corroborated by manipulating forwarding volume through another parameter set — the amount of an individual's social network that is actually influenced by a forwarded message.

# How Does the Value of Geo-Targeting Vary for Emergencies of Different Sizes and Speeds?

When events are small enough and fast enough, forwarding becomes limited by (1) the number of initial WEA recipients than can seed the forwarding process, and (2) the time available for forwarding to happen. Large, slow events, such as our major flood scenario, had many potential forwarders and time for forwarding to occur. To understand this better, we considered events and alerting zones of different sizes and speeds as a means of exploring how geo-targeting and alert diffusion interacted to influence the potential effectiveness of WEA.

With the very large, slow events, such as our major flood, we saw substantial evacuation that increased with phone compatibility and forwarding behavior. We then used a tornado scenario that had a similarly large alert zone but a much smaller emergency zone and a much quicker timeline. Here, compliance rates were smaller and did not vary by alert zone size. There was also small but noticeable out-of-zone sheltering (which may be of less concern than out-of-zone evacuation, which could cause road congestion). Returning to the flash flood results illustrated how changing the alert-zone size could dramatically effect emergency-zone evacuation. But that same scenario also demonstrated that the potential cost to this dramatic evacuation was significant over-alerting and, as implemented in the model, unnecessary out-of-zone compliance behavior.

# How Does the Interaction of Geo-Targeting and Message Diffusion Affect the Utility of Sophisticated Alerting Strategies?

One of WEA's promising characteristics is the potential to use its geo-targeting capabilities to send different alerts to people in different locations. For example, Vogt and Sorensen (1999) described an emergency response to a chemical repackaging plant in Helena, Arkansas. A key aspect of the alerting strategy was to differentiate alerts targeting those in greatest danger from those in more moderate danger. Specifically, authorities in Helena issued alerts asking those within two miles of the plant to evacuate, and residents between two and three miles to shelter in place. Despite the differential instructions, many of those in the shelter zone appeared to evacuate. WEA's geo-targeting capabilities suggest added potential for administering such differential alerting. Furthermore, WEA's relatively quick dissemination suggests that alerts could vary not only geographically but also temporally. Such capability could be helpful in implementing, for example, a staged evacuation, helping to reduce traffic problems on the roads, as has been observed in some past wide-area evacuations.

Two of our scenarios addressed these possibilities. Our hazmat scenario was modeled after the Helena, Arkansas, case. Our major flood scenario involved time-phased evacuation. Considering both sets of results, it was evident that the diffusion of alerts can dramatically limit the utility of such differentiated alerting strategies. In both scenarios, agents reacted to alerts targeting other zones. In particular, those in the hazmat inner-evacuation zone often sheltered in place (perhaps placing them at increased risk) and those in the major flood inner zone often evacuated during the earlier period when they were asked to stay in place (likely increasing the very road congestion the strategy was designed to reduce). Much of this appeared to be due to the dynamics of agent movement and the forwarding of alerts to members of agents' social networks.

# Summary

The complexity of individual behaviors in emergency alerting situations, and the way those behaviors affect the ability of alerts to serve as a protective measure during major emergencies, has been a topic of interest for decades. With the evolution of new relationships between citizens and technology — which have affected the utility of legacy alerting modes and the arrival of new options like WEA — new facets have been added to that complexity. The ability to geo-target alerts to mobile devices provides new capability, but it also raises questions about how to use geo-targeting effectively and how the interaction between targeted messages and human communication behavior will affect alert effectiveness. Our model provided a way to explore this new complexity. By distilling the different elements of alerting to a set of simplified parameters, the model allowed us to look at how changing some of those parameters could affect alerting outcomes — dialing up and down the parameters that shaped how people communicated and forwarded messages to get insights into how that behavior could become a multiplier for alerting even as it fought efforts to precisely geo-target alerts.

Extracting the relevant policy conclusions from the suite of results, the most basic bottom line is that the forwarding of messages, and the relationship between forwarding and geo-targeting, was in some ways much less important and in some other ways much more important than we thought.

Our initial framing of the study considered forwarding as a potential threat to the value of geo-targeting. This perspective considers forwarding as producing a practical limit on geo-targeting precision. The simulations certainly showed this — as would be expected — but the relationship was more nuanced than we initially thought. Very tight geo-targeting will always be compromised some by forwarding, but if very few agents are in the alert zone (i.e., are in or near the emergency zone) that compromise will be small, and hence should not hurt the value of investing in geo-targeting. Diffusion is likely of greatest concern for intermediate-sized events, where the population to be alerted is reasonably large, yet the geographic area is small enough to suggest value in precise geo-targeting. As the number of people being alerted in our simulations went up, the amount of potential forwarders also went up, and the ability to target precisely went down — though, depending on the scenario (e.g., how quickly it evolved), that could matter to differing extents.

That said, instead of considering forwarding as a threat to geo-targeting, it might be more valuable to think of forwarding as a compliance enhancer. Significant forwarding could increase the total effect of alerting a great deal. In that sense, forwarding converts the set of individuals' social networks into a new mass alerting channel of its own. In this way, forwarding or communication among individuals about the alert, here via such electronic means as social media, is simply one more type of the communication that has always occurred during milling (Committee on Public Response to Alerts and Warnings Using Social Media, 2013) before citizens make the decision to comply with the alert. The price of this effect is significant increases in out-of-emergency-zone alerting and potentially unnecessary action in response.

The model also taught us that *precise geo-targeting* can have different meanings when we consider how agents are dynamic, rather than thinking about geo-targeted alerts being aimed at a set of agents that is

sitting still “waiting for the alert to arrive.” We may want to get the message not only to people who are *in the emergency zone* but also to those *contemplating entering* the emergency zone.

Our results demonstrate quite clearly that forwarding threatens the value of trying to deliver different messages to different geographic areas in an effort to either provide messages relevant to individual risk areas or to guide population behavior in ways designed to enable more-effective response (i.e., time-phased evacuation). The dynamics of forwarding and compliance mean that success in these sorts of differential alerting efforts would require more-nuanced communication strategies (e.g., telling later staged-evacuation zones to “shelter for now”).

# Limitations and Potential Future Directions

Perhaps the most important caution to be placed on these results derives from the very nature of such a modeling exercise, which by design is an abstraction from reality. Agent-based modeling is especially strong in incorporating complex and dynamic elements, such as geographic movement and communication across social networks, which are difficult to measure empirically (on any large scale) and largely infeasible using more-direct mathematical approaches. All models are abstractions, however, and aspects of reality may not have been completely captured (e.g., the current model does not capture non-verbatim forwarding of alerts).

As an abstraction from reality, it is important to remain clear that the model itself relies on the estimates of the parameters that serve as its basis. Parameter estimates from the literature were stronger for some sets of parameters than others. We drew, wherever possible, on existing literature, supplemented by logical analysis and experimental manipulation of parameters.

The results also depended on the scenario specifications. Again, literature, exiting data and logical analysis were used to select four diverse scenarios. As suggested with the flash flood scenario, however, differential placement of the emergency zone would likely have resulted in highlighting different dynamics.

# Conclusion

The geo-targeting of alert messages to mobile devices represents a significant new capability for emergency managers, compared with simple alerting that blankets a large geographic area with a warning message. Complexities of human behavior — including both message forwarding and movement — mean that the use of geo-targeting is not as simple as just restricting the transmission of an alert to the smallest area at risk from an emergency event.

Although the potential for message spread is not the threat to the value of geo-targeting that might immediately be assumed — particularly for small events where the number of people at risk, and therefore the number of potential forwarders, is small — its use requires due consideration to ensure that emergency messages are actually targeted to the populations that need to receive them. This can produce the need for decisions about how large outside the actual footprint of the emergency the message should be transmitted, for cases where the movement of the population and the potential speed of the event put a premium on alerting individuals so they never enter the hazard zone versus warning individuals who are already there to evacuate. This could pose a difficult choice for alerting originators. An event occurring in an area where the population is largely stationary might justify a smaller alerting footprint than an identical emergency occurring in an area where people are rapidly moving toward or through the hazard area.

If the transmission of alerts through the population is viewed less as leakage from a geo-targeted area and simply as a key part of milling behavior in today’s electronic age, however, then the forwarding and sharing of alert information through social networks takes on an entirely different meaning. From this perspective, the goal is not to reduce forwarding — through message design or public education — but to reduce the cost of that forwarding in unnecessary compliance. Ongoing discussions about including more information beyond the current 90-character WEA message (Communications Security, Reliability and Interoperability Council, 2014) could be a route to do so, where more characters and the inclusion of hyperlinks or other data in messages could minimize the negative effects of forwarding while maintaining its potential value as a compliance enhancer.

In sum, the ability to transmit messages to smaller areas — which is indeed a major technical jump in emergency alerting capability — requires similar innovation in policy and practice to ensure that emergency managers are outfitted with the best understanding and tools to make the best choices regarding the use of geo-targeted alerts during the critical and time-limited decisions made in the warning phase of natural, technological or other emergency incidents.

# 1. Introduction

Emergency alerts and warnings are an element of emergency plans designed to provide information that members of the public can use to protect themselves. Alerts and warnings are used in diverse events, from slower-moving and forecastable events, such as hurricanes and some floods, to faster-moving or no-warning events, such as tornadoes or terrorist attacks. Whatever the event, the general goals of emergency alerting are similar: successfully transmit information to the potentially affected population, have that message spread or diffuse to the people who need it, and do so far enough before the incident occurs that they can act accordingly (Sorensen, 2000).

Emergency alerting has remained a mainstay of U.S. preparedness efforts for decades, including such tools as audible tornado sirens and the ubiquitous radio and television alerts that use the distinctive tone to grab listener or viewer attention.[1] But changes in technology and the way people consume media raised concerns that these legacy alert systems were no longer sufficient to achieve the goals of alerting (U.S. Government Accountability Office, 2007). For example, new devices provide additional alternatives to radio, and technologies allow ready recording of television for viewing at a later time. This means that the likelihood of reaching people in real time over those media has fallen. At the extreme, so-called cord cutters — households getting their media from streaming Internet sites rather than broadcast or cable media — might never receive emergency messages transmitted through such legacy channels.

A key part of these changes in media consumption is the rise of mobile devices, which has put cell phones and smart devices in the hands of increasing percentages of the population (comScore, 2015). These devices are the gateway for much of current media consumption and they have moved people away from traditional channels that could be used to transmit emergency alerts. Because their owners almost always have these devices in their possession, they provide an alternative channel — potentially even better than the legacy channels they are replacing — to reach members of the public with emergency alert messages. As these devices are always at least somewhat location aware, they can provide a way to transmit alert information relevant to the position where a person is at the time the emergency is occurring, rather than his or her home address or other less precise location information. To take advantage of this potential, the Wireless Emergency Alerts (WEA) system was implemented in 2012 to deliver short (up to 90 text characters) emergency messages to individual mobile devices in a designated warning area (Federal Communications Commission, 2014).[2] By delivering the message to

[1] See Moore, 2010, for a brief overview of the history of warning efforts in the United States.

[2] The website *Weather Alerts* (http://weatheralerts.alertblogger.com/) lists alert texts from WEA. An example flash flood WEA alert is: "Flash Flood Warning this area til 6:00 PM CDT. Avoid flood areas. Check local media. –NWS." An example tornado warning WEA alert is: "Tornado Warning in this area til 1:30 PM EDT. Take shelter now. Check local media. –NWS."

the mobile device, the goal is to get messages to individuals faster, where they are, and through a mode that will facilitate their paying attention and acting on the information.

WEA is designed to be a new addition to the toolbox for emergency alerting in the United States — not to replace other channels for transmitting information to the public. The full range of channels through which emergency alerts are intended to travel is shown in Figure 1.1. The figure shows the overall architecture of the Integrated Public Alert and Warning System (IPAWS), in which WEA is displayed second from the top among alert-dissemination channels. As a result, considerations about WEA and its effectiveness have to be addressed in the context of this overall alerting system, where it may play a critical role (by making it possible to deliver information to mobile devices) but is not the sole channel for disseminating a warning to the public.

*Figure 1.1. Integrated Public Alert and Warning System Architecture*

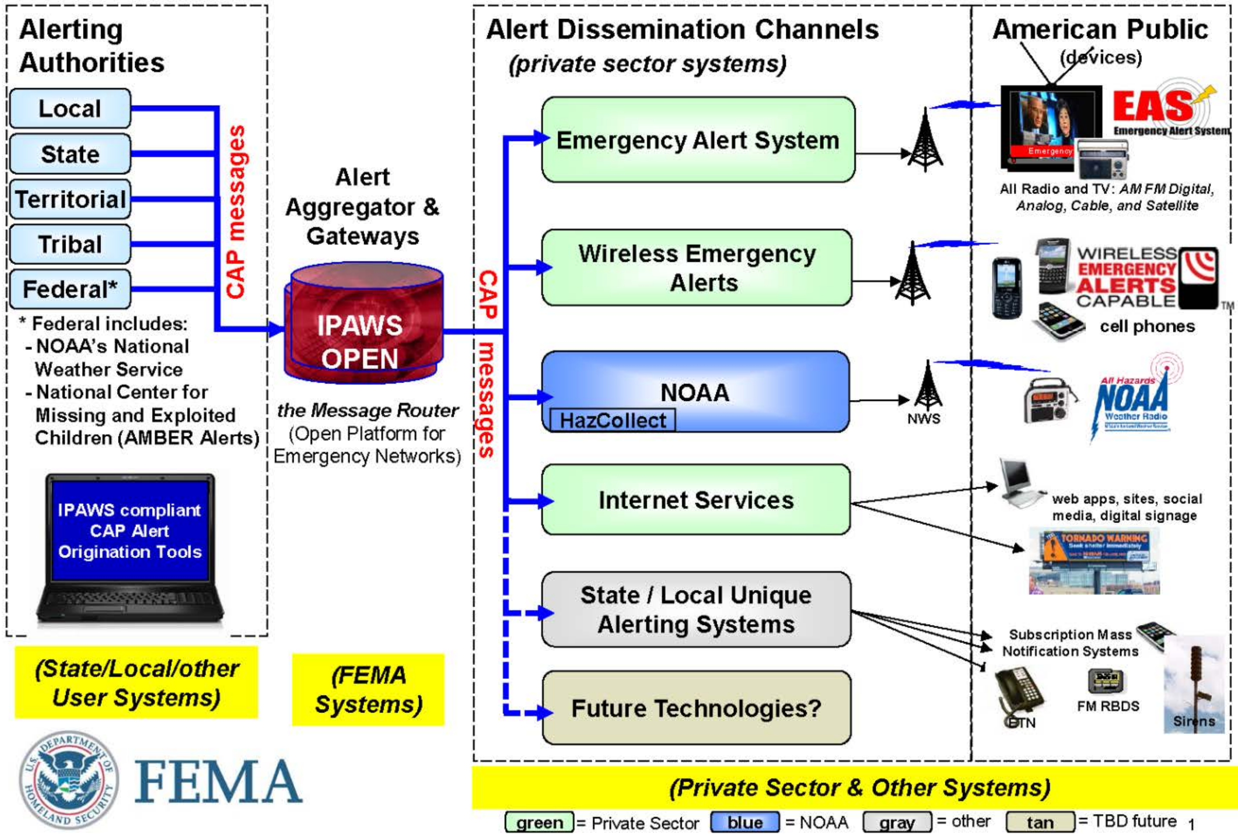


*NOTE: NOAA* = National Oceanic and Atmospheric Administration; CAP = Common Alerting Protocol; RBDS = *Radio Broadcast Data System*; TBD = to be determined.
*SOURCE: Federal Emergency Management Agency, 2015.*

## 1.1. Geo-Targeting Emergency Alerts

One of the promising features of delivery of alerts to modern mobile devices is the potential to specifically target alert information to the individuals that need it most.[3] For example, when a tornado warning is transmitted through legacy alerting channels, warning of the storm will be transmitted to a much larger area than will ultimately be affected by the incident — since, for example, television station coverage can reach over relatively large areas (Sutter and Erikson, 2010). Narrowing the alert area — termed geo-targeting the alert — offers a number of potential benefits. First, such targeting could increase the perceived salience of the alert to a receiver. Rather than a general alert that may be more likely to be ignored, a narrowly targeted message might increase the chance a recipient will pay attention and comply with alert direction. Second, narrow targeting could reduce the number of people who receive an alert who are not, in fact, at risk. This over-alerting is a particular concern for WEA as unnecessary messages delivered to mobile devices (with a distinctive tone intended to grab the recipients attention) could be viewed as a sufficient irritation for them to opt out of the system — essentially disconnecting themselves from this new alerting channel (Department of Homeland Security [DHS], 2013a). Furthermore, since WEA messages may only convey geographic information about the event through their geo-targeting, over-alerting raises the potential for those not at risk believing they actually are at risk (and behaving accordingly). Finally, high resolution geo-targeting could also make it possible to use emergency alerting in more nuanced ways, delivering different alerts to different "parts" of an emergency area to better communicate different risk levels or more effective protective action for subpopulations in the area. The rise of geo-targeting as a feature of alerting is a an alternative (and perhaps complementary) strategy to a focus on broad diffusion of alert messages — trying to ensure that the people who need an alert get it by blanketing everyone — to a focused application of alerting where one goal is seeking to limit diffusion of the message outside the targeted area.

## 1.2. Diffusion of Geo-Targeted Alerts

Technology could, in principle, make it possible to geo-target alert messages at a very specific level — e.g., delivering them to mobile subscribers within a single cell tower's footprint, or even narrower by taking advantage of technologies resident on the mobile device itself (Committee on Geotargeted Disaster Alerts and Warnings, 2013). Other factors can pull in a different direction, however, promoting the diffusion of messages beyond the area where they are delivered initially. In a connected age, individual people can be broadcasters of information, sending it to friends (e.g., through a phone call or email) or more broadly on the Internet (Vieweg et al., 2010; Gao et al., 2014; Sutton et al., 2014). Examples of the forwarding of alerts by individuals can be readily found in social media channels, such as Twitter (Figure 1.2), but forwarding could include phone, text, email, and other forms of social media or other mechanisms.

[3] See Committee on Geotargeted Disaster Alerts and Warnings, 2013, for a review.

*Figure 1.2. Examples of the Diffusion of WEA's Alerts by Individuals via Social Media*

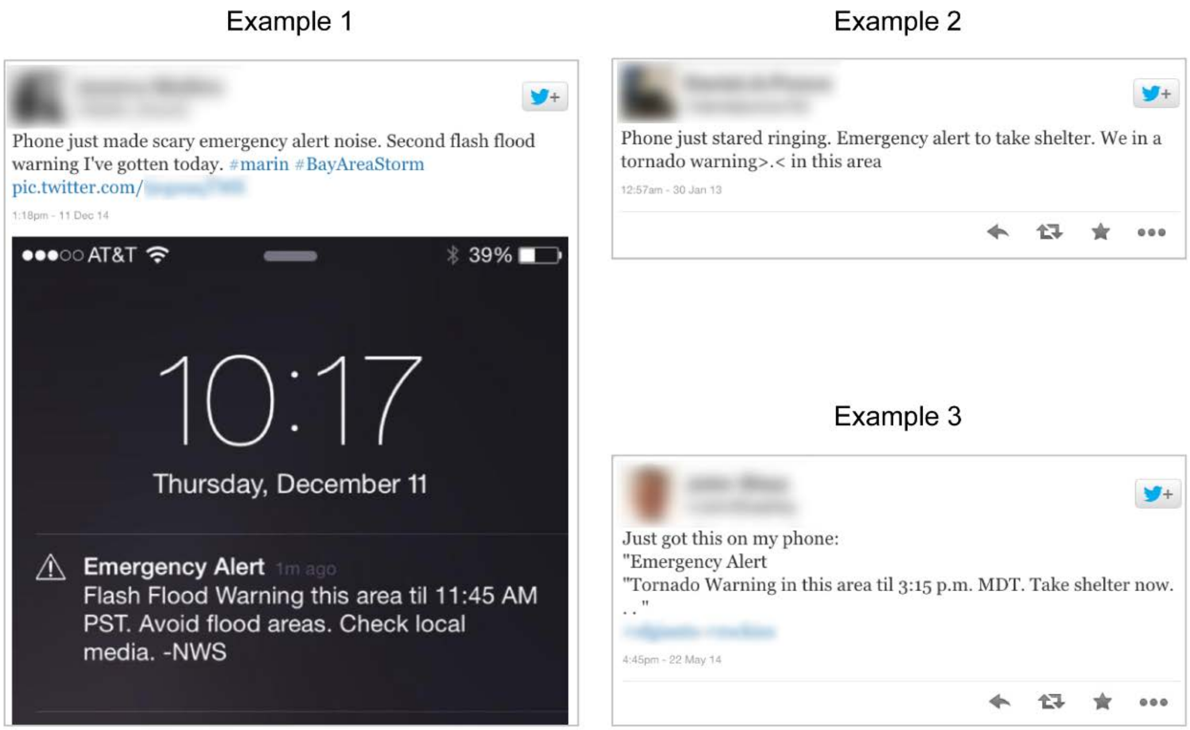


— Technology also has made it possible for individual cities or entities within a community (e.g., schools, colleges or universities) to have their own alerting mechanisms (Figure 1.3). Subscribers to these channels receive alerts transmitted from the city or others, providing another channel to obtain information relevant to their areas.

Although such mechanisms provide alternative ways for individuals to receive the information they need on channels they pay attention to, they also provide ways for a geo-targeted alert that was designed for one area to "leak out" and diffuse more generally. This potentially reduces the benefits of the focused delivery of messages to small areas. Alerting — particularly geo-targeted alerting — therefore includes different forces in opposition, pushing for both the broad dissemination of messages and the tight focus of messages in pursuit of different objectives. Understanding the value of geo-targeting and making decisions about it requires analyzing how these different forces combine to produce the net overall effect of targeting at different levels of resolution.

*Figure 1.3. Examples of the Diffusion of Emergency Alerts by Institutional Forwarders via Social Media*

Example 1


Example 2

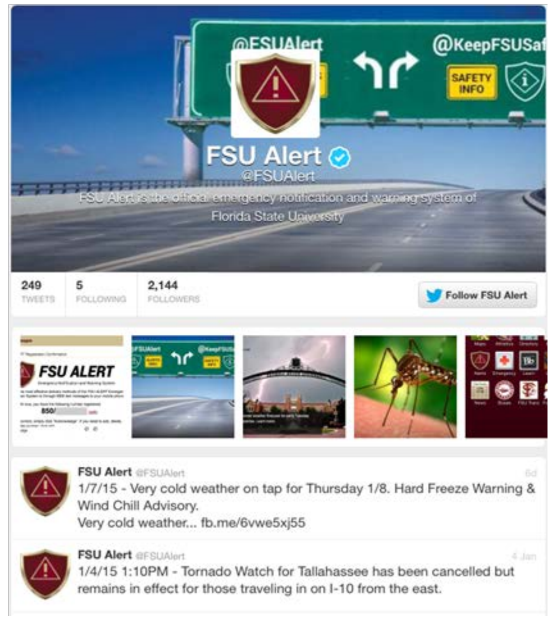


# 1.3. About This Study and Report

This project took on the fundamental question of how important diffusion behavior is for understanding the value of geo-targeting WEA messages. The study used agent-based modeling (ABM) to examine the diffusion of alerts within an area populated by individual recipients who move from place to place in the course of their daily activities. The use of ABM made it possible to simulate increases and decreases in the magnitude of different behaviors and processes, to see how this changed the effects of diffusion on geo-targeted alerts. Such modeling also makes it possible to explore dynamic behaviors that are sufficiently complex that they are difficult to solve mathematically — while still maintaining simplicity in the design of the model itself so the behaviors can be understood.

The research, and therefore this report, sought to answer three main questions:

- How does the interaction of geo-targeting and message-diffusion behavior affect the overall effectiveness of alerting?
- How do the effects of both geo-targeting and message diffusion vary for emergencies (and therefore alerting requirements) of different sizes?
- How does the interaction of geo-targeting and message diffusion affect the ability to use sophisticated alerting strategies, such as staging in time or delivering different emergency instructions to populations in different areas of risk?

Answers to these three questions make up the heart of this report, with each addressed separately in Chapters Four through Six.

Preceding that discussion, two chapters set the stage by describing the modeling effort that was used to answer the question. Chapter Two describes the modeling of emergency alerting, drawing heavily on the deep literature on alerting and human behavior during emergency situations. Chapter Three describes the broader landscape of the modeling effort, including both the human and physical geography used in the simulations and the specific emergency scenarios that were used to explore the different issues associated with geo-targeting and message diffusion. Chapter Seven concludes and explores the policy implications of the study's results.

# 2. Building a Model of Emergency Alerting and Diffusion

Due to the importance of warning as a strategy for protecting lives and property in emergency and disaster situations, the processes associated with warning and the variables that shape its effectiveness have been the focus of study for many years.[4] Seeking to build an understanding of everything from how groups of people behave in such situations to the implications for how warnings are communicated and received by individuals, sociologists, risk analysts and researchers from other fields have used methods ranging from the collection of data during and after major events to designing laboratory experiments to tease out subtle differences in human behavior (Bean et al., 2015). This deep literature provided the foundation for our efforts to build a model of individual human behavior relating to emergency alerting and diffusion. In this chapter, we sketch that model and summarize the available literature that we used to inform both how the model was designed and values of the many parameters it contained. These parameters defined how individuals in our simulations responded to and acted when alerted to an evolving emergency situation.

Because we were interested in understanding behavior by individuals during an emergency situation, affected by how an alert was received and diffused among a larger population, the study used an ABM. An ABM is a computer simulation where a population is represented by individual agents (in our model, representing people) that interact and make individual choices about their behavior based on a set of variables or rules. In our model, the agents could move around geographically, since understanding the effect of individuals moving in and out of areas where geo-targeted alerts were sent was the central aim of our work.

For understanding alerting and diffusion behavior,[5] we needed to understand and represent in the model a chain of events shown schematically in Figure 2.1. For individual agents, that chain of events was:

(1) receipt and understanding of an alert message, whether from WEA or another alerting channel;
(2) the decision by the agent to share or not share that alert with others, the central mechanism for alert diffusion in the model; and
(3) the decision by the agent whether or not to act on the alert and take whatever protective action it recommended.

[4] For example, see Quarantelli, 1982; Mileti and Sorensen, 1990; and Bean et al., 2015, which is a comprehensive annotated bibliography of this literature.
[5] The non–alert-related features of the model, including agent social networks, geography and movement, are described in the next chapter.

*Figure 2.1. Schematic Illustrating of Alert Stages for Modeling*

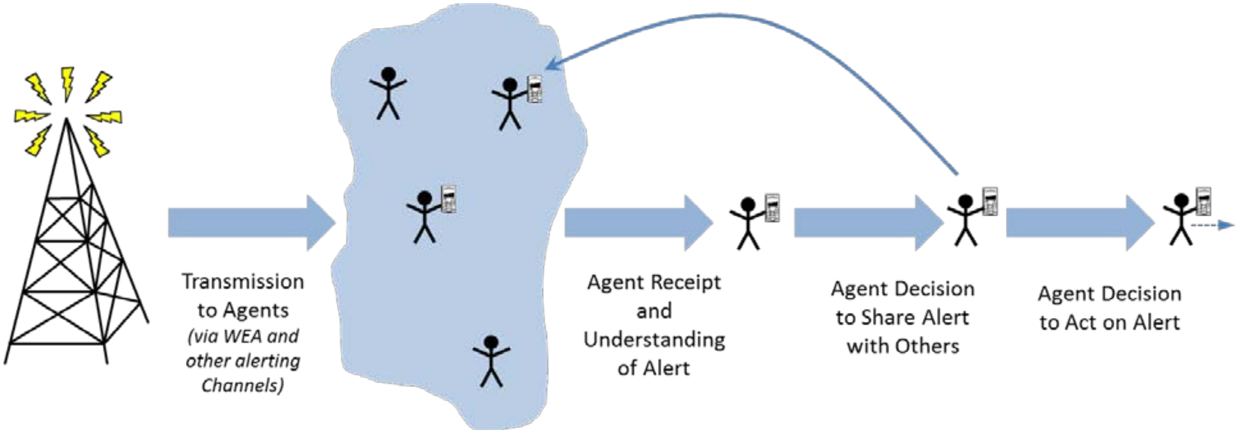


For modeling, the requirement was to both represent each of these steps or behaviors at an appropriate degree of simplification — since models, by definition, must significantly simplify reality — and make estimates of the various parameters and variables needed to define the behavior of the agents at each step. The remainder of this chapter explains how each element was represented for modeling and describes both the estimates made for each of the parameters and the support for those estimates in the available literature. We first describe the model geography and the agents themselves, then the transmission channels for alerting, and finally the agent behavior associated with the receipt, understanding, action and sharing of alerts.

# 2.1. Model Geography

Because the concept of geo-targeting alerts is inherently connected to geography, the starting point for our modeling was to (a) define a map within which our agents would move and interact and (b) provide the geographic reference points for the simulated emergencies and warning systems. Our model geography (shown in Figure 2.2) is loosely based on Dubuque, Iowa, including its network of major roads.[6] Dubuque was chosen because it was a relatively modest-sized city (with a population of approximately 58,000 in 2013) that would support a reasonable number of agents that could be simulated within computational limits (U.S. Census Bureau, 2015). Dubuque includes a river with the potential to flood, major railroads that could carry hazardous materials, and a major bridge that provides at least one natural choke point that could affect agent movement. It is important to note, however, that Dubuque was merely used as a starting point to provide a realistic foundation to the geographic model. We are not actually modeling Dubuque, since characteristics of our model (e.g., the absence of smaller roads in the network) do not reflect the city as it actually exists. In Figure 2.2 the

[6] The Dubuque road network is derived from Esri's StreetMap database, which contains detailed shapefiles of streets and highways in the United States and Canada. We limited the road network to major roads, motorways and highways to reduce the number of nodes (intersections) and edges (streets) in the graph. This allowed us to reduce the computational complexity needed to route the agents in the simulation.

boundaries of the simulated city are shown in green, waterways are identified in blue and major roads are shown in black. The yellow arrows depict agents within the simulation.

*Figure 2.2. Map of Simulated Area*

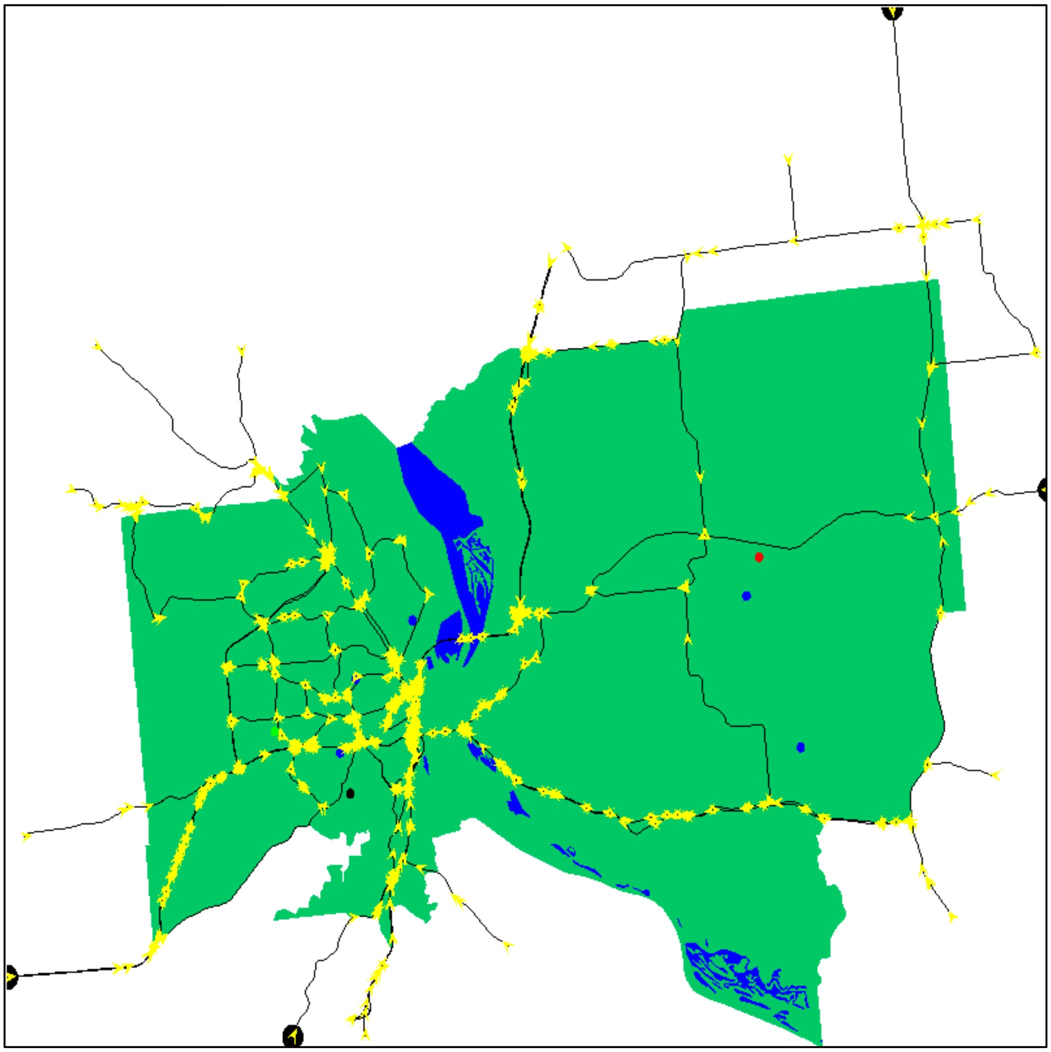

*NOTE: The city boundaries are green, waterways are blue and major roads are black. The yellow arrows depict agents. Source: Tele Atlas.*[7]

# 2.2. The Agent Population

Our model contains 10,000 agents who move about the region in the course of routine activities. Agent home and work locations were determined by using the Census 2000 Special Tabulation 64. This tabulation contains counts of workers in each census tract by tract of residence. (Here *worker* is defined as a person 16 years old or older who was employed during the last week of March 2000.) These counts

[7] Maps throughout this book were created using ArcGIS® software by Esri. ArcGIS® and ArcMap™ are the intellectual property of Esri and are used herein under license. Copyright © Esri. All rights reserved. For more information about Esri® software, please visit www.esri.com.

were used to do two probabilistic draws for each agent, selecting a home and work location. For simplicity of modeling, all agents begin at intersections of the road network — their "home" locations.[8] Our 10,000-agent model is a smaller population than the actual population of the city of Dubuque, with approximately one agent for every five actual residents.

Because one of the central behaviors we are interested in regarding alert diffusion is alert transmission, a key characteristic of our agents is how and how much they are connected to one another through social networks. In terms of the agents' social networks, we considered three different categories of connections where communications (i.e., forwarded alerts) might have different levels of influence on the recipients' behavior. We crafted ranges of connection numbers based on available literature on both individual social networks and communications during disasters (Hamill and Gilbert, 2010; Roberts et al., 2009; McCarty et al., 2001):

- *Family* — estimated to be in the range of four to five connections, on average;
- *Close social network* — defined as an intermediate-level group of friends over which the individual has some influence; we used a range of six to 15 connections for this group; and
- *Broad social network* — a much larger group of connections, perhaps only accessible through social media, over which the individual has modest influence; we used a range of 150 to 300 connections for this group.

From the perspective of examining how alerts diffuse through the population, the value used for the average number of connections among agents (and the attendant values that determine whether agents in the social network will pay attention to and act on a forwarded message) is an extremely important variable. Though our anchor points for social network sizes in the model were these values from the literature, in the course of the simulation, this variable became one that we varied to explore how changes in linkages among agents — and the resulting amount of message traffic of agents forwarding alerts to one another — affected outcomes. As a result, we ran cases where average *total agent social network size* was set at:

- 216 (corresponding to a network the size of the full broad social network described above, which was then divided into family, close and broad subsets);
- 42 (an intermediate case that essentially carved out a large percentage of the agent's broad network as uninfluenced by forwarded messages, with the remainder subdivided as above); and
- 6 (a microcase where the ability of individual agents to affect the behavior of others was quite small; the case was also subdivided into the three categories).[9]

Networks of each of these average sizes were generated with the Watts-Strogatz algorithm, with each of the values as the median degree and the rewiring probability set to 0.20 (Watts and Strogatz, 1998). This

[8] As a result of this decision, as will be described in more detail in the next chapter, on the emergency simulations, the model was run for 20 simulated minutes to allow agents to begin traveling and leave their home locations before any emergency alert was issued.

[9] Within these categories, 2.3 percent were probabilistically assigned to be in family, 4.7 percent in close social network, and 93.0 percent in broad social networks.

method was chosen because it produces “small world” networks with family-like groupings of agents. Given the total agent population of 10,000 in our simulated city, these median connectivity parameters, as well as the network sizes listed above, would result in the median agent connected to approximately 2 percent (216 of 10,000), 0.4 percent (42 of 10,000), and 0.06 percent (six of 10,000) of the other agents in the simulation, respectively.

## 2.3. Transmission of Alerts to Agents

Today, emergency alerts are transmitted to individual citizens in a variety of ways. Analogous to the performance of the legacy emergency alert system, various media providers have roles in transmitting emergency information, including the general media (e.g., television and radio broadcast, cable and satellite stations), specialized media (e.g., the Weather Channel) and non–English-language broadcast and cable stations. Radio stations transmit emergency alerts, including commercial radio and specialized radio systems (e.g., NOAA Weather Radio).

Web tools provide additional transmission options. These include:

- Social media channels connected with existing media sources (e.g., a television station’s Twitter feed).
- Specific applications to deliver alerts, subscription services (e.g., email or text alerts sent from local government or school systems).
- Dedicated websites designed to convey alert information in forms and formats useful for users (e.g., supporting individuals who are traveling by aggregating hazard data on their planned route).

Individuals using social media and direct communication from one individual to another are yet another route for alert information to spread laterally through the population, in contrast to paths that seek to reach broad subsets of the population simultaneously.[10]

Each way that alert information can be conveyed from an alert originator to citizens acts as a separate path, where one citizen might receive alerts through one or more path(s) in the course of an emergency. WEA is designed as an additional path with specific characteristics — not seeking to reach every member of an area simultaneously (as with radio or television), not aimed at a specific subscriber list (as with some web apps or text alert systems), but designed to quickly reach the population in a well-defined geographic area. WEA is intended not to replace the many other alerting mechanisms but to operate in concert with them (see Figure 1.1).

[10] Person-to-person transmission was shown in some case studies to have been a major pathway for warning information to be transmitted in a population. In a study of a tornado on a college campus, communication from others was the first warning received by one in five people and was a significant influence on their decision to comply (Sherman-Morris, 2010). In their case study, Vogt and Sorensen (1999) report even higher percentages of individuals receiving their first warning from family members, relatives, neighbors, friends, or coworkers.

Based on our review of the literature and existing alerting channels, we broke down the wide variety of systems and options into four classes of alerting effects, referred to here as *channels*, which concisely represent a large number of specific current and potential alerting mechanisms.

- The first channel is WEA itself, which is kept by itself both because of its unique characteristics and because it is the focus of this research. The WEA channel delivers an alert to everyone with a compatible device in one geo-targeted area. The WEA channel conveys limited information about the hazard and recommended protective action, without explicit information on the specific location of the incident (since this is conveyed through the geo-targeting of the alert itself). The size of the geographic area covered by the WEA channel was varied in our simulations to show the consequences of different levels of geo-targeting; any agents in the area at the time the alert was issued (or who subsequently entered the area) were sent the WEA message.
- The second channel, media, delivers an alert indiscriminately to the entire region. Media-delivered alerts are presumed to have the potential to hit everyone in that area. Media-delivered alerts contain the location of the incident, allowing agents to separately decide whether the information in the alert applies to them and whether they should take protective action.
- The third channel, technical forwarding, includes tools that deliver alerts to one targeted group or group within an area. This includes, for example, the Twitter feed of a local emergency-management agency or text messages sent to parents by their children's schools. Technical forwarding is presumed to be subscription based, so it is constrained to the list of subscribers. Those subscribers can be scattered across the entire city (e.g., for the emergency-management Twitter feed) or could be geographically constrained (e.g., parents close to a single school). Like media messages, these alerts contain both information on the location of the incident and recommended protective action.
- The final channel is individual forwarding. A citizen within the model, upon receiving an alert, has the potential to forward that alert to those in the citizen's personal social network. This forwarding could include phone calls, texts, emails, social media (e.g., Facebook, Twitter) or other mechanisms. The forwarded message contains whatever information was contained in the original message (that is, a forwarded WEA message will not contain information about the location of the event, but forwarded warnings from other sources would include it).[11]

As implemented in the model, a singular *alerting agent* (who can be considered the originator-IPAWS end of the distribution chain) sends the alert to the media and technical forwarders (who then pass the alert along to citizens), as well as transmits the alert directly to citizens through the WEA channel. Individual citizens then may forward alerts to those in their own personal social networks.

[11] Individual forwarders could add additional information when forwarding an alert. For the model presented here, we made the simplifying assumption of verbatim forwarding.

## 2.4. Agent Alert Behavior

Within the model, citizen agents move about the road network within a simulated area. In the absence of an alert, they are either idle or traveling to or from work. Once an alert is transmitted, agents have a chance to receive the alert through one of the four channels. Once received, the agents decide (1) whether or not to comply with the message directive, and (2) whether or not to forward the message to those in their personal social networks. As will be described in more detail in the next chapter, the message directive either said to evacuate (move to one of several evacuation points at the edge of the simulated area) or shelter in place (halt where the agent is at the point where he or she makes the decision to comply with the alert).

These events and agent decisions are affected by several agent characteristics, represented as parameters in the model. Phone compatibility and opting out limit WEA's reach, specifically. Individual situations may vary how likely it is that an individual receives an alert or the extent to which that receipt is delayed. And once agents receive an alert, it is not guaranteed that they will either understand or act on that alert. The probability of understanding an alert may be fundamentally different for non-English speakers or those with sensory disabilities. People may differ in their motivation to comply with alert guidance, which is reflected in the likelihood of compliance and any delay in compliance. Similarly, people may differ in their likelihood to forward an alert that they received. Finally, these parameters may vary across individuals, the alerting channel, the nature of the alert event and other contextual factors.

Broadly speaking, these parameters correspond to the sequence of steps that Mileti and colleagues proposed for how individuals respond to an alert or warning (Mileti and Sorensen, 1990; Mileti and Peek, 2000):

1. hearing the warning;
2. understanding it;
3. believing that the warning is credible and correct;
4. personalizing the warning;
5. milling behavior, including confirming the warning and seeing whether others are acting on it; and
6. acting.

When compared with the parameters in our model, step 1 broadly corresponds to the probability and timing of alert receipt; step 2 corresponds to the probability of understanding; steps 3, 4 and 6 are reflected in the probability and timing of action, broadly speaking; and steps 4 and 5 are aspects reflected in the probability of forwarding and to whom an individual forwards.[12] The following subsections describe the modeling of each of these steps in turn.

---

[12] Our model components are largely equivalent to those used by Morrow, Stoddard, and Elm, 2014, in their model of recipient trust of WEA messages.

### 2.4.1. Agent Receipt and Understanding of Alert

In the model, whether or not an agent received a message through a specific channel was modeled as a simple probability, with the probability of receipt varying across the different channel types. Table 2.1 summarizes the probability of receipt parameters for the classes of channels, with explanations and references to the relevant literature.

*Table 2.1. Parameters Determining the Probability of Agents' Receipt of Alerts*

| Channel | Probability of Receipt | Notes | Sources |
|---|---|---|---|
| WEA | Conditional on possession of a WEA-capable device:<br>Average of 60 percent<br>For individual agents, drawn from uniform distribution, 30–90 percent | The cases explored for the possession of compatible device by 30 percent of population and 70 percent of the population.<br>The main barrier to receiving an alert is an in-progress call at the time of transmission or being in an area without a wireless signal. This estimate is based on a medium level of network call traffic.<br>Agents who have made a choice to comply with or ignore an alert become immune to future WEA message, reflecting the fact that most mobile devices only display the first receipt of a specific alert from WEA. This does not apply to forwarded alerts (see below).<br>As is the case in reality, agents may opt out of receiving WEA messages, which will result in non-receipt even if they have a compatible mobile device. In our simulations, we set the opt-out rate very low, at 1 percent of the agents. | DHS, 2013b; DHS, 2013c; Jagtman, 2010 |
| Media | Average of 19 percent<br>For individual agents, drawn from uniform distribution, 5–33 percent | | Sherman-Morris, 2010; Mitchell et al., 2005; Vogt and Sorensen, 1999 |
| Technical forwarding | Conditional on being a subscriber:<br>Average of 60 percent<br>For individual agents, drawn from uniform distribution, 30–90 percent | The subscription rate was done differently for two variations of technical forwarding:<br>• For the technical forwarders who cover the entire area, a 20-percent subscription rate was used across all agents<br>• For geographically concentrated technical forwarders (e.g., in the vicinity of a school), 20 percent of the population within 2.5 kilometers subscribed | Sherman-Morris, 2010 |
| Individual forwarding | | | |

| Channel | Probability of Receipt | Notes | Sources |
|---|---|---|---|
| Family | Conditional on being a recipient ofa forwarded message:<br>Average of 90 percent<br>For individual agents, drawn from uniform distribution, 80–100 percent | There was relatively little specific basis for making a detailed estimate. This was estimated relative to (and significantly higher than) WEA due to the relationship. | Jagtman, 2010; Parker<br>and Handmer,<br>1998 |
| Close social network | Conditional on being a recipient of a forwarded message:<br>Average of 40 percent<br>For individual agents, drawn from uniform distribution, 20–60 percent | There was relatively little specific basis for making a detailed estimate. This was estimated relative to (and significantly lower than) family forwarding receipt due to the relationship. We set the same value for receipt between close and broad social network on the assumption that such forwarding would occur through the same technological pathway (e.g., social media) and consequently should be similar. | |
| Broad social network | Conditional on being a recipient of a forwarded message:<br>Average of 40 percent<br>For individual agents, drawn from uniform distribution, 20–60 percent | There was relatively little specific basis for making a detailed estimate. This was estimated relative to (and significantly lower than) family forwarding receipt due to the relationship. We set the same value for receipt between close and broad social network on the assumption that such forwarding would occur through the same technological pathway (e.g., social media) and consequently should be similar. | |

*NOTE: Since all probabilities of receipt were modeled as simple percentage chances, our model did not include the real-life behavior of warning channels having geographic zones with different probabilities of receipt (e.g., cell system dead zones, areas with reduced media reception).*

For different alerting channels, technical and other differences result in delays before issued alerts are received by the target population. Some of those delays stem from the systems that carry the alerts — e.g., congestion in wireless systems when standard text messaging systems are used for alerts. Others result from human behavior — e.g., a recipient who has a media broadcast on but does not notice when an alert is initially broadcast because he or she is not paying attention. In our initial specification of the model, we set three different ranges for receipt delays for different sets of alert channels, based on information from literature sources. WEA was the most rapid channel, with receipt at 20 seconds (DHS, 2013b). Both technical and individual forwarding channels were slower and had a wider range, with receipt between one and 10 minutes after transmission (Pries, Hobfeld, and Tran-Gia, 2006; Bambenek and Klus, 2008; Traynor, 2010). Media was the slowest, with a range from one to 25 minutes (Mersham, 2010).

Because of the emergency scenarios chosen for simulation, these delay parameters resulted in the modeled emergency alerts having very limited effects, which — while potentially realistic — prevented the exploration of the range of potential effects of message forwarding. As a result, in later trials, we diverged more considerably from reality and carried out simulations both with this parameter eliminated (all messages were delivered immediately) and a more moderate case where we reduced only the delays for WEA (immediate) and interpersonal forwarding (two to three minutes).

Even if an agent receives an alert message, there is no certainty that the message will be understood. Specific literature on the level of understanding of short alert messages is very limited, but it can be assumed that there will be some level of misunderstanding by recipients simply due to the short length of WEA messages — which, at a maximum of 90 characters, are shorter than Short Message Service (SMS) text messages can be (DHS, 2014). Beyond that potential, however, there are also certain populations with additional issues that could increase the likelihood that a short text alert is not understood. These populations include individuals with hearing and vision limitations (which occur in 3.6 percent and 2.2 percent of the population, respectively; U.S. Census Bureau, 2010) and non-English speakers or individuals who read English at "below basic" levels (groups which are estimated to make up approximately 15 percent of the population; National Center for Education Statistics, 2006). Other demographic characteristics have been identified that affect comprehension in individual emergency warning scenarios, which also suggests the value of including variance in the estimates of perceived comprehension when modeling (Brotzge and Donner, 2013).[13] For the model, we drew on this literature to represent the probability of understanding in a parametric way, focused on the relationship between different alert channels — i.e., different alert channels would have more or less of a chance that recipients would understand messages relative to other channels in the simulation (Table 2.2). To capture the fact that some populations may have greater difficulty understanding short text warnings in English, we divided the population into a low-understanding group (15 percent) and a high-understanding group (85 percent), assigned as a random draw for each agent in the simulation.

*Table 2.2. Parameters Determining the Probability of Agents' Understanding of Alerts*

| Channel | Probability of Understanding | Notes |
|---|---|---|
| WEA | High understanding:<br>Average of 90 percent, drawn from uniform distribution, 80–100 percent<br><br>Low understanding:<br>Average of 30 percent, drawn from uniform distribution, 20–40 percent | This is our best guess for the level of understanding that provided a probability for misunderstanding a message in the entire population, with a considerably higher probability for the low-understanding portion of the population. |
| Media | Average of 90 percent<br>For individual agents:<br>Drawn from uniform distribution, 80–100 percent | We assumed that agents who received a message from a media source would be using a source that addressed their specific understanding needs (e.g., non–English-language media for non-English speakers). |
| Technical forwarding | High understanding:<br>Average of 90 percent, drawn from uniform distribution, 80–100 percent<br><br>Low understanding:<br>Average of 30 percent, drawn from uniform distribution 20–40 percent | This was assigned the same as WEA, assuming similar understanding concerns for likely text message or email-based technical forwarding channels. |

[13] Note that there is a wider literature about the effect of multiple alert sources on comprehension, as well as trust in a message, which affects compliance (Mileti and Darlington, 1995; Wood and Weisman, 2003; Benavides and Arlikatti, 2010).Studies have shown that individuals are more likely to understand alerts with local information (Berry, 1999; Brotzge and Donner, 2013). Our modeling included multiple alert sources in only a modest way, and since we were not examining message content, this literature was less applicable to what we were doing here.

| Channel | Probability of Understanding | Notes |
|---|---|---|
| Individual forwarding | | |
| Family | Average of 95 percent<br><br>For individual agents, drawn from uniform distribution, 85–100 percent | We assumed that communication from a family member would have a higher probability of understanding than WEA given the ability of family to match recipient information needs and adapt the message appropriately for those needs. |
| Close social network | Average of 85 percent<br><br>For individual agents, drawn from uniform distribution, 75–95 percent | We assumed a modest reduction in the average understandability from WEA given the communication medium (e.g., individual postings on social media). |
| Broad social network | High understanding: Average 90 percent, drawn from uniform distribution 80–100 percent.<br><br>Low understanding: Average 30 percent, drawn from uniform distribution 20–40 percent. | We assigned this the same as WEA, assuming similar understanding concerns for the likely social media–based postings. |

### 2.4.2. Agent Decisions to Act on Alerts

Even if individuals receive and understand an alert, research has shown that their acting upon the alert is far from a certainty. Compliance rates with alerts have been measured for different types of emergency scenarios, showing significant variation. Rates are quite high for those at high risk — e.g., approximately 60 percent in a hurricane, 70 percent in various floods and essentially full compliance in scenarios such as chemical accidents (Rogers et al., 1990; Mitchell et al., 2005; Archibald and McNeil, 2012). In other cases, compliance can be lower — e.g., in a survey across 12 Louisiana parishes, those saying that they would evacuate from a hurricane ranged from 27 to 52 percent (Southeast Louisiana Hurricane Taskforce et al., 2005). Significant variance has even been observed within the same threat type; — in one study, measures of compliance across a wide range of tornado warnings varied from 31 to 89 percent (Sutter and Erickson, 2010). Studies have also shown significant demographic differences in compliance with warnings, which argues for including agent-to-agent variation in compliance probabilities in modeling efforts (Donner, Rodriguez, and Diaz, 2012; Vogt and Sorenson, 1999, Sherman-Morris, 2010).

Literature points to a number of phenomena that shape an individual's decision to comply with emergency alerts. A significant *normalcy bias* — where individuals' baseline assumption is that an alert does not apply to them if they do not already feel at risk — implies a low baseline for alert compliance (Gow et al., 2009; Rice et al., 2010). That baseline may be overcome, however, by individuals seeking additional information about the situation, talking to others and directly gathering information (e.g., looking out the window at the oncoming storm; Vogt and Sorenson, 1999). Assuming that the individuals find the alert credible and believe it,[14] this process of concluding that the alert applies to them has been termed *personalizing* (Mileti and Peek, 2000). A range of factors appear to increase the personalization of risk delivered in a warning, including the inclusion of personally relevant information (Drabek, 1999); geographic proximity (Brotzge and Donner, 2013; Sutton et al., 2014); smaller warning areas (Nagele and Trainor, 2012), particularly relevant for geo-targeted alerts; and the reinforcement of the message by personal contacts (Mitchem, 2003; Burnside, Miller, and Rivera, 2007), particularly

[14] From a credibility perspective, official sources appear to be especially influential for recipients (Durange et al., 2014; Freberg, 2012).

relevant for considering individual forwarding of messages. As was the case for understanding the alert, the probability of acting on an alert rises as individuals receive more copies of it — i.e., the credibility, salience or information transmitted to the agent accumulates with receipt of the warning from different channels or over time (Quarantelli, 1982; Casteel and Downing, 2013).

Although much of noncompliance with emergency alerts involves individuals simply ignoring them, in some cases, it includes action different from that recommended in the alert. For example, compliance may be lower in cases where individuals have been directed to shelter in place rather than evacuate (Sorensen, Shumpert, and Vogt, 2004; Vogt and Sorensen, 1999), and a large number of characteristics can affect individuals' decision one way or the other (Liu, Murray-Tuite, and Schweitzer, 2012). Furthermore, in some cases, individuals converge rather than evacuate, seeking loved ones or pets before evacuating as a single unit (Vogt and Sorensen, 1999; Eisenman et al., 2007; Durange et al., 2014; Liu, Murray-Tuite, and Schweitzer, 2012; Schultz et al., 2010). In our model, we defined probabilities of compliance for different alerting channels based on this general literature, but with a focus on placing probabilities' relative values appropriately given these challenges (e.g., that there would be a higher probability of compliance for a warning forwarded by a family member than from a general alerting channel). To include the potential for convergence behavior, the technical forwarder we designed to simulate messages forwarded from a neighborhood school (where subscribers were individuals living in close proximity to the school) produced convergence. Table 2.3 summarizes the parameters used for the probability of compliance for each alerting channel. Some available literature supports the idea that different alerting channels will have differential compliance rates. In survey data from a tornado, Mitchum (2003) observed that individuals receiving warning from the National Weather Service Weather Radio network or from a friend or relative (i.e., our individual forwarding pathways) were more likely to perceive the threat of the event, which produced different behaviors. As will be described in the modeling experiments we explored multiple levels of compliance to WEA messages when exploring the different potential effects of geo-targeting.

*Table 2.3. Parameters Determining the Probability of Agent Compliance with Alerts*

| Channel | Probability of Compliance | Notes |
|---|---|---|
| WEA | Average of 20 percent<br>For individual agents: drawn from uniform distribution, 10–30 percent | Because of the wide variation in reported compliance values across different emergency incidents, we selected a single average value for WEA compliance (with a distribution around it) that was used for all scenarios. This average value, 20 percent, was at the low end of reported compliance values in some scenarios (described in the text). Having compliance at a low value provided room for assigning values to other alert modes that were viewed as likely to produce higher compliance. |
| Media | Average: 15 percent<br>For individual agents, drawn from uniform distribution, 0–30 percent | We assumed this to be lower than WEA. |
| Technical forwarding | Average of 20 percent<br>For individual agents drawn from uniform distribution, 10-30 percent | We assumed this to be comparable to WEA. |
| Individual forwarding | | |
| Family | Average of 50 percent<br>For individual agents, drawn from | We assumed to be more than twice WEA. |

| Channel | Probability of Compliance | Notes |
|---|---|---|
| | uniform distribution, 30–70 percent | |
| Close social network | Average of 30 percent<br>For individual agents, drawn from uniform distribution, 10–50 percent | We assumed this to be slightly higher than WEA. |
| Broad social network | Average of 10 percent<br>For individual agents, drawn from uniform distribution, 0–20 percent | We assumed this to be much lower than WEA. |

In addition to the initial decision to comply with alert directions, the *time before an individual acts* is also an important factor. While some delay before action may have a limited effect in a slow-moving emergency event (e.g., a long pre-hurricane evacuation), delay in a fast-moving emergency (e.g., a tornado) can be equivalent to receiving no warning at all. As was the case with the probability of compliance, considerable variation has been observed in the timing of response to warnings in disaster events. For example, after a train derailment, almost 60 percent of the warned population evacuated within 30 minutes (Sorensen, 1991). But slower-moving events have shown *S*-shaped curves in evacuation timing; some evacuate very quickly, a larger fraction of the population do so over an intermediate time period and a straggling group takes longer. Such a curve was reported by Vogt and Sorensen (1999) for a hazardous-material event. In this case, approximately 10 percent of the evacuees left within 20 minutes of receiving the warning, but there were significantly longer delays before the later-evacuating groups left, reaching and then exceeding an hour of delay. The literature has shown that cues that confirm risk (e.g., events where recipients of an alert can visually see an approaching hazard) can hasten evacuation (Vogt and Sorensen, 1999).

Delay before compliance is related to what has been termed *milling behavior* (Committee on Public Response to Alerts and Warnings Using Social Media [CPRAWSM], 2013), where recipients of an alert seek out additional information from other sources, look at the reactions of other nearby people to the alert information and participate in other pre-evacuation activities. Part of this milling behavior includes sharing the alert information with others (our modeling of which is discussed in the next section) as part of seeking their views of the risk and appropriate actions. In our initial modeling efforts, we included a delay before compliance (milling time) of zero to 15 minutes (drawn from a uniform distribution) for all alert channels. Additional messages received by the agents reduced the time before action (by multiplying the time by 0.9 whenever an additional warning message was received). As was the case with our modeled delay before messages were *received* by the agents, this delay before *compliance* meant that few warning effects were observed in many of our scenarios, limiting our ability to examine message-diffusion effects. In later simulations (as described in Chapters Four, Five and Six), this time was reduced zero, effectively consolidating the delay into the receipt-delay parameters.

### 2.4.3. Agents' Decisions to Share Alerts with Others

The central behavior of interest in this study is the diffusion of messages, which is driven by agents' decisions to share the alerts they receive (from whatever source) with others. We divided the groups that information may be shared with into three classes: family, close social network and broad social network. There were increasing numbers of individuals in each class. There is extensive documentation of such diffusion in the literature, with data providing breakdowns of the source of warning for individuals in specific incidents. For example, of the warned population in a number of events, 22 to 46 percent said that they received a warning from either friends or family (with friends closer to the low end of that range and family higher; Parker, Priest, and Tapsell, 2009; Parker, Tunstall, and McCarthy, 2007; Mitchell et al., 2005; Vogt and Sorenson, 1999; Sherman-Morris, 2010; Legates and Biddle, 1999). Analyses of individual communication behavior *during* disasters showed a spread of information by individuals involved in an emergency; of the increase in communications that occurred in an emergency by people contacted by individuals within the emergency region, 25 percent of those communications went to individuals within the emergency region (i.e., no geographic diffusion) while 70 percent was to other individuals outside the region (i.e., geographic diffusion; Gao et al., 2014).

Large percentages of individuals (50 to 80 percent) also expressed a willingness to post about disasters on social media, potentially enabling broader diffusion (Chae et al., 2014; CPRAWSM, 2013). More specific studies of social media behavior in emergencies have also identified key high volume individuals who forward more information about the event and examined the increased retweeting probability of emergency related tweets during events (Hughes and Palen, 2009; Vieweg et al., 2010).

Beyond establishing the existence of the behavior, available literature provided only limited insight into how to estimate the probability of individuals forwarding alert information received during emergencies. For example, Chew and Eysenbach (2010) documented an increase in tweets using the term *H1N1* between May and December 2009. This work did not specifically examine WEA. This likelihood, however, almost certainly differs across message channels.  For example, it is less likely than an individual contacted directly by a member of their immediate family will forward that message onward to others, while someone messaged about a hazard situation on social media by a friend might pass it on. In the absence of data to support more detailed estimates, we divided the alert channels into three groups:

- No forwarding — messages from family members had zero chance of being forwarded onward by agents;
- Five percent chance of forwarding — messages from broad social networks, the media and technical forwarders fell into this category; and
- 10 percent chance of forwarding — messages from close social networks and WEA fell into this category.

For individual agents, the probability of forwarding was varied slightly (plus or minus 20 percent of the average value) to provide some heterogeneity across the agent population. In later experiments, we used forwarding probability for WEA as a manipulated variable, turning it off and increasing it to double or triple the baseline rate to explore changes in behavior. Note that, in our simulation, we are only capturing the spread of the alert from agent to agent through individual forwarding behavior — we do not explicitly model the potential for the spread of incorrect or misinformation via these routes which has been raised as a concern in the literature (Sutton, Palen, and Shklovski, 2008).

A relationship may exist between the type of forwarding behavior we are considering here and the process of milling before acting on an alert, though such a relationship has not been established in the literature:

> Although many first responders believe that social media have given them less control over the warning process, informal dissemination of messages has always played an important role in the warning process. Indeed, one might conjecture that the inherently social nature of social media might help reduce milling time, but whether this is true and under what conditions is an open research question. (CPRAWSM, 2013; Sutton et al., 2014)[15]

Figure 2.3 shows an example of social media forwarding activity which, when coupled with the subsequent discussion, clearly resembles the social interaction and confirmation that goes on during milling and delay before compliance with traditional warning methods.

*Figure 2.3. Forward Warning/Social Media Exchange as Milling Behavior and Warning Confirmation*

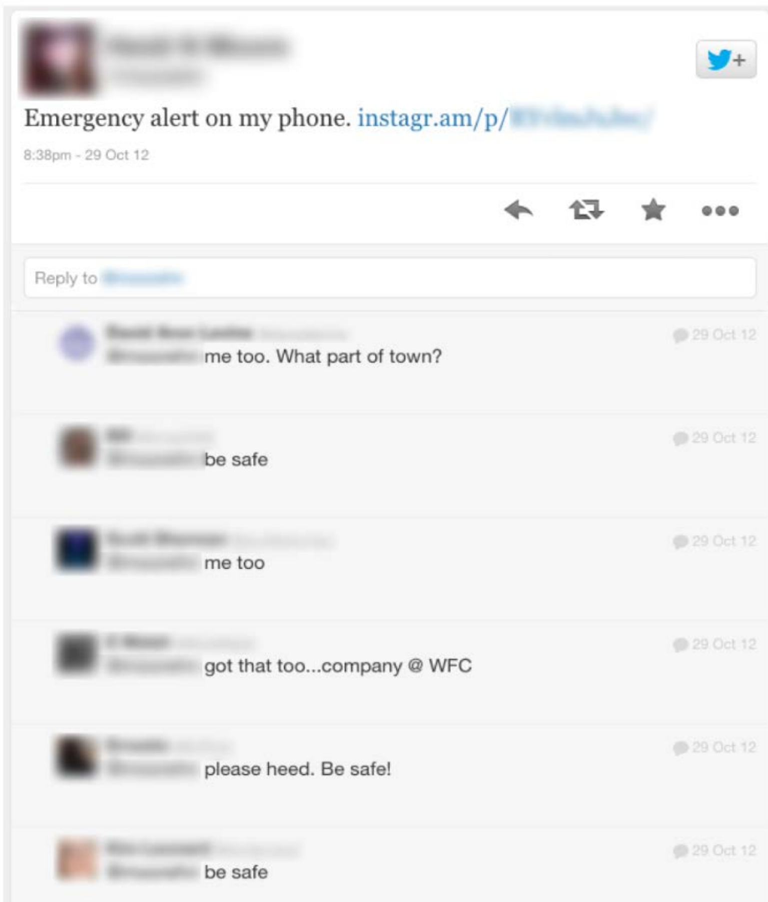


[15] In their WEA-focused modeling work, Morrow, Stoddard II and Elm (2014) include "confirmation by social media" as one contributor to individuals believing, and therefore acting, upon alerts.

## 2.5. Summary

Because of the role emergency alerting has played as a component of emergency preparedness for decades, there has been a significant amount of research on the process of individual receipt, understanding and action upon warning information. More recently, with the increasing role of mobile and social media in individual communications, literature on their particular role in disasters has been published. Data collected from past events provides a basis for making parameter estimates for modeling purposes, and for defining ranges to vary those estimates over the agent population in our simulation. The availability of information across the different parameter sets does vary, however. The most useful data was available on parameters surrounding agent receipt and understanding of messages. While information from surveys in past disasters provided a wide variety of values for compliance, that variability led us to a more relative strategy — defining a compliance value for the WEA alert channel and logically placing the values for other channels above or below it. And while individuals clearly share and forward warning messages, data were insufficient to conclusively estimate the likelihood of forwarding, leading us to simply define an approximate set of values. As is frequently the case for modeling efforts, such limitations provide a reason for exploratory analysis within the parameter space, some of which was done as part of the analyses reported in the following chapters.

# 3. Defining the Emergencies and WEA Performance in the Study

Beyond the structure of the model itself, which defined geography and agent behavior, two other elements were required for the study: the emergency scenarios in which warning and agent behavior would be tested and the measures against which outcomes (and changes in those outcomes as a result of alert diffusion) would be assessed.

## 3.1. Modeled Emergency Scenarios

Due to the wide variation that exists in types and characteristics of emergency and disaster scenarios, many possible scenarios could have provided the basis for the study. They range from very large to very small, with hazards originating from such sources as weather, geological events, technological accidents and security risks (e.g., a terrorist attack). To provide a starting point, we examined data provided by IPAWS at DHS describing the WEA messages that had been issued to date. The graph of that data is shown in Figure 3.1, highlighting the high frequency of flash flood and tornado warnings in WEA alerting.

*Figure 3.1. Frequency of WEA Messages by Event Type, as of October 2014*

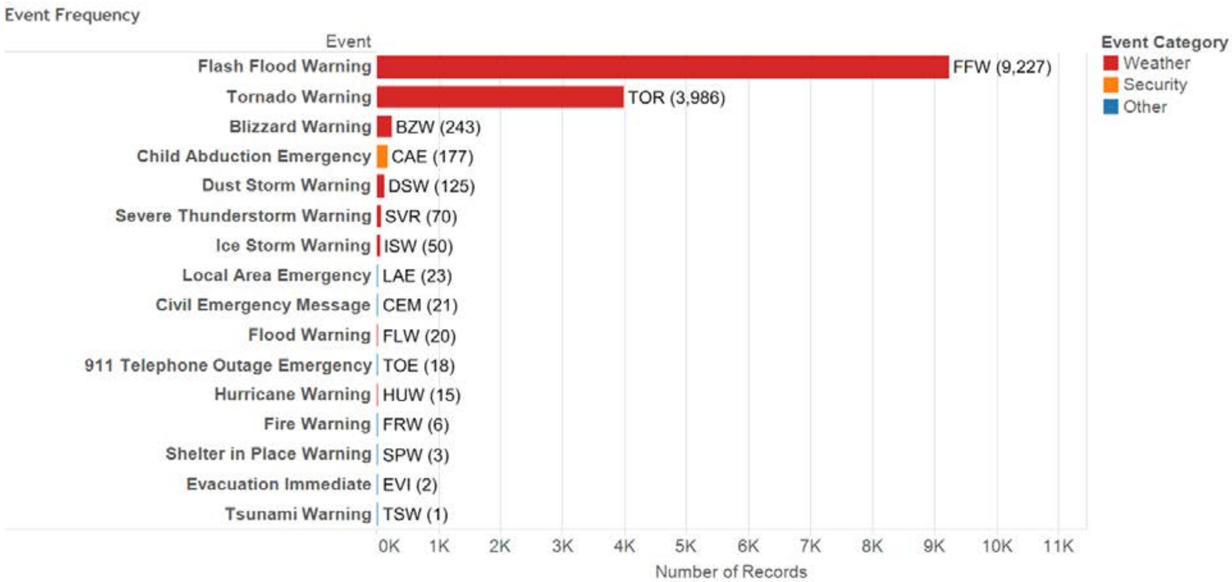


To better represent the characteristics of these different emergency events beyond their frequency, we qualitatively binned them into small, medium and large events (based on rough estimates of the size of the emergencies themselves) and the average length of the alerts (which was included in the IPAWS data set.) The resulting two-dimensional representation (where the sizes of the circles represent the frequency data from Figure 3.1) is shown in Figure 3.2. Based on information from the literature we analyzed in the model design and parameterization, we added one type of emergency that was not included in the IPAWS data set (a release of hazardous materials) because an example was available in the literature (Vogt and Sorensen, 1999). This provided the opportunity to examine a scenario where different alerts were sent to different geographic areas.

*Figure 3.2. Qualitative Representation of Emergency Characteristics Based on IPAWS Warning Data*

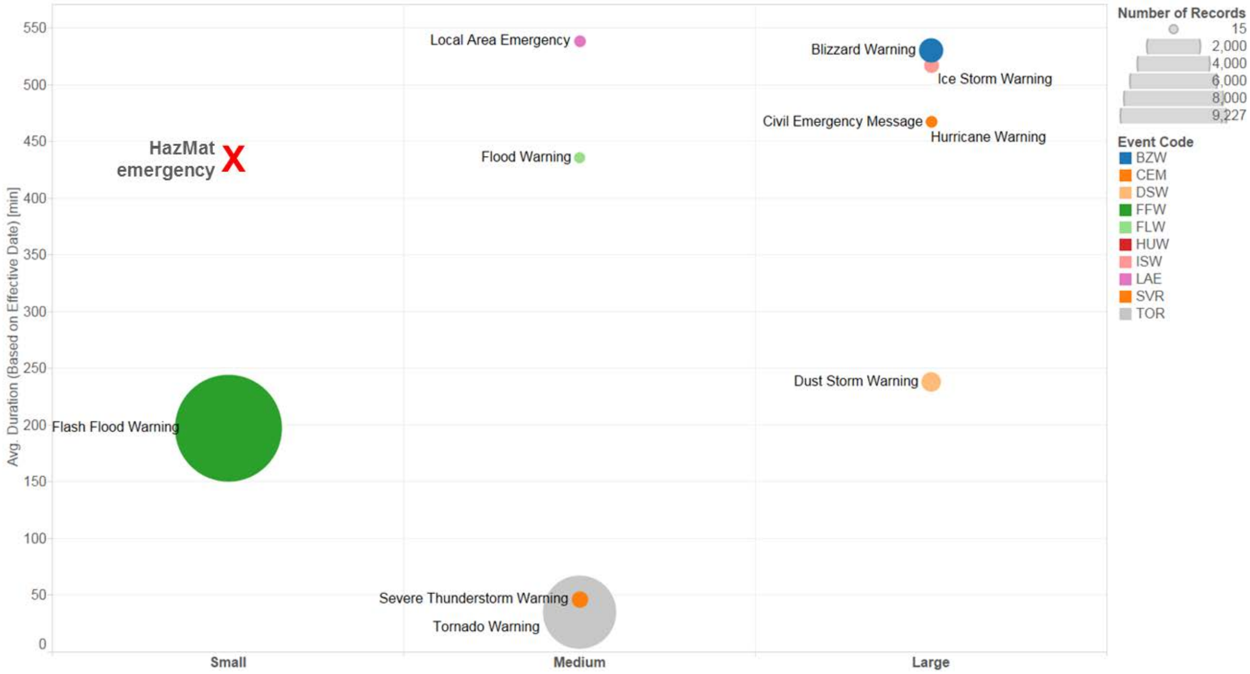


*NOTE: The sizes of the circles represent the frequency data from Figure 3.1.*

For our simulations, we used four scenarios; this provided a large enough set to cover a number of emergency characteristics while remaining practical, given the scope of the effort. The scenarios were:

1. *Flash flood* — a frequent scenario that has been the focus of large numbers of WEA messages, covering a small area and with a short timeline, where individuals at risk would be directed to evacuate the affected area;
2. *Tornado* — also a frequent focus of past WEA messages, but covering a larger area and with a potentially very short alert-to-emergency timeline, where individuals at risk would be directed to shelter in place until the emergency ended;
3. *Hazmat plume* — a small-sized event, but one that could involve multiple hazard areas where different protective actions (evacuation versus sheltering in place) would be most appropriate; and
4. *Major flood* — a rarer but larger event, potentially affecting a very large population and therefore requiring a significant multistage evacuation over a longer emergency timeline.

These four scenarios provided a set of emergencies with distinct characteristics that would make it possible to test different elements of alerting strategy and performance. The scenarios also provided cases where the ability to geo-target might facilitate more-complex responses (e.g., in a phased evacuation or to deliver different protective-action directions to agents in distinct hazard areas). The scenarios were explicitly selected to focus on cases that could be modeled below the county level, which was the previous level of targeting resolution for WEA and many legacy emergency alerting systems and where precise geo-targeting was most relevant and interesting.

In the model, the scenarios began in a similar way and then diverged in both content and timeline based on their specific characteristics. Figure 3.3 provides a schematic representation of the scenario timelines, which we will refer back to as we discuss each scenario in turn. After its initiation, each scenario was allowed to run for 20 simulated minutes to allow the agents to begin moving about their daily business before an alert was issued.

*Figure 3.3. Scenario Timelines*

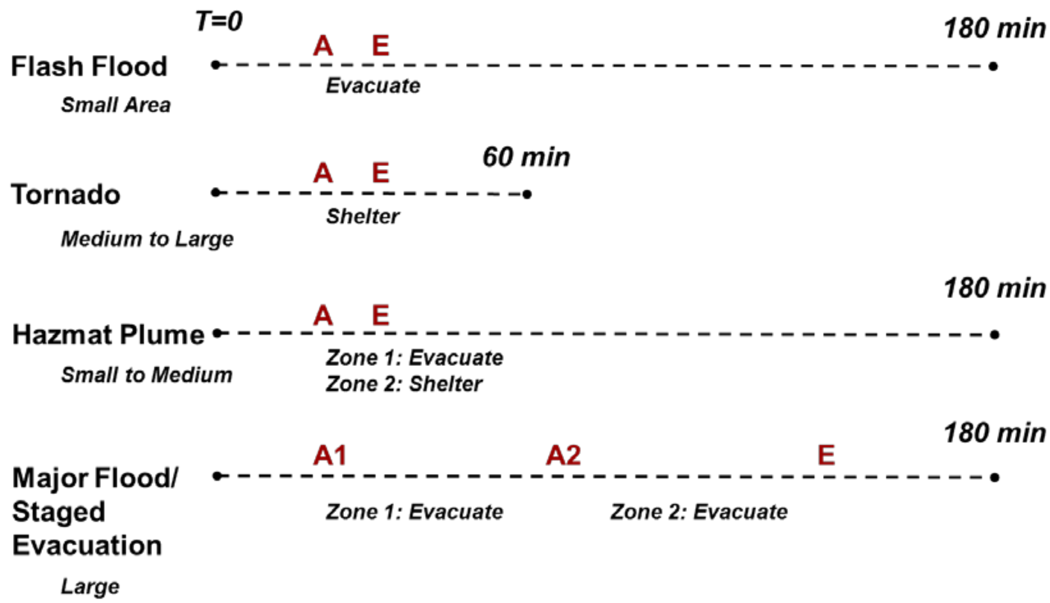


NOTE: T = time in minutes; A = time point of alert(s) in scenario; E = time point of emergency in scenario.

### 3.1.1. Flash Flood

Our flash flood scenario occurred across a major roadway in the northern half of the simulated city (Figure 3.4). The left side of Figure 3.4 shows the small square of the emergency area itself (the red rectangle), where the flood was defined to occur. The right side of Figure 3.4 shows both the emergency area and a larger WEA alerting area (the blue rectangle) for cases where the alert message was delivered beyond the edges of the emergency itself. In the flash flood scenario, the emergency alert was disseminated 20 simulated minutes (identified as *A* in Figure 3.3) into the scenario, through all the relevant alerting channels. As described, some channels covered the entire area (e.g., media), some only covered subsets of subscribed agents and WEA was geographically restricted. To assess the effects of targeting, cases were run where the alert was delivered only to the emergency area (Figure 3.4, left) and to a larger alerting area (Figure 3.4, right). In this simulation, the lead time before the emergency was short, beginning 10 minutes after initiation (*E* in Figure 3.3). The simulation was run to 180 simulated minutes, allowing the agents to continue to move and respond to the alert and subsequent forwarding among the population.

*Figure 3.4. Location of Simulated Flash Flood Emergency and Alerting Areas*

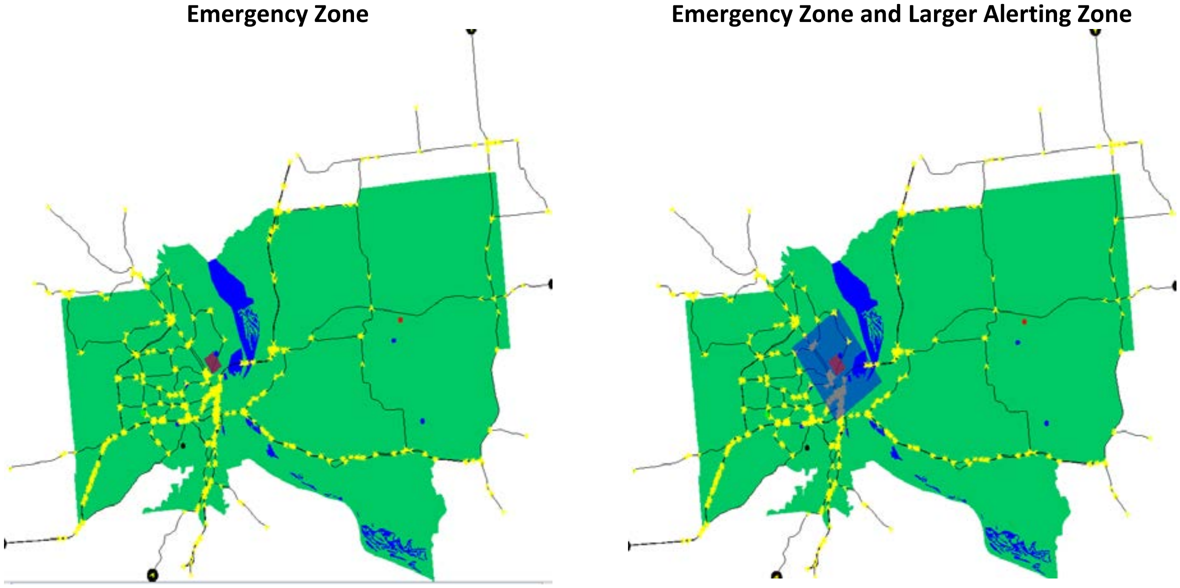


*NOTE: The left side shows the small square of the emergency area itself (red rectangle), where the flood was defined to occur. The right side shows both the emergency area (red rectangle) and a larger WEA alerting area (blue rectangle) for cases where the alert message was delivered beyond the edges of the emergency itself. The yellow arrows depict the agents. Source: Tele Atlas.*

### 3.1.2. Tornado

Our tornado scenario occurred in a small strip through the upper left of the simulated city, with larger alerting zones (the blue areas in Figure 3.5) beyond the narrow tornado track (the red line in Figure 3.5). As a simplifying assumption, the path of the tornado was presumed to be largely known and a straight line. Alerting was always larger than the emergency zone, given the inherent uncertainty that exists in tornado tracks. To examine differences in geo-targeting, we used two different alerting zones that differed in size and the fraction of the population they covered. The track of the tornado crossed several roadways, and agents would only be affected by the emergency as they traveled back and forth on those roads. The timeline of the emergency was similar to the flash flood, with the alert (*A* in Figure 3.3) at 20 simulated minutes into the simulation and the emergency occurring at the 30-minute mark (*E* in Figure 3.3). Unlike the flash flood scenario, in this case, agents were instructed to shelter in place, which, for the purpose of the simulation, meant halting in the place where they made the decision to comply. As a result, with the alerting zone extending beyond the actual tornado track, alerting would generally cause agents to never enter the emergency zone, since they would shelter somewhere in the alert area. Because of the short timeline of tornadoes, the simulation was kept to a total of 60 minutes rather than the longer runs for the other emergencies.

*Figure 3.5. Location of Simulated Tornado Emergency and Alerting Areas*

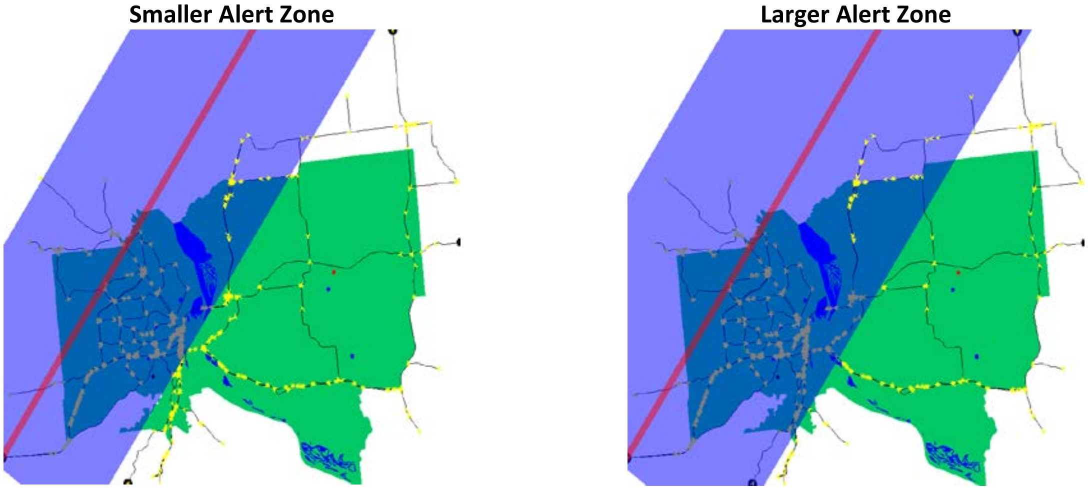


*NOTE: The blue areas represent the larger alerting zones beyond the narrow tornado track, shown by the red line. The yellow arrows depict the agents. Source: Tele Atlas.*

### 3.1.3. Hazmat

Our hazardous material scenario was approximately modeled after a chemical plant explosion described by Sorensen and Vogt (1999). It was placed in the same location as the flash flood scenario (Figures 3.4 and 3.6). When the alert was issued 20 minutes into the simulation (*A* in Figure 3.3), individuals in the small inner zone were directed to evacuate. In a larger outer zone, agents were directed to shelter in place. As in the previous scenarios, the emergency occurred at the 30-minute mark in the simulation (*E* in Figure 3.3). In this case, the delivery of two different messages created the potential to cross-contaminate directions, with agents traversing the zones or receiving forwarded messages with the incorrect direction for their positions. As a result, this scenario represented a test of one possible application of precise geo-targeting, as well as an exploration of how message diffusion could cause a breakdown in the delivery of different messages to different hazard areas.

*Figure 3.6. Location of Simulated Hazmat Inner and Outer Evacuation Zones*

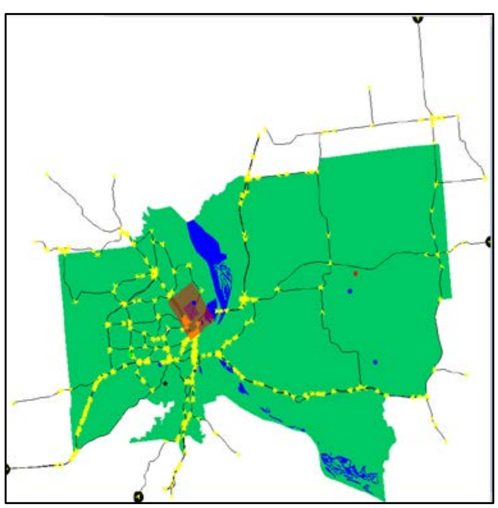

*NOTE: The small red rectangle shows the inner evacuation zone and the larger orange rectangle shows the shelter-in-place zone. Source: Tele Atlas.*

### 3.1.4. Major Flood

Our major flood scenario was a larger-scale event covering a substantial percentage of the area and population of the simulated city. Framed as a flood affecting a significant portion of the city, the directions given to agents were to evacuate in a phased way, in an effort to produce a more orderly evacuation. One group (in the outer region, the red rectangle in Figure 3.7) was directed to evacuate in an alert disseminated at minute 20 (labeled *A1* in Figure 3.3), and the second group (from the inner, blue region) was directed to evacuate in an alert sent one hour later, at minute 80 (labeled *A2* in Figure 3.3). The flood occurred 140 minutes into the simulation (*E* in Figure 3.3)

*Figure 3.7. Location of Simulated Major Flood Inner and Outer Evacuation Zones*

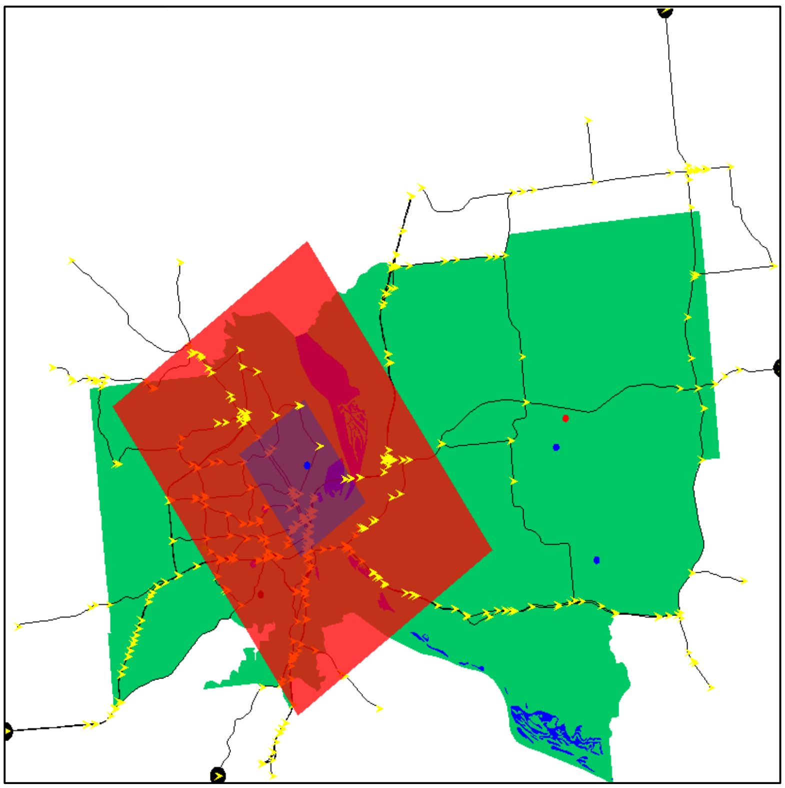

*NOTE: The group in the outer region, the red rectangle, was directed to evacuate in an alert disseminated at minute 20. The group from the inner region, the blue rectangle, was directed to evacuate in an alert sent one hour later. The yellow arrows depict the agents. Source: Tele Atlas.*

## 3.2. Measures of WEA Performance Drawn from the Model to Assess Alerting and the Effects of Geo-Targeting

Generally considering WEA performance, and more specifically the role of geo-targeting, we took into consideration four basic concepts.

First, improving geo-targeting increases the ability to deliver a precise alert to very specific groups. If the alert is transmitted outside the area where it is supposed to be — by whatever mechanism — then its precision is reduced.  One of the goals of the study was to understand this leakage. The first basic leakage mechanism, therefore, is alerting areas wider than the emergency zone. This was certainly the case when alerts could only be sent at the county level, and this should be improved with greater geo-targeting. Hence, *specificity* is reflected in the proportion of agents receiving the alert inside versus outside the emergency zone.

Within IPAWS, however, WEA does not exist in isolation, so the receipt of an alert from WEA outside the zone is not simple leakage into a population that has no other means of being alerted. A policy goal of IPAWS is to get the message out to the full population at risk, and understanding the role WEA plays in the bigger picture is important. Because of this, we explicitly modeled other alerting modes that function similar to the real world, operate on different timescales and have differential penetration. Within this larger picture, the second performance metric is *coverage*, which is reflected in the number of agents receiving a WEA message, inside and outside the emergency zone.

Our third performance metric is *alert response*, or the proportion of agents within the emergency zone following directions (in our simulations, either to evacuate or shelter in place). Getting people to safety during the period before an emergency is the main focus of alerting. The goal is to minimize the number of agents at risk during the time period when the emergency is projected to strike (e.g., still in a flood zone, in a tornado zone but not sheltered). This is a key outcome that intentionally combines multiple behaviors, pushing people out of the zone and pulling them into the zone (e.g., through convergence or transit). Hence, a key question asks how increased forwarding affects alert response.

Finally, because of geo-targeting, the very fact that an individual receives an alert provides its recipient with evidence of its relevance organically. Forwarded alerts from WEA, when they leak outside the geographic area where they are supposed to go, become divorced from this contextual information. A recipient of a forwarded alert outside the emergency zone, therefore, has a chance of acting in error. This might have consequences for that individual (inconvenience) or for others (e.g., unnecessary evacuation by people who then clog the roads). Hence, we are also concerned with *unnecessary alert response*, or the proportion of agents outside the emergency zone following directions to evacuate or shelter.

## 3.3. Analytic Strategy

Chapters Four through Six present the main substantive results from the study. In these chapters, we take a policy-driven approach, with each chapter addressing a specific policy-research question. As some results are more relevant to these questions, and relevant results cut across scenarios and specific analyses, we organized the presentation to best discuss each question rather than scenario by scenario. Scenario-by-scenario results are presented in Appendix C, for those who are interested.

Notably, the analysis that follows emphasized alert response and unnecessary alert response, as these represent the bottom line for our research questions. We went about operationalizing these outcomes in two ways. First, for evacuation-based scenarios, we traced the number of agents in the emergency zone over the course of the simulation. If WEA and other alert channels are effective, then alerting and forwarding should reduce the number of agents in the zone. For sheltering scenarios, agents stopped moving, so this metric was not applicable. We also measured response by assessing agent goals. Goals to evacuate or shelter were sticky in our model — once adopted, the goals remained with the agent for the rest of the simulation. Hence, we assessed how many agents adopted goals consistent with the scenario (either evacuation or shelter) between the time that alerts were first transmitted and the projected time of the emergency.

Because WEA was a well-defined process within our model, specificity and coverage (our first two measures) effectively acted as checks on whether WEA was implemented correctly. Since this is of secondary interest, we present these results, along with other scenario-specific results, in Appendix C.

# 4. How Do Geo-Targeting and Alert Diffusion Affect Alerting Effectiveness?

In principle, WEA messages should reach those who need to be alerted, but not those who do not need to be. Geo-targeting has the potential to serve this goal by reducing over-alerting, and the diffusion of a geo-targeted alert has the potential to water down such gains. Many prominent arguments for geo-targeting emphasize the benefits of limiting over-alerting, including the reduction in both the annoyance caused by unnecessary alerts and subsequent opt-out behavior. There is also concern that over-alerting, especially with short WEA messages that convey limited information, could cause unnecessary compliance, such as evacuation or sheltering in place, by those outside an emergency zone.

This chapter presents results bearing on alerting effectiveness and how that effectiveness is affected by geo-targeting and message diffusion. We start with a set of simulation results using realistic, first-approximation parameters, drawing from the literature where possible (and described in detail in Chapter Two). Our initial focus is on the flash flood scenario, which is both very small geographically and very quick moving. We find limited effectiveness of the WEA messages, especially within a small, quick event. Increasing the size of the alerting zone to extend beyond the emergency zone and the forwarding of alerts by citizens (two key diffusion mechanisms) had very little effect on this result. As a result, the location and context (e.g., timing) of an event are critical to understanding the ideal targets of an alert.

We then provide a contrasting set of flash flood cases, where we reduced a key set of parameters: delay of alert receipt and delay of alert compliance. Effectively, this increased the churn of forwarding by allowing more for iterations. With greatly reduced delays, we see marked increases in message effectiveness within the emergency zone.

To better understand the generality of these results, we then contrast the flash flood case to a larger and longer event — a major flood. We find that, with the increase in the number of agents in the emergency zone (and the number alerted via WEA) and the more protracted event, WEA forwarding has a distinct impact on evacuation behavior, even with the longer delay parameters.

We conclude with a discussion of how forwarding can act as an alerting magnifier, amplifying the number of alerts received and increasing compliance over what would otherwise be much lower baseline levels.

## 4.1. Baseline Dynamics Within the Flash Flood Scenario

Figure 4.1 provides an exploded view of the emergency and alert zones for the flash flood scenario.

*Figure 4.1. Flash Flood Scenario, with Emergency Zone (red) and Larger Alert Zone (blue)*

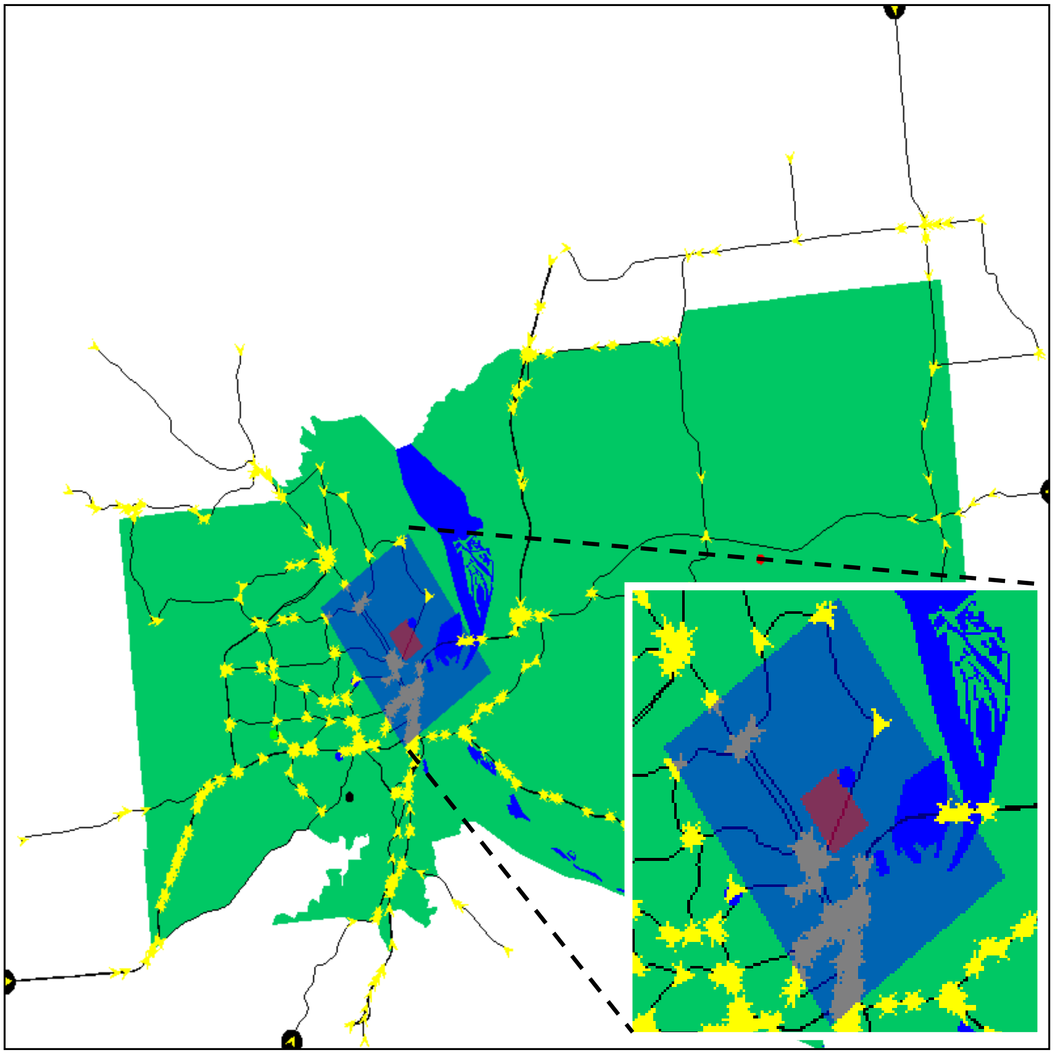

*NOTE: The yellow arrows depict the agents. Source: Tele Atlas.*

In interpreting the following results, it is important to consider not just the size of the zone but its location and context. As can be seen in the inset for Figure 4.1, the simulated flash flood occurred in a flood plain along a thoroughfare street. As described in Chapter Two, the agents "lived" at intersections, and this is where they started their days. Since there were no intersections within the emergency zone, the choice of zone location is analogous to affecting a freeway away from inhabited areas, and agents will be traversing this zone in the course of their day. To visualize normal-day commuting patterns, Figure 4.2 presents a baseline case in which no alert was issued. Displayed is the number of agents in the potential emergency zone, over time (measured as minutes since the start of the simulation). Before alerts were issued, we let the simulation run for 20 minutes, which allowed agents some time to start moving away from their home intersections and around the road network. The figure is scaled to start (on the left-hand side) at this 20-minute point.

*Figure 4.2. Number of Agents in the Flash Flood Emergency Zone over Time with No Alert*

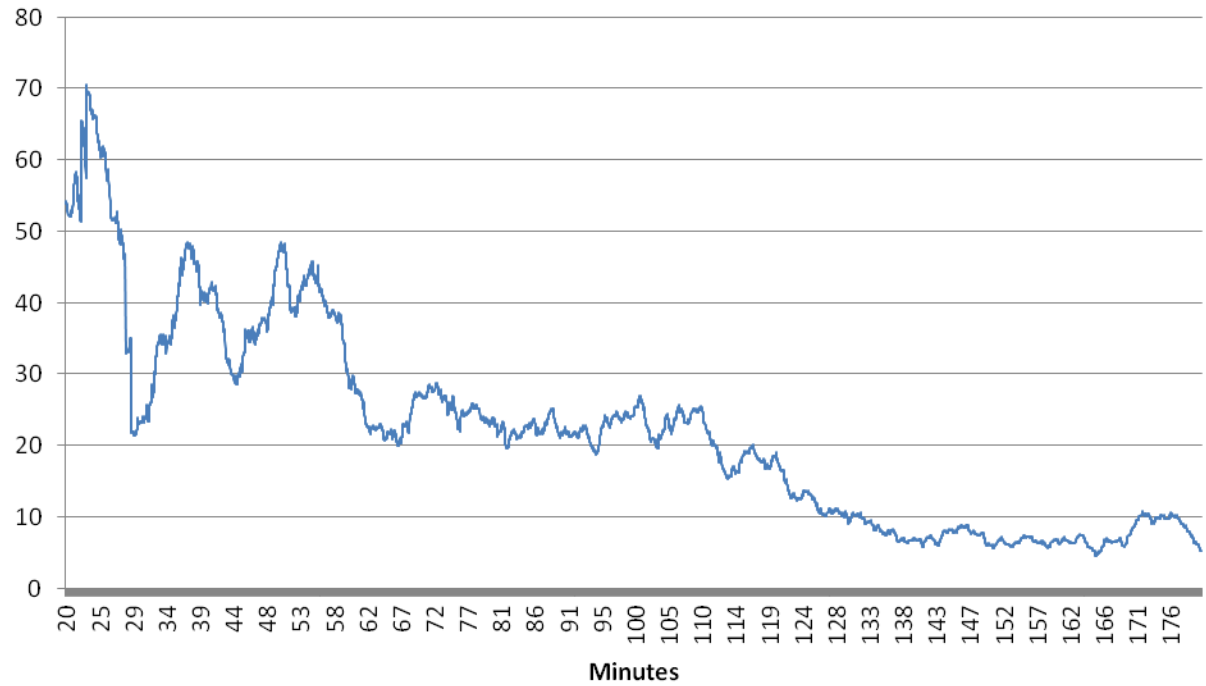


As noted above and as can be seen in Figure 4.1, there were no intersections in the emergency zone, and no agents started in the zone (at time zero, which not shown). As the simulation progressed, agents started traversing the zone as part of their daily activities. The trace presented in Figure 4.2 shows a natural ebb and flow of agents in the zone, with an initial peak, a trough, a few more minor peaks and an eventual tailing off. It is this trace against which we compare our alerting cases within the flash flood scenario.

## 4.2. Flash Flood Simulations with Realistic Delay Parameters

Figure 4.3 overlays onto this no-alert case four different emergency alert cases, including an alert with no WEA forwarding, an alert with normal WEA forwarding (reflecting our best-guess parameters, based on the literature), and cases with double and triple WEA forwarding. The upper panel of Figure 4.3 represents cases where 30 percent of agents had WEA-compatible devices, and the lower panel represents a future world where 70 percent of devices are WEA compatible. Both presume an alert zone that extends beyond the emergency zone, as this presents the greater opportunity to observe a WEA effect (results for simulations with the smaller alert zone, where alert and emergency zones correspond, are presented with other supplementary analyses in Appendix C).

*Figure 4.3. Number of Agents in the Flash Flood Emergency Zone over Time, by WEA Compatibility and WEA Forwarding Rate, in Alert Zone Larger Than Emergency Zone*

**30 Percent Phone Compatibility**

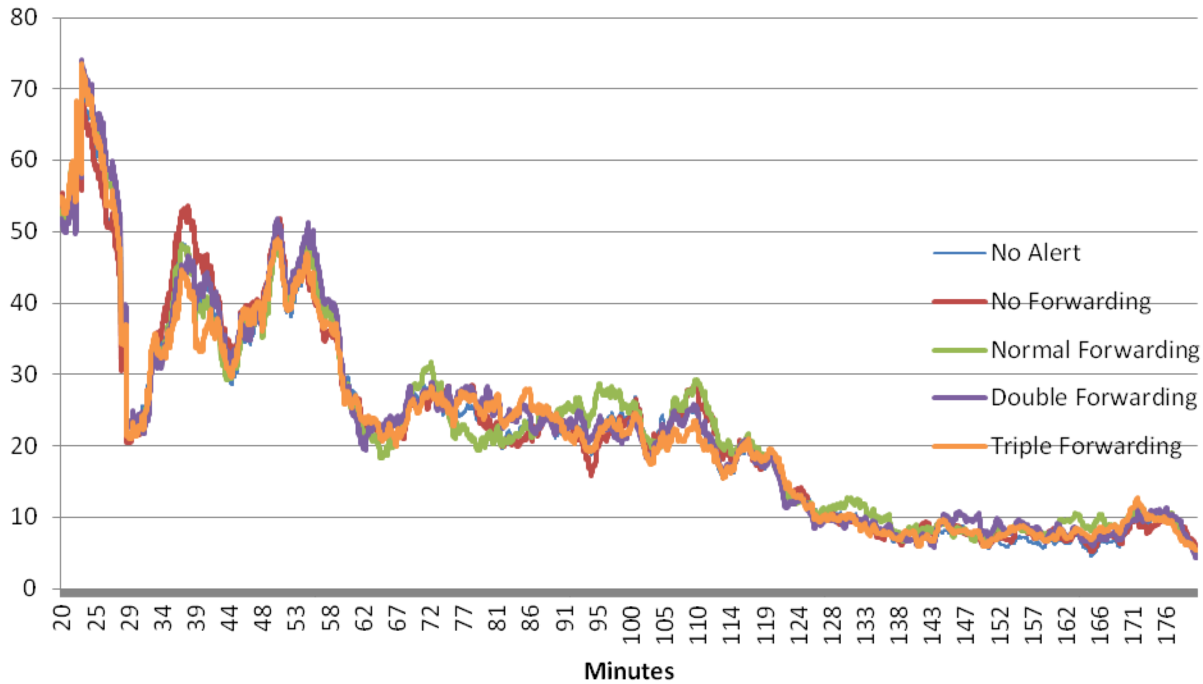


**70 Percent Phone Compatibility**

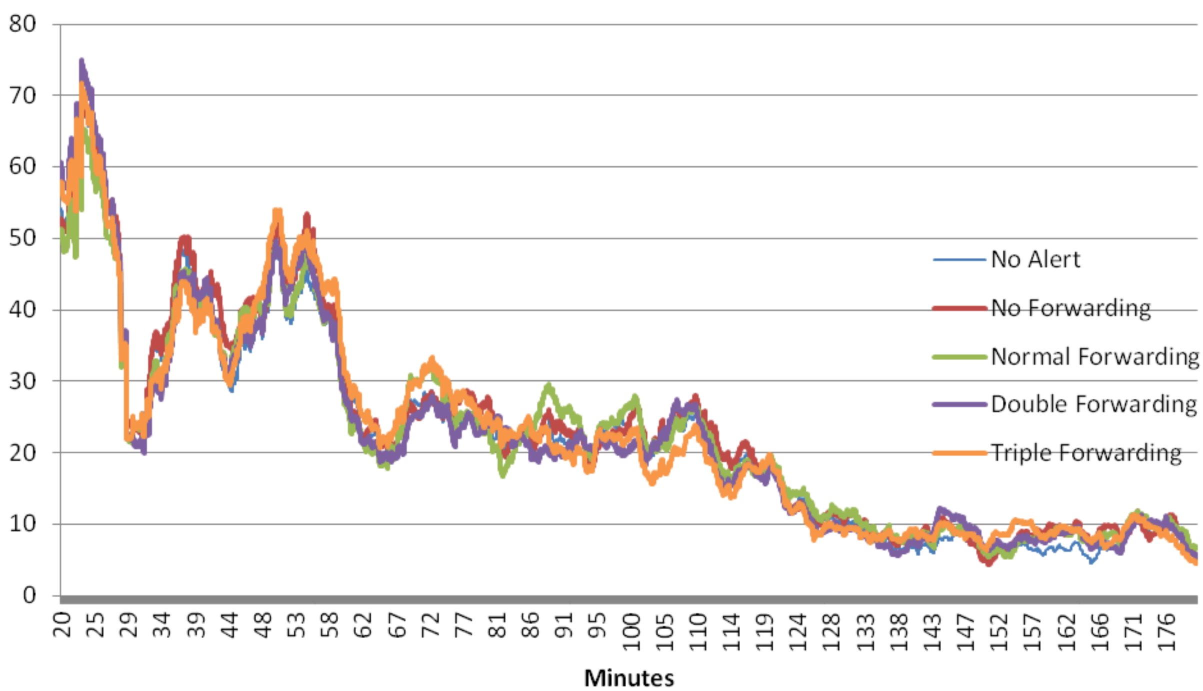

If WEA or WEA forwarding is to have a measurable effect, we would expect to see the curves for the alerting cases to be below those of the no-alert case, indicating that fewer people were in the emergency zone as a result of WEA or WEA forwarding. This is noticeably not the case in Figure 4.3, with each of the curves essentially falling on the no-alert line (with some noise). In other words, we see no noticeable evacuation behavior.

## 4.3. The Effect of Modestly Reducing Delays to Receipt and Compliance

Based on these results, we chose to explore our delay parameters (both for the receipt of alerts and for compliance with alerts) as a fulcrum of exploring the role of forwarding. Decreasing delays will increase the opportunity for and the number of iterations of alert forwarding. Consequently, we see this as a forwarding-volume manipulation, rather than as simply one specifically about delays.

Figure 4.4 presents graphs analogous to those in Figure 4.3, again focusing on cases where the alerting zone is larger than the emergency zone, but where we removed delays in compliance and reduced delays in receipt. Specifically, for the latter, we reduced WEA receipt to no delay, reduced delays in interpersonal forwarding receipt to two to three minutes, and kept delays at their default for media and technical forwarding.

As can be seen in Figure 4.4, WEA and WEA forwarding had only minimal effects, at best. In the time between alert (20 minutes) and the projected time of the emergency (30 minutes), the curves are indistinguishable from each other and from the baseline no-alert case. Only in the 70-percent compatibility cases, and only after at least an hour had elapsed, do we see any divergence of the curves. Here, greater forwarding results in fewer agents in the emergency zone, as would be expected, but this effect is quite modest and occurs well after the time the simulated flash flood began.[16] Still, this effect suggests that, given enough time, WEA and WEA forwarding might have an effect. Increasing the rate at which forwarding occurred might increase the effectiveness of the WEA message.

[16] Put another way, there would be an observable effect for an event where a warning could be issued a longer time before the emergency.

*Figure 4.4. Number of Agents in the Flash Flood Emergency Zone over Time, After Reduced Delay in Compliance or WEA Receipt, in Alert Zone Larger Than Emergency Zone*

**30 Percent Phone Compatibility**

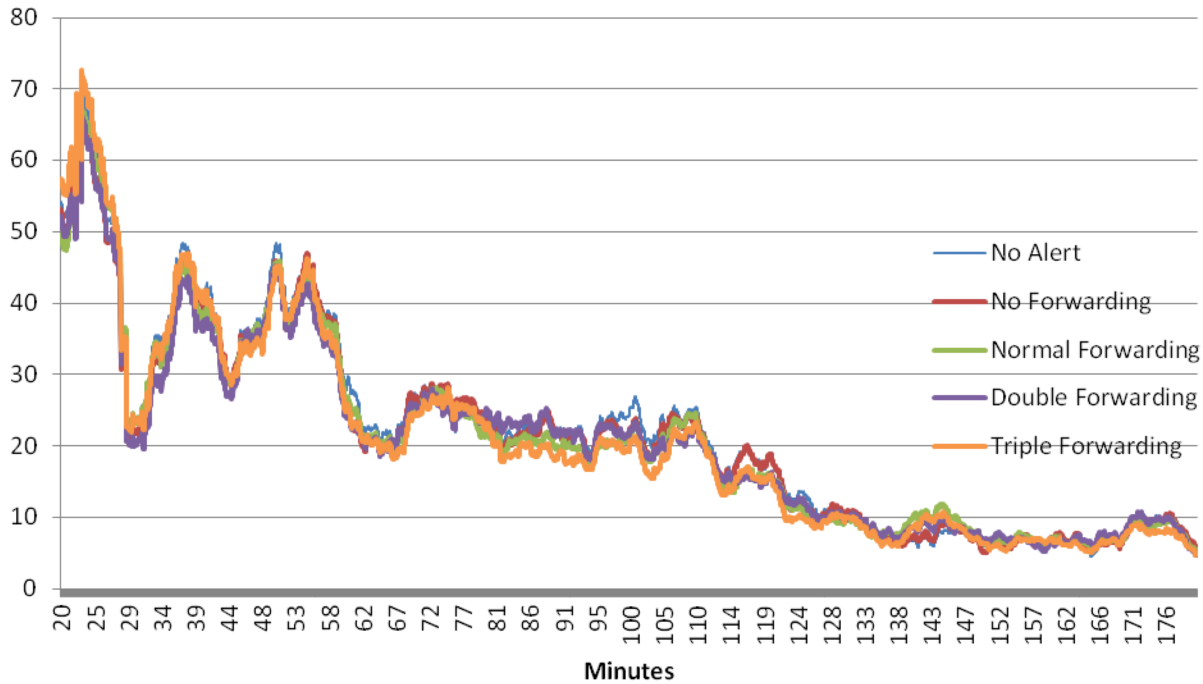


**70 Percent Phone Compatibility**

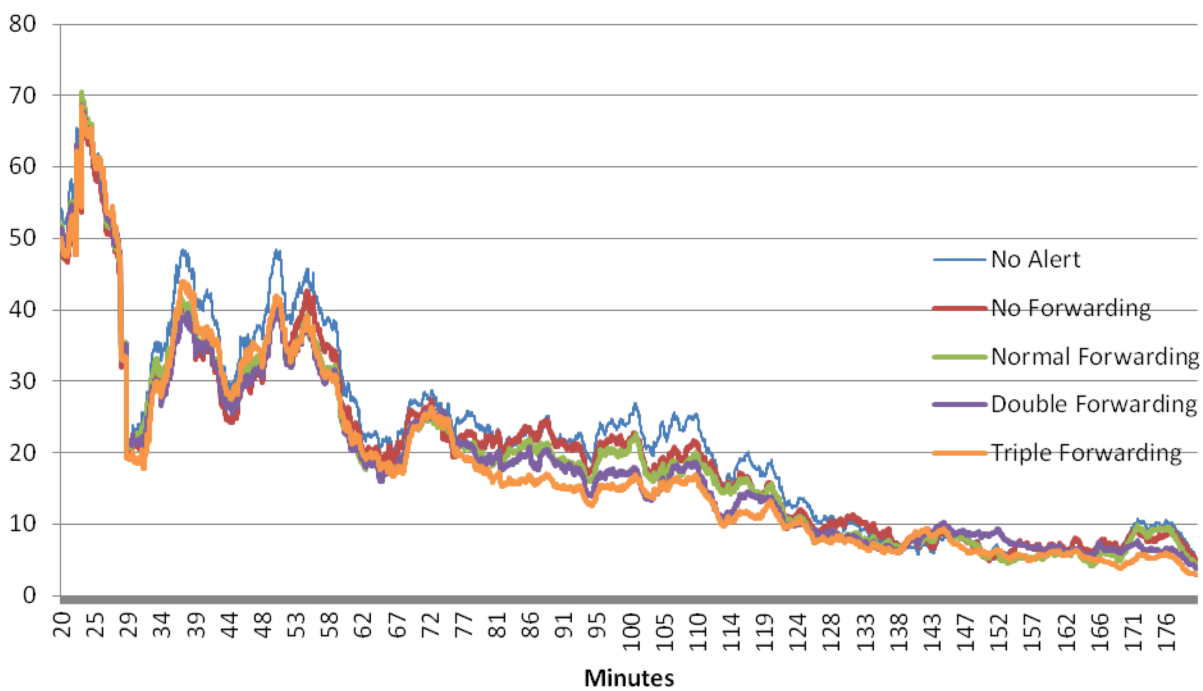

## 4.4. The Effect of Greatly Reducing Delays to Receipt and Compliance

The data in Figure 4.4 suggested the promise of forwarding as a means of increasing compliance. To take this logic to its limit, we next considered an extreme case. Figure 4.5 also presents graphs analogous to those in Figure 4.3, but we removed all delays in compliance and receipt. It is important to note that this is an intentionally unrealistic set of cases, but one that provides a useful contrast to the results presented earlier.

When delays in receipt and compliance were dramatically reduced, as in Figure 4.4, there was a remarkable clearing of the emergency zone, as can be seen in all the lines with alerting markedly below the no-alert case. With both alerting and WEA forwarding, the lines go to near zero by about the 35-minute mark. In the case with alerting but no WEA forwarding, some agents did stay in the emergency zone, but that figure was lower than in the no-alert case. This latter result is instructive; as it suggests a strong role of forwarding in increasing the in-zone effectiveness of a WEA message (recall that the manipulation of forwarding depicted in Figure 4.4 was *only* on the forwarding of WEA messages, without altering the forwarding of other types of alerts).

*Figure 4.5. Number of Agents in the Flash Flood Emergency Zone over Time, with No Delay in Compliance or WEA Receipt, in Alert Zone Larger Than Emergency Zone*

**30 Percent Phone Compatibility**

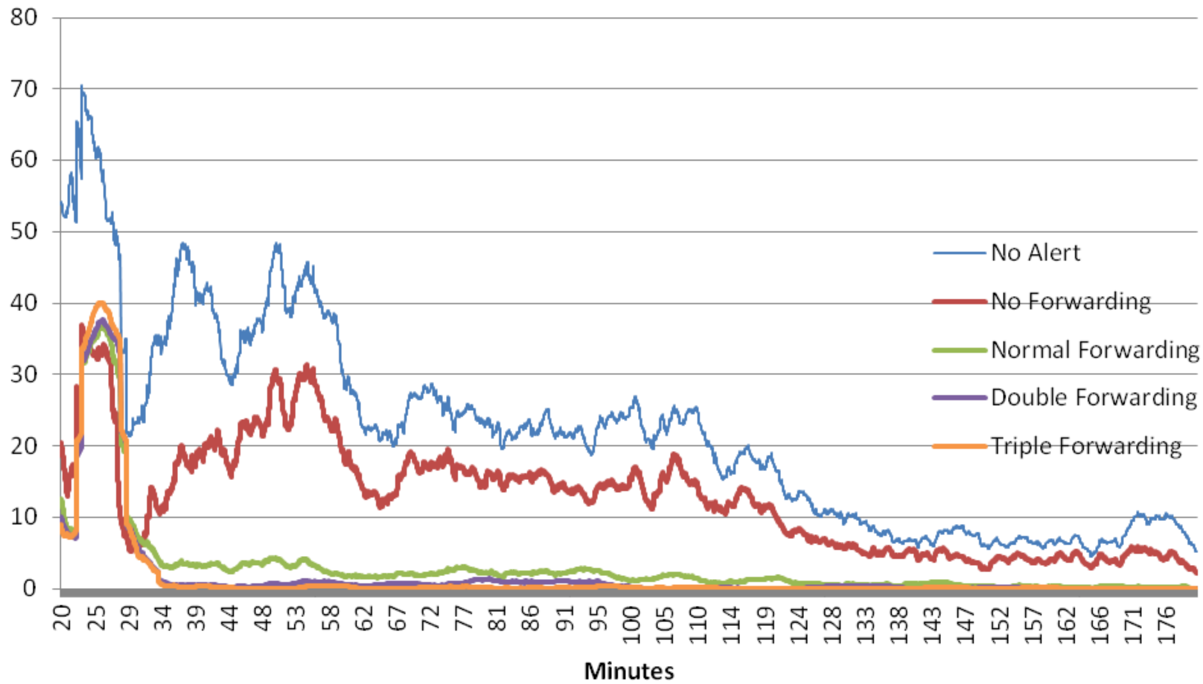


**70 Percent Phone Compatibility**

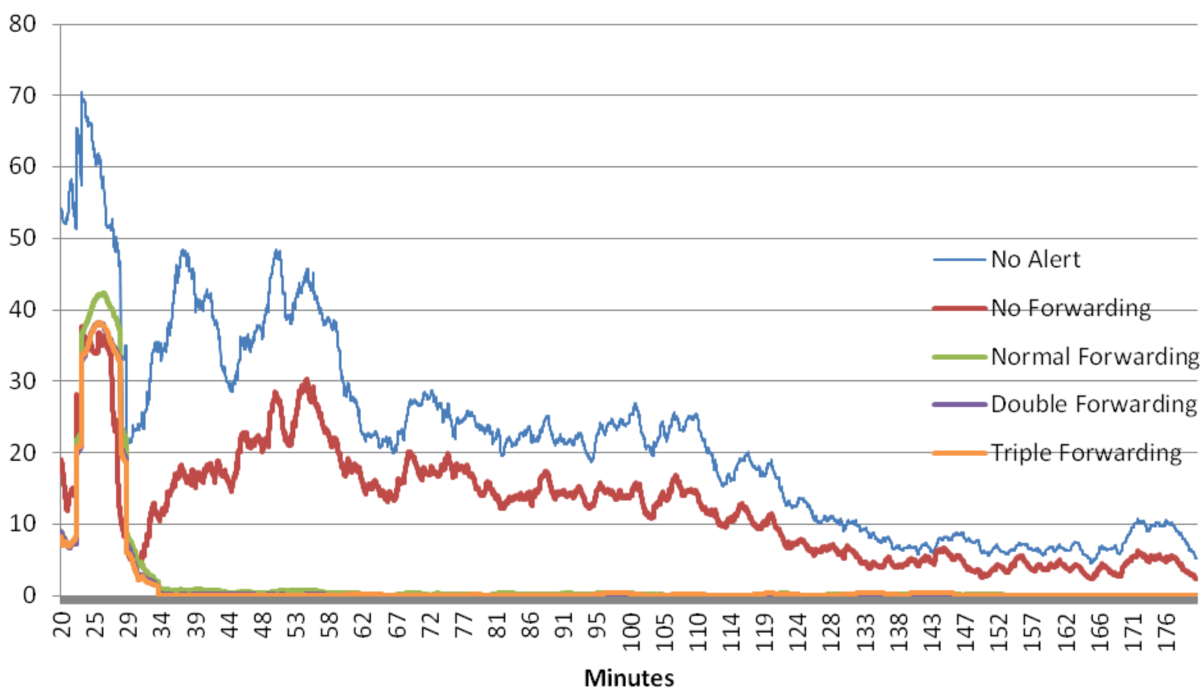

Figure 4.6 shows the same reduced-delay cases, but with an alerting zone that exactly corresponds to the emergency zone (i.e., perfect geo-targeting). The effects here are still pronounced but less extreme than those in Figure 4.5, showing how these parameters combined to increase or decrease the effectiveness of the WEA message. When device compatibility is near current levels (i.e., 30 percent), alerting with reduced delays does decrease the number of agents in the emergency zone, but the effect of forwarding is not noticeable until after the 60-minute mark. With increased device compatibility, these forwarding effects happen earlier and more strongly. Because of the smaller footprint of the alerting zone, fewer agents were in or entered the zone and received the WEA message, representing fewer potential “seeds” to initiate WEA message forwarding.

It is worth noting that the no-forwarding cases (the red lines) in Figures 4.5 and 4.6, which *only* differ in the size of the WEA alerting zone, are virtually identical. This suggests that the initial WEA message was not having much of an effect. Instead, it appears as if the main in-zone effect in this scenario was due to forwarding of WEA messages.

*Figure 4.6. Number of Agents in the Flash Flood Emergency Zone over Time, with No Delay in Compliance or WEA Receipt, in Alert Zone Corresponding to Emergency Zone*

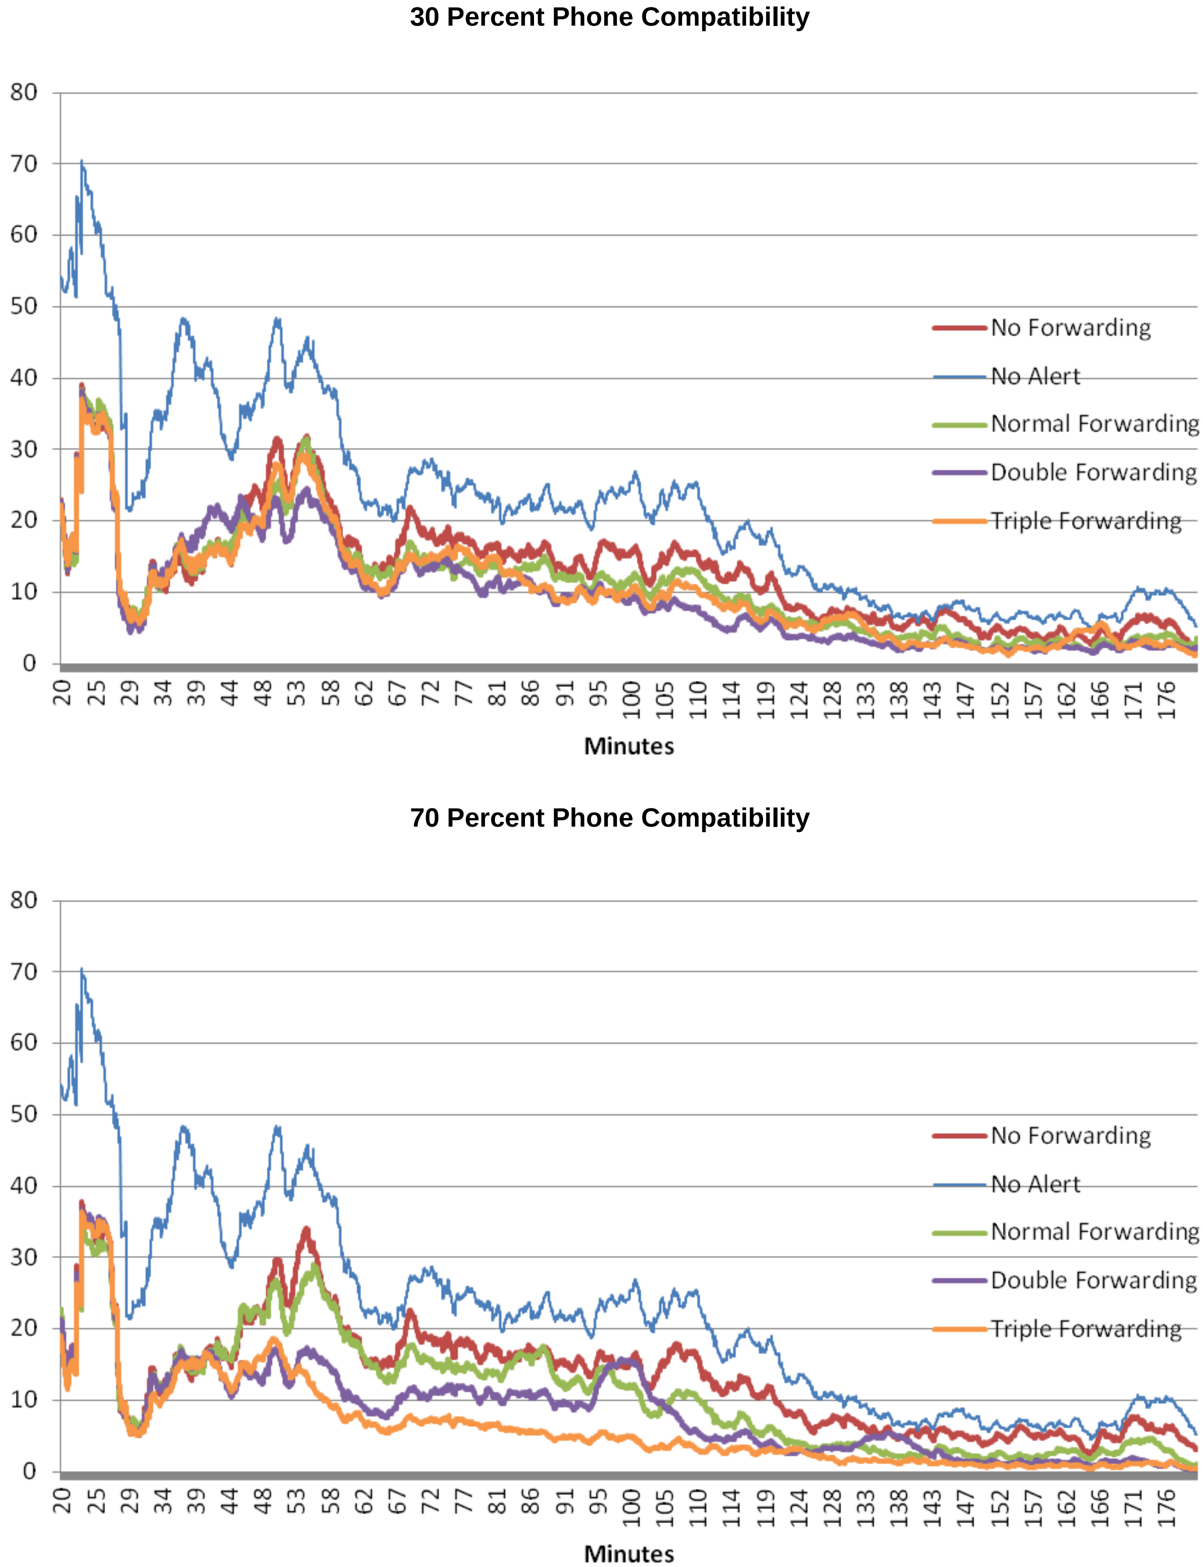

## 4.5. The Effect of Greater In-Zone Population and a Longer Event

The major flood scenario, depicted in Figures 4.7 and 4.8, presents a useful comparison to the flash flood. Whereas a flash flood is very small and very fast, a major flood is, relatively speaking, very large and very slow. This means a much larger starting population of potential WEA recipients, but also the chance for more forwarding and many more forwarding iterations.

Figure 4.7 displays the number of agents in the combined, inner and outer major flood emergency zones, returning to the more realistic delay parameters. There is small but distinct differentiation across cases, such that more WEA forwarding led to fewer agents in the combined emergency zones. Most of this differentiation is evident in the inner zone, which may reflect the fact that those in the inner zone needed to traverse the outer zone as they evacuated (something that we will return to in Chapter Six). Figure 4.8 displays these data, but it uses the reduced (but not eliminated) delay parameters. Here the differentiation is more substantial. What this means is that, with enough people and time, WEA forwarding makes up for an initially weak response, even in the presence of delays in the system.

*Figure 4.7. Number of Agents in the Major Flood Emergency Zone over Time, Default Delays*

**30 Percent Phone Compatibility**

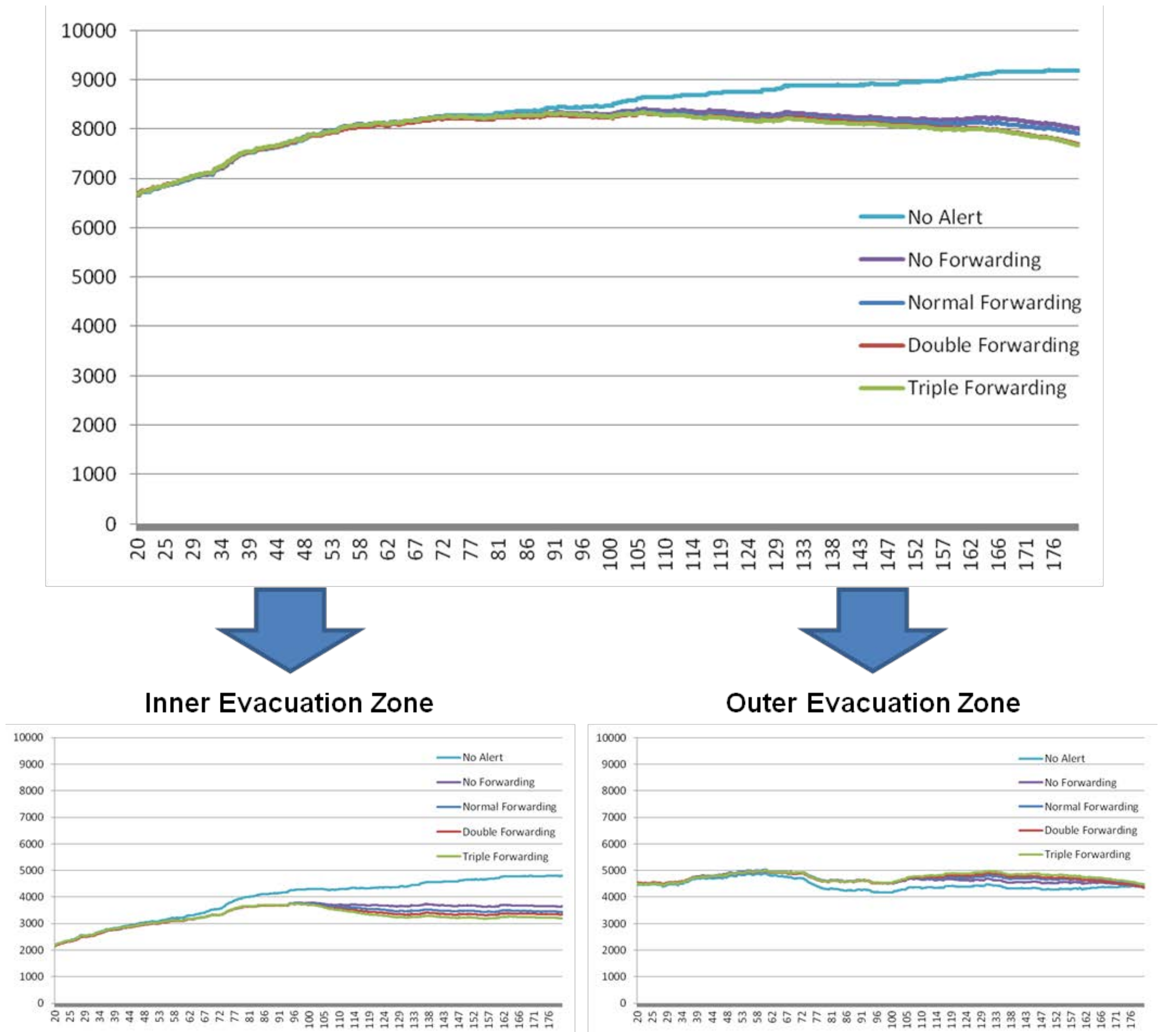

**70 Percent Phone Compatibility**

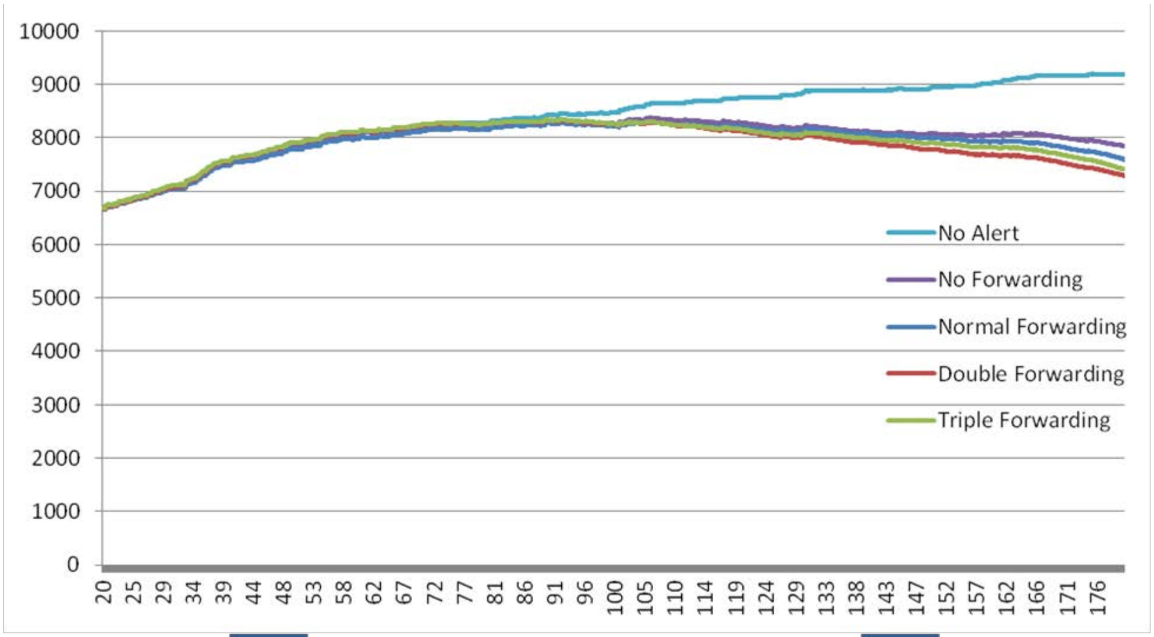


**Inner Evacuation Zone**

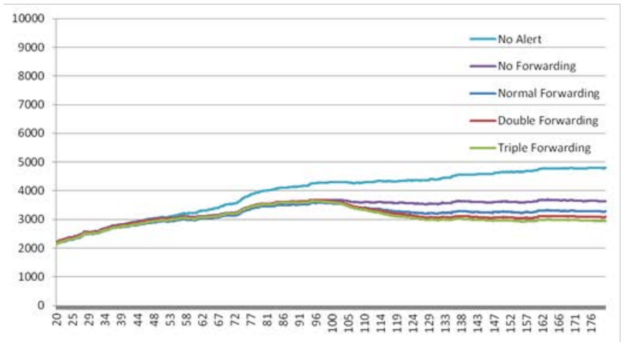


**Outer Evacuation Zone**

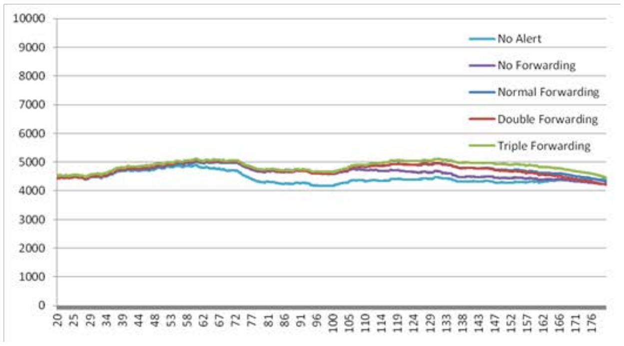

*Figure 4.8. Number of Agents in the Major Flood Emergency Zone over Time, Reduced Delays*

**30 Percent Phone Compatibility**

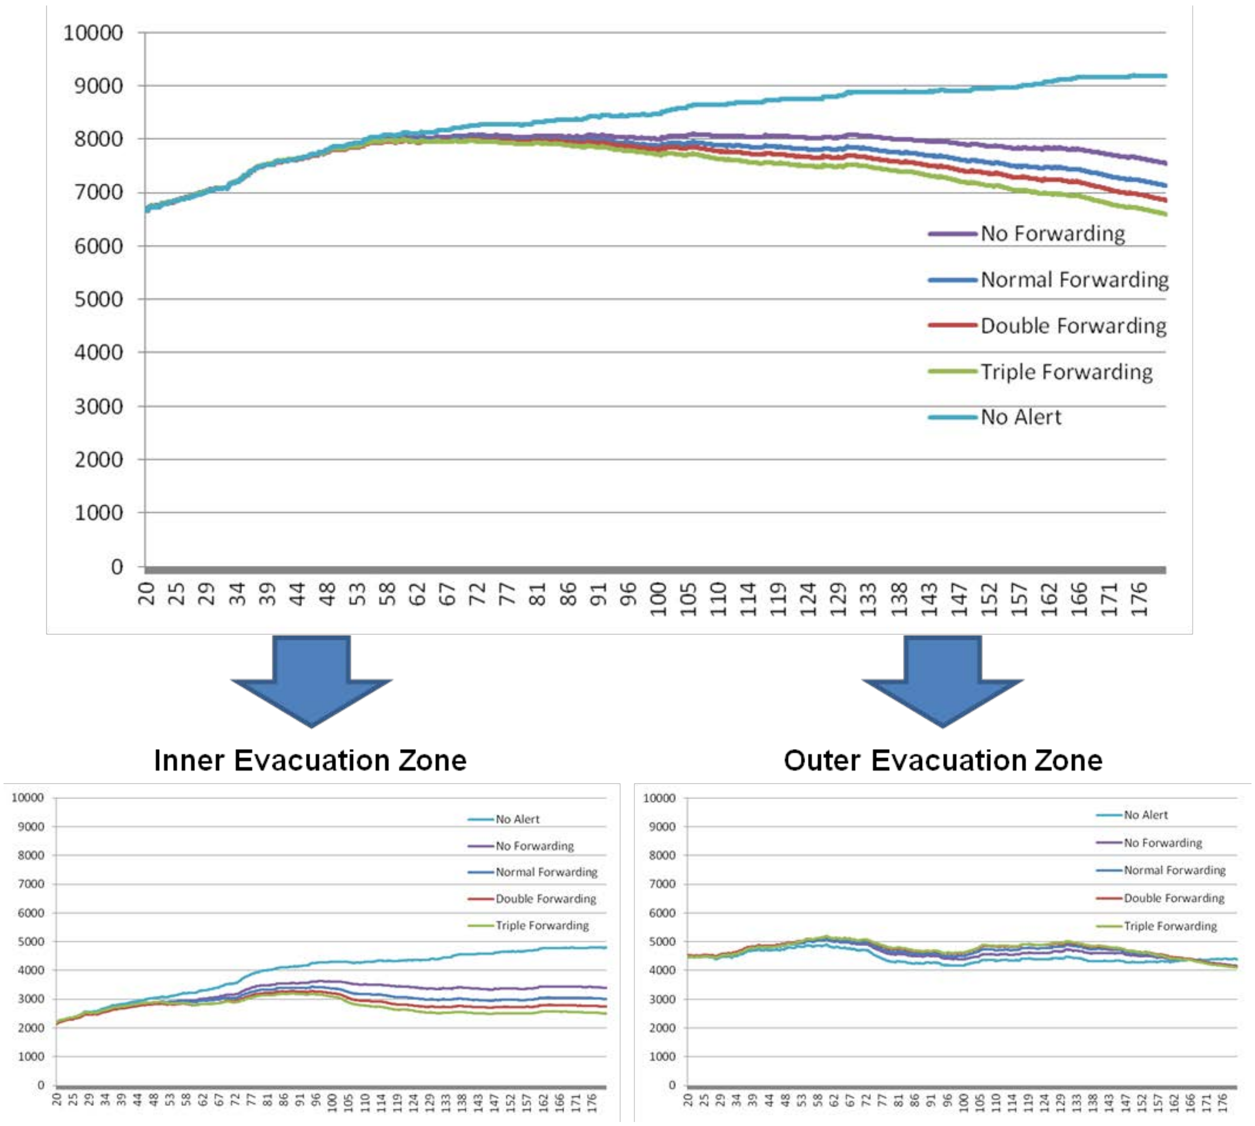

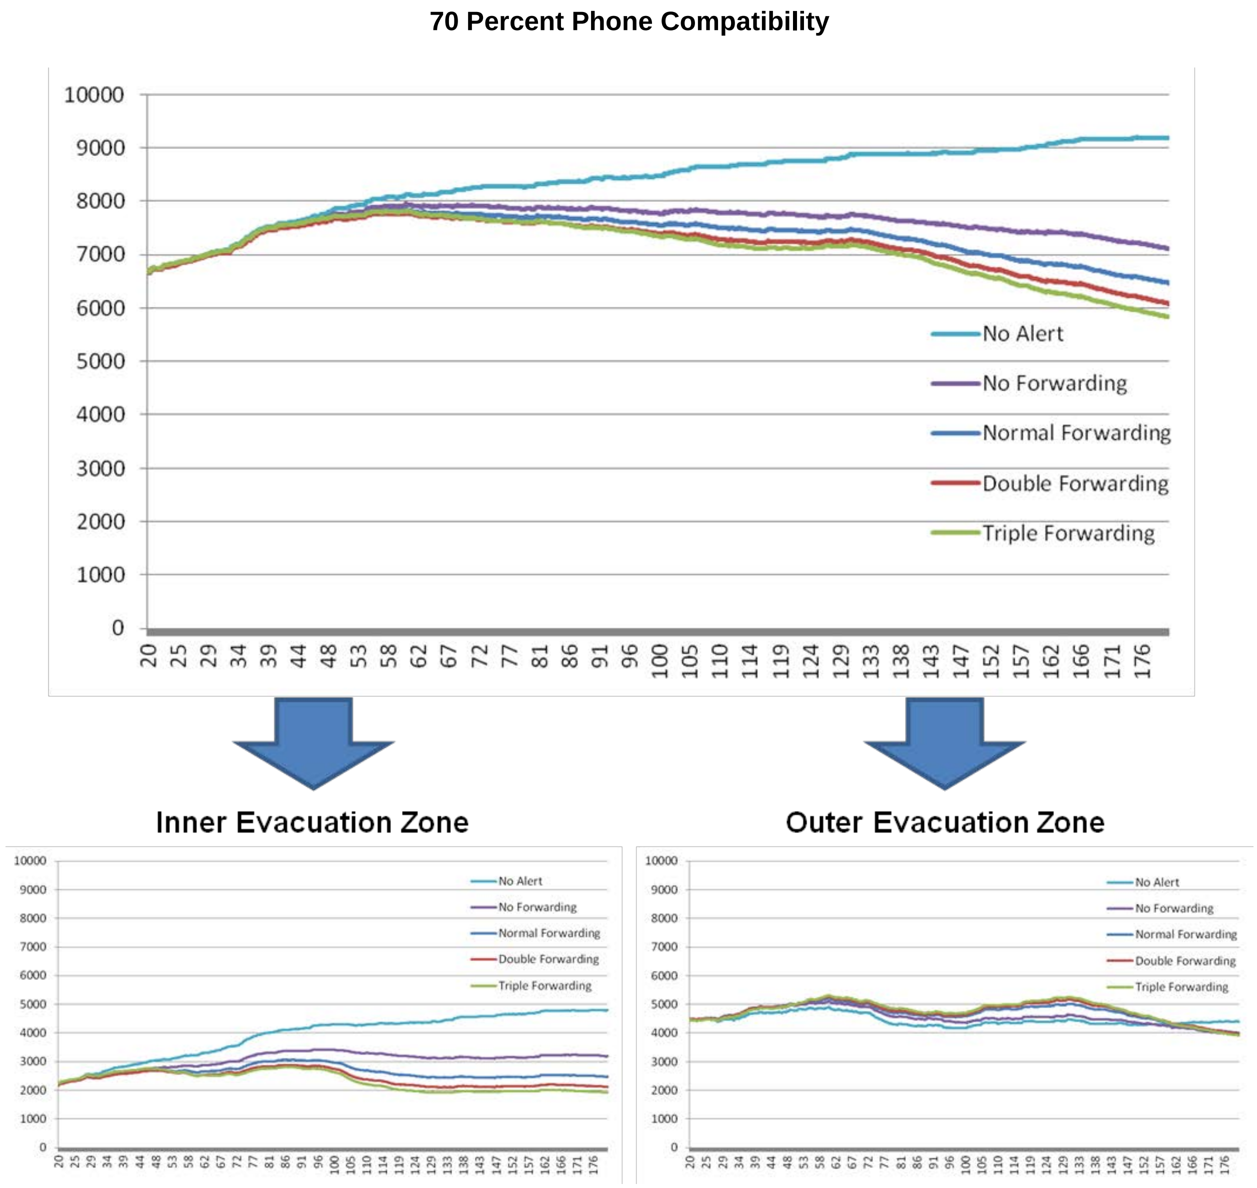


# 4.6. Forwarding as an Amplifier of Alert Receipt and Compliance

To better understand these results, it is instructive to return to the string of parameters that influence (and in many ways limit) agent response to WEA. The model and the parameters that came from the literature clearly demonstrate one of the challenges for alerting mechanisms. The multistep process of people receiving, understanding and acting upon alerts (introduced in Chapter Two) creates multiple opportunities for the system to fail, which can greatly reduce the effectiveness of alerting.

Figure 4.9 summarizes how the parameters in our model (for simplicity, we focus on just the point estimates and midpoints of ranges described in Chapter Two) conspired to reduce WEA effectiveness. The upper blue box represents 100 percent of possible WEA recipients, including those who were starting out or entering the alert area during the time between the simulated alert and the time when the event was projected to occur. Of these, only 30 percent had a compatible device, and 99 percent had not opted out. Of those with a compatible device who had not opted out, 60 percent received the alert, and of those, 81 percent understood it (the product of low- and high-understanding groups with different rates). Among those who understood the message, only 20 percent complied with it, resulting in less than 3 percent of the target audience actually complying with the alert and attempting to evacuate.

*Figure 4.9. WEA Parameters Limit Effectiveness, but Forwarding Can Make Up for It*

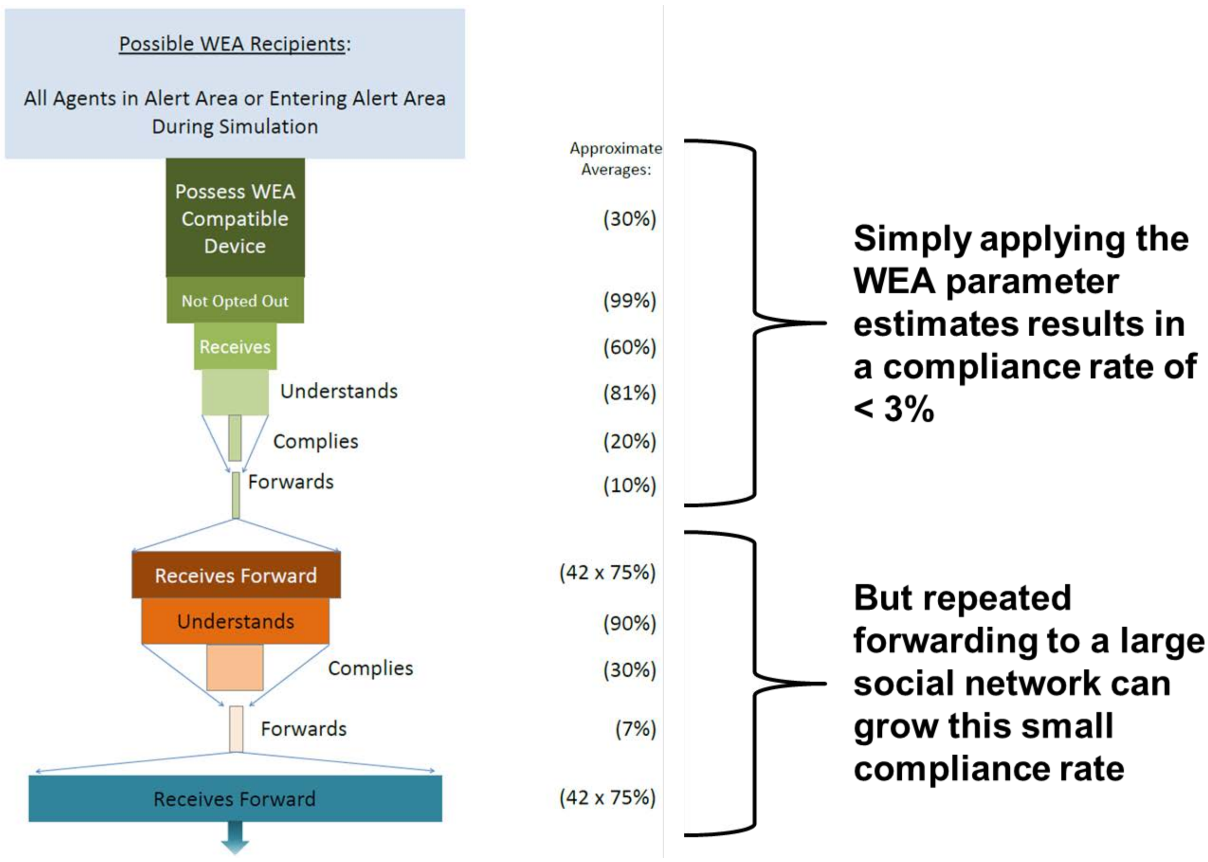


With no forwarding, this is the end of the story, especially with a small, quick event that leaves very few agents trying to evacuate. Assuming a 10 percent forwarding rate (among those who receive and understand the alert), with each agent forwarding the alert to 42 people in his or her social network, there is the potential to grow this pool. Assuming that people are more likely to receive, understand and comply with a message that comes from a friend or family member, the number complying after the first round of forwarding will be larger than the number complying with the initial WEA message. Because this first generation of forward recipients can subsequently forward, this pool of recipients (compliers) will continue to grow across generations of forwarding.[17]

[17] This path is similar to a logic described in Sutton et al. (2014) (p. 4 and references therein that have been removed), based on analysis on Twitter use: “Not only are new audiences exposed to the tweet when others

The analysis in Figure 4.9 is a simplified representation of our model. The figure does not consider heterogeneity in agent characteristics (which were built into the model). It also does not account for the fact that, across generations of forwarding, some growing portion of the agents receiving the alert would already be complying (and would not be eligible to convert to compliance). That said, the overarching point regarding the compounding of the parameters in the model highlights the real-world challenges faced by WEA. The logic in Figure 4.9 also strongly suggests considering social forwarding as its own mass-alerting channel, or at a minimum, an amplifier for existing channels.

To explore this behavior further, we conducted an experiment to step through different amounts of forwarding traffic and its effect on behavior. In the structure of the model, three main sets of relevant parameters could be used to do so: (1) the number of people to which each agent is connected, and therefore to whom he or she could forward; (2) the probability that individuals will forward a received message; and (3) what a recipient of the forwarded message will do upon receipt, including whether he or she will forward it again to someone else. Social network density has a strong basis in the literature, and therefore provides a relatively well-defined way to modulate the overall amount of forwarding behavior. This can be thought of as the subset of each person's social network, where a forwarded message has any probability of producing a change in behavior; this sort of influence might be strongest for those socially closest to the sender.

To investigate this, we conducted a small-scale experiment across the three levels of social network described in Chapter Two. A small social network (where each agent is linked to a median of six other agents) emphasizes relatively closer connections. A medium-sized social network (where each agent is linked to a median of 42 other agents) expands the influence beyond the family and close networks, and suggests a moderate amount of influence on a more extended social network. A large social network (where each agent is linked to a median of 216 other agents) is a limiting case, where influence is extended through all of a person's contacts.

Figure 4.10 presents results from this small-scale experiment within the flash flood scenario. We focus on the case with 70 percent phone compatibility because it reveals the greatest contrast among the various social network sizes even with the realistic delay parameters. Rather than looking at how many agents were in the zone (see Figure 4.6), here we focus on the agents who were evacuating. Each of the graphs in Figure 4.10 displays the cases with the smaller alert zone (exactly corresponding to the emergency zone) on the left and the cases where the alert zone is larger than the emergency zone on the right. The bars also break out by the probability of forwarding a WEA message and whether the agent was ever inside the emergency zone.

retransmit the message, but some users may be re-exposed to the tweet. Multiple exposures to messages have been linked to more confidence in [their] veracity . . . , which can lead to further sharing. . . . Indeed, repeated exposures from multiple network ties are often a prerequisite for the spread of information through networks."

*Figure 4.10. Percentage Evacuating in Flash Flood Scenario, by Social Network Size, Alerting Zone Size and Location*

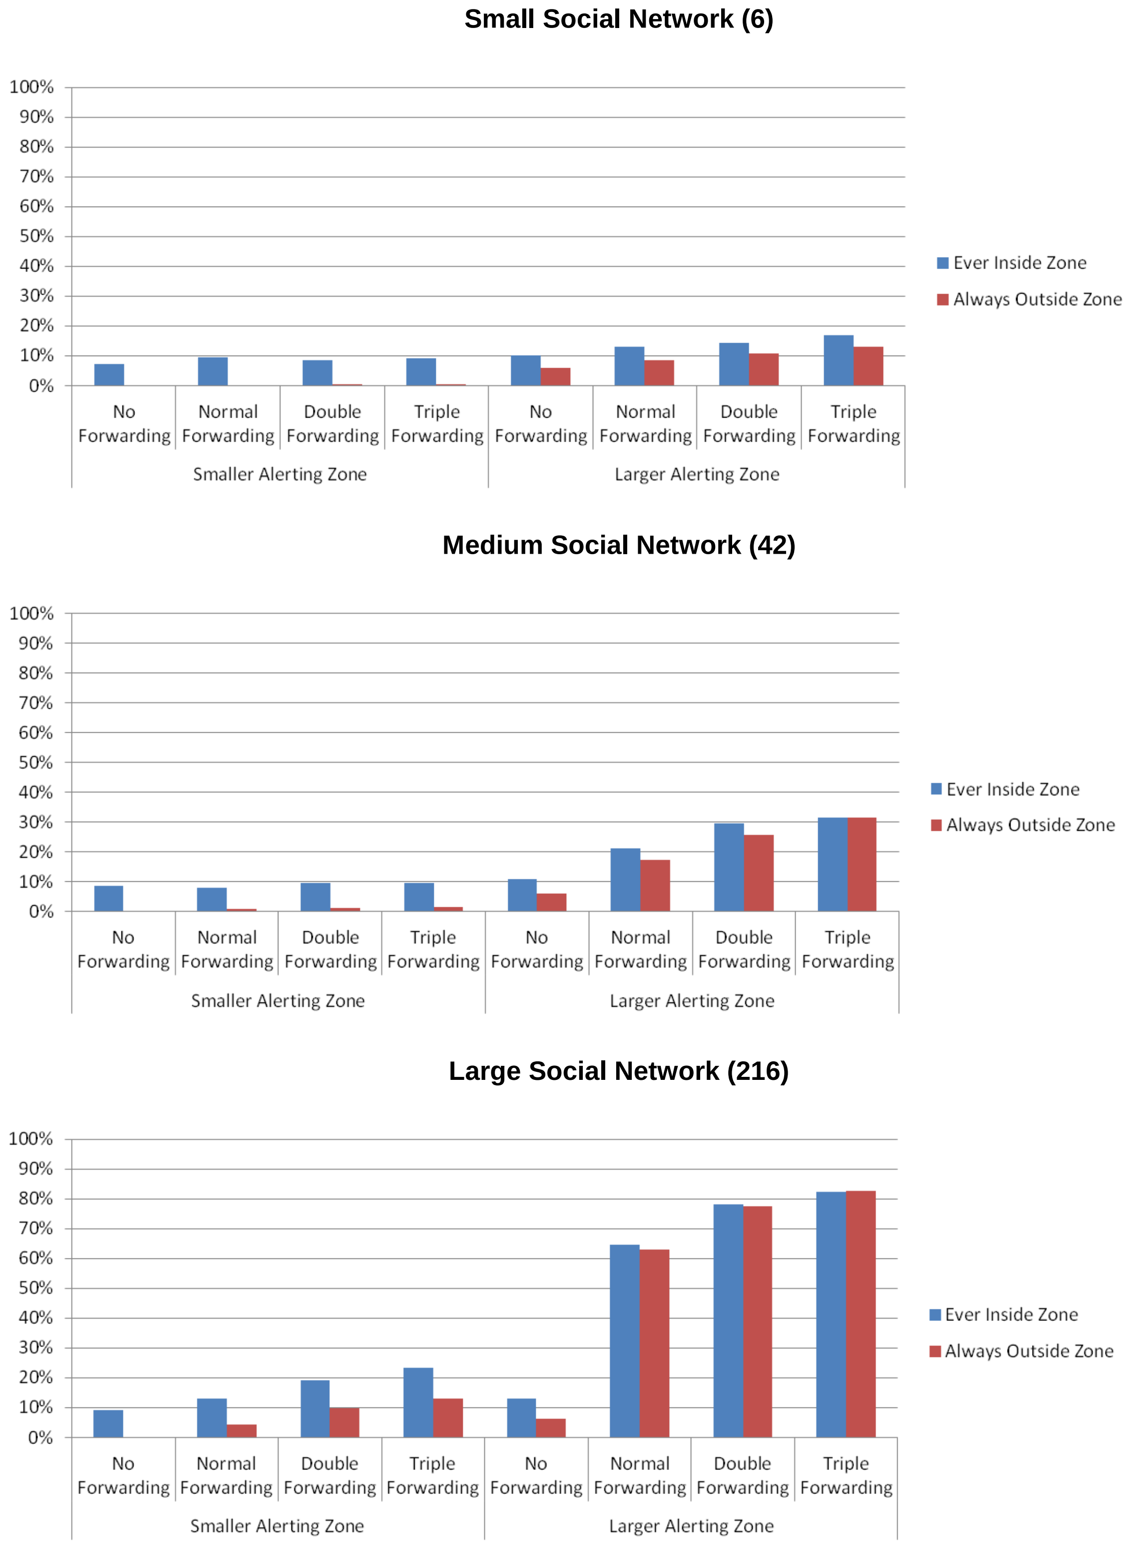


*NOTE: The results focus on cases with 70 percent phone compatibility and realistic delay parameters. Each case is the average of five runs of the model.*

Figure 4.10 shows the dramatic influence of the size of an individual's social network. When the agent really only influences the few people closest to him or her (top panel), there is relatively little evacuation, and it is relatively uninfluenced by either alert-zone size or likelihood of forwarding. When the agent has extremely wide influence (bottom panel), we see huge effects, such that with a larger alert zone and any forwarding whatsoever, we see perhaps unrealistic levels of evacuation behavior. Furthermore, this is true even among those who never enter the emergency zone. The medium social network size (center panel) shows a more modest amount of evacuation behavior, with a substantial forwarding effect in the larger alert zone.

In this small-scale experiment, we observe greater alert-forwarding effects with larger social networks. Put another way, increasing social network density can have a large effect — when network members are attentive and might comply with forwarded alerts. Therefore, it matters not only how many people the average individual forwards to, but also the number of those people who are behaviorally influenced. Other parameters influencing forwarding and compliance would presumably have similar effects.

# 4.7. Summary

In this chapter, we used our model to examine how geo-targeting and alert diffusion affect alert effectiveness. Considering first the flash flood scenario, which is both very small and very quick moving, we found no effect of alerting or alert forwarding using our first-guess parameter set. We then explored the potential of greater alert forwarding, implemented through reducing our parameters' delay in receipt and compliance. A modest decrease showed the potential for alerting and alert forwarding to influence the evacuation of the emergency zone, but the effects were not realized until too late — after the hypothetical emergency time. A more extreme reduction in these delay parameters showed dramatic increases in clearing the emergency zone. A logical analysis of our parameters helped illustrate how our parameters combine to imply very limited initial compliance with a WEA message, but also how forwarding can amplify small initial effects. This explanation was corroborated by manipulating forwarding volume through another parameter set — the amount of an individual's social network that is actually influenced by a forwarded message.

# 5. How Does the Value of Geo-Targeting Vary for Emergencies of Different Sizes and Speeds?

As we saw in Chapter Four, when events are small and fast enough, forwarding becomes limited by (1) the number of initial WEA recipients that can seed the forwarding process, and (2) the time available for forwarding to happen. Large, slow events, such as our major flood scenario, offer many potential forwarders and ample time for forwarding to occur. In this chapter, we consider events and alerting zones of different sizes and speeds as a means of exploring how geo-targeting and alert diffusion interact to influence the potential effectiveness of WEA.

## 5.1. Large, Slow Events Have Many Forwarders with Plenty of Time

Returning to our major flood scenario, we considered a very large area of effect; much of the population either started in or entered the emergency zone during the time between alert(s) and emergency. Figures 5.1 presents a somewhat different view on agent reactions. Rather than looking at how many agents were in the zone (see Figures 4.7 and 4.8), we focus on how many agents were evacuating. We also bring to bear another feature of the major flood scenario, two alert zones, inner and outer, which were told to evacuate at different times. The larger outer zone was initially told to evacuate. This was intended to help clear the roads for the inner zone, which was told to evacuate an hour later. Here we focus specifically on the first time period, to consider diffusion with larger alert zones. In Chapter Six we will return to this scenario to focus specifically on the role of differential instructions.

*Figure 5.1. Percentage Evacuating in the Major Flood Scenario, Between First and Second Alerts, for Agents Starting in the Inner Zone (top) and Outer Zone (bottom)*

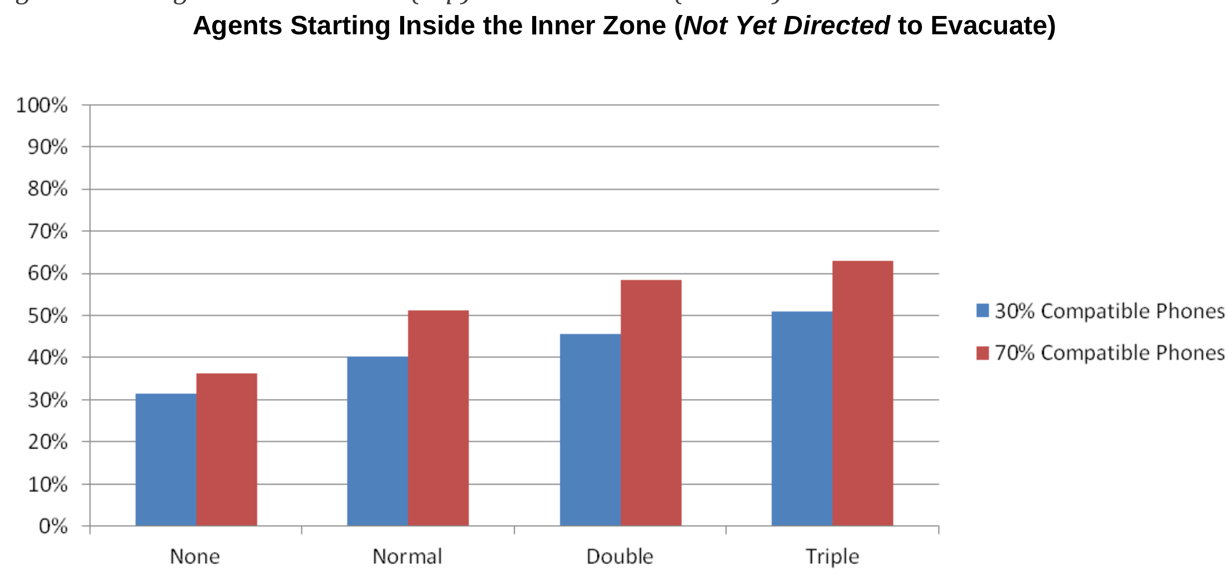

**Agents Starting Inside the Outer Zone (*Directed* to Evacuate)**

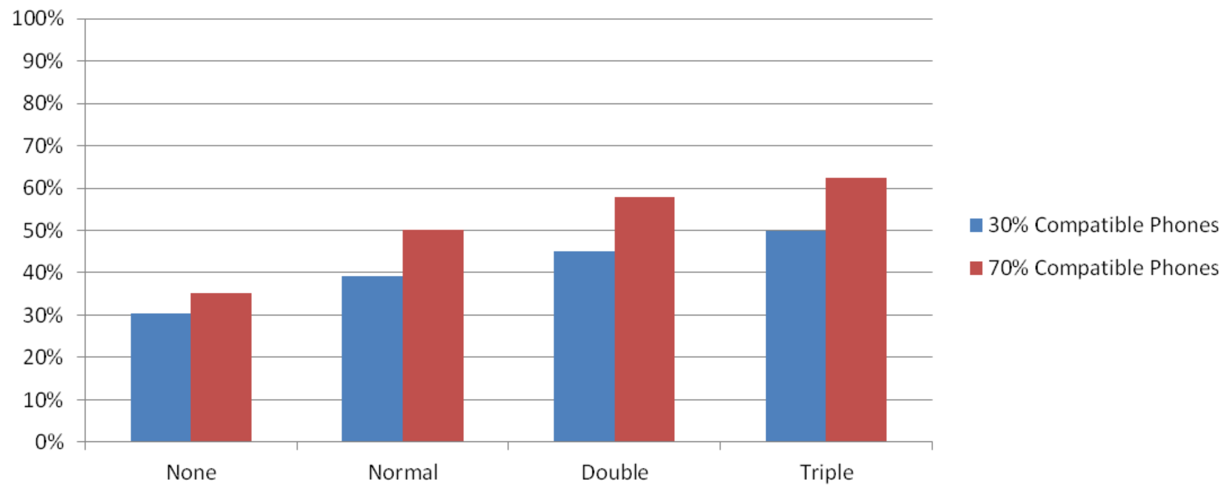


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases. The results reflect the reduced (but not eliminated) delay parameters. Similar results with default parameters are included in Appendix C.*

The bars in Figure 5.1 show the percentage of agents choosing to evacuate. Looking first at the bars on the left-hand side of the two panels, it is apparent that a large number chose to evacuate, even without forwarding. This is a direct effect of the larger zone and a larger number of agents receiving the alert directly. Moving from left to right, it is also apparent that WEA forwarding drove the protective-action rate upward, as would be expected due to the large number of initial forwarders within the emergency zone. Note that, as emphasized in Chapter Four, this is the case even though the default, realistic delay parameters were in place. Finally, phone compatibility (the blue versus red bars) appeared to have a modest but constant positive effect. The fact that the two figures are quite similar, however, may indicate that with so many forwarders and so much time, targeting almost necessarily broke down. Specifically, there was just as much evacuation happening in the inner zone (which was expected to hold in place since people had not yet been directed to evacuate) as in the outer zone. In Chapter Six, we return to this discussion, further considering the implications for differential alerting strategies.

## 5.2. Tornadoes Require Very Tight Timelines, but the Alert Zones Are Large Enough to Result in Significant Diffusion

Recall that the tornado scenario is fast moving, and whereas the actual emergency zone is a narrow strip, the uncertainty surrounding the precise path requires a much wider alerting zone. As displayed in Figure 5.2, even the smaller alerting zone was much larger than the emergency zone. The tornado scenario also differs from our other scenarios in that the sole directive is to shelter in place (operationalized as agents stopping where they are).

*Figure 5.2. Tornado Scenario, with Emergency Zone (red), Smaller Alert Zone (blue, left) and Larger Alert Zone (blue, right)*

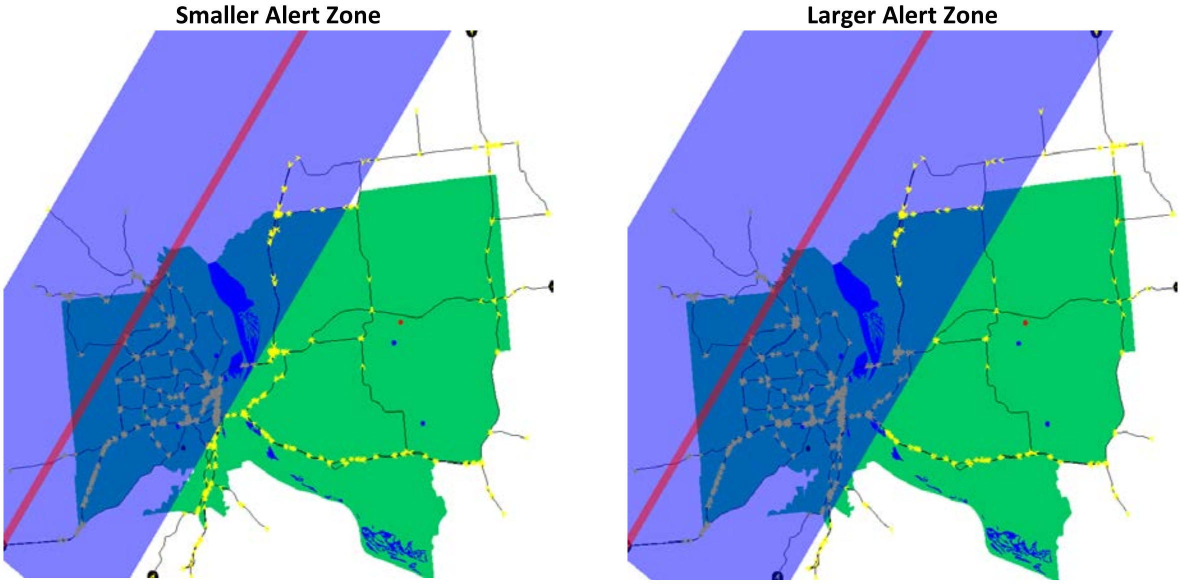


*NOTE: The yellow arrows depict the agents. Source: Tele Atlas.*

As before, we compared those within the emergency zone with those outside the emergency zone. Because the tornado emergency zone was so narrow and the few roads crossing it did so rather directly, the probability that anyone was in the emergency zone for a substantial amount of time was very small. Due to this, there was relatively little opportunity for individuals in the zone to receive multiple forwards of a WEA message. The upper panel of Figure 5.3 displays the percentages of those starting inside the emergency zone who actively attempted to shelter in place. The number of agents in this class was relatively small (approximately 50), since the emergency zone was so constricted, but even within this class, the proportion was small (ranging between approximately 4 and 13 percent). This may be because of the inherent limitations of WEA compliance (highlighted in Chapter Four). Furthermore, because agents traversed the zone quickly, many may have exited before passing the compliance threshold. The different bars in the upper panel of Figure 5.3 show a relatively constant boost to sheltering with increased phone compatibility, as would be expected. There were no strong effects, however, of either alert-zone size (both alerting zones were quite large compared with the emergency zone) or increased forwarding (which may have been the case because the incident timeline provided too little time for forwarding to have a substantial effect).

*Figure 5.3. Percentage Sheltering in the Tornado Scenario*

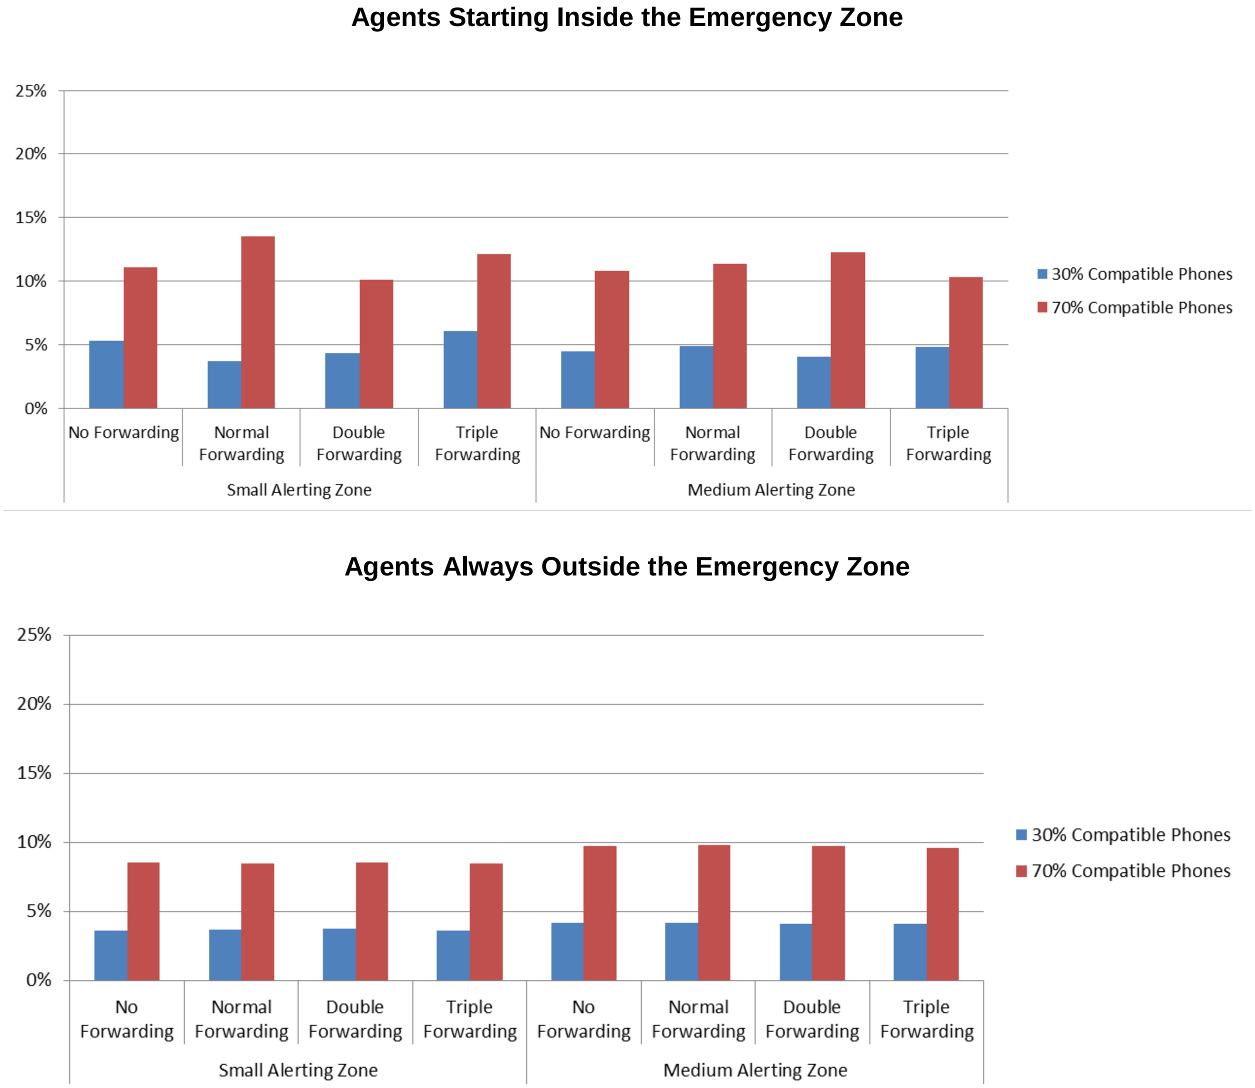


The lower panel of Figure 5.3 shows the results for those who started outside the emergency zone and stayed outside it. Within this class of agents, many also sheltered in place, presumably because they received a WEA message and did not realize that they were outside the true emergency zone, but this number was somewhat lower than those who started in the emergency zone. Otherwise, the lower panel looks remarkably similar to the upper panel. This unnecessary sheltering may not be entirely negative, since the downside consequences may be lower — and sheltering by some agents outside the emergency zone almost certainly prevented them from entering the emergency zone. Specifically, while it may be potentially inconvenient for the person, unnecessary sheltering does not carry the negative consequence of road congestion that unnecessary evacuation does.

Note that Figure 5.3 excludes the small group (also about 50 agents) who started outside the emergency zone but entered at some point during the simulation. This group showed almost complete noncompliance (averages at or near zero for all cases). We believe that this was caused by a set of factors. First, if an agent received an alert while outside the emergency zone (but perhaps heading toward it), he or she could either comply (and stop in place, consequently never entering the emergency zone) or ignore the message (consequently becoming immune to future WEA messages). This means that of those entering the emergency zone, the number who had not previously received and ignored it would be quite a small number. The additional agents who received a forwarded WEA message would also be a small number. This analysis also points out that the effect of interest here may not be the number of agents entering the emergency zone, but the number who would have entered but did not. This includes those who sheltered in place outside the zone and those who turned around, rather than enter the zone.

## 5.3. With Smaller Events, the Alert Area Can Provide Insight into Event-Size Effects

The flash flood scenario highlights zone size in a different way, by considering smaller versus larger alerting zones with reduced delay parameters. Recall that, there was very little effect of WEA or WEA forwarding with our best-guess and modestly reduced delay parameters. We did start to observe effects once we reduced those delays more substantially (and therefore increased forwarding traffic). Figure 5.4 summarizes the key results from Figures 4.5 and 4.6, focusing on normal forwarding and 30 percent phone compatibility cases. The figure contrasts an alerting zone that corresponds to the emergency zone (red) against one that is larger than the emergency zone (green). As before, the blue line represents the no-alert case. Focusing on just these two cases, we see clearly that moving from alerting only those in the emergency zone to including those in the larger alert area resulted in dramatically fewer agents in the emergency zone. Because these two cases differ only in the size of the alerting area, this effect is primarily due to agents outside the emergency zone, who may have entered instead avoiding the zone.

*Figure 5.4. Effect of Alert-Zone Size for Flash Flood on Number of Agents in the Emergency Zone, No Delays in Compliance or Receipt*

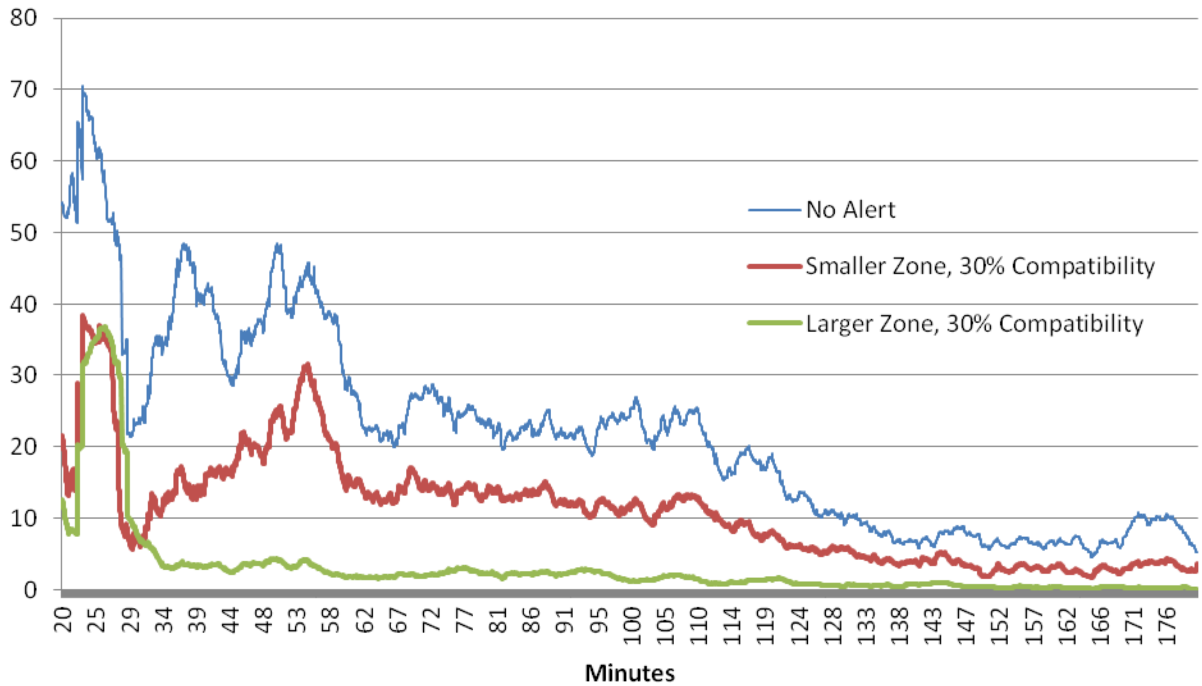


*NOTE: The figure considers only cases with normal forwarding and 30 percent phone compatibility.*

# 5.4. Considering the Role of Leakage Outside an Emergency Zone

These results highlight that when asking “how small is too small,” it is important not to think solely geographically. Diffusion outside the pure hazard zone may be problematic in some ways but positive in others. If we think primarily in terms of alert specificity, people outside the emergency zone should not get the alert. When agents are dynamic, however, their trajectories may be as important as their locations. We will likely want to alert individuals on their way into the hazard zone, and the faster they are going and the longer the compliance lead times, the further out we have to alert them, enabling them to take protective action. It is useful to return to the map inset in Figure 4.1, which illustrates the problem. At that moment within the simulation, no agents (yellow arrows in the figure) were in the emergency zone. Many, however, were immediately outside and several were likely to enter. By alerting an area larger than the emergency zone, these proximal agents were potentially warned away.

Figure 5.5 demonstrates the contrasting risk, however. The bars represent the percentage evacuating among those who never entered the emergency zone. This population is important, since it is never directly reflected in the earlier time trace graphs, because it never enters the emergency zone. We see that with a small, quick event and an alert area that corresponds to just the emergency zone (blue bars), not much evacuation occurred outside the zone. That said, evacuation outside the zone was clearly dependent on a delay in receipt. When there was even a modest delay in receipt (the left side of Figure 5.5), little evacuation occurred outside the emergency zone.

*Figure 5.5. Percentage Evacuating Among Those Always Outside the Flash Flood Emergency Zone, by Alert-Zone Size and Delay in Receipt*

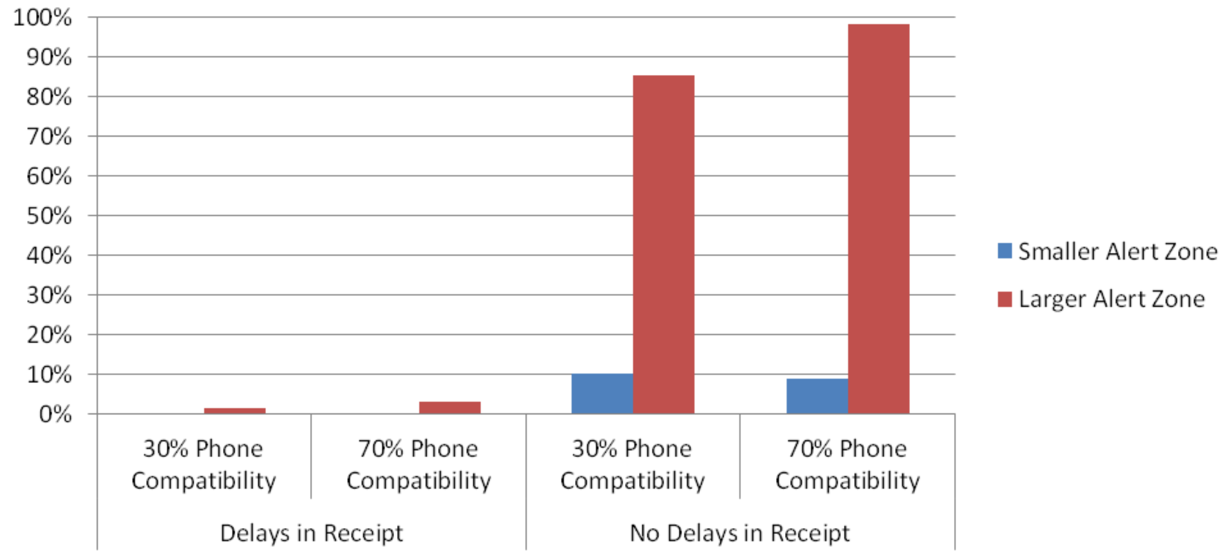


The delay in receipt itself was probably less the issue than the effect that reducing the delay had on the amount of forwarded messages. When receipt delays occurred, there was less forwarding and no or minimal evacuation outside the emergency zone. When we removed the receipt delays, forwarding was amplified, and those outside the zone started to evacuate. Specifically, it appears as if the interaction between delay in receipt and alert-zone size combined to dramatically increase the rate of evacuation outside the emergency zone (i.e., unnecessary compliance). In fact, in the hypothetical case with no receipt delays, when an alert zone was larger than the emergency zone, the rate of unnecessary evacuation approached the entire population. Clearly, this case had an unrealistic boundary condition, but it does illustrate the potential negative consequences of extreme alert diffusion.

It is also important to note how much the initial number of forwarders mattered. Even with no receipt delay, the small alert zone only had a modest number of people outside the zone evacuating. This was due to the small number of *seed forwarders*. With the medium alert zone, this number of seed forwarders was much greater, and hence the potential for initiating forwarding cascades increased.

In summary, whereas we may be concerned about alert diffusion because of the annoyance and unnecessary compliance it may cause, some leakage outside the emergency zone is not really leakage. In particular, understanding the traffic patterns is not simply about understanding the effectiveness of the alert; it is also about the people you want to receive the alert. For example, we may want to target everyone driving toward the emergency, but not those driving away from it. The longer the delay of receipt and the longer it takes people to respond, the further away from the hazard you want to alert people to give them time to avoid the emergency zone. In cases similar to what we simulated, where the emergency happens on a thoroughfare, this is especially important.

## 5.5. Summary

In this chapter, we examined the effects of emergency size and speed on the value of geo-targeting. With the very large, slow events, such as our major flood, we saw substantial evacuation that increased with phone compatibility and forwarding behavior. The tornado had a similarly large alert zone, but a much smaller emergency zone and a much quicker timeline. Here, compliance rates were smaller and did not vary by alert-zone size. There was also small but noticeable out-of-zone sheltering (which may be of less concern than out-of-zone evacuation, which could cause road congestion). Returning to the flash flood results illustrated how changing the alert-zone size can have dramatic effects on emergency-zone evacuation. That same scenario also demonstrated that the potential cost to this dramatic evacuation (and the related results in Chapter Four) was significant over-alerting and, as implemented in the model, unnecessary out-of-zone compliance behavior.

# 6. How Does the Interaction of Geo-Targeting and Message Diffusion Affect the Utility of Sophisticated Alerting Strategies?

One of WEA's promising characteristics is the potential to use its geo-targeting capabilities to send different alerts to people in different locations. For example, Vogt and Sorensen (1999) described an emergency response to a chemical repackaging plant in Helena, Arkansas. A key aspect of the alerting strategy was to differentiate the alerts targeting those in greatest danger from those in more moderate danger. Specifically, authorities in Helena issued alerts asking those within two miles of the plant to evacuate, and residents between two and three miles to shelter in place. Despite the differential instructions, many of those in the shelter zone appeared to evacuate. WEA's geo-targeting capabilities offer added potential for administering such differential alerting. Furthermore, WEA's relatively quick dissemination suggests that alerts could vary not only geographically but also temporally. Such capability could be helpful in implementing a staged evacuation, which would reduce traffic problems, as has been observed in some past evacuations in wide areas. For example, during the evacuation before Hurricane Rita in 2005, the traffic congestion from the number of people evacuating caused major problems (Litman, 2006), and the vast majority of the fatalities associated with the event were from the evacuation rather than the hurricane (Zachria and Patel, 2006).

## 6.1. Using WEA to Differentiate Evacuation and Sheltering Areas

Our hazmat scenario was modeled after the Helena, Arkansas, case. In our scenario, agents in an inner zone (1 km by 0.75 km) were told to evacuate, while those in an outer zone (3 km by 2.25 km) were told to shelter in place. Separate WEA messages were sent to each zone. Figure 6.1 illustrates the geography of this scenario.

*Figure 6.1. Hazmat Scenario, with Inner Evacuation Zone (red) and Outer Shelter-in-Place Zone (orange)*

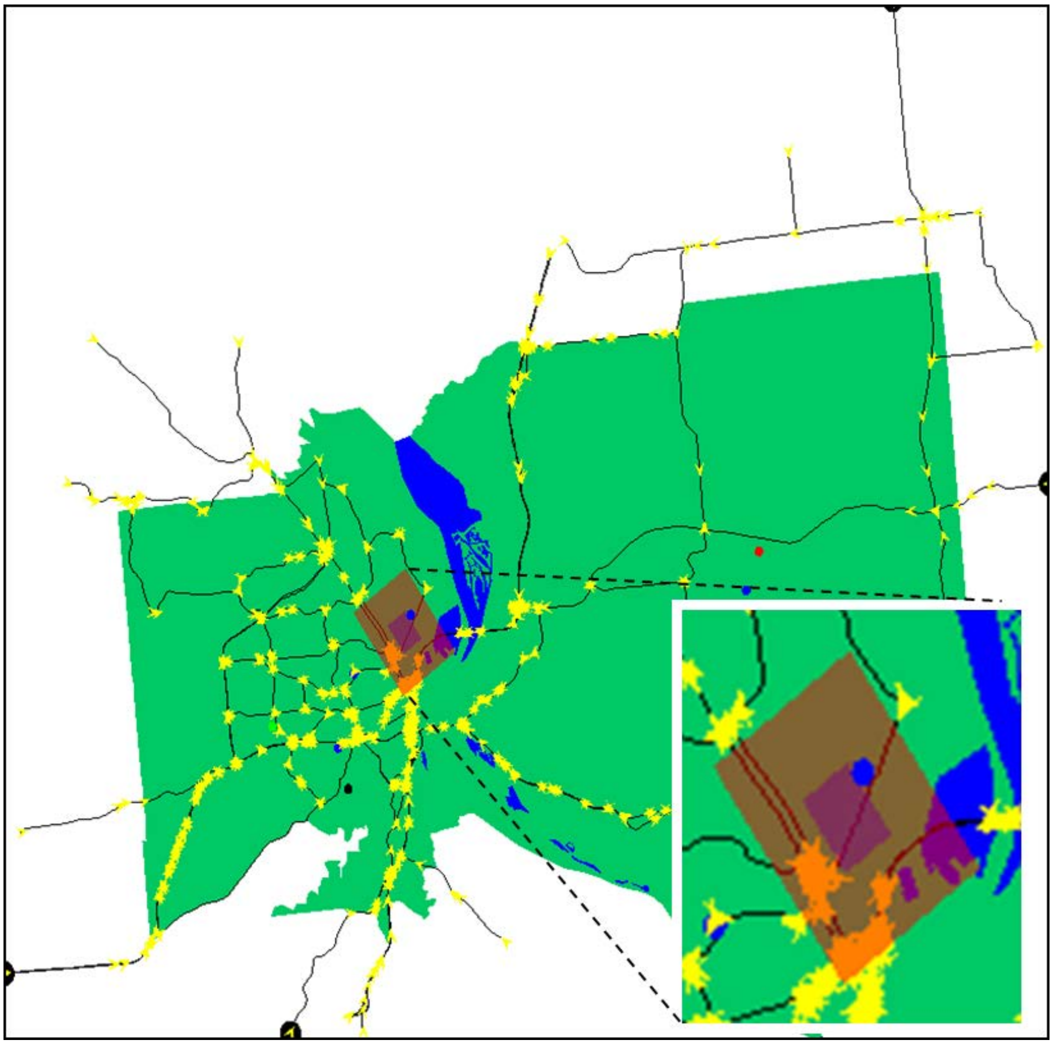

*NOTE: The yellow arrows depict the agents. Source: Tele Atlas.*

Figure 6.2 presents the evacuation behavior of agents in each of the two zones with the reduced delay parameters. As would be hoped, a greater percentage of agents were evacuating in the inner evacuation zone than in the outer shelter-in-place zone, but some in the outer zone did evacuate. And whereas increased phone compatibility (red bars) showed consistently greater evacuation, any forwarding effects (moving from left to right in the graphs) were modest at best.

*Figure 6.2. Percentage Evacuating in the Hazmat Scenario, for Both Inner Evacuation Zone (top) and Outer Shelter-in-Place Zone (bottom)*

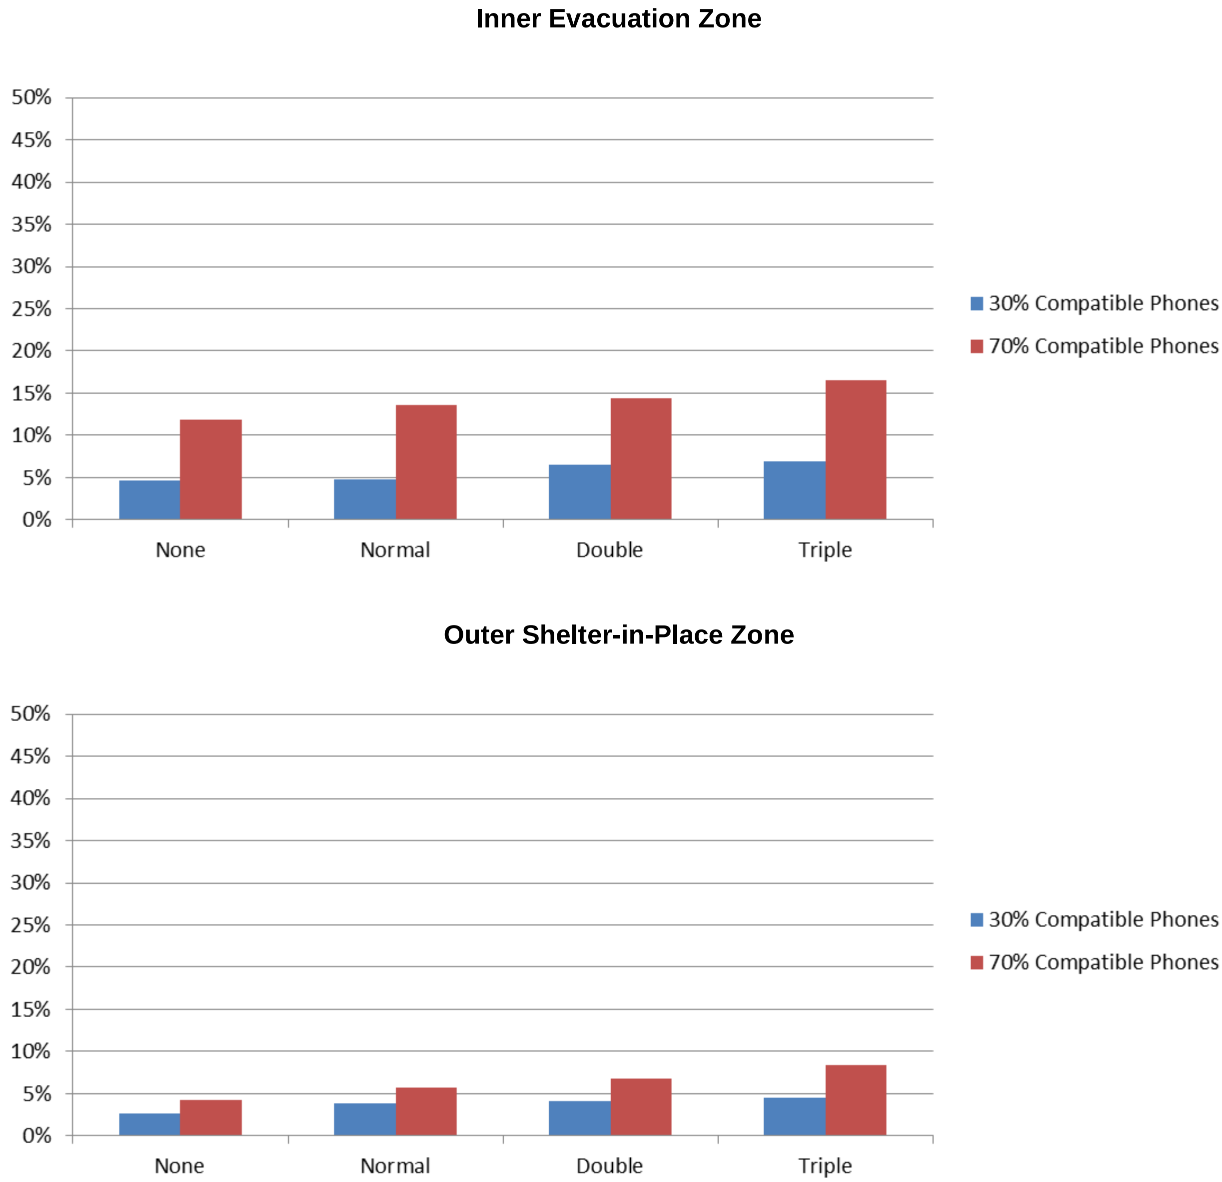


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases.*

Figure 6.3, in contrast, presents the sheltering behavior of those in the inner and outer zones. Recall that the outer zone, where the directive was to shelter in place, was much larger, so that instruction was received by more agents. Here we saw more sheltering in the outer zone, but only by a small percentage of the agents. There was substantial sheltering in the inner evacuation zone, however.[18] There is a

[18] In our simulation, we did not include differential levels of compliance for evacuation versus shelter-in-place directions, which have been observed in actual emergency events. As a result, agents had the same probabilities of compliance with either type of direction.

strong forwarding effect, which suggests that those in the inner zone were unnecessarily (and perhaps dangerously) sheltering because they mistook a forwarded WEA message as applying to them.

*Figure 6.3. Percentage Sheltering in the Hazmat Scenario*

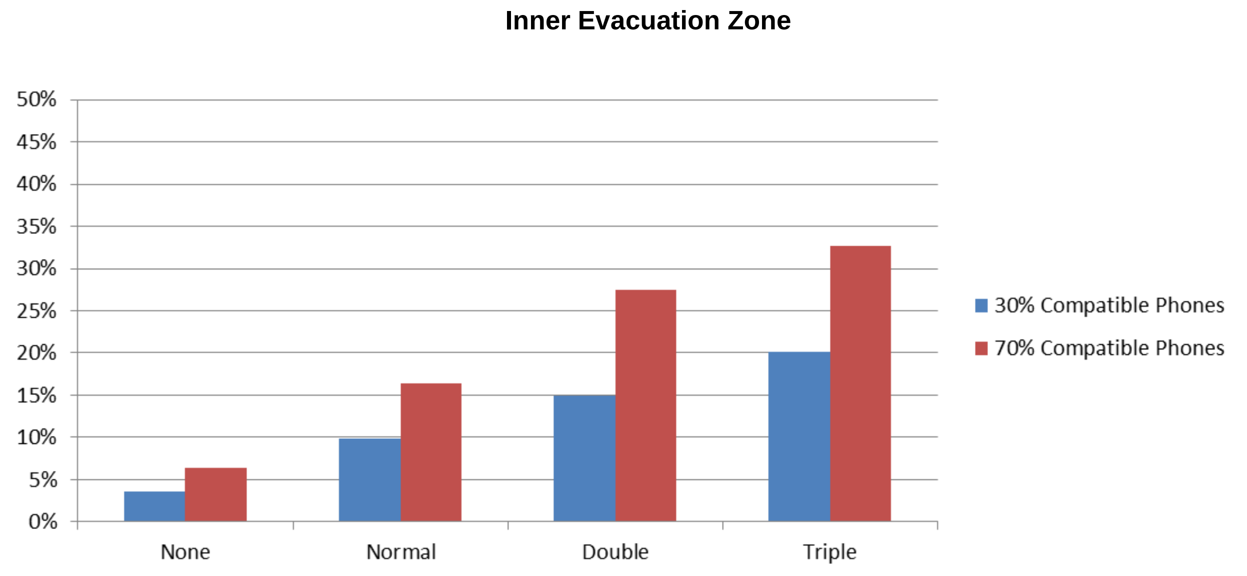


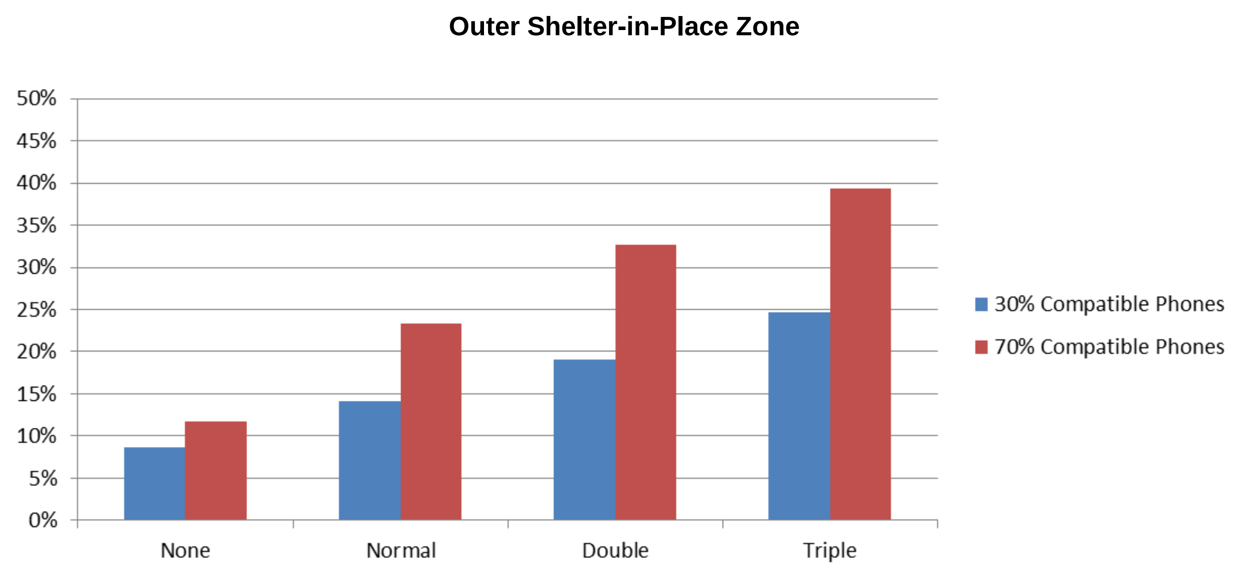


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases.*

# 6.2. Using WEA to Implement a Staged Evacuation

Our major flood scenario, depicted in Figure 6.4, also consisted of an inner and an outer zone. Rather than differential actions, agents in the two zones were told to evacuate, but at different times. Those in the outer zone were told to evacuate first and those in the inner zone second. The major flood scenario also involved a very large overall alert zone, covering much of the population, and it had a more protracted alerting period, spanning two simulated hours.

*Figure 6.4. Major Flood Scenario, with Inner (blue) and Outer (red) Evacuation Zones*

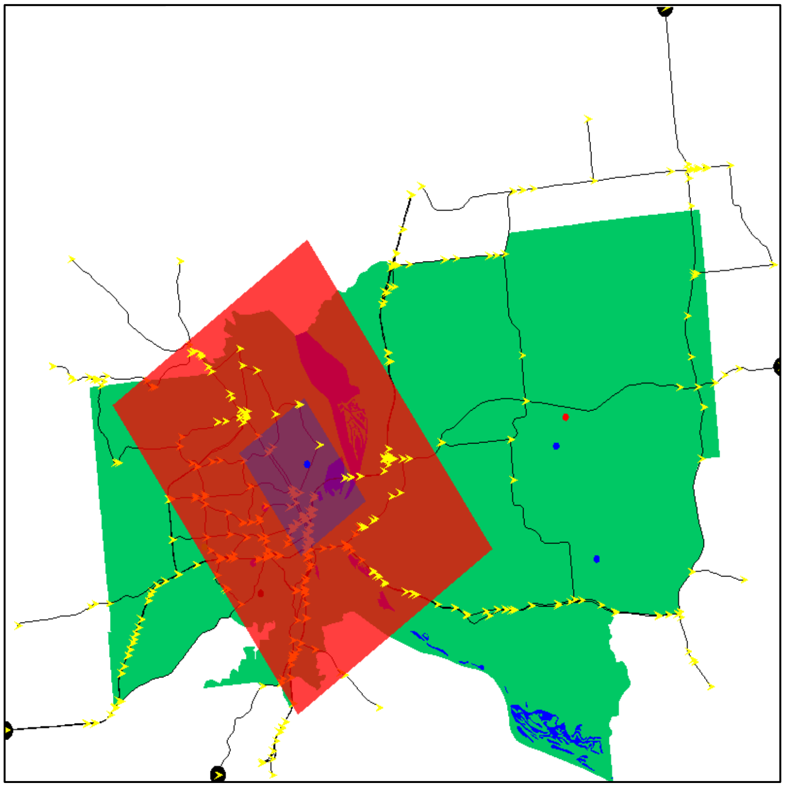

*NOTE: The yellow arrows depict the agents. Source: Tele Atlas.*

Figures 6.5 and 6.6 divide the results according to two separate time periods. Figure 6.5 replicates Figure 5.1, but we focus our discussion on differential alerting (rather than size, speed and diffusion). The figure looks at the interval between the first alert (telling those in the outer zone to evacuate) and the issuance of the second alert, and it focuses on those agents in each zone at the time of the first alert. If the staged evacuation proceeded as planned (and ignoring the potential for agent movement from zone to zone), only agents in the outer zone would evacuate during the first interval. Figure 6.6 focuses on the interval between the second alert (telling those in the inner zone to evacuate) and the projected time of the emergency, and it emphasizes those agents in each zone at the time of the second alert. Each interval lasts one hour. Each figure displays the percentage of agents who evacuated, compared with those who started the period in the inner zone (upper panel) and those who started the period in the outer zone (bottom panel).

*Figure 6.5. Percentage Evacuating in the Major Flood Scenario, Between First and Second Alerts, for Agents Starting in the Inner Zone (top) and Outer Zone (bottom).*

**Agents Starting Inside the Inner Zone (*Not Yet Directed* to Evacuate)**

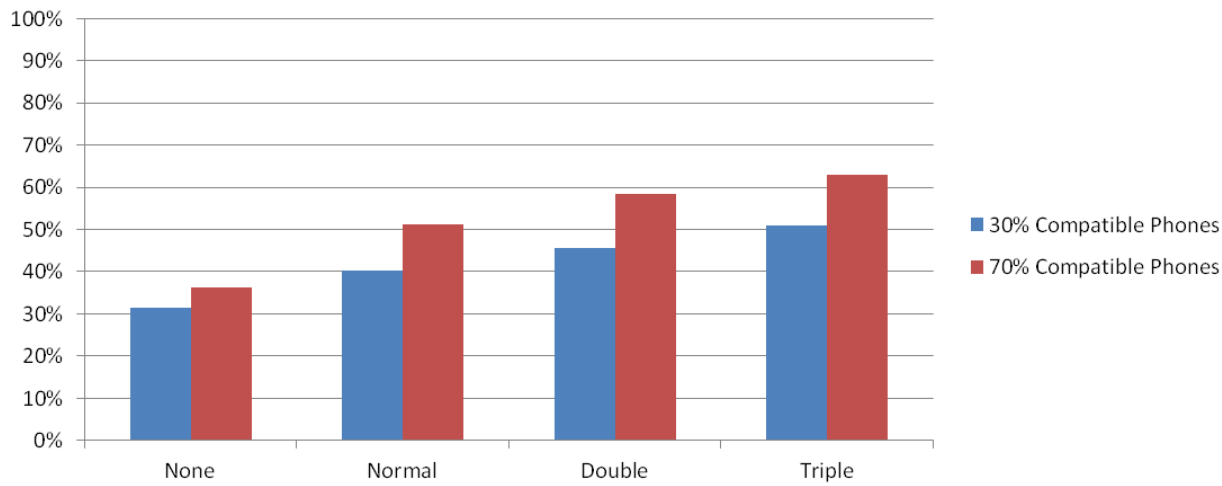


**Agents Starting Inside the Outer Zone (*Directed* to Evacuate)**

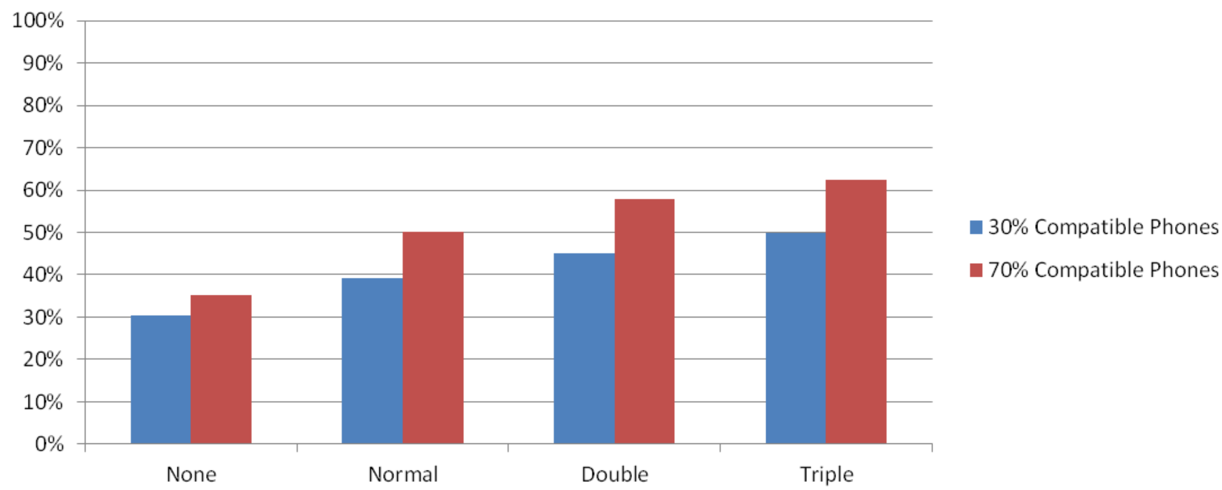


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases. The results reflect the reduced (but not eliminated) delay parameters. Similar results with default parameters are included in Appendix C.*

What we see is *very little differentiation* in evacuation rates across the two zones, and indeed across the two time periods. In particular, many agents in the inner zone evacuated during the first phase (Figure 6.5, top) when they should not have. When considering Figure 6.5, it appears as if the initial WEA was not the central driver of evacuation behavior, or else we would see substantial differentiation between the top and bottom panels. Instead, it was *the forwarding of the WEA message* that really had the effect, and we do indeed see substantial forwarding effects moving from left to right in Figure 6.5. Because of the large number of agents in the initial evacuation zone, there was a large pool of potential message forwarders, which created the potential for a large amount of forwarding traffic circulating in the agent population during the time between the first and second alert messages.

One noteworthy result in Figure 6.5 is the number of agents in the inner zone who evacuated even in the absence of any WEA forwarding. This is likely a direct consequence of the slowness of the event and the movement of agents around the region. Many of those starting in the inner zone traveled out of that zone during the hour between the first and second alerts. Because the inner zone was completely surrounded by the outer zone, these agents moved into an area where they could receive the WEA message directly — i.e., the natural movement of the agents in the model itself was a central contributor to the strategy of phased evacuation breaking down.

The time between the second alert and the emergency (Figure 6.6) showed more differentiation, with the outer zone having a greater percentage evacuating than the inner zone. The transmission of the second WEA message — given the large amount of WEA forwarding after the first alert — would be expected to have a relatively modest effect on the message and forwarding traffic in the population of agents. The large amount of forwarding of the first WEA message would mean that many, if not all of the agents, would have received either an evacuation WEA (which they either complied with or ignored) or a forwarded evacuation message. The second WEA aimed at the agents in the inner zone would therefore be an incremental addition on top of that message traffic (with a smaller number of agents in the inner zone to receive the WEA and begin forwarding that second WEA anew). As a result, the populations of agents observed between the second alert and the emergency represent both the effects of the first alert (with its massive forwarding) and the second alert.

The differentiation in Figure 6.6 is therefore the result of a number of factors and likely driven more by the effects of the *first* versus the *second* alerts. At this point in the simulation, some agents (in both zones) would have received and rejected the WEA message when they passed through the outer, earlier evacuation zone (and hence became immune to the alert). For those agents, the transmission of a second WEA message (if they were in or entered the inner, second evacuation zone) would have no effect, though the forwarding of the message might. A significant number of agents would also have already left the zone through evacuation. As noted, those evacuating the inner zone necessarily traversed the outer zone, so they may have been counted there in Figure 6.6;any traffic effects could hold agents who were evacuating in the outer zone, increasing the apparent percentage of the population evacuating in that zone. Finally, the outer zone was much more likely to have outside agents (those initially in neither zone) enter during the course of the period, in a sense replenishing the pool of potential evacuees.

*Figure 6.6. Percentage Evacuating in the Major Flood Scenario, Between Second Alert and Emergency, for Agents Starting in the Inner Zone (top) and Outer Zone (bottom) at the Point the Second Alert Is Issued.*

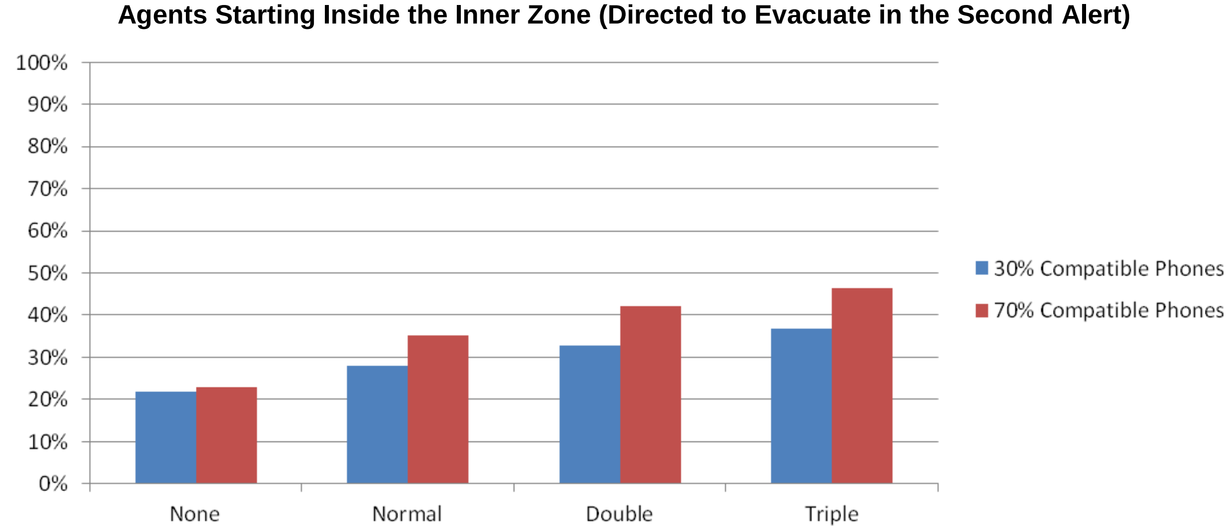


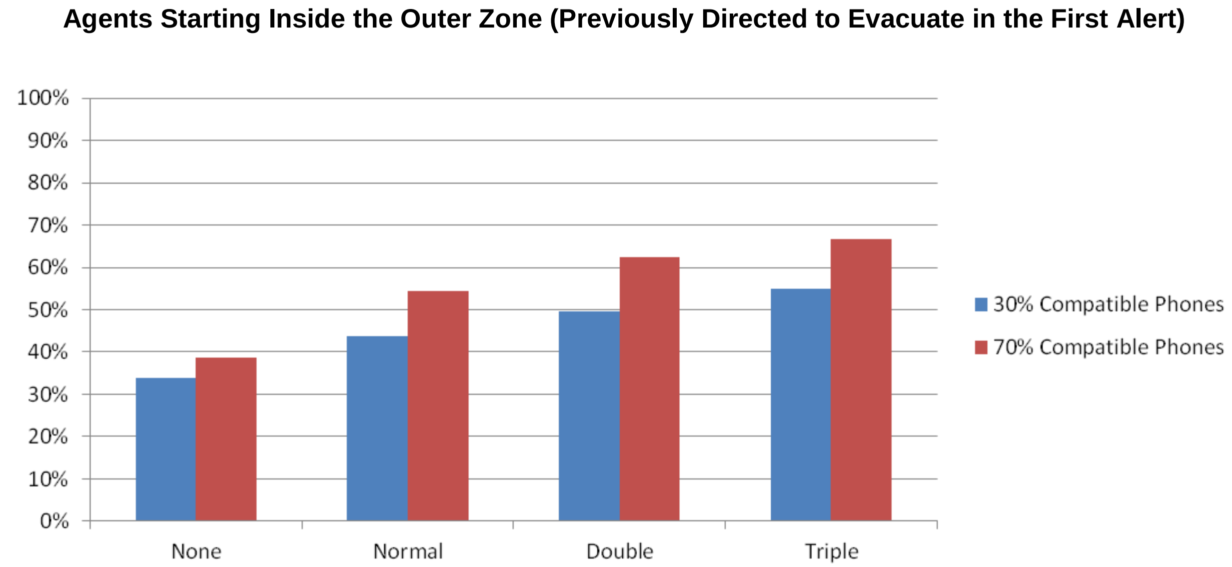


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases. The results reflect the reduced (but not eliminated) delay parameters. Similar results with default parameters are included in Appendix C.*

# 6.3. Summary

This chapter examined how the combined effects of geo-targeting and diffusion affect the ability to implement more-sophisticated emergency responses, such as different instructions for different zones (as in the hazmat scenario) or the same instruction at different times for different zones (as in the major flood scenario). Considering both sets of results, it is evident that the diffusion of alerts can dramatically limit the utility of such differentiated alerting strategies. In both scenarios, agents reacted to alerts targeting other zones. In particular, those in the hazmat scenario's inner evacuation zone often sheltered in place (perhaps placing them at increased risk), and those in the major flood scenario's inner zone often evacuated during the earlier period when they were asked so stay in place (likely increasing the very road congestion the strategy was designed to reduce).

# 7. Conclusions and Policy Implications

The complexity of individual behaviors in emergency alerting situations, and the way those behaviors affect the ability of alerts to serve as a protective measure during major emergencies, has been a topic of interest for decades. The evolution of new relationships between citizens and technology has affected the utility of legacy alerting modes, and the arrival of new options, such as WEA, has added new facets to that complexity. The ability to geo-target alerts to mobile devices provides new capability, but it also raises questions about how to use geo-targeting effectively, and how the interaction of targeted messages with human communication behavior will affect the effectiveness of alerts. Our model provided a way to explore this new complexity. By distilling the different elements of alerting to a set of simplified parameters, the model allowed us to look at how changing some of those parameters could affect alerting outcomes — dialing up and down the parameters that shaped how citizens communicated and forwarded messages to get insights into how that behavior could become a multiplier for alerting even as it fought efforts to precisely geo-target alerts.

After extracting the policy-relevant conclusions from the suite of results, it is evident that the relationship of forwarding to geo-targeting, was in some ways much less important and in some ways much more important than we thought. In this final chapter, we consider some of the more interesting implications of this conclusion.

## 7.1. Forwarding as a Compliance Enhancer, Rather Than Simply a Limit on Geo-Targeting

Our initial framing of the study considered forwarding as a potential threat to the value of geo-targeting. This perspective assumes that forwarding produces a practical limit on geo-targeting precision. The simulations certainly showed this, but the relationship was a lot more nuanced than we initially thought. Our results suggest that very tight geo-targeting will always be somewhat compromised by forwarding, but if very few agents are in the alert zone, that compromise will be small. From a policy perspective, if we ask whether forwarding is a reason not to invest in geo-targeting at very tight resolution, then the answer is no. This is simply because in those cases where one was trying to alert very few agents, there were very few people to forward the message in the first place.

This means that this effect is of greatest concern for our intermediate-sized scenarios and cases, where the population to be alerted is reasonably large, yet the geographic area is small enough to suggest value in precise geo-targeting. As the number of people being alerted in our simulations increased, the amount of potential forwarders also increased, and the ability to target precisely decreased — though, depending on the scenario, that could matter to differing extents.

This result was also obviously sensitive to the amount of forwarding traffic that occurred, as demonstrated in the experiment that varied the size of the influenced social network. As forwarding traffic decreases (in the case of our experiment, by cutting back the size of agents' affected social networks), forwarding effects become less important.

Instead of considering forwarding as a threat to geo-targeting, it might be more valuable to think of forwarding as a compliance enhancer (presuming that the information in the forwarded alert remains intact and correct). Our model parameters assumed relatively low rates of net compliance for many channels, based on available literature. As we showed in Chapter Four, under those conditions, significant forwarding (with our assumed increased compliance with messages received through some forwarding channels) could increase the total effect of alerting a great deal. In that sense, forwarding converts the set of individuals' social networks into a new mass alerting channel of its own. In this way, forwarding and communication among individuals about the alert, here via electronic means, such as social media, are simply one more type of the communication that has always occurred during milling before citizens make the decision to comply with the alert. Taking this view, narrowing geo-targeting to the greatest extent possible becomes less of an issue, since doing so would minimize potential forwarders of the message. Furthermore, the various steps in Figure 4.9 become potential targets for policy intervention — not just to increase individual recipients' understanding and compliance with a message but to increase their effectiveness as a forwarder to communicate the message to others.

The "price" of this effect is a significant increase in out-of-emergency-zone alerting and potentially unnecessary action in response. This price cannot be ignored — one reason for increasing the specificity of the targeting of WEA is the hypothesis that individuals, annoyed by alerts that are not relevant to them, will opt out of the system, damaging its future value as an alerting channel (DHS, 2013a). Forwarding may be of less concern, however, than *direct* over-alerting by WEA. While direct over-alerting could trigger opting out of WEA, it seems unlikely that people would opt out of WEA because their family or friends forwarded them a WEA message.

Furthermore, this role of forwarding as a compliance enhancer appears to become more important in emergencies where the alerting zones are larger and the timescales are longer. For large events — assuming there is no attempt at differential messaging, such as the phased evacuation — the central goal is to broadly disseminate a message and have people act on it in the time available before the event. In that case, demand for precision in geo-targeting is not an issue, and there is little downside for messages escaping the area (again, presuming that alerts are forwarded intact). As a result, the value of forwarding messages to improve compliance would dominate the cost of leakage by increasing the percentage taking action to protect themselves, among the very large number of citizens who could be affected by such an event.

# 7.2. Considering What "Ideal" Targeting Might Look Like

The model also taught us that *precise geo-targeting* can have different meanings when we consider how agents are dynamic, rather than thinking about geo-targeted alerts being aimed at a set of agents that is sitting still, waiting for the alert to arrive. This observation comes directly out of the fact that the flash flood scenario was situated on a thoroughfare where agents were traversing but not stopping. Considering Figure 4.1 again, it is worthwhile to ask which agents (or alternately, what region) would be ideally targeted by a WEA message? This would certainly include sending the message to people who are *in the emergency zone*, to trigger a goal to evacuate. But, as highlighted in that scenario, restricting the alert to only the emergency zone fails to warn agents driving toward the flood. Consequently, ideal targeting would also likely include getting the message to people *contemplating entering* the emergency zone, to convince them to turn around and not enter. This would be different in situations where alerting largely targets stationary individuals; for example, in a case where agents are all motionless (such as a residential area at night), leakage from the hazard zone would not have these positive effects. This suggests context specificity as to how large a geo-targeted zone should be based not just on the size of the emergency but also on the likely direction and speed of the transportation processes that would move people in and out of the zone during the emergency. These observations regarding ideal targeting also moderate what might otherwise be a fairly pessimistic view on geo-targeting for smaller emergencies. We found that alerting has little effect, at least on a quick timescale, but considering the need to *keep people away from a zone* (versus solely alerting individuals who are in the zone to promptly evacuate) would argue for increasing the size of the geo-targeted area and therefore increasing the effect of alerting.

# 7.3. Forwarding Clearly Threatens the Value of Differential Alerting

Our results demonstrate quite clearly that forwarding threatens the value of trying to deliver different messages to different geographic areas in an effort to either provide messages relevant to individual risk areas or to guide population behavior in ways designed to enable a more effective response (i.e., time-phased evacuation). In our model, compliance with forwarded messages was higher than with the direct messages. Hence, forwarding, and the potential for spreading forwarded messages outside their initially targeted area, could rapidly overwhelm the desired differential behavior from people complying with the original message. We observed this both in the major flood (phased-evacuation) scenario and in the hazmat scenario, where different populations were directed to shelter or evacuate. The breakdown that we saw occurred even in the absence of some of the differential compliance behavior that has been observed in some real-life emergencies, including the often-higher compliance rates with evacuation compared with the shelter-in-place direction.

The dynamic movement of agents was an additional challenge for differential alerting strategies. In the simulation, individuals who started in one zone but did not comply with the alert might move into the other zone — again demonstrating that alerting strategies for populations in motion are more difficult to frame than if the population is viewed as static.

The dynamics of forwarding and compliance mean that success in these sorts of differential alerting efforts would require more-nuanced communication strategies. A more nuanced strategy would have to try to persuade people — such as in the inner evacuation area of the major flood scenario — that it really is in their interest to wait to evacuate until they receive their own direction to do so. It is unclear (but worth considering) how a “shelter for now” or “shelter until 12:00 a.m.” message would work in such situations (DHS, 2014). The fact that WEA messages are restricted to no more than 90 characters limits the amount of such information that could be included in an alert, however.

# 7.4. Limitations and Potential Future Directions

Perhaps the most important caution to be placed on these results derives from the very nature of such a modeling exercise, which is an abstraction of reality. The model was designed to create an analytic environment that captures key elements of real-world dynamics. Agent-based modeling is especially strong at incorporating complex and dynamic elements, such as geographic movement and communication across social networks, which are difficult to measure empirically (on any large scale) and largely intractable using more-direct mathematical approaches. Arguably, this approach allowed us to make observations that would not have been feasible with other approaches. Most notably, the strong result that the forwarding of messages has the potential to be the dominant force in alert dissemination was only observable through the examination of social-communication dynamics. All models are abstractions, however, and aspects of reality might not have been completely captured. For example, the model included just one or two WEA messages being sent per event, whereas in reality, multiple alerts might be sent. Similarly, the model presumed verbatim forwarding of alerts, without the modification, commentary or augmentation that may occur in reality (and as exemplified in Figure 1.2).

A set of questions outside the scope of the current project, but perhaps of interest for future research, has to do with potential emergent technologies. A simulation platform, such as the one implemented in this research, could be used to explore the effects of adding additional WEA capabilities. For example, one could envision a future smart phone app that provides the user with some indication of the perceived value of an alert, drawing perhaps on temporal lag or network distance (i.e., the number of forwards it took to receive the alert). Our 70 percent phone compatibility cases were one step in the direction of considering future technology, but one focused on incremental gains rather than fundamentally new capabilities.

It is important to remain clear that the model itself relies on the estimates of the parameters that serve as its basis. Parameter estimates from the literature were stronger for some sets of parameters than others. Wherever possible, we drew from the literature, supplementing with logical analysis to make our assumptions. Furthermore, in many of our experiments, we strategically focused on parameters (such as the size of the social network) for which the literature was richest. In other cases, we systematically swept a range of values to capture uncertainty, to examine potential future changes in the world, and to target key elements of our research questions. This was the case with varying WEA phone compatibility rates (which are likely to increase with time) and WEA forwarding rates (which are both relatively uncertain and critical given the research question we were taking on). Other questions, which may be of interest to some audiences, were outside the scope of this project. For example, our assumed level of alert understanding was informed by U.S. census figures on sensory disabilities and English-language proficiency. Other factors could also lead to greater or lesser understanding, such as prior hazard exposure, and could be explored in future research.

The results also depended on the scenario specifications. As described in Chapter Three, our main choice of scenarios was based on frequency of WEA use in real events, as well as cases from the literature (e.g., the Arkansas chemical plant explosion). Specific details were sometimes dictated by other choices or based on researcher judgment, however. For example, emergency and alert zones had to be placed based on judgment but were also bound by the abstracted geography of Dubuque, Iowa, that we were using as our simulation location. In many cases, the alternative placement of the emergency zone or consideration of another geographic area might have led to other insights. Such was the case in Chapter Four, where the chosen location of the flash flood emergency zone informed insights regarding commuter dynamics, but an alternate placement that included intersections might have highlighted other dynamics (e.g., hitting larger numbers of agents with the emergency alert message while they were at home, before they started moving around).

Finally, the modeling focused exclusively on highly localized emergencies, in accord with our emphasis on the value of more-precise geo-targeting. Even our largest event (major flood) covered a fraction of a metropolitan area. Use of WEA and WEA geo-targeting in larger events, including such issues as cross-jurisdictional coordination, would require a somewhat different approach.

# 7.5. In Conclusion

The geo-targeting of alert messages to mobile devices represents a significant new capability for emergency managers compared with simple alerting that blankets a large geographic area with a warning message. Complexities of human behavior — including both message forwarding and movement — mean that the use of geo-targeting is not as simple as restricting the transmission of an alert to the smallest area at risk from an emergency event.

Although the potential for message spread is not the threat to the value of geo-targeting that might immediately be assumed — particularly for small events where the number of people at risk, and therefore the number of potential forwarders, is small — its use requires due consideration to ensure that emergency messages are actually targeted to the populations that need to receive them. Proper targeting can require deciding how far outside the actual footprint of the emergency the message should be transmitted, for cases where the movement of the population and the potential speed of the event put a premium on alerting individuals so they never enter the hazard zone versus warning individuals who are already there to evacuate. This could pose a difficult choice for alerting originators. An event occurring in an area where the population is largely stationary might justify a smaller alerting footprint than an identical emergency occurring in an area where people are rapidly moving toward or through the hazard area.

But, as suggested above, if the transmission of alerts through the population is viewed less as leakage from a geo-targeted area and simply as a key part of milling behavior in today’s electronic age, then the forwarding and sharing of alert information through social networks takes on an entirely different meaning. From this perspective, the goal is not to reduce forwarding but to reduce the cost of that forwarding in unnecessary compliance through message design or public education. Ongoing discussions about including more information beyond the current 90-character WEA message could be a route to do so; more characters and the inclusion of hyperlinks or other data in messages could minimize the negative effects of forwarding while maintaining its potential value as a compliance enhancer.

In sum, the ability to transmit messages to smaller areas — which is indeed a major technical jump in emergency alerting capability — requires similar innovation in policy and practice to ensure that emergency managers are outfitted with the understanding and tools needed to make the best time-limited and high-stakes decisions during the warning phase of natural, technological or other emergency incidents.

# References


Archibald, E., & McNeil, S. (2012). Learning from traffic data collected before, during and after a hurricane. *IATSS Research*, *36*(1), 1–10.

Bambenek, J., & Klus, A. (2008). Do emergency text messaging systems put students in more danger? *Educause Quarterly*, *31*(3), 12–16.

Bean, H., Sutton, J., Liu, B. F., Madden, S., Wood, M. M., & Mileti, D. (2015). The study of mobile public warning messages: A research review and agenda. *Review of Communication*, *15*, 60–80.

Benavides, A., & Arlikatti, S. (2010). The role of the Spanish-language media in disaster warning dissemination: An examination of the emergency alert system. *Journal of Spanish Language Media*, *3*, 41–58.

Berry, L. (1999) C*yclone Rona: February 11, 1999 evacuation of Caravonica and Lake Placid report.* Townsville City, Australia: Centre for Disaster Studies, James Cook University.

Brotzge, J., & Donner, W. (2013). The tornado warning process: A review of current research, challenges, and opportunities. *Bulletin of the American Meteorological Society*, *94*(11), 1715–1733.

Burnside, R., Miller, D. S., & Rivera, J. D. (2007). The impact of information and risk perception on the hurricane evacuation decision-making of Greater New Orleans residents. *Sociological Spectrum*, *27*(6), 727–740.

Casteel, M. A., & Downing, J. R. (2013). How individuals process NWS weather warning messages on their cell phones. *Weather, Climate, and Society*, *5*(3), 254–265.

CGDAW — *See* Committee on Geotargeted Disaster Alerts and Warnings.

Chae, J., Thom, D., Jang, Y., Kim, S., Ertl, T., & Ebert, D. S. (2014). Public behavior response analysis in disaster events utilizing visual analytics of microblog data. *Computers & Graphics*, *38*, 51–60.

Committee on Geotargeted Disaster Alerts and Warnings. (2013). *Geotargeted alerts and warnings: Report of a workshop on current knowledge and research gaps*. Washington, DC: National Academies Press.

Committee on Public Response to Alerts and Warnings Using Social Media. (2013). *Public response to alerts and warnings using social media: Report of a workshop on current knowledge and research gaps*. Washington, DC: National Academies Press.

Communications Security, Reliability and Interoperability Council. (2014, October). *Geographic targeting, message content and character limitation subgroup report*. Washington, DC: Federal Communications Commission. As of June 18, 2015:

http://transition.fcc.gov/pshs/advisory/csric4/CSRIC_CMAS_Geo-Target_Msg_Content_Msg_Len_Rpt_Final.pdf

comScore. (2015). U.S. digital future in focus 2015. Reston, VA: Adam Lella & Andrew Lipsman. As of June 22, 2015: http://www.comscore.com/Insights/Presentations-and-Whitepapers/2015/2015-US-Digital-Future-in-Focus

CPRAWSM — *See* Committee on Public Response to Alerts and Warnings Using Social Media.

CSRIC — *See* Communications Security, Reliability and Interoperability Council.

Department of Homeland Security. (2013a). *Study of integration strategy considerations for wireless emergency alerts*. Pittsburgh, PA. As of June 18, 2015: http://www.firstresponder.gov/TechnologyDocuments/Wireless%20Emergency%20Alerts%20Integration%20Strategy.pdf

Department of Homeland Security. (2013b). *Wireless emergency alerts: Computer model and simulation results*. Washington, DC. As of August 18, 2014: http://www.firstresponder.gov/TechnologyDocuments/Wireless%20Emergency%20Alerts%20Computer%20Model%20and%20Simulation%20Results.pdf

Department of Homeland Security. (2013c). *Wireless Emergency Alerts mobile penetration strategy*. Washington, DC. As of August 18, 2014: http://www.firstresponder.gov/TechnologyDocuments/Wireless%20Emergency%20Alerts%20Mobile%20Penetration%20Strategy.pdf

Department of Homeland Security. (2014). *Comprehensive testing of imminent threat public messages for mobile devices*. College Park, MD: National Consortium for the Study of Terrorism and Responses to Terrorism. As of June 18, 2015: http://www.firstresponder.gov/TechnologyDocuments/Comprehensive%20Testing%20of%20Imminent%20Threat%20Public%20Messages%20for%20Mobile%20Devices.pdf

DHS — *See* Department of Homeland Security.

Donner, W. R., Rodriguez, H., & Diaz, W. (2012). Tornado warnings in three Southern states: A qualitative analysis of public response patterns. *Journal of Homeland Security and Emergency Management*, *9*(2), Article 5.

Drabek, T. E. (1999). Understanding disaster warning responses. *The Social Science Journal*, *36*(3), 515–523.

Durage, S. W., Kattan, L., Wirasinghe, S. C., & Ruwanpura, J. Y. (2014). Evacuation behaviour of households and drivers during a tornado. *Natural Hazards*, *71*(3), 1495–1517.

Eisenman, D. P., Cordasco, K. M., Asch, S., Golden, J. F., & Glik, D. (2007). Disaster planning and risk communication with vulnerable communities: Lessons from Hurricane Katrina. *American Journal of Public Health*, *97*(Supplement 1), S109–S115.Federal Communications Commission. (2014). Wireless Emergency Alerts (WEA). As of April 30, 2015: http://www.fcc.gov/guides/wireless-emergency-alerts-wea

Federal Emergency Management Agency. (2015). Basic architecture description of the integrated public alert and warning system and its components. As of June 19, 2015: https://www.fema.gov/media-library/assets/documents/105491

Freberg, K. (2012). Intention to comply with crisis messages communicated via social media. *Public Relations Review*, *38*(3), 416–421.

Gao, Liang, Song, C., Gao, Z., Barabási, A-L., Bagrow, J. P., & Wang, D. (2014). Quantifying information flow during emergencies. *Scientific Reports*, *4*(3997), 1–6.

Gow, G. A., McGee, T., Townsend, D., Anderson, P., & Varnhagen, S. (2009). Communication technology, emergency alerts, and campus safety. *Technology and Society Magazine, IEEE*, *28*(2), 34–41.

Hamill, L., & Gilbert, N. (2010). Simulating large social networks in agent-based models: A social circle model. *Emergence: Complexity and Organization*, *12*(4), 78–94.

Hughes, A. L., & Palen, L. (2009). Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management*, *6*(3), 248–260.

Jagtman, H. M. (2010). Cell broadcast trials in the Netherlands: Using mobile phone technology for citizens' alarming. *Reliability Engineering & System Safety*, *95*(1), 18–28.

Legates, D. R., & Biddle, M. D. (1999). *Warning response and risk behavior in the Oak Grove–Birmingham, Alabama, Tornado of 08 April 1998*. Natural Hazards Center Quick Response Report 116. Boulder, CO: Natural Hazards Research and Applications Information Center.

Litman, T. (2006). Lessons from Katrina and Rita: What major disasters can teach transportation planners. *Journal of Transportation Engineering*, *132*, 11–18.

Liu, S., Murray-Tuite, P., & Schweitzer, L. (2012). Analysis of child pick-up during daily routines and for daytime no-notice evacuations. *Transportation Research Part A: Policy and Practice*, *46*(1), 48–67.

McCarty, C., Killworth, P. D., Bernard, H. R., Johnsen, E. C., & Shelley, G. A. (2001). Comparing two methods for estimating network size. *Human Organization*, *60*(1), 28–39.

Mersham, G. (2010). Social media and public information management: The September 2009 tsunami threat to New Zealand. *Media International Australia, Incorporating Culture & Policy*, (137), 130–143.

Mileti, D. S., & Darlington, J. D. (1995). Societal response to revised earthquake probabilities in the San Francisco Bay Area. *International Journal of Mass Emergencies and Disasters*, *13*(2), 119–145.

Mileti, D. S., & Peek, L. (2000). The social psychology of public response to warnings of a nuclear power plant accident. *Journal of Hazardous Materials*, *75*(2), 181–194.

Mileti, D. S., & Sorensen, J. H. (1990, August). Communication of emergency public warnings: A social science perspective and the state of the art assessment. ORNL-6609. Ok Ridge, TN: Oak Ridge National Laboratory.

Mitchell, J. T., Edmonds, A. S., Cutter, S. L., Schmidtlein, M., McCarn, R., Hodgson, M. E., & Duhé, S. (2005). *Evacuation behavior in response to the Graniteville, South Carolina, chlorine spill*. Boulder: Natural Hazards Center, University of Colrado.

Mitchem, J. D. (2003). *An analysis of the September 20, 2002, Indianapolis tornado: Public response to a tornado warning and damage assessment difficulties*. Quick Response Research Report #161. Boulder: Natural Hazards Research and Applications Information Center, University of Colorado. As of June 22, 2015: http://www.colorado.edu/hazards/qr/qr161/qr161.html

Moore, L. K. (2010, December 14). The Emergency Alert System (EAS) and all-hazard warnings. Washington, DC: Congressional Research Service, RL32527.

Morrow, T. B., Stoddard, R. J., II, & Elm, J. P. (2014, February). Wireless Emergency Alerts: Trust model simulations. Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, CMU/SEI-2013-SR-026.

Nagele, D. E., & Trainor, J. E. (2012). Geographic specificity, tornadoes, and protective action. *Weather, Climate, and Society*, *4*(2), 145–155.

National Center for Education Statistics. (2006). *A first look at the literacy of America's adults in the 21st century*. Washington, DC: U.S. Department of Education, NCES 2006-470. As of June 21, 2013: http://nces.ed.gov/naal/pdf/2006470.pdf

Parker, D. J., & Handmer, J. W. (1998). The role of unofficial flood warning systems. *Journal of Contingencies and Crisis Management*, *6*(1), 45–60.

Parker, D. J., Priest, S. J., & Tapsell, S. M. (2009). Understanding and enhancing the public's behavioural response to flood warning information. *Meteorological Applications*, *16*(1), 103–114.

Parker, D. J., Tunstall, S. M., & McCarthy, S. (2007). New insights into the benefits of flood warnings: Results from a household survey in England and Wales. *Environmental Hazards*, *7*(3), 193–210.

Pries, R., Hobfeld, T., & Tran-Gia, P. (2006, May). On the suitability of the short message service for emergency warning systems. In *Vehicular Technology Conference, 2006. VTC 2006-Spring. IEEE 63rd* (Vol. 2, pp. 991–995). Melbourne: IEEE. As of June 22, 2015: http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=1682973

Quarantelli, E. L. (1982, June). People's reactions to emergency warnings. University of Delaware Disaster Research Center, Preliminary Paper #75.

Rice, S., Trafimow, D., Hunt, G., & Keller, D. (2010). We don't dislike the alarm: We dislike complying with the alarm. *International Journal of Technology, Knowledge and Society*, *6*(1), 93–105.

Roberts, S. G., Dunbar, R. I., Pollet, T. V., & Kuppens, T. (2009). Exploring variation in active network size: Constraints and ego characteristics. *Social Networks*, *31*(2), 138–146.

Rogers, G. O., Watson, A. P., Sorensen, J. H., Sharp, R. D., & Carnes, S. A. (1990, April). *Evaluating protective actions for chemical agent emergencies.* Oak Ridge, TN: Oak Ridge National Laboratory, ORNL-6615.

Schultz, D. M., Gruntfest, E. C., Hayden, M. H., Benight, C. C., Drobot, S., & Barnes, L. R. (2010). Decision making by Austin, Texas, residents in hypothetical tornado scenarios. *Weather, Climate, and Society*, *2*(3), 249–254.

Sherman-Morris, K. (2010). Tornado warning dissemination and response at a university campus. *Natural Hazards*, *52*(3), 623–638.

Sorensen, J. H. (1991). When shall we leave? Factors affecting the timing of evacuation departures. *International Journal of Mass Emergencies and Disasters*, *9*(2), 153–165

Sorensen, J. H. (2000). Hazard warning systems: Review of 20 years of progress. *Natural Hazards Review*, *1*(2), 119–125.

Sorensen, J. H., Shumpert, B. L., & Vogt, B. M. (2004). Planning for protective action decision making: Evacuate or shelter in place. *Journal of Hazardous Materials*, *109*(1–3), 1–11.

Southeast Louisiana Hurricane Taskforce; Survey Research Center, UNO; Center for Hazards Assessment, Response and Technology (CHART), UNO; and Department of Geography, UNO. (2005). Citizen hurricane evacuation behavior in southeastern Louisiana: A twelve parish survey. University of New Orleans, Survey Research Center Publications, Paper 10. As of June 23, 2015: http://scholarworks.uno.edu/src_pubs/10

Sutter, D., & Erickson, S. (2010). The time cost of tornado warnings and the savings with storm-based warnings. *Weather, Climate, and Society*, *2*(2), 103–112.

Sutton, J., Palen, L., & Shklovski, I. (2008). Backchannels on the front lines: Emergent uses of social media in the 2007 Southern California wildfires. *Proceedings of the 5th International ISCRAM Conference — Washington, DC, USA, May 2008*. Eds. F. Fiedrich & B. Van de Walle.

Sutton, J., Spiro, E. S., Johnson, B., Fitzhugh, S., Gibson, B., & Butts, C. T. (2014). Warning tweets: Serial transmission of messages during the warning phase of a disaster event. *Information, Communication & Society*, *17*(6), 765–787. As of June 23, 2015: http://dx.doi.org/10.1080/1369118X.2013.862561

Traynor, P. (2010). Characterizing the security implications of third-party emergency alert systems over cellular text messaging services. In Sushil Jajodia & Jianying Zhou (Eds.), *Security and Privacy in Communication Networks* (pp. 125–143). Berlin: Springer.

U.S. Census Bureau. (2010). *2010 American community survey, 1-year and 5-year estimates*. Washington, DC.

U.S. Census Bureau. (2015, May 28). State & county QuickFacts: Dubuque (city), Iowa. As of June 23, 2015: http://quickfacts.census.gov/qfd/states/19/1922395.html

U.S. Government Accountability Office. (2007, March). Emergency preparedness: Current emergency alert system has limitations, and development of a new integrated system will be challenging. Washington, DC, GAO-07-411.

Vieweg, S., Hughes, A. L., Starbird, K., & Palen, L. (2010, April 10–15). Microblogging during two natural hazards events: What Twitter may contribute to situational awareness. CHI 2010: Crisis Informatics, Atlanta, Georgia.

Vogt, B. M., & Sorensen, J. H. (1999). *Description of survey data regarding the chemical repackaging plant accident West Helena, Arkansas*. Oak Ridge, TN: Oak Ridge National Laboratory, ORNL/TM-13722.

Watts, D. J. & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature*, *393*(4), 440–442.

Wood, V. T., & Weisman, R. A. (2003). A hole in the weather warning system: Improving access to hazardous weather information for deaf and hard of hearing people. *Bulletin of the American Meteorological Society*, *84*(2), 187–194.

Zachria, A., & B. Patel. (2006). Deaths related to Hurricane Rita and mass evacuation. *Chest*, *130*(4), (meeting abstracts), 124S-c–124S.

# Appendix A. Model Summary

## A.1. Purpose

The model in this project uses the Multi-Agent Platform for Interactive Environments (MAGPIE) framework, written in Java. MAGPIE itself is built atop the Multi-Agent Simulator of Neighborhoods (MASON) agent-based modeling tool kit developed at George Mason University. MAGPIE represents agents as a set of finite state machines (FSMs) — one each for navigation/movement; sensing/collecting information from the environment; communication; and command and control/decision making. Message passing between the FSMs provides a synchronized and coherent view of the world for the agent and implements the actions resulting from decision making. The model examines the diffusion of alerts within an area populated by individual recipients who move through a simulated road network as they execute their daily routines. Parameters controlling agent behavior and alert propagation were varied to study their effects on the diffusion of geo-targeted alerts.

## A.2. Entities

The model contained two primary entities: the emergency and the citizen. The emergency agent was simple in design and had one primary function: generate a signaling message with details about himself or herself to send to forwarding authorities at the appointed time. The citizens were more complex. Citizens moved through the model with a combination of steering behaviors to determine the desired direction and an acceleration allowance that determined how much the agent could change his or her speed in each time step of the model. At the beginning of a model run, agents moved from their initial positions (residence locations) toward their work locations.[19] Agents made choices about their behavior by choosing among a set of goals executed in the command and control FSM.

## A.3. Processes

The process of communications within the model was triggered by the emergency agent — a single agent who corresponds to the government alert originator. The model ran in time steps of one second.

Once the model started, each agent could receive communications through the various channels and send out communications (by forwarding to their social networks). The communications FSM maintained an inbox and outbox structure, which managed messages received and transmitted by each agent. Information included in communications updated the agent's assessment of what he or she should do, determining whether the agent should change from pursuing the baseline goal of "go to work" to another available goal.

---

[19] For simplicity in structuring the model, the only reason that agents move routinely about the model is transit from their homes to work locations. Work could be defined broadly across the agent population to include other daily activities as well, however.

In all the modeled scenarios, the earliest point at which this happened was 20 minutes (or 1,200 time steps) into the simulation. At the point the alert was issued, the model defined one or more polygons on the simulated emergency area (the footprint of the event that would result in harm to the agents) and the alert area(s) where geo-targeted Wireless Emergency Alerts (WEA) messages were sent. WEA messages were sent to all agents inside the geo-targeted area at the point when the alert was issued and to all other agents who subsequently entered the alert area (i.e., as they crossed into the alert polygon, they received the WEA message). As is the case for real-life WEA messages, once an agent has received the message once, he or she does not receive any additional direct (i.e., non-forwarded) versions of that WEA alert for the remainder of the simulation (equivalent to the agent's phone not displaying a message that it has already displayed).

# A.4. Design Concepts

### A.4.1. Adaptation

Citizen agents had a defined set of goals that governed their behavior: idle, work, shelter and evacuate.

When the agent received a warning message through some of the simulated channels (media, technical forwarding or a forwarded message from one of those channels), geographic information was communicated about the location of the emergency area, and the agent had the opportunity to compare his or her location with that area and decide whether the alert applied (and if it did not, to ignore the message). For WEA messages (or forwarded WEA messages), that geographic information was included only implicitly by the geo-targeting of the message itself; all recipients assumed that a received WEA or WEA-forwarded message was relevant for them.

Having applied the probabilities described in the body of the text with respect to receipt of, understanding of and action on a message, the agent's command and control FSM would determine whether the agent should (1) communicate the message by forwarding it to the agent's social network, and (2) change the agent's behavior. This change would be from its nonemergency state of either being idle (stationary at a destination) or going to work to either a *safety* goal (where the agent is executing the direction provided in the alert) or a *family safety* goal (where the agent is converging on the emergency area because a family member is present there). The family safety goal was only associated with subscribers to the technical forwarder, identified in the simulation as a school; the only agents who adopted that goal equated to parents going to their children's schools when notified of an emergency at that location. Once an agent changed his or her goal to either safety (i.e., evacuating or sheltering in place, depending on the direction the agent received) or family safety, that goal is *sticky* — i.e., the agent would not change again until the end of the simulation. The agent's decisions to comply with the message (i.e., change his or her goal) and forward the message were separate, so an agent could perform one or both actions upon receiving an alert.

### A.4.2. Objectives

Each agent's chief objective was to work toward the highest-ranked goal after all had been evaluated. The shelter and evacuate goals, once decided on, were permanently fixed as the agent's objective for the duration of the simulation. Although no special action was required for the shelter goal, the evacuation goal was considered met when the agent reached a designated evacuation node on the road network.

### A.4.3. Interaction

Citizen agents interacted with other entities in the environment through messages sent via defined communication channels that modeled broadcast media, such applications as Facebook and Twitter, and Short Message Service (SMS) and cellular communication.

### A.4.4. Stochasticity

Probabilistic draws were used throughout the model. At initialization, draws were used in assigning agent positions on the road network and in building the social network. During the model run, draws were used to determine the delay in message reception, whether or not it is understood, and whether it is forwarded via a specific channel.  Compliance to emergency directives was also determined probabilistically.

### A.4.5. Observation

Data generated by the model captured entity position, message traffic and goal execution.

The simulation recorded information about message traffic and agent position that made it possible to calculate measures associated with alert receipt (e.g., which agents received the WEA message based on where they were, which agents received different types of other messages, forwarding behavior) and agent risk (e.g., which agents were in the emergency zone at the time of the alert, how many were in the zone at the time the emergency occurred).

## A.5. Initialization

The model was initiated with agents placed at road intersections, proportional to population residence data from the 2000 U.S. Census for Dubuque, Iowa. Specifically, agent home and work locations were determined by using the Census 2000 Special Tabulation 64. This tabulation contains counts of workers in each census tract by tract of residence. (Here *worker* is defined as a person 16 years old or over who was employed during the last week of March 2000.) These counts were used to perform two probabilistic draws for each agent, selecting a home and work location. Agent social networks were generated at the beginning of the model run, using the Watts-Strogatz method, with 6, 42 or 216 as the median degree and the rewiring probability set to 0.20 (Watts and Strogatz, 1998). Initially, all agents were to prioritize the goal of “go to work,” since no emergency messages or other information existed at that point that would cause them to prioritize other goals.

# Appendix B. Validation Analysis of Alert Receipt by Alerting Channel and Forwarding Behavior

As much of our study design centered on the effects of forwarded alerts, we conducted a series of validation exercises to ensure that the model worked as expected. We sought to check for coherence in the model results, rather than estimate "on average" behavior, and the exercises are therefore based on single model runs (i.e., no replication). The validation analyses were conducted within the flash flood scenario, with an alerting zone size that was larger than the emergency zone, as well as a 70 percent phone compatibility rate (to provide greater opportunity to observe forwarding behavior).

Our strategy was to isolate each of the four channels that might be forwarded: government technical, media, school technical and Wireless Emergency Alerts (WEA). We turned on each channel individually (i.e., when examining each channel, all other channels were turned off) and then varied whether agent forwarding was turned on or off. This four-by-two set of possibilities resulted in eight model runs.

Table B.1 presents the number of agents not receiving and receiving the WEA message directly, based on the eight conditions above and whether the agent was ever in the emergency zone. For ease of interpretation, and because the focus is on WEA, we averaged across the four non-WEA channels. No agents received the WEA message outside the WEA channel. Furthermore, forwarding had no effect on direct WEA receipt. The number of agents receiving the WEA alert directly through the WEA channel also appropriately matched model parameters. Across the two WEA runs, the average observed in-zone receipt rate was approximately 73 percent. This is very close to the 70 percent phone compatibility rate. In principle, this should be reduced by the 60 percent receipt-probability parameter. Given that in-zone agents had multiple opportunities to receive a WEA message, and might have been in the zone for a period of time, these results seem reasonable. The outside-zone receipt rate was approximately 49 percent. This is harder to judge, since it depends on how many agents entered the alert zone (which is larger than the emergency zone).

*Table B.1. Validation Results for Receipt of WEA Message*

| Channel | Forwarding | In Emergency Zone? | No Receipt | Receipt |
|---|---|---|---|---|
| WEA | Off | Yes | 143 | 369 |
| WEA | Off | No | 4,821 | 4,667 |
| Other | Off | Yes | 481 | 0 |
| Other | Off | No | 9,519 | 0 |
| WEA | On | Yes | 119 | 341 |
| WEA | On | No | 4,921 | 4,619 |
| Other | On | Yes | 468 | 0 |
| Other | On | No | 9,530 | 0 |

Table B.2 captures the receipt of alerts through any channel (including interpersonal forwarding). Here we break out by all channels, as receipt for those is nonzero. As would be expected, there were positive numbers across the board. The receipt rates for all channels except WEA were similar inside and outside the emergency zone. The values here, by and large, appear reasonable. For example, media had a 100 percent subscription rate and a probability of receipt that varied from 5 percent to 33 percent. With forwarding turned off, the rate of receipt was about 10 percent, which, while lower than the midpoint, was within the expected range. Rates of receipt were greatly amplified when forwarding is turned on. This was especially true for WEA, and is in contrast to Table B.1.

*Table B.2. Validation Results for Receipt of Any Alert*

| Channel | Forwarding | In Emergency Zone? | No Receipt | Receipt |
|---|---|---|---|---|
| Government technical | Off | Yes | 387 | 71 |
| Government technical | Off | No | 8,322 | 1,220 |
| Media | Off | Yes | 444 | 47 |
| Media | Off | No | 8,538 | 971 |
| School technical | Off | Yes | 460 | 34 |
| School technical | Off | No | 8,954 | 552 |
| WEA | Off | Yes | 143 | 369 |
| WEA | Off | No | 4,821 | 4,667 |
| Government technical | On | Yes | 295 | 166 |
| Government technical | On | No | 5,967 | 3,572 |
| Media | On | Yes | 120 | 351 |
| Media | On | No | 2,806 | 6,723 |
| School technical | On | Yes | 391 | 86 |
| School technical | On | No | 7,774 | 1,749 |
| WEA | On | Yes | 15 | 445 |
| WEA | On | No | 629 | 8,911 |

# Appendix C. General Results by Scenario

This appendix presents scenario-by-scenario results, instead of using the policy-question format of Chapters Four through Six. A similar format is used for each scenario, leaving out inapplicable analyses (e.g., tracking the number of agents in the emergency zone has little meaning in the shelter-in-place tornado scenario) and supplementing analyses where useful. The general format is to present results as follows:

1) Wireless Emergency Alerts (WEA) specificity: the percentage of agents receiving the WEA message, both inside and outside the emergency zone;
2) WEA coverage: the number of agents receiving the WEA message, both inside and outside the emergency zone;
3) alert response: the number of agents in the emergency zone over time, where applicable, and the percentage of agents with evacuation or shelter goals, from among those inside the emergency zone; and
4) unnecessary alert response: the percentage of agents with evacuation or shelter goals, from among those outside the emergency zone.

For each set of analyses, we specify any assumptions (e.g., delay of receipt).

## C.1. Flash Flood Scenario

### C.1.1. Default Delay Parameters

The following results, presented in Figures C.1 through C.3, assumed the default delay parameters, as specified in Chapter Two. With these parameters, however, we observed no switching to evacuation goals until after the time of the emergency was passed. Because of this, we do no present results on goals here.

Figures C.1 and C.2 present the baseline results for two of our performance metrics, specificity and coverage. Figure C.1 shows the percentage of agents receiving the WEA message (directly from the originating agent) inside and outside the emergency zone, broken out by size of the alerting area and phone-compatibility rate. As expected, with an alerting zone that was just the emergency zone, a substantial proportion of agents inside the emergency zone received the WEA message over the course of the simulation, but no agents outside of the zone receive the WEA alert. With an alert zone larger than the emergency zone, we saw agents outside the zone receiving the WEA alert, as expected. Rates inside the zone approached the phone-compatibility rates.

*Figure C.1. Flash Flood: Percentage Receiving the WEA Message, Default Delay Parameters*

**30 Percent Phone Compatibility**

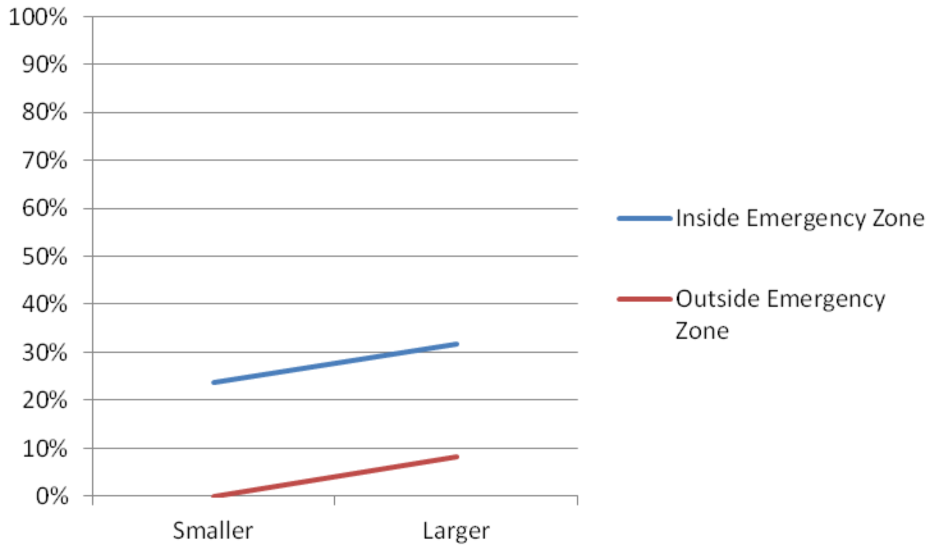


**70 Percent Phone Compatibility**

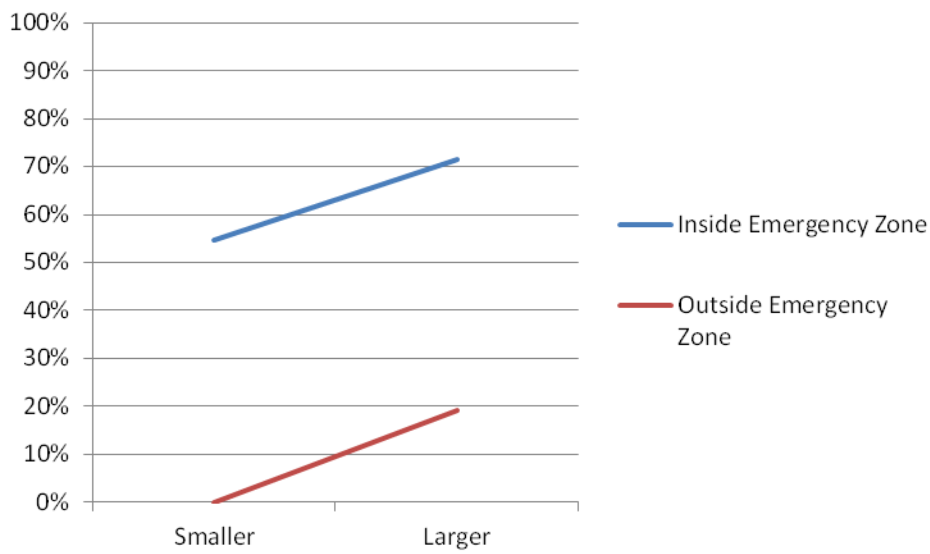


*NOTE: Smaller alert zones (left) correspond to the emergency zone, whereas larger alert zones (right) extend beyond the emergency zone.*

Figure C.2 shows a different view of the same data. Coverage is reflected in Figure C.2 by the number of agents receiving the WEA message. In contrast to Figure C.1, we can see in Figure C.2 that even though the percentage of agents receiving the alert was larger inside the emergency zone, the sheer numbers were much greater outside the emergency zone, since the total number in the zone was small.

The results in Figures C.1 and C.2 are replicated in other runs and scenarios, listed below, and act as a basic validity check on how the WEA messages were implemented in the model.

*Figure C.2. Flash Flood: Number Receiving the WEA Message, Default Delay Parameters*

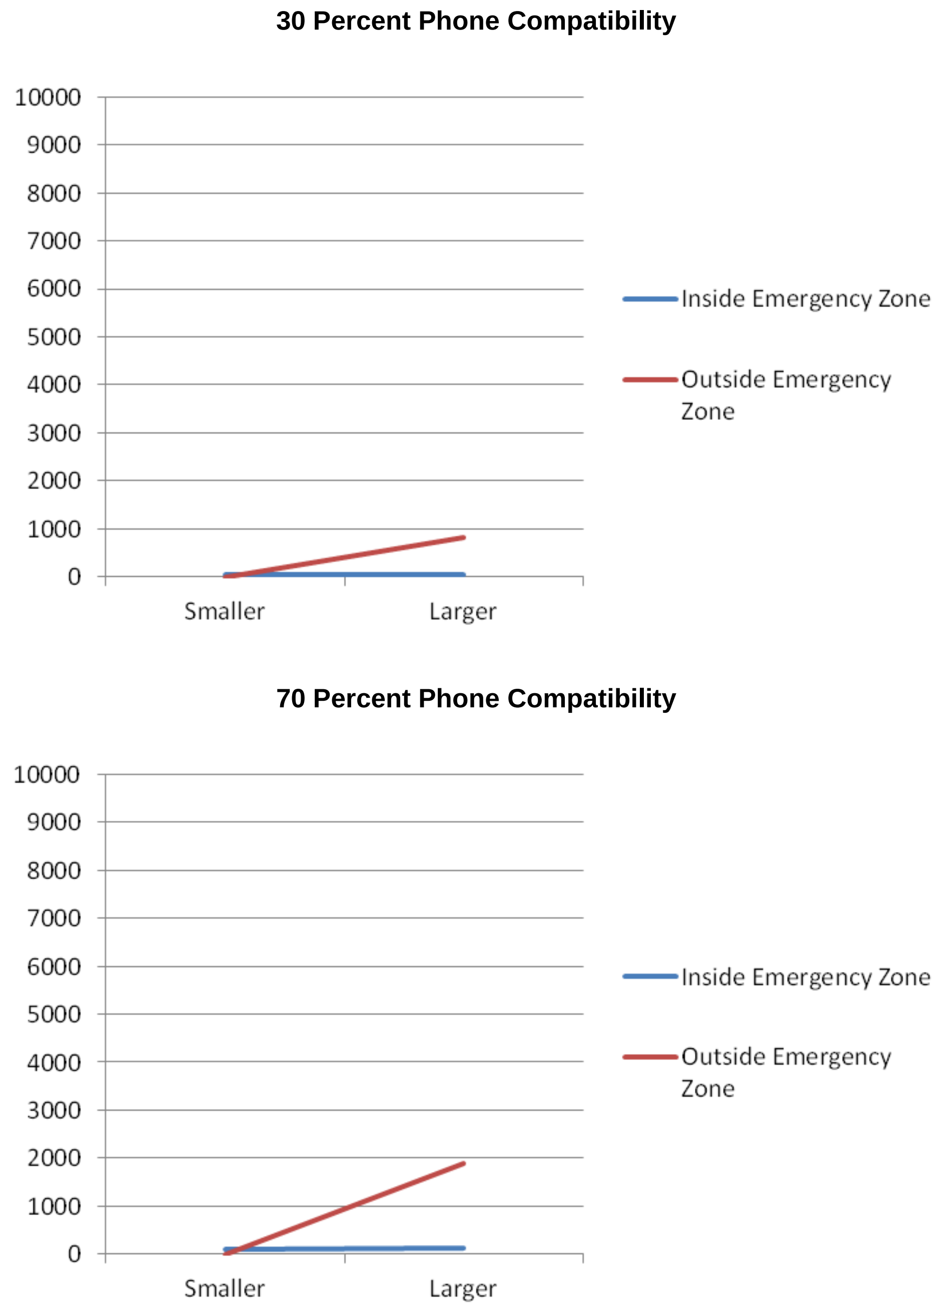


*NOTE: Smaller alert zones (left) correspond to the emergency zone, whereas larger alert zones (right) extend beyond the emergency zone.*

*Figure C.3. Number of Agents in the Flash Flood Emergency Zone over Time, Smaller Alert Zone, Default Delay Parameters*

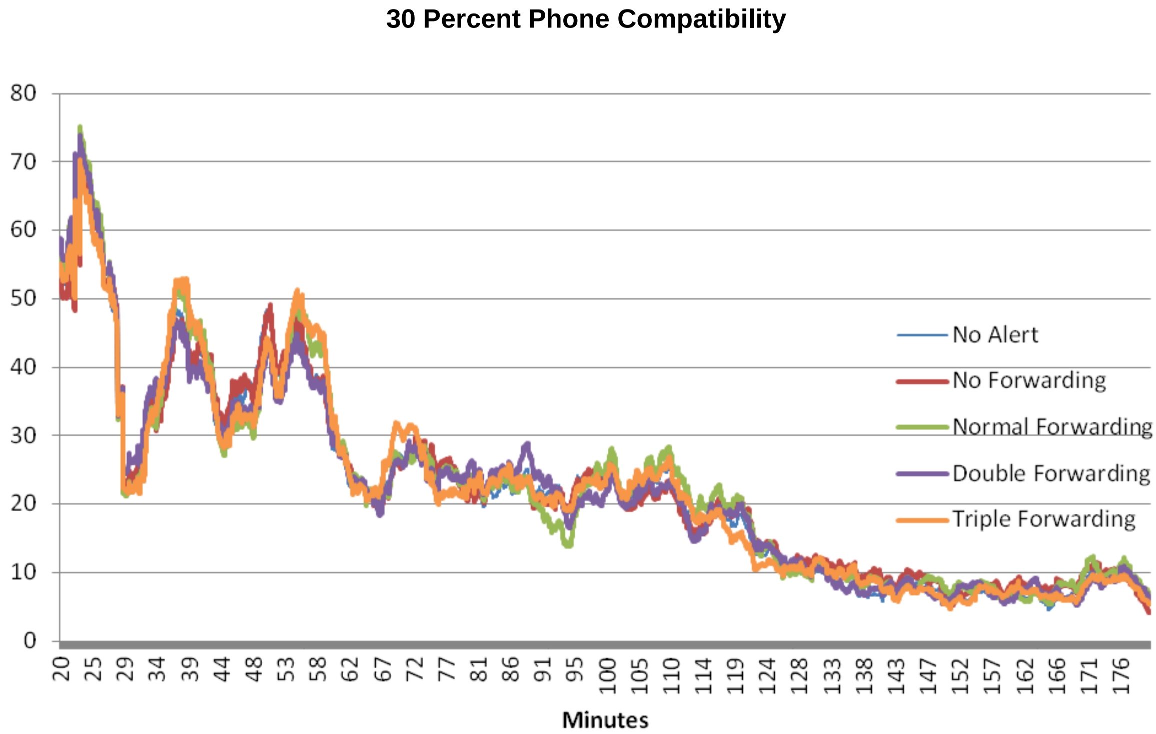


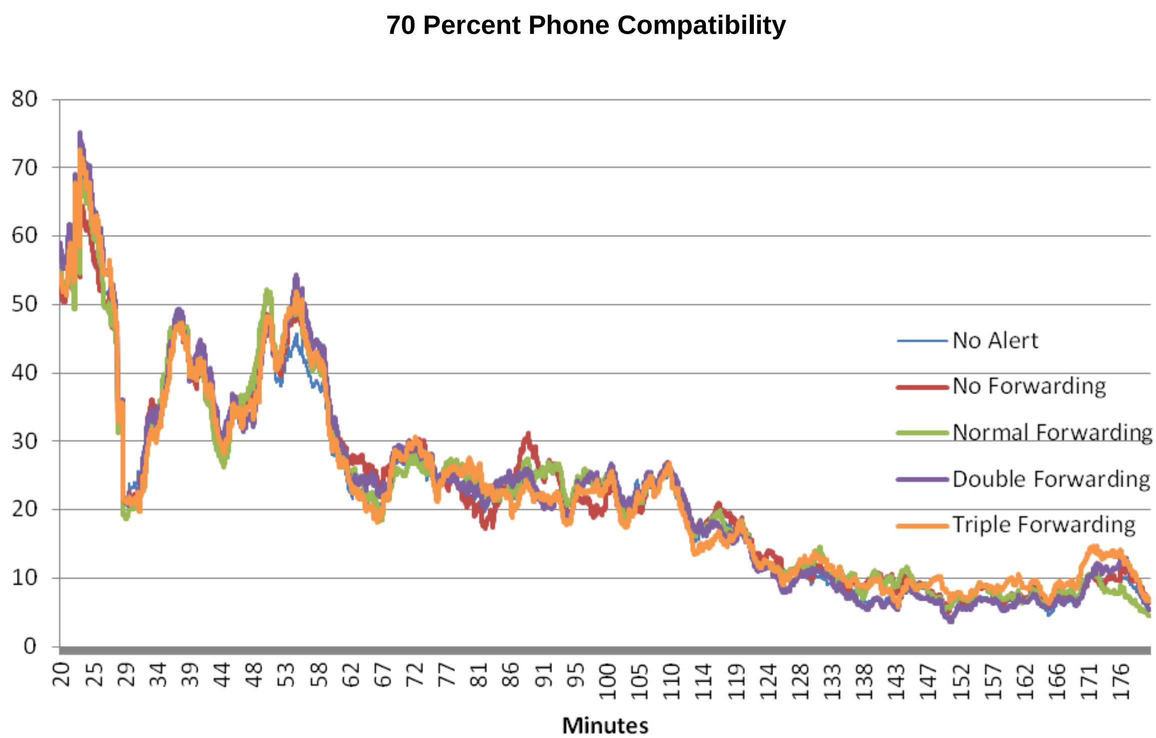


*NOTE: Results for larger alert zone appear in Figure 4.3.*

### C.1.2. Reduced Delay Parameters

Figures C.4 to C.7 present the results for the reduced delay parameters, as described in Chapter Four. Specifically, there was no delay in compliance. There was also no delay in receipt of WEA messages, and delay in receipt of interpersonal forwarding was reduced to two to three minutes. Delays in receipt for technical forwarding and media remained at their default levels. The percentage of agents with an evacuation goal among those outside the zone is presented in Figure 5.5.

*Figure C.4. Flash Flood, Percentage Receiving the WEA Message, Reduced Delay Parameters*

**30 Percent Phone Compatibility**

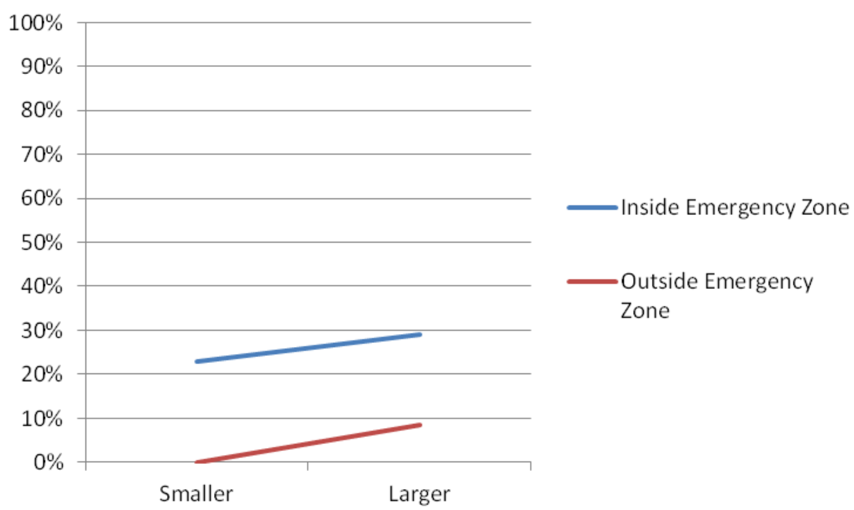


**70 Percent Phone Compatibility**

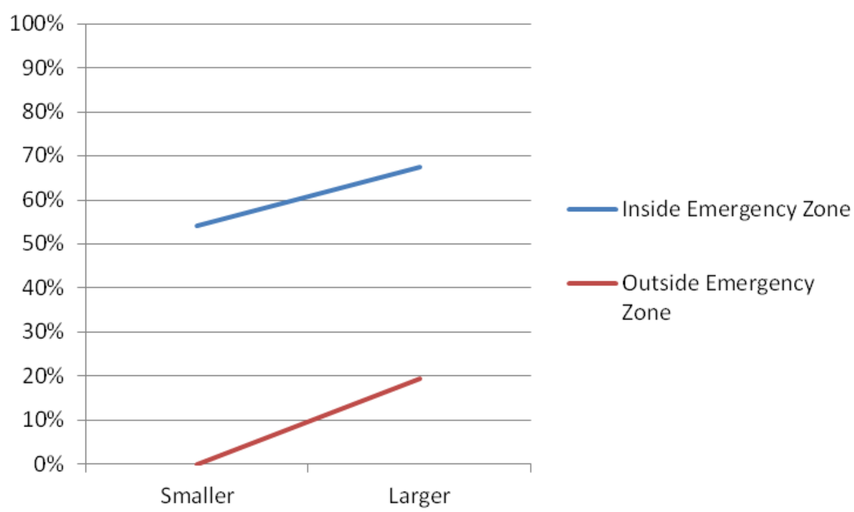


*NOTE: Smaller alert zones (left) correspond to the emergency zone, whereas larger alert zones (right) extend beyond the emergency zone.*

*Figure C.5. Flash Flood, Number Receiving the WEA Message, Reduced Delay Parameters*

**30 Percent Phone Compatibility**

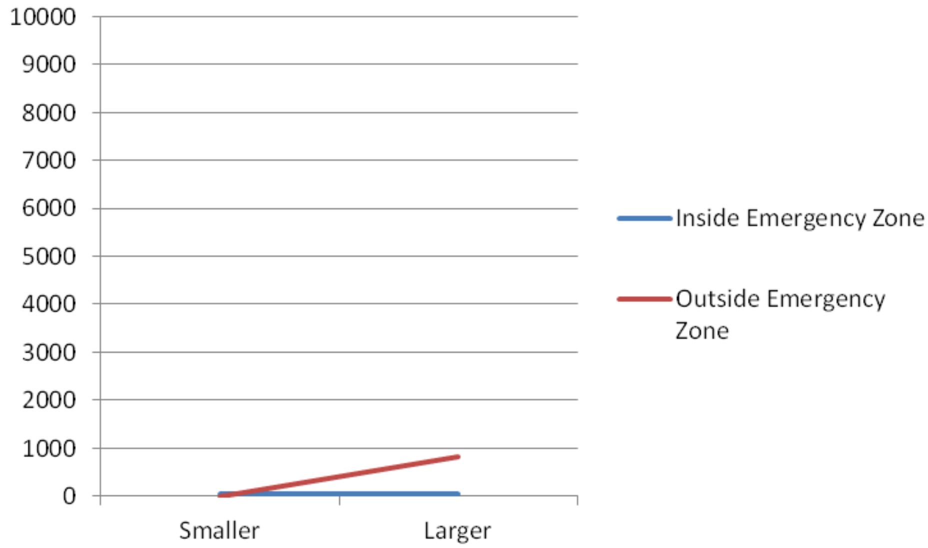


**70 Percent Phone Compatibility**

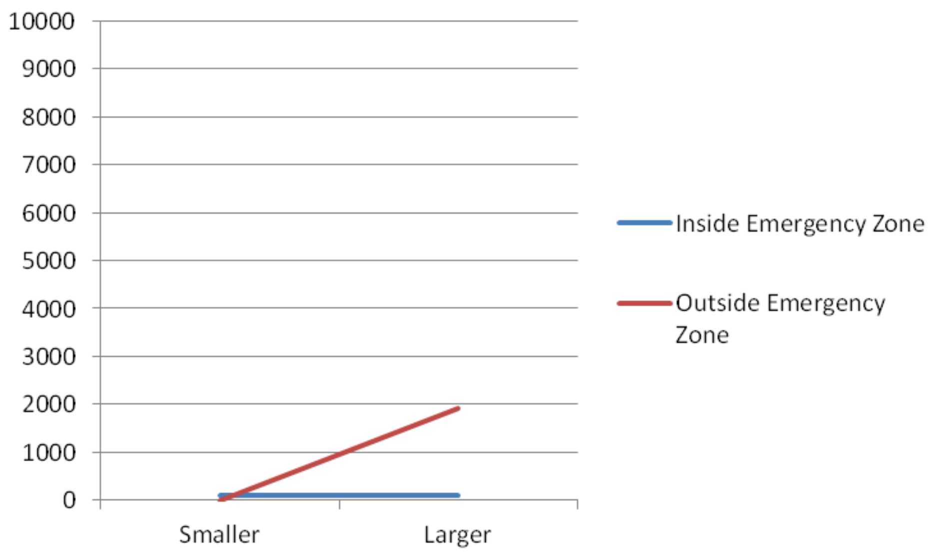


*NOTE: Smaller alert zones (left) correspond to the emergency zone, whereas larger alert zones (right) extend beyond the emergency zone.*

*Figure C.6. Number of Agents in the Flash Flood Emergency Zone over Time, Reduced Delay Parameters, Smaller Alert Zone*

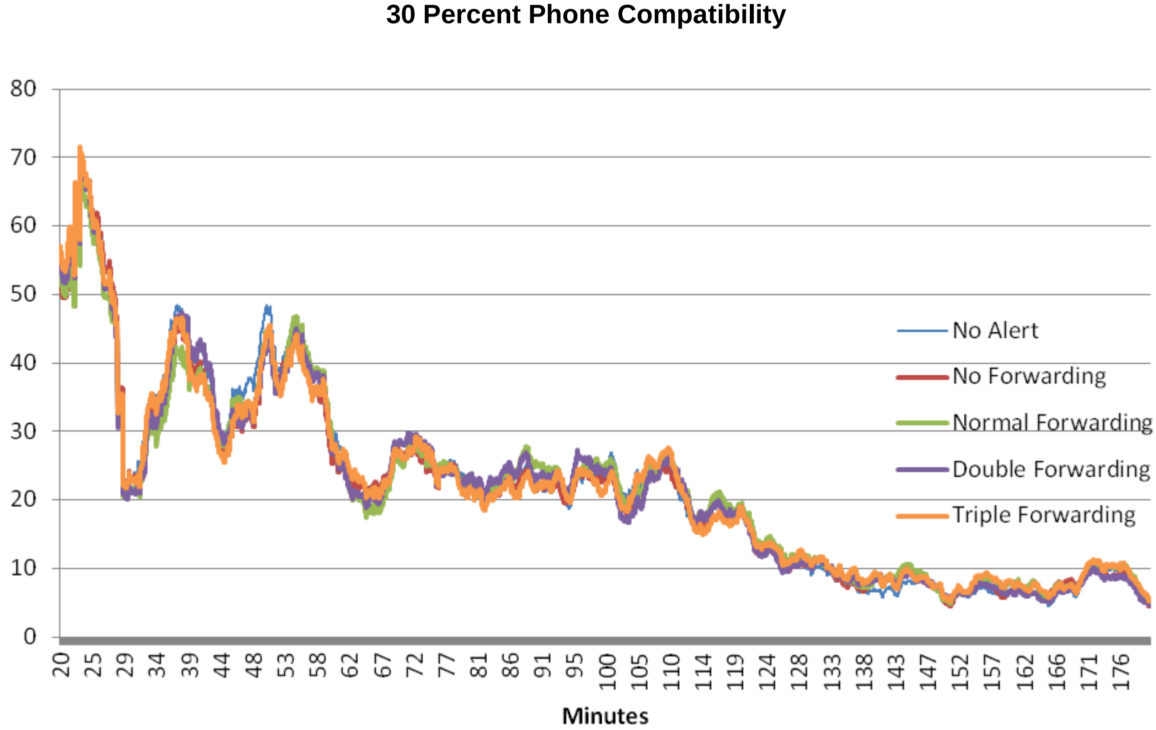


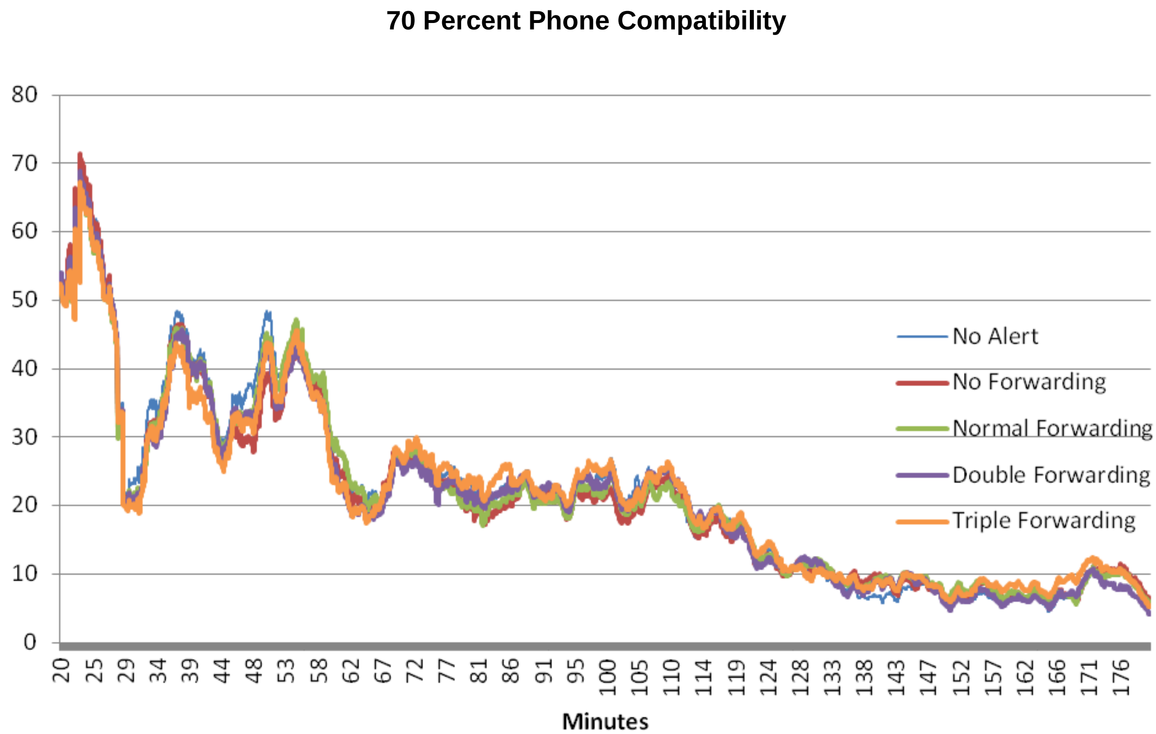


*NOTE: The larger zone is presented in Figure 4.4.*

*Figure C.7. Flash Flood: Percentage of Agents with an Evacuation Goal, from Among Those Inside the Emergency Zone at the Time of the Alert, Reduced Delay Parameters*

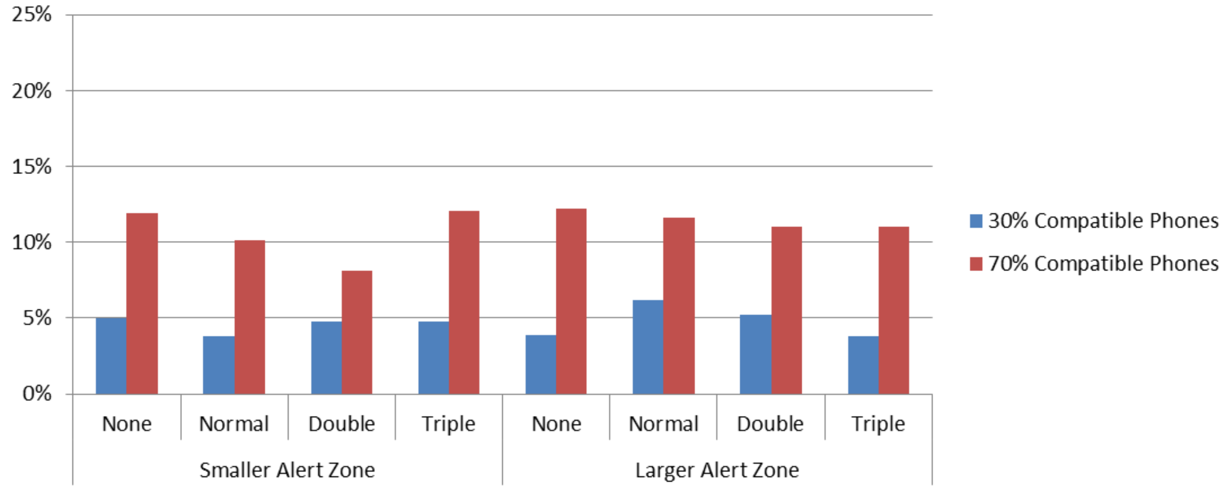


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases.*

*Figure C.8. Flash Flood: Percentage of Agents with an Evacuation Goal, from Among Those Entering the Emergency Zone Between Times of Alert and Emergency, Reduced Delay Parameters*

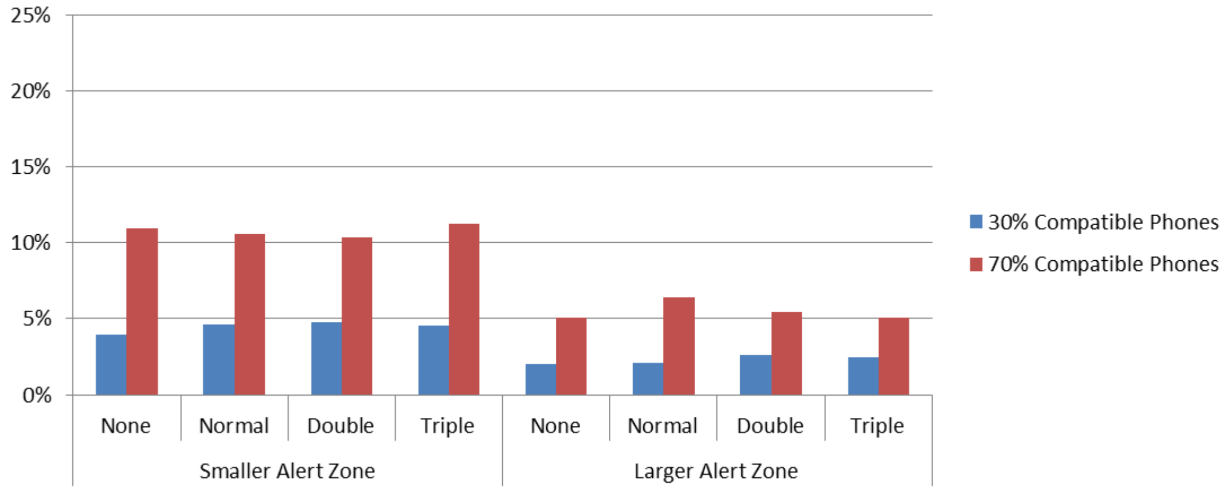


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases.*

### C.1.3. No Delay Parameters

Figures C.9 through C.12 present the results for no delay parameters, as described in Chapter Four. Specifically, there was no delay in compliance or receipt for any channel. Agents in the emergency zone over time are presented in Figures 4.5 and 4.6, whereas the goals of those always outside the zone are presented in Figure 5.5; those results are not replicated here.

*Figure C.9. Flash Flood: Percentage Receiving the WEA Message, No Delay Parameters*

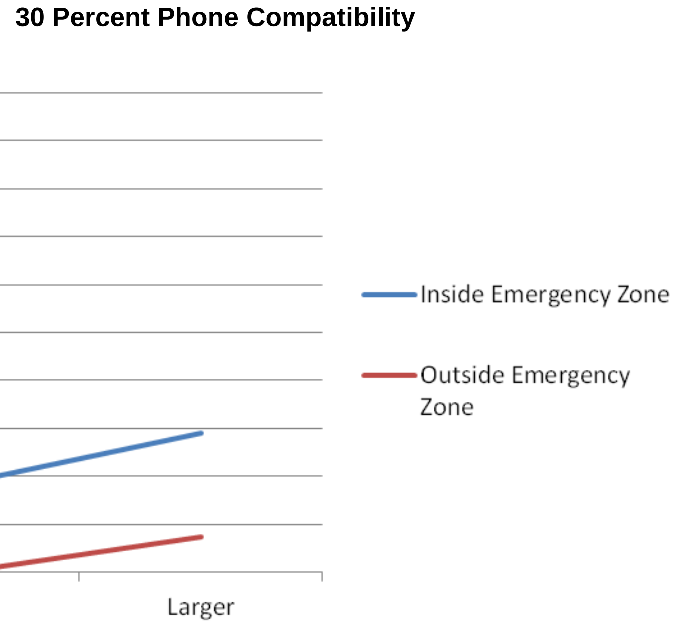


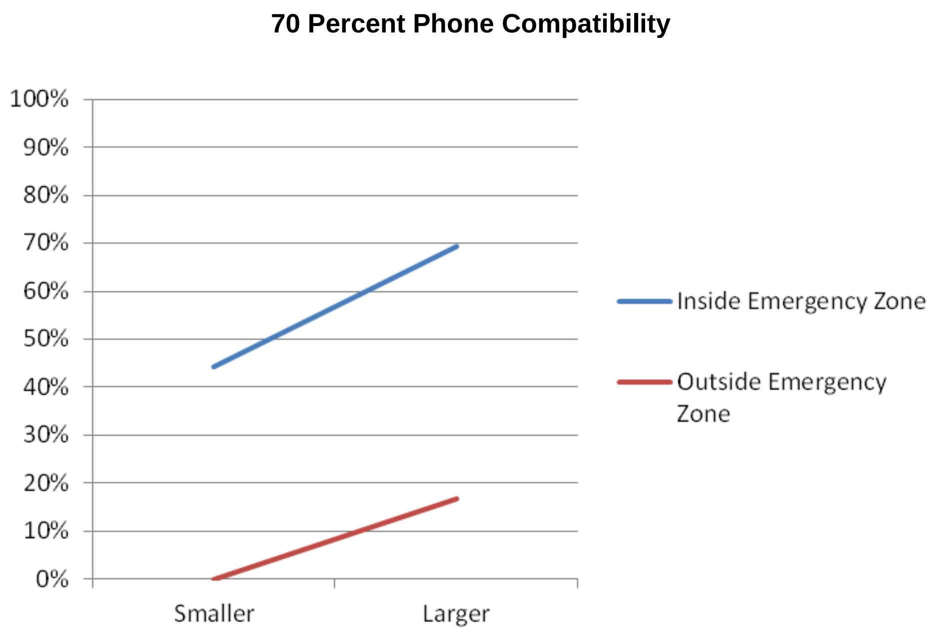


*NOTE: Smaller alert zones (left) correspond to the emergency zone, whereas larger alert zones (right) extend beyond the emergency zone.*

*Figure C.10. Flash Flood: Number Receiving the WEA Message, No Delay Parameters*

**30 Percent Phone Compatibility**

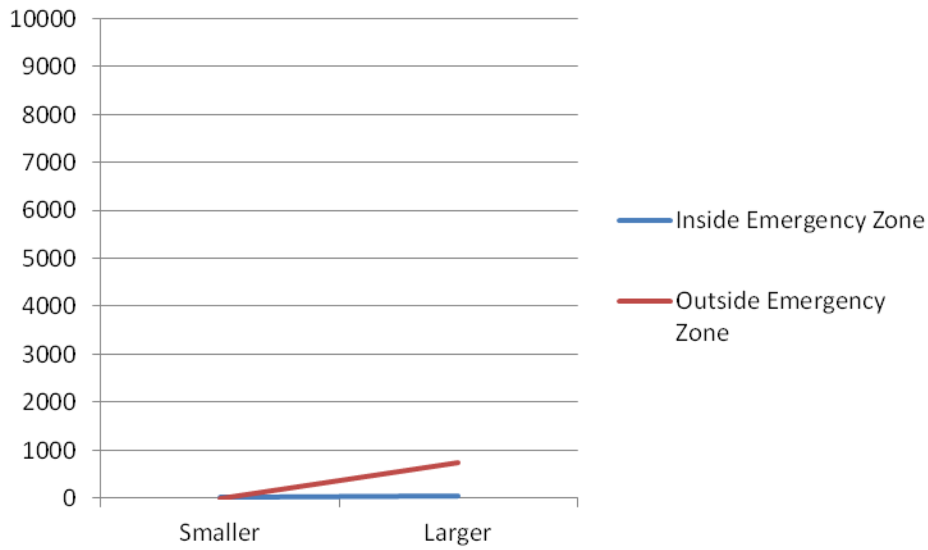


**70 Percent Phone Compatibility**

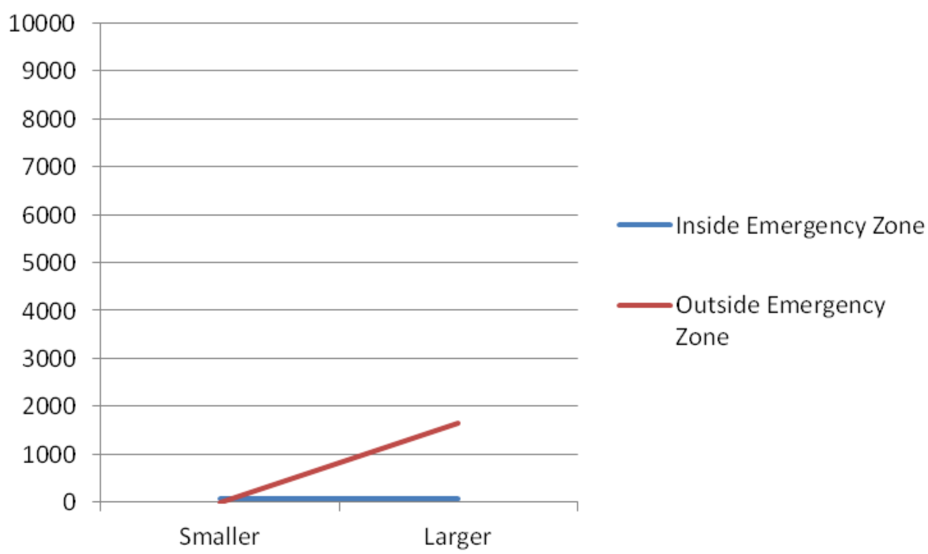


*NOTE: Smaller alert zones (left) correspond to the emergency zone, whereas larger alert zones (right) extend beyond the emergency zone.*

*Figure C.11. Flash Flood: Percentage of Agents with an Evacuation Goal, from Among Those Agents Inside the Emergency Zone at the Time of the Alert, No Delay Parameters*

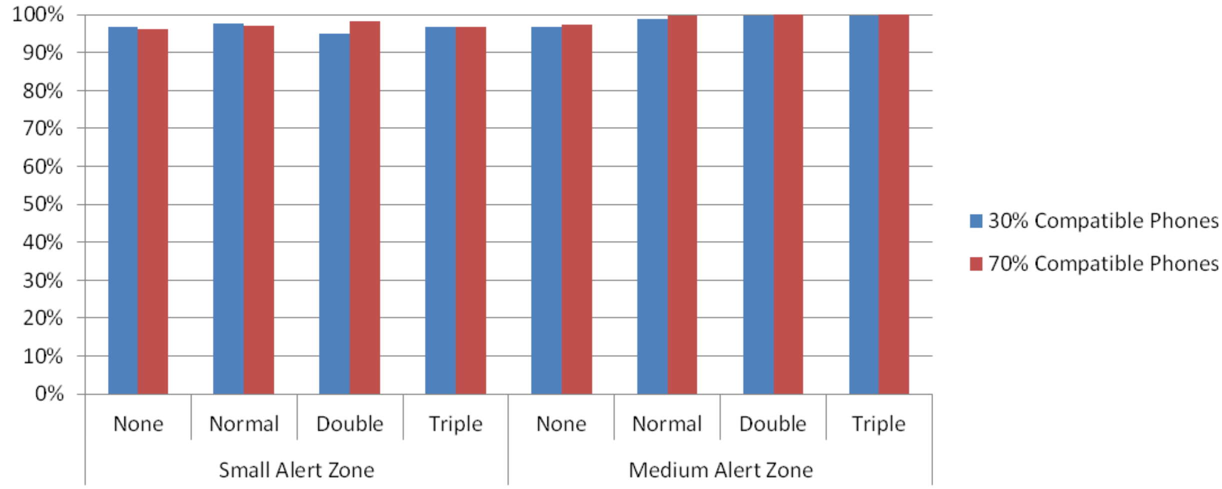


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases.*

*Figure C.12: Flash Flood, Percentage of Agents with an Evacuation Goal, from Among Those Agents Entering the Emergency Zone Between Times of Alert and Emergency, No Delay Parameters*

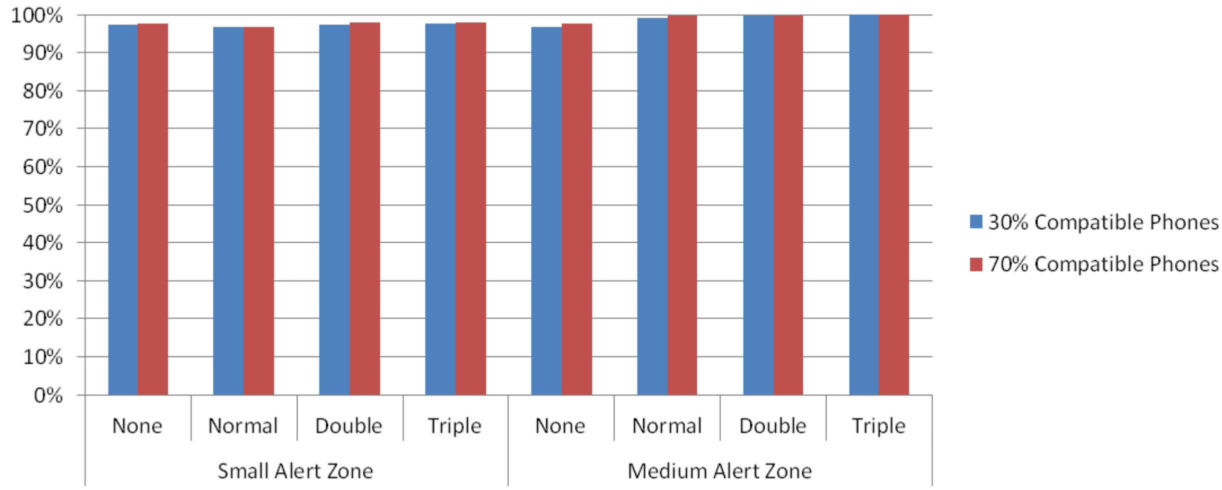


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases.*

# C.2. Tornado Scenario

We present in Figures C.13 through C.16 the percentage and number of agents receiving WEA for both the simulations, with default delay and reduced delay. As with the flash flood scenario, with the default delay parameters, we saw no agents switching to the shelter goal between the alert and emergency. Because of this, we do not present the goal results here for default delay. The data about agent goals by location are provided in Figures 5.2 and 5.3.

For each of the following sections, we do not produce the trace of the number of agents in the zone over time. This metric is less obviously relevant, given that the shelter-in-place directive in this scenario caused agents to stop where they were.

### C.2.1. Default Delay Parameters

*Figure C.13. Tornado: Percentage Receiving the WEA Message, Default Delay Parameters*

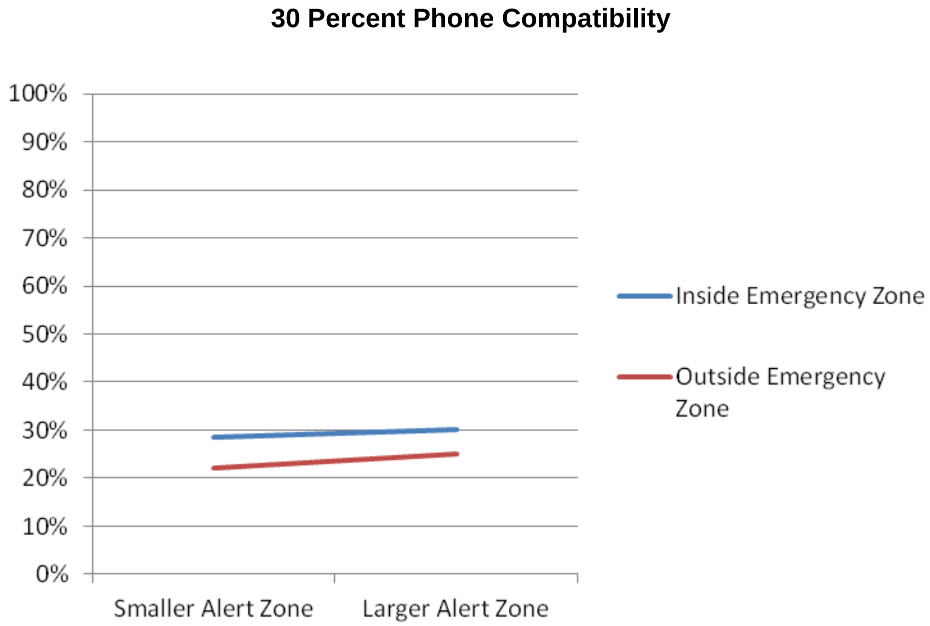


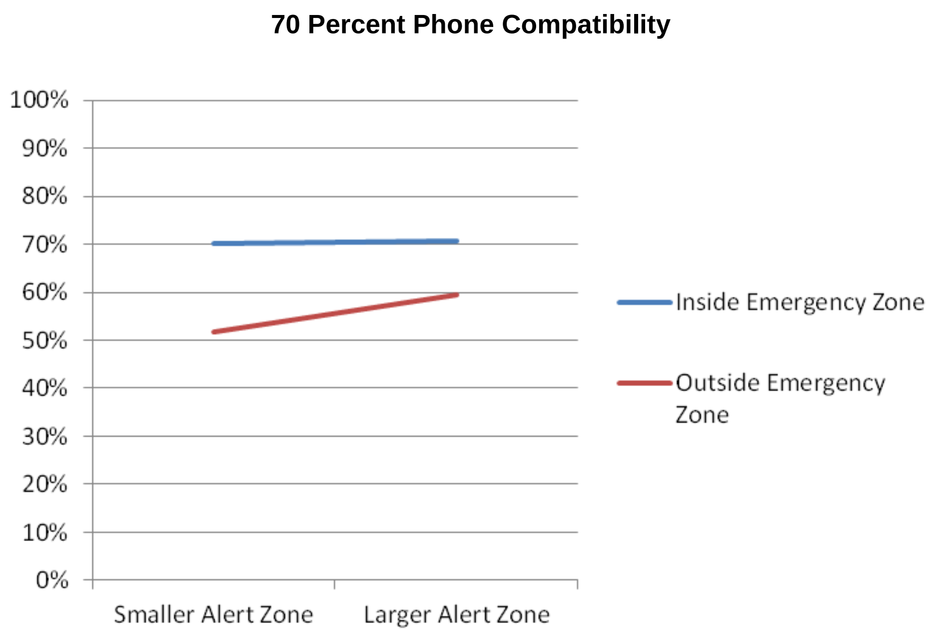

*Figure C.14. Tornado: Number Receiving the WEA Message, Default Delay Parameters*

**30 Percent Phone Compatibility**

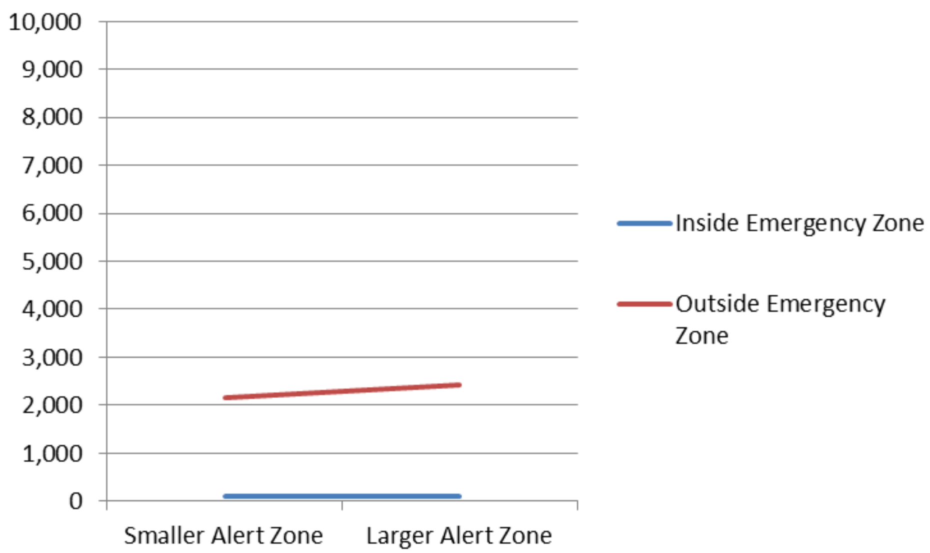


**70 Percent Phone Compatibility**

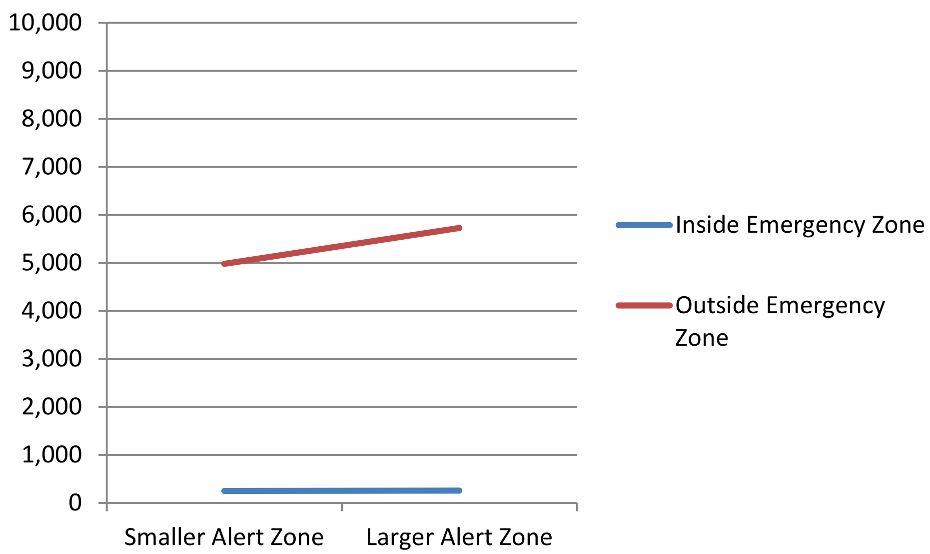

### C.2.2. Reduced Delay Parameters

*Figure C.15. Tornado: Percentage Receiving the WEA Message, Reduced Delay Parameters*

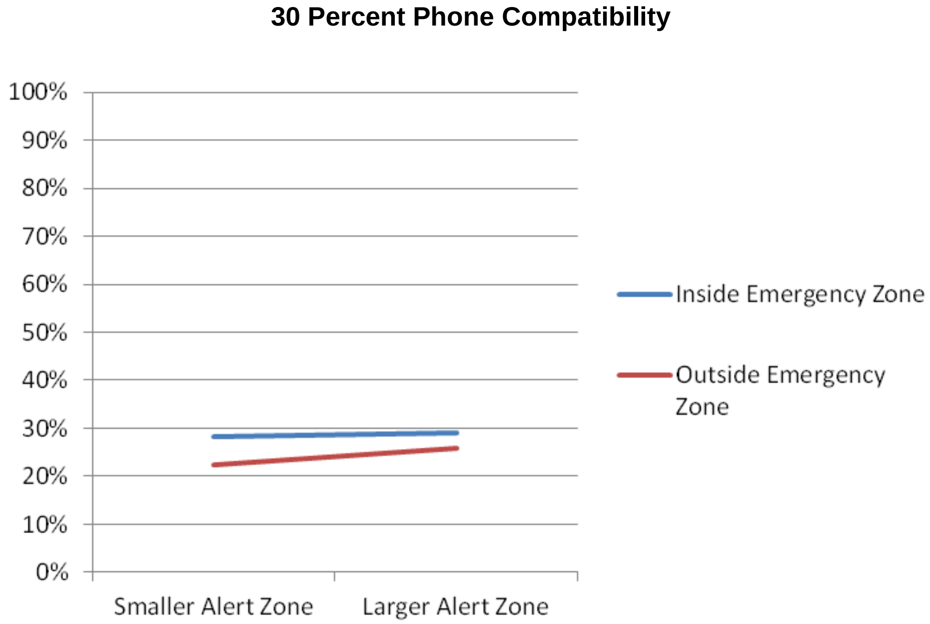


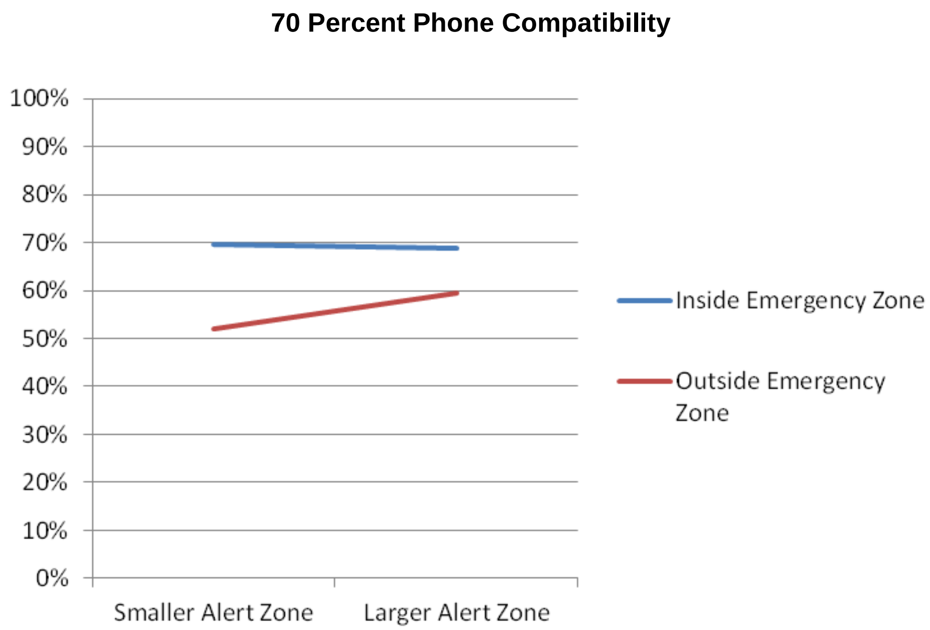

*Figure C.16. Tornado: Number Receiving the WEA Message, Reduced Delay Parameters*

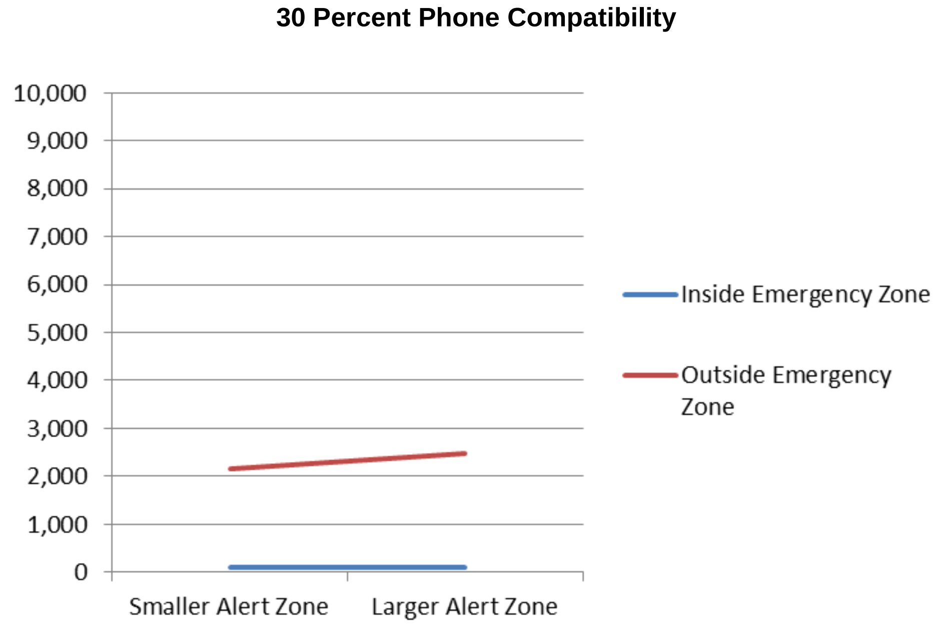


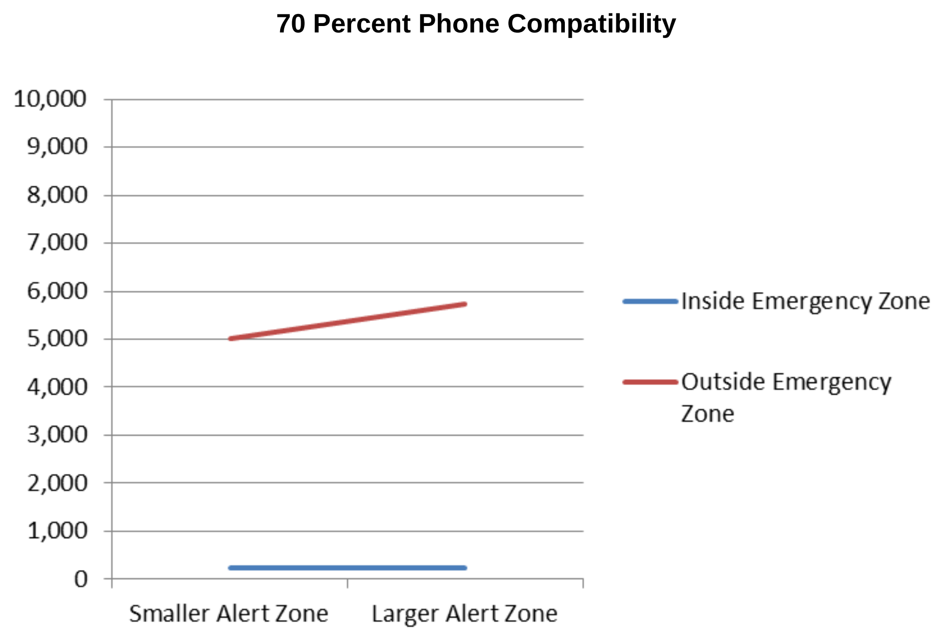


# C.3. Hazmat Scenario

Because there was only one alerting zone size, and because the combined zone maps onto the same area as the flash flood scenario, the percentage and number of agents receiving WEA message provide no additional information and are not presented.

### C.3.1. Default Delay Parameters

As with the flash flood and tornado scenarios, with default delay parameters, we did not observe any agents switching to an evacuation or shelter goal before the time of the emergency. Because of this, we do not present the percentage evacuating and sheltering. Figure C.17 presents the number agents in the emergency zone over time, both overall and broken out by evacuation and shelter zones, assuming the default delay parameters.

*Figure C.17. Number of Agents in the Hazmat Emergency Zone over Time, Default Delay Parameters*

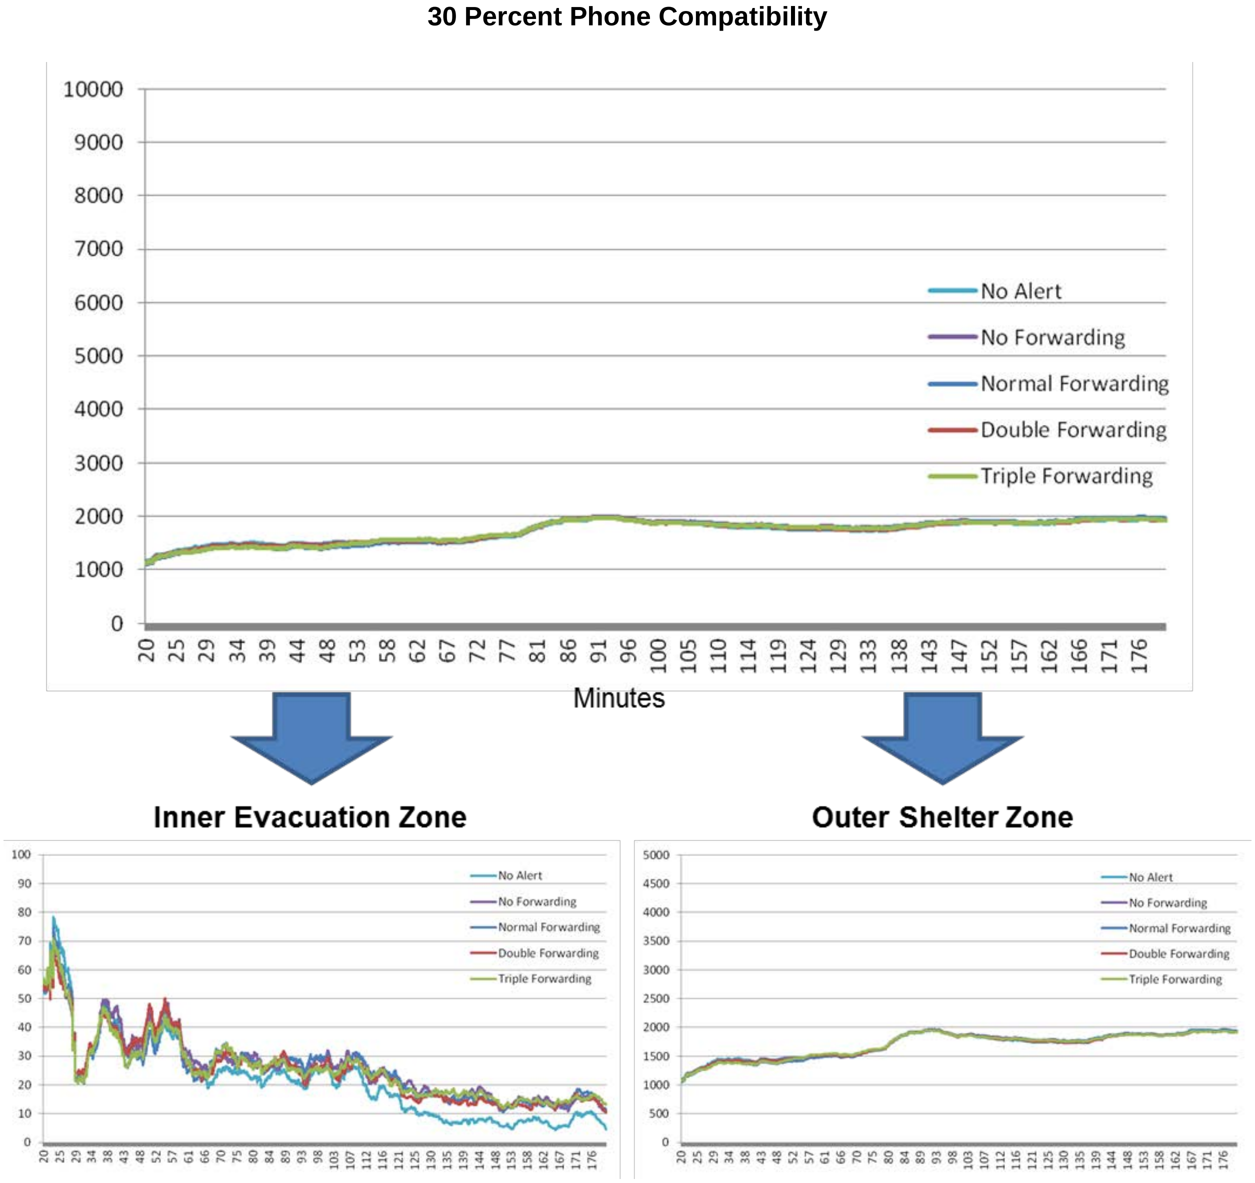

**70 Percent Phone Compatibility**

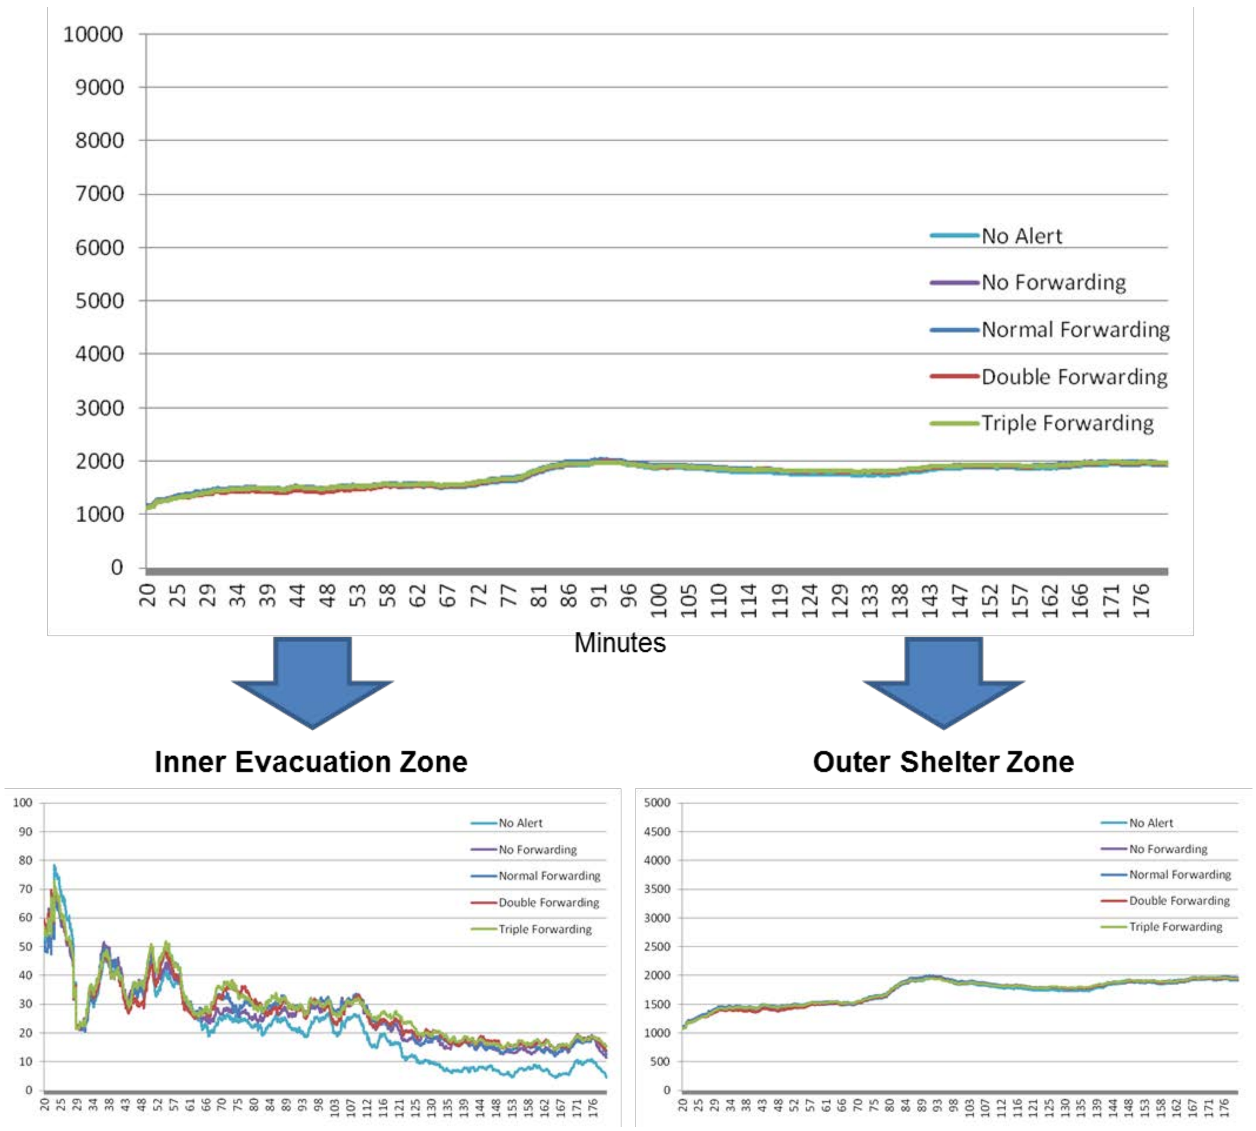


### C.3.2. Reduced Delay Parameters

The percentage of agents with evacuation and shelter goals, by zone, is presented in Figures 6.2 and 6.3. Figure C.18 presents the number agents in the emergency zone over time, both overall and broken out by evacuation and shelter zones, assuming the reduced delay parameters.

*Figure C.18. Number of Agents in the Hazmat Emergency Zone over Time, Reduced Delay Parameters*

**30 Percent Phone Compatibility**

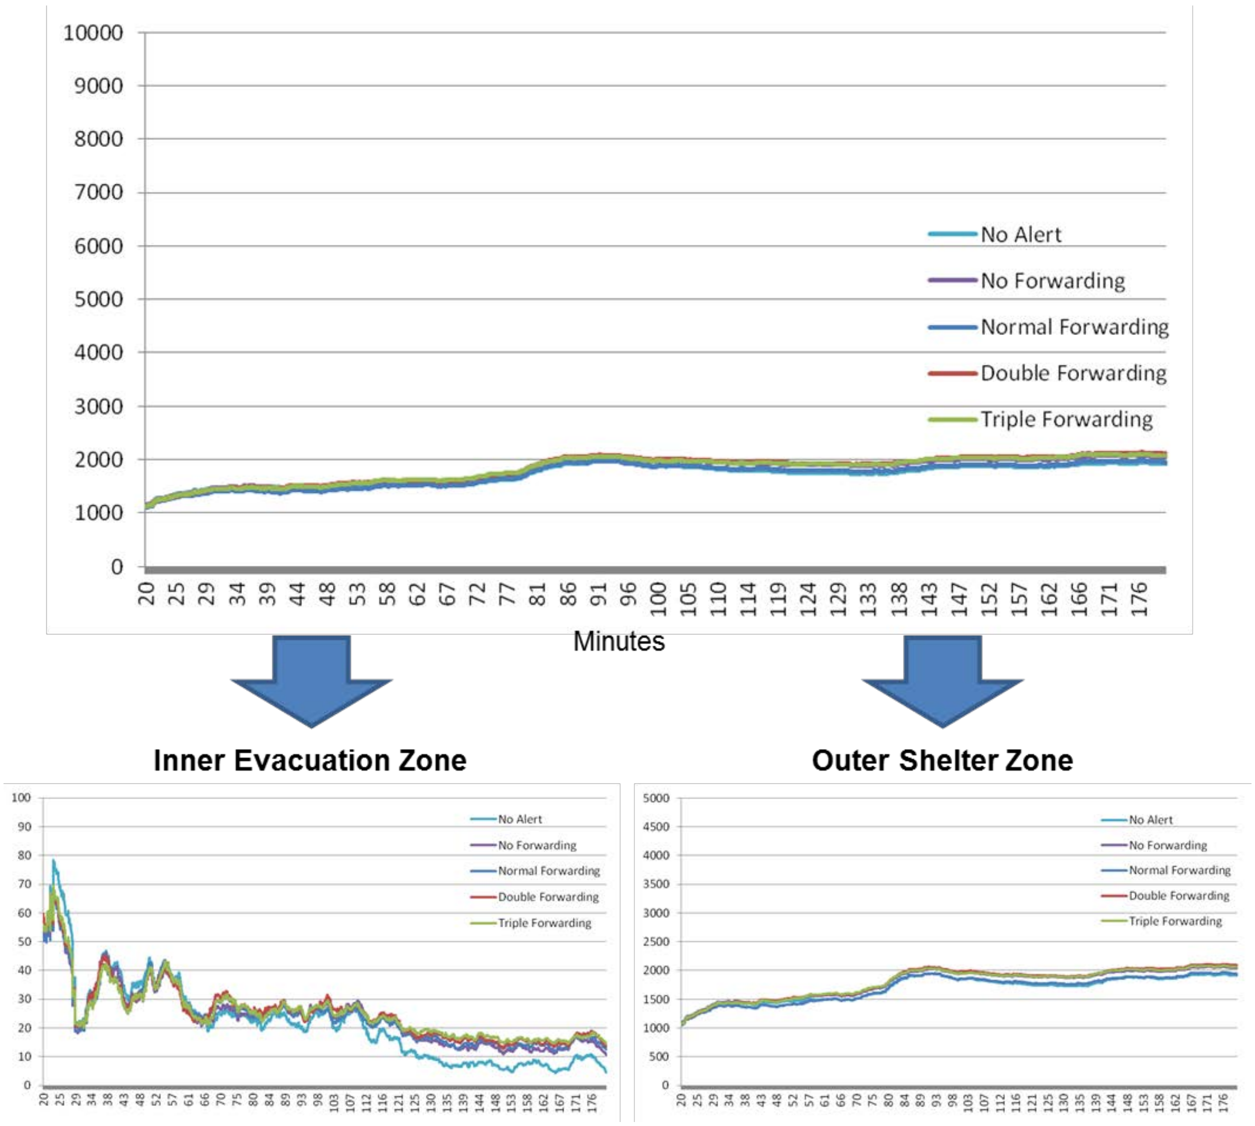

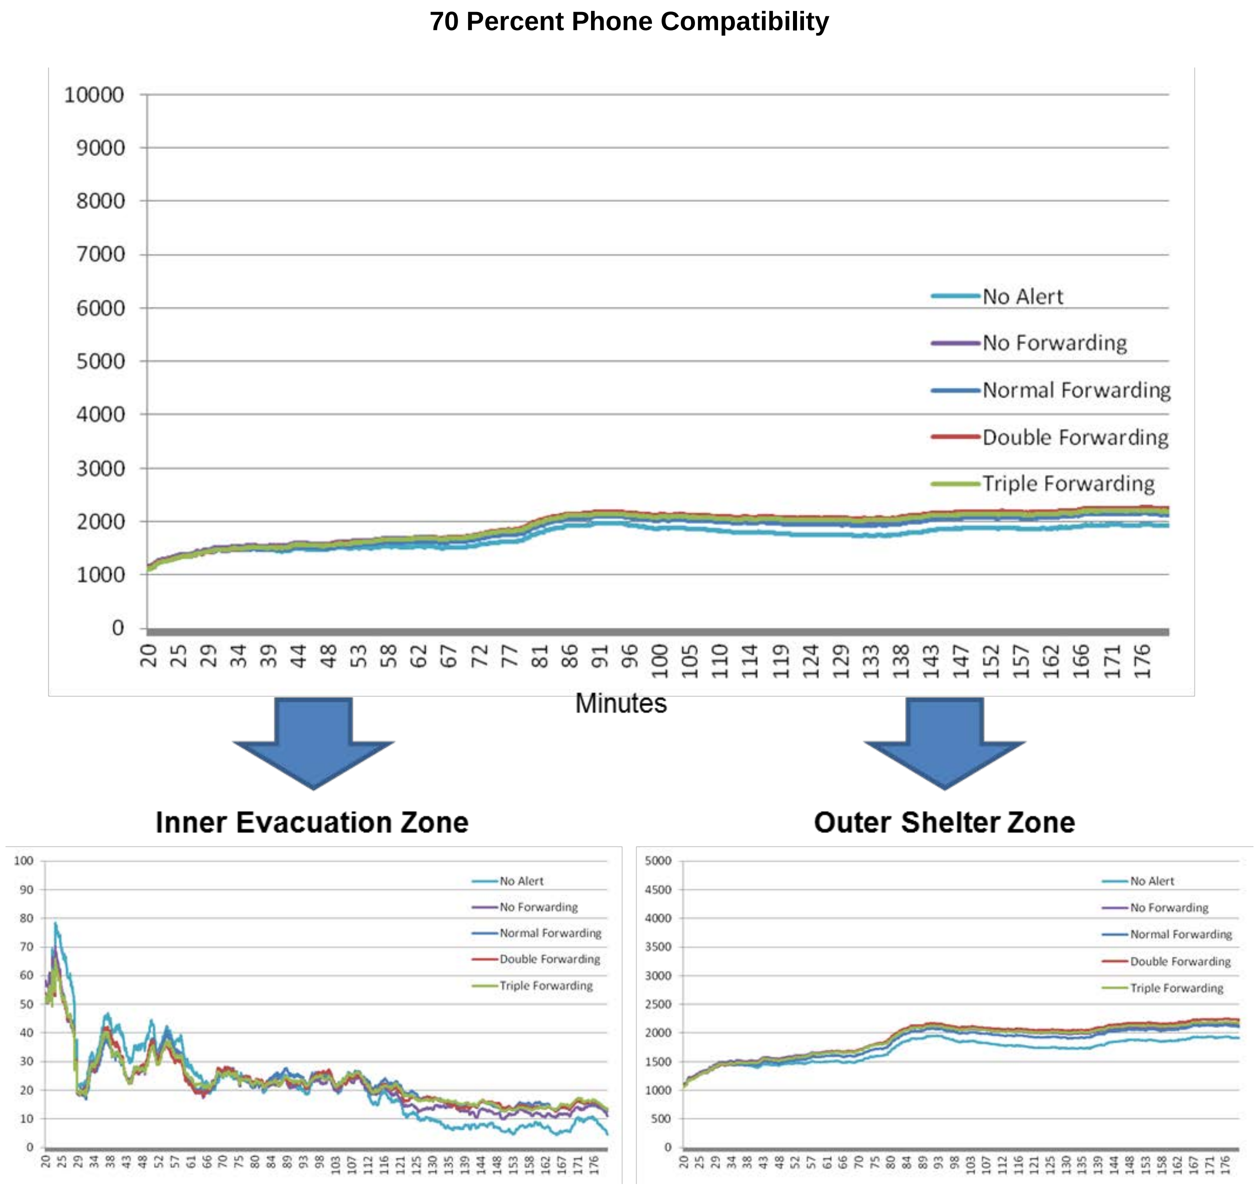


# C.4. Major Flood Scenario

The major flood scenario consisted of two zones; agents in each zone were alerted and told to evacuate at different times (outer zone first, then inner zone). There was no reason to expect differential WEA receipt across these zones and the alerting area did not extend beyond these zones; therefore, data about specificity and coverage are is not as relevant in this scenario.

## C.4.1. Default Delay Parameters

The number of agents in the zone is presented in Figures 4.7. Goal data are presented in Figures C.19 and C.20.

*Figure C.19. Major Flood: Percentage Evacuating Between First and Second Alerts, Default Delay Parameters*

**Agents Starting Inside the Inner Zone (*Not Yet Directed* to Evacuate)**

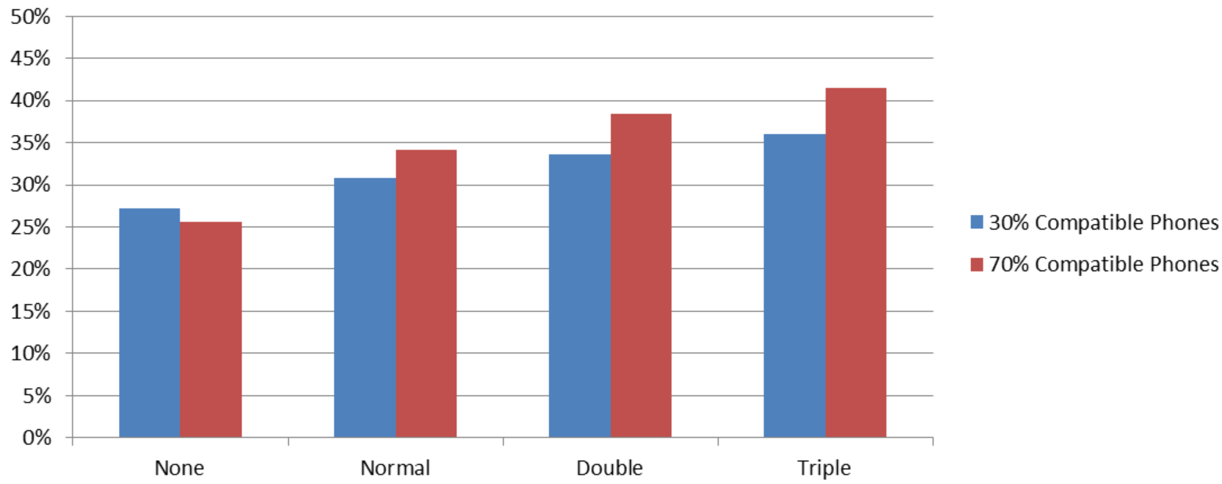


**Agents Starting Inside the Outer Zone (*Directed* to Evacuate)**

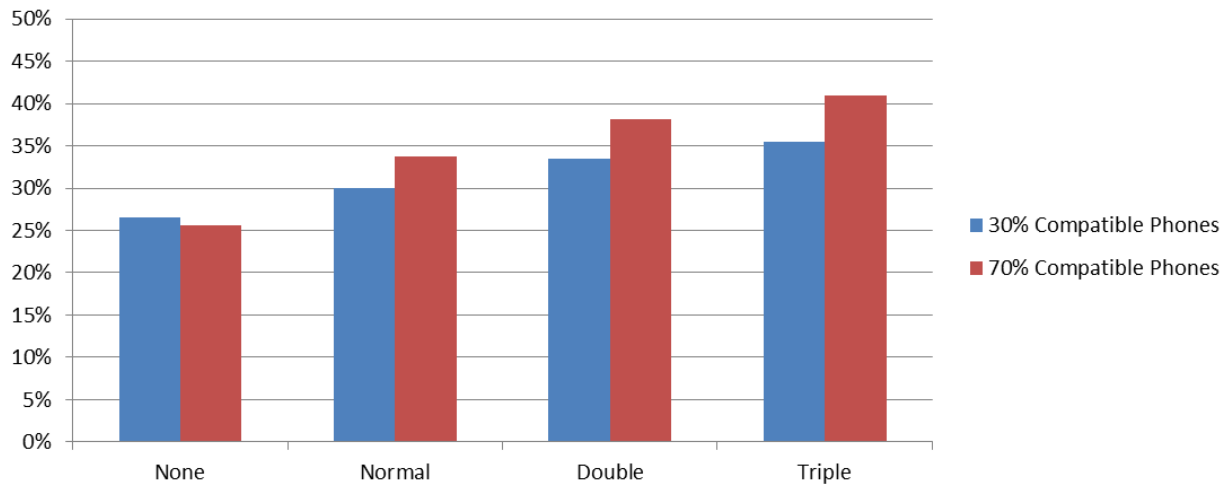


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases.*

*Figure C.20. Major Flood, Percentage Evacuating Between Second Alert and Emergency, Default Delay Parameters*

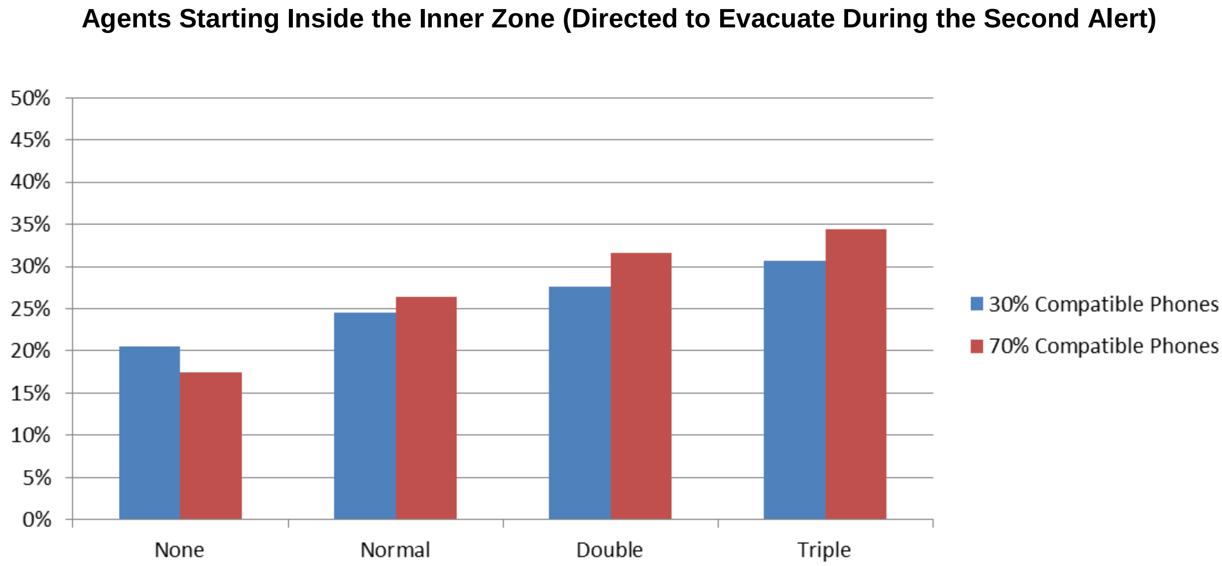


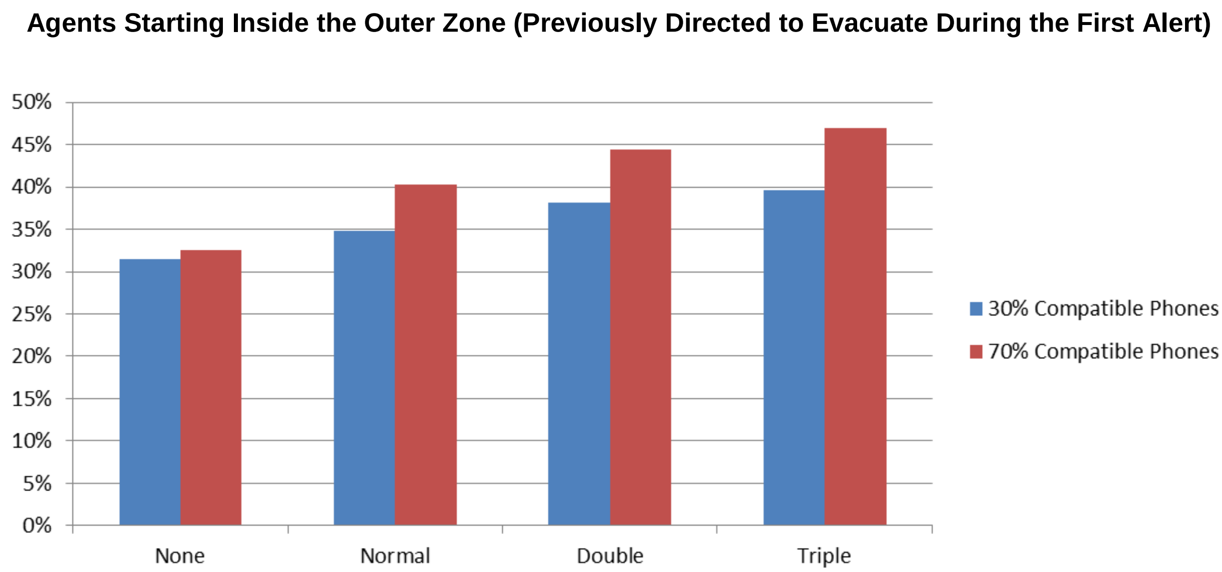


*NOTE: Bars distinguish among no, normal, double or triple forwarding cases.*

### C.4.2. Reduced Delay Parameters

The number of agents in the combined alerting zones is presented in Figure 4.8. The percentage of agents with an evacuation goal, broken down by zone and time period, is presented in Figures 6.5 and 6.6.

# Appendix D. Acronyms List

| | |
|---|---|
| ABM | Agent-Based Modeling |
| CPRAWSM | Committee on Public Response to Alerts and Warnings Using Social Media |
| FSM | Finite State Machine |
| IPAWS | Integrated Public Alert and Warning System |
| MAGPIE | Multi-Agent Platform for Interactive Environments |
| MASON | Multi-Agent Simulator of Neighborhoods |
| NOAA | National Oceanic and Atmospheric Administration |
| SMS | Short Message Service |
| WEA | Wireless Emergency Alerts |