The Impact of Educational Interventions on Real & Stylized Cities

Brian R. Hirshman, Alaiksandr Birukou*, Michael K. Martin,
Michael W. Bigrigg, and Kathleen M. Carley

July 10, 2008
CMU-ISR-08-114

Institute for Software Research
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Center for the Computational Analysis of Social and Organizational Systems
CASOS technical report.

*Alaiksandr Birukou is a Ph.D. Student at the University of Trento.

Abstract

Different cities and different populations react differently to various types of educational interventions due to small differences in population, socio-demographic traits, intervention media type, and a variety of other factors. The social simulation tool Construct was used to evaluate the effectiveness of several educational interventions on one stylized and four real cities to examine the types of behavior that could be observed given these small variations in initial conditions. This technical report describes the experimental design, parameters, and setup for the Construct educational intervention virtual experiments, then reports the results obtained when the experiments were run. The results from the stylized city and real cities are compared, and general similarities and conclusions are discussed.

Table of Contents

List of Figures

List of Tables

1   Introduction & Motivation

Recent success in modeling educational interventions has lead to follow-up questions about information access. For instance, if individuals are illiterate, how effective is a print advertisement in modifying behavior? Additionally, if the majority of the high-risk population lacks Internet access, how effective will a sophisticated web site be? Construct, a multi-agent simulation platform developed by the CASOS center, has previously been used to ask questions about the relative strengths of various educational interventions [1]. In order to understand the effects of these physical and cognitive barriers to information access, a variety of virtual experiments were conducted using Construct to understand the magnitude of these effects. This technical report describes the setup for these experiments, , the changes made to Construct to facilitate these experiments, and the results of a series of experiments performed on one stylized and several real cities.

The remainder of this technical report is organized as follows. Section 1 outlines the virtual experiment – the problem being modeled, the types of societal features modified, the information access parameters used, and other important global parameter settings. Section 32 discusses the kinds of modifications made for each virtual experiment, describing the interventions and information access mechanisms varied over the course of the runs. Section 4 presents high-level outlines of the changes made to the Construct code in order to support the information access virtual experiments. Section 5 describes the socio-demographic distributions and experiments run for a stylized city, while Section 6 presents the results of this work. Section 7 describes four real cities – Hartford CT, San Diego CA, Orlando FL, and Kansas City MO – and the parameters used to model and run virtual experiments on them; Section 8 presents the results of this set of experiments and compares them to the stylized cities. Section 9 concludes.

The modeling, simulation, and analysis contained in this report contains information regarding the demographics of the population for the purpose of describing how the members of the population are most likely to interact; consistent with the underlying well validated social science theories as embedded in the Construct simulation model. For real cities, the demographic data originates from the US Census. For stylized cities, the demographic data is consistent with the type of data extractable from the US Census. For both real and stylized cities information on access constraints, such as literacy and readership, comes from a variety of sources as described in the technical report "Information Access in Construct" [2]. It is recognized that government agencies do not and cannot use information regarding race, creed, or gender as discriminators in their decision making.

2   Experiment Parameters

2.1   Construct Parameters

Construct is a social network analysis tool which examines the evolution of networks and the processes by which information moves around a social network [3][4][5]. It seeks to combine social network analysis, a field that has typically sought to describe static networks, with an understanding of information diffusion to create a simulation environment that models network change over time. Construct is an agent-based model, which means that individual actors are the sources of decisions in the model. Agents in Construct interact via homophily -- the principle of

"like attracts like" -- which leads similar agents to interact, share knowledge, and become more similar to each other. As a simulation engine, Construct draws from a wide variety of fields: social network research for analytic techniques, sociology for its interaction mechanisms, psychology for agent choices and decision procedures, computer science for its algorithms and artificial intelligence as the glue that binds the various ideas together. Past work with Construct has sought to compare and contrast different interventions on a social network, to examine belief diffusion and information propagation among people, to study the effect of different network topologies on diffusion rates, and several other topics [1].

Additional information about the Construct system is available in the references of this paper. The Construct system itself is freely downloadable from the CASOS website, http://www.casos.cs.cmu.edu/projects/construct.

2.2 Modeled Scenario

The overall suite of experiments run, of which this experiment was just one part, sought to model the effects of different cognitive and access constraints on individuals who could interact with an educational intervention [1][2]. The educational intervention sought to encourage agents to perform in a specific way. In the taxpaying domain in which this problem was conceived, the interventions can be thought of those seeking to dissuade individuals from participating in an illegal tax scheme, to encourage them to take a legitimate tax credit, or to help them their file tax forms in a correct fashion. The prototypical model for this problem, a model explored in previous work, was the Earned Income Tax Credit (EITC) model: what is the most appropriate educational intervention or combination of educational interventions which will maximize legitimate behavior and simultaneously minimize fraud [6]? The generalization of this model, as explored in this technical report, seeks to expand the educational intervention model for use with other sorts of educational interventions in other domains; this has lead to a slight shift in terminology from previous work and publications [1]. The low-income focus of the EITC model, however, has been preserved since low-income individuals have been found to be most likely to suffer from lack of information access [7][8][9]. This allowed the literacy and access mechanism modifications to have a larger impact on the society as compared to tax credits or interventions used for other taxpaying subpopulations.

2.3 Knowledge and Belief

As has been done in previous work, the action examined was modeled as an action to be deterred. The action was modeled using two components: "knowledge" and "belief" [1]. Knowledge represented the facts necessary to perform the action; without sufficient knowledge, agents would be ignorant of the action and would not perform it even if they believed that they should. Belief was a function of three parts: a sequence of facts, similar to the knowledge component described previously; the belief an agent held in the previous time period; and a weighted function of the beliefs of agents in the surrounding social network. Table 1 presents the types of facts associated with knowledge and belief, the number of facts associated with each, and their interaction and transmission weights. Knowledge was broken down into two components. A single bit was used to represent whether the agents knew that the action existed. If an agent knew this bit, then it was said to "know of" the action in this work as well as previous experimental work. This was done to model the spread of the action in the society, and was given a very high weight in determining the transmission of information between agents. An agent who knew this bit was able to perform the action; any agent lacking this bid was unable to

2

Table 1: Numbers and weights on key facts

Fact Type		Characteristics		
		number of bits	interaction weight	transmission weight
Action exists	knowledge	1	1x	3x
How to perform action	knowledge	6	1x	3x
Action is right	belief	3	1x	3x
Action not right	belief	4	1x	3x
Action is worth doing	belief	3	1x	3x
Action not worth doing	belief	3	1x	3x
General knowledge	other	500	1x	1x

perform the action. Since many more agents could know that the action existed even while not knowing how to take the action, this fact was often the most widely known of the facts related to the modeled behavior.

As can be seen in the first column of Table 1, six bits were used to represent whether the agents knew enough information about the action in order to perform the action. These bits were often referred to as the "know how" bits in this work as well as previous work. In order to know how to take the action, agents needed to know at least three of the six bits. Agents could learn any three of the six facts, meaning that two agents could perform the action without having any overlapping information between them. This allowed an agent with slightly different knowledge of the action to make a decision to perform the action independently of any other agent, and also ensured that agents could have some variability in the composition of their knowledge when making the decision. All six of the facts were weighted equally with respect to transmission weight; no particular how-to fact was necessarily favored over any other.

The second requirement for taking the action was to have sufficient belief. There were two beliefs in the model: belief that the action was right and therefore a legitimate thing to do, and belief that the action was worthwhile and a good thing to do (a factor which could be held independent of whether agents felt it was right). An agent held a belief if sum of its current knowledge, prior beliefs, and social influence was above a certain threshold. To take the action, an agent only needed to be above this threshold; any additional belief beyond this threshold was superfluous.

The belief that the action was "right" was modeled with seven facts, three facts associated with the idea that the action was right and four with the idea that the action was not. Each fact was weighted equally, so belief knowledge component was a linear function of the number of facts known related to pro or con belief. Thus, the setup of three pro-action (action is right) and four anti-action (action is not right) belief facts ensures that an agent who knows all facts would believe that the action would not be right. This also helped to guarantee that the "not right" beliefs would be passed slightly more frequently than the "is right" belief. When communicating, an agent would be more likely to select a fact suggesting that the action was not right in comparison to a fact suggesting that the action was right, since it was likely that they knew more of them.

The "worth doing" belief was modeled with six facts, three of which suggested that performing the action was a good idea and three of which suggested that performing the action was a bad idea. Again, each fact was weighted equally so the belief knowledge component was a linear function of the number of known belief facts. Unlike the is-right belief, there was no difference in the number of facts associated with the positive and negative worth-doing belief.

Table 2: Initial distribution of key facts

Fact Type	Chance that individual fact is known by		
	human agent	promoter	seminar
Action exists (1 fact)	0%	100%	100%
Action know-how (6 facts)	0%	100%	100%
Believe right (3 facts)	1%	100%	100%
Believe not right (4 facts)	5%	0%	0%
Believe worth doing (3 facts)	1%	100%	100%
Believe not worth doing (3 facts)	5%	0%	0%
General knowledge (500 facts)	20%	2%	2%

This meant that, if an agent knew all the facts associated with the worth-doing belief, it would have no opinion, not an anti-activity opinion.

For an agent to hold a belief, it would often need to know more facts relating to one of the belief statements than it knew for the other. For instance, an agent who knew one is-right fact but zero is-not-right facts would believe the action was right. If the agent then learned two is-not-right facts, it would hold the is-right belief for a time, since the influence of the past belief would linger and the agent's belief would not be modified instantaneously. After a short time, however, the agent would believe that the action was not right because the sum of two is-not-right facts is greater than the one is-right fact. These effects could be temped by social influence, which might have a large or small effect on the agent depending both on the influentialness of one agent and the influenceability of the other. If the agent's local social network held the is-right belief, and the agent was very influenceable, it would be possible for the agent to persist in its older belief even though it held knowledge to the contrary.

In order to perform the action, an agent had to either believe that the action was right or that the action was worth doing. If the agent held the is-right belief but not worth-doing belief, it exemplified the case – actually observed in the real world on a not-infrequent basis – where an agent did not mean to be deceptive but was genuinely misinformed as to whether the action was right. The agent believed that its action was correct, even though there may have been information that it was not aware of which could have, in a perfect world, convinced it otherwise. On the other hand, if the is-right belief was not held but the worth-doing belief was, this represented the case in which the agent intentionally attempted to commit fraud. Agents who did not believe the behavior was right and who did not believe that it was worth doing– by far the majority group in the simulated population – did not try to perform the action even if they had sufficient knowledge. These agents had been successfully deterred from the behavior.

Additionally, a total of five hundred "social facts" were used in the experiment, facts which had no bearing on whether agents performed the action. These facts served to increase or decrease the homophily between two agents and helped agents select interaction partners. Due to the preponderance of these general knowledge facts, they were exchanged in the vast majority of interactions. As the simulation progressed, agents would exchange these facts, gradually becoming more similar to some agents and (relatively speaking) less similar to others. While the social facts did not directly affect the number of agents who performed the action, they indirectly helped to define the interaction partners of an agent and therefore greatly influenced agent sources of information.

2.4 Simulated Society

The experimental society was comprised of three thousand human agents for all three of the virtual experiments performed. While the socio-demographic characteristics of the individual agents depended on the experiment, a large number of other factors about the agents in the society were held constant.

The first column of Table 1 describes the knowledge associated with the human agents. The human agents did not have any of the knowledge facts initially, so they had to learn these facts via interaction with an outside source such as a promoter. Human agents, however, did have initial beliefs. On simulation start, agents have a 5% chance of knowing each of the anti-action facts (that the action was not right and that it was not worth doing), and a 1% chance of knowing each of the pro-action facts (that the action was right and that the action was worth doing). This created a society in which a small minority of agents initially held pro-action beliefs, a larger number held an opposing belief, and the majority of the agents did not have any belief at all and were primarily swayed by the beliefs of agents comprising their surrounding social network.

Additionally, each of the human agents had a 20% chance of knowing each of the five hundred social facts, meaning that each of the agents had about a hundred social facts in order to jumpstart agent-agent differentiation and homophily. These facts were selected from a random uniform distribution. Due to small perturbations between the types of facts known by each agent, agents were more similar to some agents than to others. These small differences would influence interaction, and would help determine interaction patterns.

The experiment was designed such that the action was propagated through the society by a single "promoter" agent, an agent who entered the society at the beginning of the first simulated period and attempted to convince as many human agents as possible to perform the action. The promoter had an associated companion, a "seminar" agent, which was active less frequently but could deliver a more powerful message to its attendees. This allowed the promoter to, from time to time, deliver a more powerful message to human agents who chose pay attention to it.

Both the promoter and the seminar had full knowledge of the action and how to perform the action. When no intervention was present, the promoter and seminar served as the sole source for all knowledge about the action. The promoter and the seminar knew all the is-right facts and the worth doing facts in an attempt to convince as many agents as possible to perform the action.

The simulation ran for one hundred and four time periods. This represented one year with two interaction periods per week for the simulated society. Human agents were active every time period, choosing one potential interaction partner to communicate with and receiving communication from one another agent. The promoter was active only every other period, but could initiate communication with several human agents and receive communication from several other human agents during that time. The seminar was active only four times during the entire simulation, and could not contact other agents and had to wait for human agents to contact it.

2.5 Network Topology

The human agents in the virtual experiment were all linked together via an underlying social network. This allowed one agent to interact with a subset of the other human agents in the network but not with all of the possible agents. However, these subsets were not all alike. Some slices of the agent population could form small, highly interconnected cliques, while other slices

could have links that spanned different parts of the networks, while still others could have a jumble of assorted edges. The underlying network topology helped to define this type of structure, though it was up to the individual agents to interact according to homophily and build an actual communication structure on top of the possible communication framework provided by the network topology.

The density of the human agent to human agent social network was about .0133. This meant that each of the three thousand human agents had about forty neighbors. An independent social network was created for each replication, as drawn from the underlying distribution of social network topologies, to ensure that the results observed were not biased by the specific layout of any one individual social network.

The network topology used in this research was a uniform random graph [10]. The uniform random network was a random network of three thousand nodes, created by selecting 1.33% of the possible edges from the fully-connected graph and creating the social network from them. The uniform random network was generated in a manner consistent with the generator suggested by Erdos and Renyi: edges were selected randomly from a complete graph until the requisite density was achieved [10]. No preferential attachment, triadic closure, or clustering algorithm was used in this process, but the edges were selected to ensure symmetry (if an ego could contact an alter then the alter could contact the ego). This process created a network in which the mean number of neighbors was forty, the average distance was empirically determined to be 2.5, and the diameter was three. While it has been shown that Erdos-Renyi graphs do not capture some of the subtleties in real-world social networks, they provide a reasonably accurate first approximation of them.

While it was not performed for this particular work, past research using Construct has also investigated different types of underlying social networks. For instance, the paper "Societal Factors as Moderators of Intervention Strength," a paper draft submitted to IEEE SMC, suggests that different underlying social networks may lead to different observed results [1]. Much of the work described there is present in this version of Construct, including the descriptions of the simulation setup, knowledge distribution descriptions, and interaction sphere descriptions. One noticeable change, however, is the fact that the number of agents in this experiment has been increased to three thousand (from two thousand as described in that paper); this has resulted in scaling of other simulation parameters.

2.6   Other Parameters

While Section 2.3 described the facts and beliefs used in the simulation, there were a number of subtleties with respect to how they were modeled and used.

First, when an agent attempted to compute its similarity with another agent, it would use a feature called "transactive memory" to compute its similarity score with another agent. Transactive memory represents second-order information about a fact; for example transactive memory could represent that an ego was aware that an alter knew a specific fact [4][11]. Thus, agents computed knowledge similarity using their perception of the agent's knowledge, not the agent's knowledge itself. Similarly, agents used a special form of belief, belief transactive memory, when computing the belief level of the surrounding community; unlike fact transactive memory, which was binary (either an agent knew it or it did not), belief transactive memory was real-valued and therefore more nuanced. In this way, transactive memory served as a mental

6

Table 3: Distribution of key facts for interventions

Fact Type	Intervention				
	Ad campaign	Website	Call center	Radio ad	Mailing
Action exists (1 fact)	100%	100%	100%	100%	100%
Action know-how (6 facts)	10%	33%	10%	10%	10%
Believe right (3 facts)	0%	0%	0%	0%	0%
Believe not right (4 facts)	33%	100%	100%	33%	33%
Believe worth doing (3 facts)	0%	0%	0%	0%	0%
Believe not worth doing (3 facts)	33%	100%	100%	33%	33%
General knowledge (500 facts)	10%	2%	5%	10%	10%

model of another agent's knowledge and can have many of the features seen in a realistic mental model. The transactive memory perception could be inaccurate, incomplete, or incorrect, depending on the situation -- alter agents could learn, evolve, or change beliefs without an ego agent knowing about it. An ego's transactive memory of an alter agent was only updated if an agent sent a specific fact to the alter agent, or if the ego agent learned (possibly from a third party) that the alter knew a particular fact. About twenty percent of all messages sent between agents contained exclusively transactive memory to help keep this information up-to-date.

Second, knowledge and belief were represented slightly differently in the simulation. Knowledge was represented as binary facts: either the agent knew a fact, or the agent did not. Agents could not partially learn a fact, nor could they partially tell the fact to another agent; communication was all or none. Forgetting was not enabled, so agents retained any information that they learned over the course of the experiment. In contrast, belief was allowed to vary between agreement, disagreement, and no information along a real-number scale. Agents could change their beliefs by learning new information, or by finding out that an alter agents in their social network had changed their beliefs and updating their belief accordingly. When determining whether an agent performed an action, however, a hard cutoff was established, effectively turning the real-valued belief into a binary yes-or-know decision. Beliefs stored in beliefs transactive memory were stored as either zeros, ones (agreement), or negative ones (disagreement) to summarize agent belief; the exact nuances were not stored by individual agents as they computed the beliefs in their society. A side effect of this implementation of binary transactive memory was a reduction in the memory footprint of Construct, a factor which resulted in faster completion of experiments and allowed more instances of an experiment to run simultaneously on the same machine.

Lastly, facts were not cognitively equivalent to the agents in the simulation. As can be seen in the second and third columns of Table 3, agents had different interaction and transmission weights for particular facts. In this experiment, all facts – knowledge facts, belief facts, and general facts -- had an interaction weight of one. This meant that agents would consider all facts equally when evaluating a potential interaction partner's knowledge. Since agents often held vastly more general knowledge facts than they did knowledge facts, this lead to situations in which human agents were selecting interaction partners primarily based upon their general knowledge similarity – the action-related knowledge facts played only a minor role. However, the fact transmission weight for action-related facts was much higher: in order to over-sample for conversations of interest for facts related to the action, the action-related fact transmission weight was set to three. This ensured that agents were more likely to pass information related to the action in a given interaction (if they knew it), but the passing of this type of information was still relatively infrequent in the general society. The net effect of these weights was to ensure that

Table 4: Percentage of population that could interact with an intervention

Property	Intervention				
	Ad campaign	Website	Call center	Radio ad	Mailing
Percent of society	*	*	50%	90%	50%

* affected by the access mechanisms discussed in Section 3.2

general knowledge homophily largely drove agent-agent interaction, but even when knowledge homophily was dominant agents would be more likely to transmit action-related facts.

3   Virtual Experiment Setup

In the experiments performed, two classes of parameters were varied. The first parameter varied was the number and type of interventions present in the society, discussed in detail in Section 3.1. These represented the attempt of an outside entity to influence behavior in a society by a series of educational interventions to warn at-risk agents about the action. The second parameter varied was the information access method that was active and is further explained in Section 3.2. These information access methods could weaken the intervention in various fashions, by restricting the number of agents who could interact with the intervention or cutting down the amount of information learned from the intervention if interaction occurred.

3.1   Interventions

In addition to the society of human agents and the promoter, there were also up to five other intervention agents. These interventions were representations of print advertisements, web pages, call centers, radio advertisements, and postal mailings. Interventions had differing amounts of knowledge and belief facts available for transmission; however, all were similar in the sense that they could not initiate communication with a human agent and had to be contacted by an informed party before spreading their information. The knowledge holdings of the individual intervention agents are described in Table 3. The fraction of agents who could communicate with the interventions is highlighted in Table 4.

- The advertisement was meant to provide a small amount of knowledge and belief, while also containing a large amount of general knowledge information to make agents more likely to examine it. This meant that the intervention could have a small impact on a variety of agents, and could serve to dissuade agents before they ever learned about the action. The advertisement had the action exists bit, had a ten percent chance of knowing each of the six know-how facts, and had a thirty percent chance of knowing each of the anti-action belief facts (for both the is-right and worth-doing beliefs). The advertisement could only send a short message, but it had about ten percent of all the social facts in the society – the highest of any of the interventions. The percentage of agents who could interact with the advertisement was dependent on newspaper readership when active; when it was inactive, the percentage was set to 100%.

- The web site was designed to provide a large amount of belief information by proving a large number of reasons why the action was not right and not worth doing. In doing so, however, it could potentially be scraped for knowledge information. The web site knew the action existed and had a thirty percent chance of knowing each of the knowledge facts necessary to complete the action. It knew all of the anti-action belief facts in order to convey a long, strong anti-behavior message. The interaction sphere of the web site was

dependent on the Internet access parameter as described in Section 3.22.5 when it was inactive it was set to 100% as a comparison case.

- The call center was designed to answer agent questions with regard to the action. It knew less information about the action than the web site, but had the same percentage of the anti-action belief facts. It had more social knowledge then the web site, meaning that agents were more likely to communicate with it than with the web site. However, it could not send as powerful a message in one communication as the web site in a single interaction, and also was not accessible by all the agents in the simulation.

- The radio advertisement was very much like the print advertisement: it was designed to provide a small amount of information and belief while reaching a large number of agents. It too could only send a short message. Unlike the advertisement, it was not affected by the literacy and access mechanisms described in Section 3.2.

- The postal mailing was designed to represent an official mail intervention conveying information about the action with the intent of deterring at-risk agents from performing the behavior. It had the same information content as the advertisement, but it had a different interaction mechanism than the other interventions. When the mailing agent was active, it could send a message to some fraction of the 50% of the agents in the society. For the next four time periods, the mail message resided in the agent's "mailbox". The agent had a certain probability of checking their mail during that period and learning the information contained in the mailing. Agents who read the mailing absorbed some of the information contained in the mail message; agents who did not were unaffected by the intervention.

These five intervention types lead to the development of twelve intervention cases used in the experiment. The first, a baseline case, was a case with none of the interventions present. This was used to measure the amount of activity occurring if the promoter was allowed to spread the behavior unimpeded. The next five cases were single-intervention cases: the advertisement alone, the web site alone, the call center alone, the radio advertisement alone, and the mailing alone. Six combination cases were also tested. Three of these cases were combinations with the web site: the advertisement and the web site, the radio advertisement and the web site, and the mailing and the web. The remaining three interventions were super-bundles containing more complex interventions – the mailing, radio ad, and web site; the call center, mailing, radio ad, and web site; and all of the interventions together (print ad, call center, mailing, radio ad, and web site).

3.2   Information Access

In addition to the intervention modifications, three different types of information access mechanisms were used: literacy, Internet access, and newspaper readership. These access mechanisms affected whether agents could interact with a specific intervention (if the intervention was present).

- The literacy mechanism affected all interventions that required reading textual material. This meant that the print advertisement, the web site, and the postal mailing. When this information access parameter was enabled, illiterate agents could still access these interventions; however, the agents did not receive the full message (and could suffer from message distortions). Literate agents were unaffected.

9

- The Internet access mechanism affected all interventions that required web access, namely the web page. When this information access parameter was enabled, agents lacking web access could not communicate with the web site intervention at all. Agents with Internet access were unaffected.

- The readership mechanism affected all interventions that required newspaper access, namely the print advertisement (which was assumed to appear in a print newspaper or associated magazine). When this information access parameter was enabled, agents lacking newspaper subscriptions could not communicate with the advertisement intervention at all. Agents with subscriptions were unaffected.

To determine which agents were literate, had Internet access, or obtained a newspaper or a regular basis, data was gathered from a number of places including the National Center for Statistics, the Pew Research center, and Newspaper Association of America. Results were gathered for aggregate population groups, including breakdowns by gender, age, race, education, and income. The changes necessary to support these interventions are described in Section 4. Additional changes on the formulae used in the computation, as well as the validation numbers themselves, are available in the technical report "Modeling Information Access in Construct" [2].

The virtual experiments in this research used five different levels of information access. The baseline case had all of the mechanisms disabled for comparison with earlier experiments. Three cases enabled each of the mechanisms separately to try and understand the individual effect of each of these mechanisms. Finally, a combination case enabled all three access mechanisms, thereby attempting to model processes in the real world as realistically as possible.

4 Changes to Construct

In support of this project, two major changes were made. The first change was to add literacy, access, and readership information access functionality to Construct. The second change was a mechanism to automatically generate literacy, access, and readership rates for individual agents from aggregate data. Additional details about the changes to Construct, as well as a description of data gathered to calibrate the literacy and access mechanisms, is available in the CASOS technical report "Modeling Information Access in Construct" [2].

While the overall mechanisms for adding literacy, web access, and readership were very similar, as all three seek to limit the effectiveness of information access, there are some subtle differences between them. Each mechanism was only invoked if an intervention agent had a specific parameter set; for instance, a web site had an "access_required" parameter in order to require potential interaction partners to have Internet access. A similar parameter was used for the newspaper setting and readership. The Internet access and readership parameters prevented interaction between an agent and an intervention if the agent did not have the requisite attribute set. If a human agent did not have Internet access, the intervention agent was removed from the human agent's interaction sphere. This meant that there could be no direct contact between the agent and the intervention, though information could still be conveyed between the two via intermediate agents. This mechanism was relatively simple to implement, and was done by applying a post-processor on top of the Construct initialization system in order to ensure that only specific agents were removed from the interaction spheres of the intervention.

Unlike the web and readership access parameters, the literacy parameter did not prevent interaction between an agent and an intervention. Instead, it cut the amount of information conveyed between an illiterate agent and an intervention requiring literacy. While Internet access and newspaper readership were representations of physical connections, which required either a printed newspaper or a computer, the concept of literacy was less concrete. Since other research has shown that illiterate agents can gain information from printed materials even if they are not able to understand everything printed there, a more nuanced mechanism was used for literacy than for the other access mechanisms.

For this simulation, the process was modeled as follows: If an illiterate human agent chose to interact with an intervention requiring literacy, the intervention would choose a message to send to the agent exactly as it would for a literate agent. Then, each bit in the message had a 50% chance of being removed to represent the lack of comprehension by the illiterate agent. Following this, each remaining bit would have a 50% chance of being replaced to represent miscomprehension. If the message component was a fact, the original fact was replaced with a randomly-chosen fact; if the message component was a belief, the belief was inverted to represent the opposite belief. In sum, then, an illiterate agent would learn about 25% of the original message, learn a set of random facts equivalent to 25% of the original message, and lose the remaining 50%. This process was implemented by inserting additional code in Construct to post-process messages sent by intervention agents.

The implementation of the information access parameters had important secondary and tertiary effects. Since the information access parameters were represented using agent attributes, they could increase or decrease similarity with other agents. Research has shown that readers and non-readers tend to self-associate; this manner of representation ensured that there was increased similarity between pairs of literate agents and decreased similarity between literate and illiterate pairs. This had the further effect of decreasing the spread of intervention information diffusion to illiterate agents and agents who lacked access to the intervention; without direct access to the information, and with decreased similarity with those who actually can learn the information, these agents were even less likely to obtain access to the information. This process did not require additional modeling in order to configure – instead, it was an emergent property that resulted from past code decisions regarding the way attributes are used to calculate agent similarity scores. Additionally, there were tertiary similarity-building effects due to literacy and information access. If two agents had access to the web site intervention, each would have the potential to interact with the web site and learn additional belief or knowledge information. If both learned the same facts, this process would make the agents slightly more similar to each other. In contrast, an illiterate agent would have no way of learning the fact directly, and thus might not benefit from this (very slight) increase in similarity. While the direct effects of these secondary or tertiary effects were not measured in the course of this research, it is important to note that such processes were actively taking place as agents were interacting, learning, and adapting to new situations.

In order to implement these literacy, Internet access, and readership parameters on a per-agent basis, it was necessary to use aggregate population-level statistics to derive a set of agents which modeled empirical data as faithfully as possible. In a web search used to ground the simulation in empirical data, it was found that information was usually aggregated by one parameter – by race, by age, or other parameter. To calculate the literacy for a single agent (for instance, the probability that an Asian American high school graduate between the ages of thirty

11

Figure 1: Socio-demographic breakdown for stylized city

and forty with an income of under twenty-five thousand dollars a year), it was necessary to combine the parameters in such a way that population demographics were matched and individual agent parameters made sense.

5   Virtual Experiment Ia: Stylized Cities

The virtual experiment performed examined several types of stylized cities in order to tease out the various effects of agents who performed the action. Five variations on a stylized city were replicated multiple times in order to create an agent database from which to sample for the number and types of agents which performed the action, who knew sufficient information to perform the action, and who believed that the action was right.

5.1   Agents Used

The stylized city used in this experiment was similar to other stylized cities used in past virtual experiments [2]. In this experiment, however, there were minor modifications to the experiment design. While agents were over-sampled to ensure that a high number of low-income and low-education agents were represented in the society – the groups that were most likely to suffer from illiteracy and a lack of Internet access – a decision was made to perform a virtual experiment separately for agents of each race. The breakdown of agents is depicted in Figure 1.

In the virtual experiment, four socio-demographic attributes were modeled as important factors for driving interaction: education, income, race, and age. A substantial amount of empirical research has indicated that these factors have been shown to be strong predictors of interaction patterns. These factors have also been shown to correlate strongly with access to

information access media, such as literacy, newspaper readership and internet access. While other virtual experiments have dealt with societies that are homogenous over one of these attributes, this experiment was performed on a diverse city with substantial variation in all four of the socio-demographic factors.

As this virtual experiment sought to explore the effect of interventions on low income individuals, several modifications were made to the stylized city in order to over-sample for low-income agents. First, the agent population was skewed predominantly toward low income individuals. Three in four agents were in the lowest income bracket, and nearly three in four were in the lowest education bracket as well. Additionally, the homophily of the promoter was increased for low income agents relative to high-income agents, since this demographic was its target audience.

When computing statistics on the society, however, only the low-income subset of the population was used. A matching algorithm was decided upon in discussions with simulation experts in order to reflect the fact that while the modeled activity was geared towards low-income individuals, high-income agents could participate at lower level of probability. In conjunction with other simulation experts, it was decided that if an agent was not in the income range or not in the education range, it would have a 50% chance of being able to performing the action, even if it knew enough about it and believed the action was right. If the agent did not match the income level or the education level, it would only have a 25% chance of being able to perform it. Agents who did not match, and therefore could not perform the action, were modeled as those likely to be flagged and caught if they performed the action. While this flag prevented the agents from performing the action directly, it did not prevent them from serve as conductors of information throughout the society; these agents could pass on any knowledge or belief fact in their communication with other agents in the society and could even communicate with the promoter if they desired.

Using these matching criteria, slightly less than 75% of the agents in the simulated society had the necessary characteristics to perform the action. Since the other 25% of the population could not perform the action, they were excluded from the analysis. All resulting percentages presented in the remainder of this document deal with these matching agents only, except where specially noted.

5.2 Experiment Design

The experiment results were first generated as a 12x5 matrix, with twelve interventions crossed by five different access mechanisms. The script creating the input deck did create the full factorial design, even cases (such as Internet access with the web site intervention) that would be no different than another case (here, the Internet access without the access intervention). A total of just over two hundred experimental cases were eventually submitted to the Condor cluster at the TeraGrid supercomputing center in West Lafayette, Indiana [13]. A total of fifty iterations were performed for each input deck for a total of 15,000 Construct iterations.

Execution of the experimental cases took about two days on the heterogeneous cluster of computers available on the TeraGrid. Since the TeraGrid had a very large number of machines, replications could be done in parallel and computation could proceed much more rapidly. For comparison, an earlier instance of the experiment was performed on local machines, including

Table 5: Modifications active in the most and least realistic cases

Intervention name	active modifications		notes
	least realistic	most realistic	
No intervention	none	none	
Advertisement	none	literacy, readership	
Call center	none	none	
Mailing	none	literacy	
Radio message	none	none	
Website	none	literacy, internet access	
Ad + web	none	all three	
Mail + web	none	literacy, internet access	
Radio + web	none	literacy, internet access	
Mail + radio + web	none	literacy, internet access	
Call + mail + radio + web	none	literacy, internet access	"all_but_ad"' in graph labels
Ad + call + mail + radio + web	None	all three	"all"' in graph labels

two dual-core, 4CPU machines with 64GB of memory. Performing 15,000 experiments locally took almost a month.

Table 5 summarizes the modeled cases which contained the most realistic results of the fifteen thousand experiments. These "most realistic" cases modeled the literacy mechanism for the advertisement, web page, and mail message; modeled the newspaper readership for the advertisement; and modeled Internet access for the web page. In some intervention combinations, all three of these mechanisms could be active such as occurred in the most realistic model of the advertisement and web site. However, only relevant interventions were affected in the combined cases: while literacy might affect both the ad and the web site, the newspaper readership modification affected only the advertisement. In three of the cases, none of the mechanisms were modeled in the most realistic case since they were not relevant.

The reader should recognize that the research on access constraints commonly find correlations between gender, income, age, race, education and literacy, newspaper readership, and internet access. These are simply correlations not causations. In this work, the authors use these correlations to characterize the population; they do not infer causal mechanisms. The modeling, simulation, and analysis contained in this report contains information regarding these correlations for the purpose of describing how the members of the population are most likely to interact; consistent with the underlying well validated social science theories as embedded in the Construct simulation model. The authors recognize that US government and IRS decisions do not and cannot use information regarding race, creed, or gender as discriminators in its decision making. As such, effort was made to ensure that the causal mechanisms did not focus on these factors.

6 Results for Stylized Cities

The most realistic cases were the advertisement with literacy and readership, the web site with both literacy and Internet access, and the combination case with literacy, Internet access, and newspaper readership modeled. The table presents the percentage of eligible agents in the society who had the requisite knowledge, had requisite belief, and who decided to take the action. Since some agents were not eligible to take the action due to the matching implementation, the agents were removed from the database and were not included in any of the analyses. The values in the table were calculated using the following formulas.

Table 6: The most and least effective intervention combinations for stylized city

	Knowledge	Belief	Behavior
Most effective	mail	all but ad	ad + web
mean	3.11%	3.07%	0.551%
std dev	± 0.56%	± 0.53%	± 0.181%
Second most	mail + radio + web	ail + radio + web	web alone
mean	3.15%	3.14%	0.551%
std dev	± 0.49%	± 0.51%	± 0.167%
Third most	radio + web	radio + web	all interventions
mean	3.17%	3.14%	0.551%
std dev	± 0.54%	± 0.53%	± 0.142%
Least effective	mail + web	no intervention	ad alone
mean	3.24%	3.58%	0.589%
std dev	± 0.62%	± 0.64%	± 0.157%

- The value MIR_{ki} was the mean percent of matching agents who met criteria k (either knew sufficient information, held the belief, or took action) when intervention i was present and the more realistic model of agent ability and access to information was used, i.e., all cognitive and information access constraints of relevance were turned on. This is the mean of means, the mean number of agents over the two hundred and fifty runs performed over all five of the societies of individual races in the set of intervention conditions i.

- The name of the most effective intervention NR_{ki}, the name of the intervention which lead to the value of MIR_{ki} observed for k when intervention bundle i was present

Many of the tables in the remainder of this document are organized to display MIR_{ki}, NR_{ki}, or both, depending upon the context. Many tables report the three most effective interventions in terms of decreasing the number of agents who performed the activity to account for the fact that the simulated data may be noisy and that different interventions may have different effects in different runs. Often, the least effective intervention is reported for comparison.

The remainder of this section, as well as the remainder of this technical report, is organized as follows. Section 6.1 discusses the general effects observed in this experiment for the societies as an average. Since the societies are broken down individually by race, the results observed here differ slightly from results observed in past work as well as work examining the effects observed in heterogeneous societies. Section 6.2 presents the high-level overview of the effects observed in this research and presents recommendations for policy and analysis.

6.1   General Effects

The experiment results were examined for trends to find best intervention combination for the society. While there was considerable variation in the effectiveness of an intervention over the course of the two hundred and fifty replication experiment, some general trends were apparent. The most effective interventions examined, as well as the number of agents who knew of the action, had sufficient belief to complete the action, and actually followed through with the action are summarized in Table 6.

While there was some variation in the amount of knowledge about the action known in the society, the effect was not statistically significant after the two hundred runs performed ($p > .37$).

The difference between the most effective intervention case, the mail intervention in which 3.11% of the matching agents in the society knew of the action, and the least effective intervention case, the mail + web case in which 3.24% of the society knew sufficient agents to perform the action, was not robust given the number of experiments run. This difference was roughly equivalent to two additional agents for every five experiments run.

The fact that knowledge did not change substantially is not particularly surprising. Since knowledge of the action was defined as the number of agents who knew at least three facts related to the action, it was not always easy for agents to enough information to participate. Interventions, designed to sway belief and not to strongly influence knowledge of the action, were unlikely to have a substantial effect on knowledge except in relatively rare situations. Nevertheless, it is important to consider ways in which interventions might or might not boost the amount of knowledgeable agents present in the society. On the one hand, the intervention did present some information to the society, so it was possible for an agent who knew only two of the six knowledge-related facts to communicate with the intervention and find out the necessary information needed to qualify as an agent who "knew about the activity". On the other hand, agents who communicated with the intervention may have previously communicated the promoter and could have already known about the action before interacting with the intervention, in which case any additional knowledge from the intervention might not have resulted in a recorded change. In a third situation, the intervention could actually decrease knowledge: if agents preemptively communicated with the intervention instead of contacting the promoter, it would be highly unlikely that the agent would learn sufficient information about how to perform the action from the intervention alone and thus lead to a potential decrease in total knowledge.

In this experiment, it was observed that additional interventions did increase the total amount of action-related facts floating about in the social network – a factor that was highly statistically significant. However, these facts were not always co-located in the same agents, meaning that there was additional diffusion of knowledge but this knowledge was less concentrated than it was in the no-intervention case. This meant that any of the there scenarios described above were possible: some agents might learn more knowledge due to the diffusion and then go on and participate, others might not change, and still others could decrease. In this experiment, all three occurred, and suggested that they were a wash in terms of the overall effects. Thus, while the amount of knowledge increased in general, the concentration of knowledge in individual agents did not increase substantially relative to either the no-intervention case or to any of the other cases.

Figure 6, included as an appendix, plots the reliability versus effectiveness for each of the intervention cases. To do so, it plots the relative change between each of the interventions and the no-intervention case. The x and y axes represent the relative change in the mean number of agents who know how to perform the action (the intervention effectiveness or magnitude) and the relative change in the standard deviation (the intervention reliability or variability), respectively. As can be seen from the plot, some interventions decrease the amount of knowledge in the society and others do not; however, all interventions decrease the amount of variability in the society. The decrease in variability initially can seem surprising, but has a logical explanation. In the absence of any of the interventions, there was no check on the number of agents who wanted to seek out the promoter and learn about the action. This lead to some scenarios in which a large number of agents sought out the promoter and learned a

16

substantial amount of action-related facts. This then contributed to a rise in the number of knowledgeable agents. In other simulation runs without interventions, there was substantially less contact with the promoter. With less promoter contact, fewer agents ever learned about the action. However, when interventions were present, such high and extremely low outliers occurred less frequently. Of the top three interventions by effectiveness, the mail + radio + web intervention was seen to be the most reliable since it had the largest decrease in standard deviation. It is worth noting that the mail + web intervention had increased variability as well as an increased number of knowledgeable agents; while some of this was due to a presence of a very high outlier in which 5.42% of the society knew enough information to perform the action, much of the increase was due to a spate of about twenty runs in which the knowledge was substantially higher than that observed for the other intervention cases.

In contrast to knowledge, where the effect observed was not statistically significant for the number of experiments performed, the effect of the interventions on belief was highly statistically significant ($p < .0001$). In the no intervention case, there was no effective deterrent against the belief information spread by the promoter. This lead a relatively large number of agents to take the action, as the only factor opposing the spread of the promoter's belief was the relatively weak bias of some of the agents in the society to oppose the promoter's belief. In the no intervention case then, an average 3.58% of the society held sufficient belief to perform the action, although there was substantial variation between the runs performed. In contrast, the three most effective interventions in reducing belief and deterring the activity were the all but ad case, in which 3.07% of the population held the promoter's belief; the mail + radio + web case, in which 3.12% of the society held the belief; and the radio + web case, in which 3.14% of the society held the belief. Indeed, all of the interventions except the advertisement and the call center lead to statistically significant decreases in belief for the general society aggregated over all incomes, ages, education levels, and other factors.

The substantial change in belief observed is also not surprising. As seen in previous work, belief decreased substantially when interventions were added[1]. The decrease in belief was due to four factors:

- First, the interventions distributed belief information directly to some of the agents in the society. While some of these agents were affected by the cognitive and access limitations modeled in the society, a substantial percentage of the agents were not. The agents who could communicate with the interventions were able to learn this belief information in messages sent from the interventions, and could lead these agents to change their belief. If the agent held a pro-action belief prior to communicating with the intervention, many of the interventions (and especially the web page intervention) could provide a substantial message which might help change the agent's belief directly.

- Second, agents who communicate with the intervention were also able to further relay this information to other agents in the society. Even if an agent did not have direct access to a powerful intervention, it was very likely that at least one of the forty-odd agents in its interaction sphere had the possibility of communicating with an intervention. Since agents could relay information learned from the intervention to other agents in the society, it was possible for agents to learn belief information second-hand or third-hand. Thus, while some of the effect of the intervention was directly related to direct contact, a

substantial portion of the effect was due to the additional power of agent-agent communication and the belief spreading that occurred naturally within the society.

- Third, the belief information spread by the intervention could decrease the similarity between some agents and the action promoter. While the interventions provided a small amount of knowledge, they also provided anti-action belief facts which could increase the differences (decrease the relative similarity) between an agent and the promoter. This could lead to fewer agents communicating with the promoter overall. The consequence of this would be a net decrease in the number of agent who held the belief about the action, as well as fewer pro-action belief facts entering the society and creating second-order effects. While this effect was not always present – in most of the cases, not much of a decrease in the number of unique agents communicating with the promoter was observed – the decrease in belief from this effect cannot be discounted.

- Last, the interventions could decrease belief very indirectly by changing the overall belief in a society. Even if an agent did not communicate with the promoter, or learn belief facts indirectly from another agent who did, it would still take the prevailing belief in its social network into account when computing its overall belief in each simulation period. In the absence of the intervention, there was a small minority of agents who held the belief mainly because other agents in their immediate social network held the belief. When the interventions were present, this result was counteracted by the anti-action belief coming from one or more sources. While there still could be pockets of agents who held the belief about the action, it was less likely that these pockets would dominate for any one particular agent. Since the belief mechanism was designed to take the beliefs of other agents into account, the mechanism could give rise to these indirect effects with potential consequences for agents who never did interact with the interventions or even recognize that other agents in their interaction sphere had interacted with the interventions.

While the interventions had an effect on belief, they also decreased the variance in belief observed in between individual runs. Figure 7 in the appendix presents a plot contrasting the change in means versus the change in standard deviations for each intervention, comparing each intervention's effectiveness relative to the no-intervention baseline case. As can be seen, all of the cases have a smaller standard deviation relative to the no intervention case, mainly because these cases lack the very low and very high outliers present when the interventions were not active. As discussed before for knowledge, when the interventions were not active there was less consistency in the number of agents who held a belief: some runs could be very low if the promoter's message was not successful initially, or it could be very high if it spread rapidly. Such swings in belief were less likely when the intervention was present: agents would be somewhat more likely to know about the action, decreasing the chance of an extremely low communication rate with the promoter and probably ensuring that belief was above some very low level; more importantly, though, it decreased the number of high cases and very high outliers in order to make the intervention more effective overall. Of the top three interventions, the mail + radio + web intervention was seen to be more reliable than the others, though the differences between the three were very small.

The number of agents who performed the action was a combined function of the number of agents who knew how to perform the action as well as held a positive belief about the action.

18

Not all the agents who knew how to perform the action became participants; some held an opposing belief, and others did not satisfy the necessary matching criteria (which usually meant that their income level was too high). In the experiments without any of the interventions active, a mean of 0.580% of the matching agents went on to participate – a small fraction of the total population of the society, but about eighteen percent of those who knew how to perform the activity and about sixteen percent of those who held the pro-activity belief.

When the interventions were compared, a statistically significant change in behavior was observed (p < .005). The two most effective interventions, the ad + web intervention and the web alone intervention, decreased the number of agents who performed the behavior to .551% of the qualifying agents in the society went on to perform the behavior. The next most effective intervention combination, the all intervention bundle, resulted in a decrease to 0.552%; following this, the radio + web bundle, the all but ad bundle, and the mail + radio + web bundle resulted in drops to participation. Graphically, these decreases are depicted on the y-axis of Figure 9 at the end of this document. In contrast, the majority of single-intervention bundles contributed to increases in behavior, the opposite effect of what the intervention was attempting to achieve. Of these, the advertisement alone performed the worst: 0.589% of the matching went on to perform the action, a nearly 3% rise over the no-intervention case and a 7% increase over the most effective intervention combination in which the same advertisement was combined with the web page.

In some ways, these results may seem surprising: the advertisement alone lead to a slight increase in the number of agents participating in the activity, but when combined with the web page it lead to fewer agents doing so. The initial conclusion might be that the web site intervention alone was the strongest, and that the effect of the web page dominated any slight increase in activity. However, the actual results suggest slightly more nuance. Consider, for instance, the radio intervention as modeled in Figure 8. When the radio intervention was modeled alone, there was a slight increase in the number of agents participating relative to the no-intervention case. When combined with the web site, the radio intervention lead to a decrease in activity, and then when the mail intervention was also included there was an observed increase in activity – even though the mail intervention lead to an decrease in activity when modeled alone. Such results might be consistent with the overall additive effects. When the call center intervention was added to create the all but ad result, however, a decrease in activity was observed, and when the advertisement was then added a still further decrease was observed. Both the call center and the advertisement increased behavior when modeled alone, but when placed in conjunction with the multiple other interventions modeled lead to an interaction which produced a decrease in behavior. While the decrease observed in the all-intervention case was not as large as the decrease observed in the web + ad case or the web site alone case, the five-intervention bundle was superior to most of the sum of its parts.

The interventions could affect behavior in four major ways: by changing the number of people who knew how to perform the action generally, by changing the number of people who held a belief about the action generally, by changing the proportion of knowledgeable agents who held the pro-action belief, or by changing the proportion of believing agents who were knowledgeable. While these four concepts were highly interrelated, the interventions could have differing effects on each of them. For instance, Figure 8 suggests that the advertisement reduced the number of agents who knew about the activity. However, the advertisement intervention had

the second-highest belief. Despite the decrease in knowledge, however, the fraction of knowledgeable agents who ended up participating rose to 18.8% of that sub-population, the highest rate observed in any of the virtual experiments performed. This suggested that decreasing the number of agents who held the belief was not sufficient to lead to a decrease in the behavior. On the other hand, while the mail + radio + web intervention helped to decrease belief in the stylized case, a total of 18.4% of the agents who held the belief went on to participate, the highest ratio observed in any of the experiments. In contrast, the lowest percentage of potential participants in terms of knowledgeable agents occurred for the all intervention case, in which only 17.3% of the agents who knew how to perform the action ended up doing so, and 16.4% of those who believed that they should perform the action actually did so in the no-intervention case.

Such analysis is not meant to suggest that interventions which raise the proportion of activity in the subpopulation are less effective. Indeed, interventions can be effective even if they increase the percentage of the sub-population that performs the action. If the intervention decreases the total number of knowledgeable agents, or decreases the number of agents who hold a specific belief, then it is possible for it to increase the percentage of people performing the behavior relative to the number who either know how to do it or hold the belief. For instance, switching from no intervention to the mail + radio + web intervention combo lead to a large decrease in the total number of agents who held the belief. At the same time, it lead to an increase in the relative number of agents who performed the activity relative to those that held the belief. This occurred because the drop in the number of agents who held the belief was sharper than the drop in the number of agents who performed the activity. While using the percentage of agents who performed the activity relative to those who had a specific characteristic can be a useful proxy, it may often not tell the full story.

6.2   Intervention Reliability

While there were differences in the mean effects for each of the interventions modeled, these fluctuations were not always consistent across runs. Indeed, in some of the cases, the standard deviations were larger or smaller than others, suggesting that some of the interventions lead to results that were closer to the mean than in others. This suggested that some of the interventions would be more likely to produce results that are reliable, as these results would be more likely to occur in the simulated society.

The x axis of Figure 6, Figure 7, and Figure 8 present these standard deviations relative to the baseline (no intervention) case. The positive numbers on the x axis are percent increases relative to the no intervention case, which suggest that the intervention is less reliable than the no-intervention case and has a larger standard deviation. The negative numbers on the x axis are percent decreases relative to the no intervention case and indicate that the replications performed for this experimental case have less deviation from the mean. As can be seen in each of the figures, most of the interventions lead to decreases in standard deviations for knowledge, belief, and activity. This suggests that the estimates obtained for the intervention cases are more reliable than those obtained for the no-intervention case, some of them substantially so. While at first this result seems unexpected – interventions add a certain degree of variability to the society, as agents may or may not communicate with them in each particular run – the effects that are actually occurring in the society have somewhat more nuance and subtlety.

In the baseline case, when there is no intervention, there is a fair degree of variability in terms of the agents which speak with the promoter. In some runs, there were relatively few agents that speak with the promoter: in one of the no-intervention runs, only 5.1% of the matching agents in the society contacted or were contacted by the promoter, resulting in only a small amount of belief and knowledge information into the society. In other runs, nearly double that number was contacted, resulting in much larger numbers of agents knowing how to perform the activity and holding the belief. This variability was due to differences in initial social knowledge, differences in the underlying social network, differences in communication patterns and messages conveyed during the virtual experiment, and even randomness in terms of message content. Thus, some runs could end up with very low knowledge and very low belief, while others could end up with a large amount of knowledge and a large amount of belief. The lack of moderating factors could yield results with a relatively large standard deviation as well as a fair number of outlying agents.

Interventions, when present, could also serve as sources of knowledge and belief in the society. Interventions could increase the amount of knowledge in the society by providing it directly to agents, or they could decrease the amount of knowledge in the society by deterring communication with the promoter. Interventions could also decrease belief by infusing additional anti-activity belief information. All three of these factors served to decrease the number of high and low outliers observed. In terms of knowledge, as represented in Figure 6 differences could be seen in terms of variability because of the removal of the high outliers. Indeed, a substantial portion of the drop that occurred because of the elimination and reduction of high-knowledge outliers. Similarly, the decreases observed in Figure 8 was partially associated with the fact that most interventions had replications in which the highest percentage of agents who held the belief was much less than the 5.4% observed as the highest percentage in the no-intervention case. All of these together helped to contribute to a situation in which the size of the standard deviations from mean behavior decreased as additional interventions were added, some of which had mean estimates that were nearly 20-30% more reliable than the no-intervention case.

7 Virtual Experiment Ib: Real Cities

The virtual experiment performed examined four real cities in order to determine the various effects of agents who performed the action. Four real cities were replicated multiple times in order to understand how the effects on stylized cities were affected by the presence of actual US populations derived from PUMS data. These cities were Hartford, CT; San Diego, CA; Orlando, FL, and Kansas City, MO.

7.1 Agents Used

Unlike past virtual experiments, which used a stylized city and over-sampled to ensure that a large number of the agents in the population were able to participate in the activity, the real cities were created by examining PUMS data gathered from the US census [14]. Additionally, unlike in past experiments where the distribution of agent attributes was done independently without respect to any of the other parameters, the real cities had a more accurate representation of a city population at a per-agent level.

Figure 2, Figure 3, Figure 4, and Figure 5 show the agent breakdowns for the four cities used. As can be seen from the diagrams, the actual aggregate differences between the cities were

relatively minimal. The largest difference between the cities was in terms of race: Hartford was the most homogeneous city, with nearly three quarters being of the same race, while San Diego was the most racially diverse with nearly half the population being non-white. There were also small differences in terms of education levels and income levels among the cities, but most socio-demographic slices for a city were within three percentage points of those for any other city. Nevertheless, these aggregate statistics do not present the true differences between the cities. Some of the cities had closer ties between education and income than other cities, which meant that the agents of the same income level were more likely to communicate among themselves due to the additional education-based socio-demographic similarity. In cities in which agents were heterogeneous by socio-demographic attribute, it was substantially less likely to have agents talking to agents in different slices of the different pie charts – agents with different incomes would communicate based on the increased similarity due to identical education levels, for example. Such interactions could lead to differing outcomes for different types of interventions: as low income agents were statistically more likely to lack access to web page interventions, cities that had low income but more-highly-educated agents would be more likely to have interaction between low- and other income-level agents, leading to increased effectiveness for the intervention among such groups.

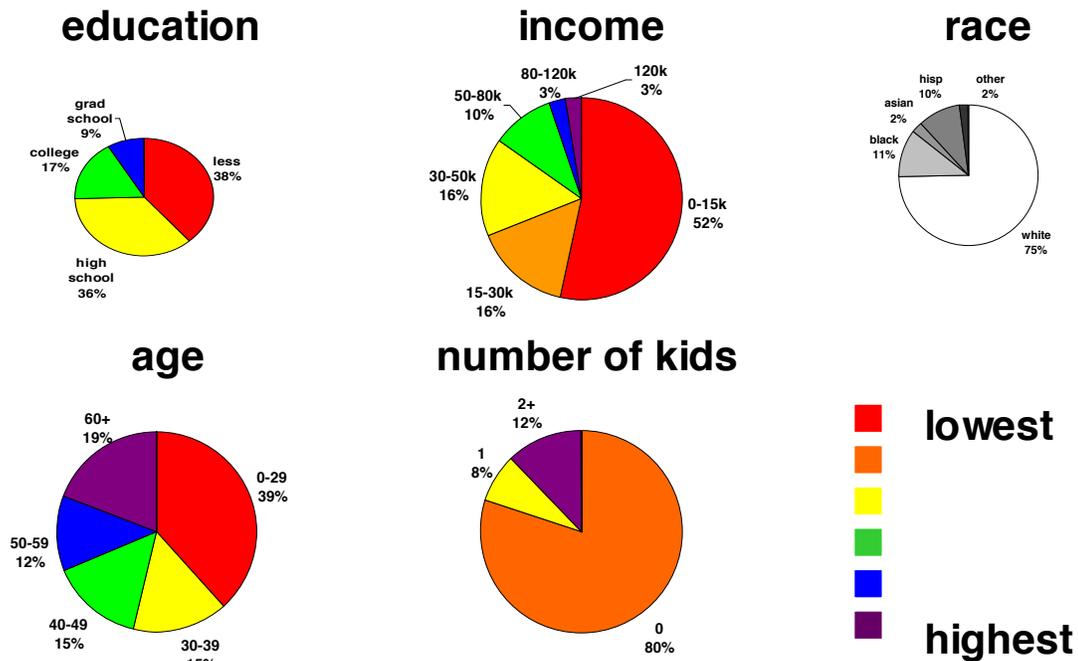Figure 2: Socio-demographic breakdown of Hartford, CT

Figure 3: Socio-demographic breakdown for San Diego, CA
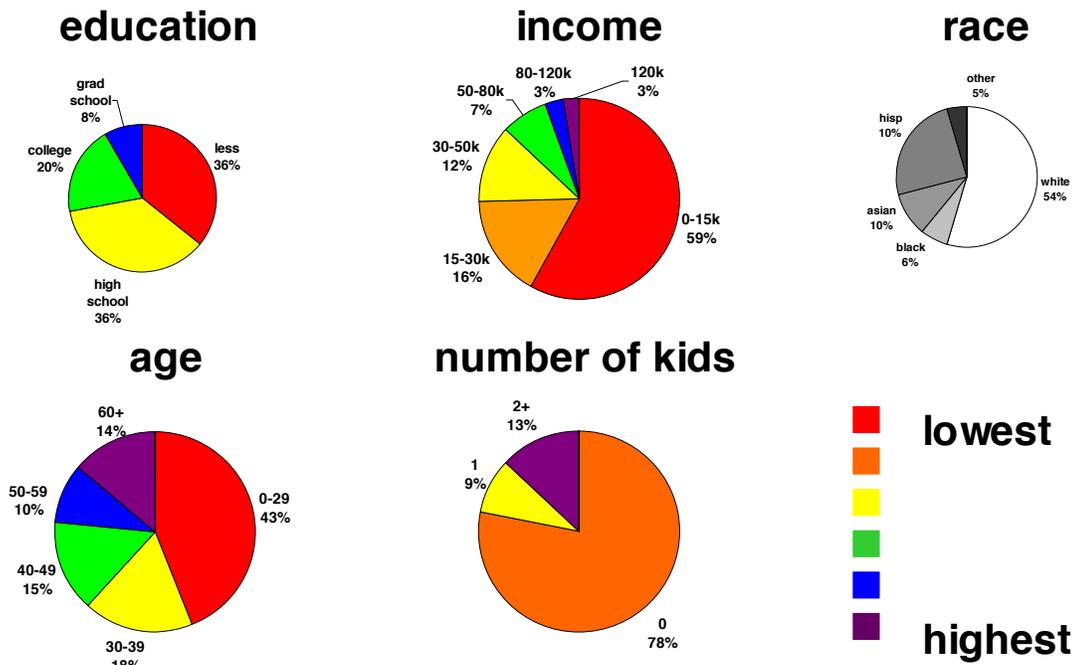

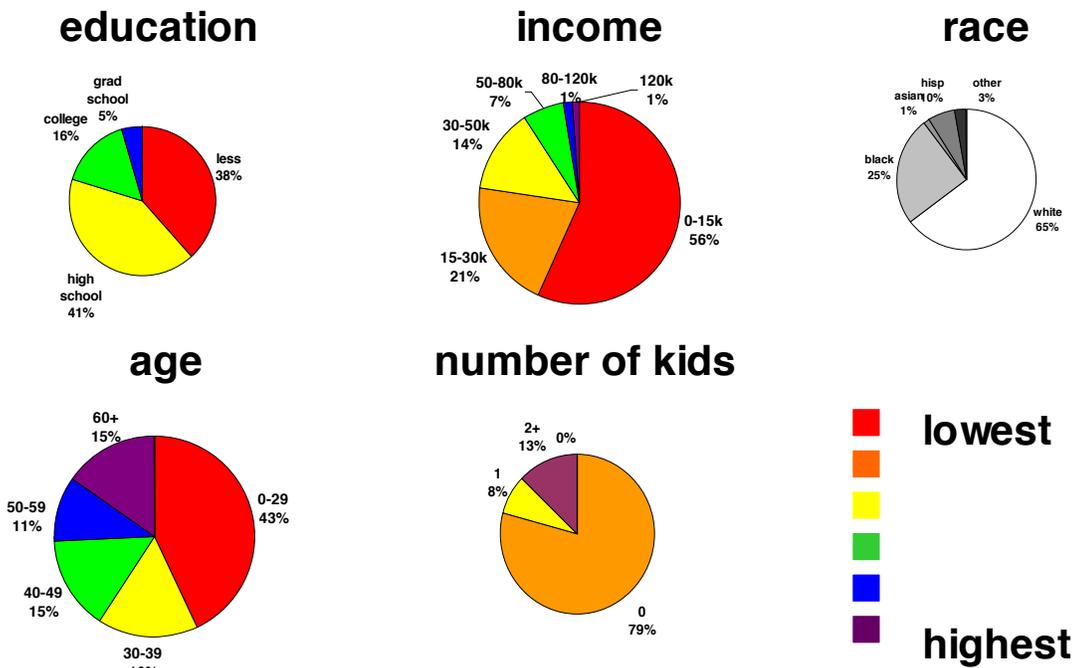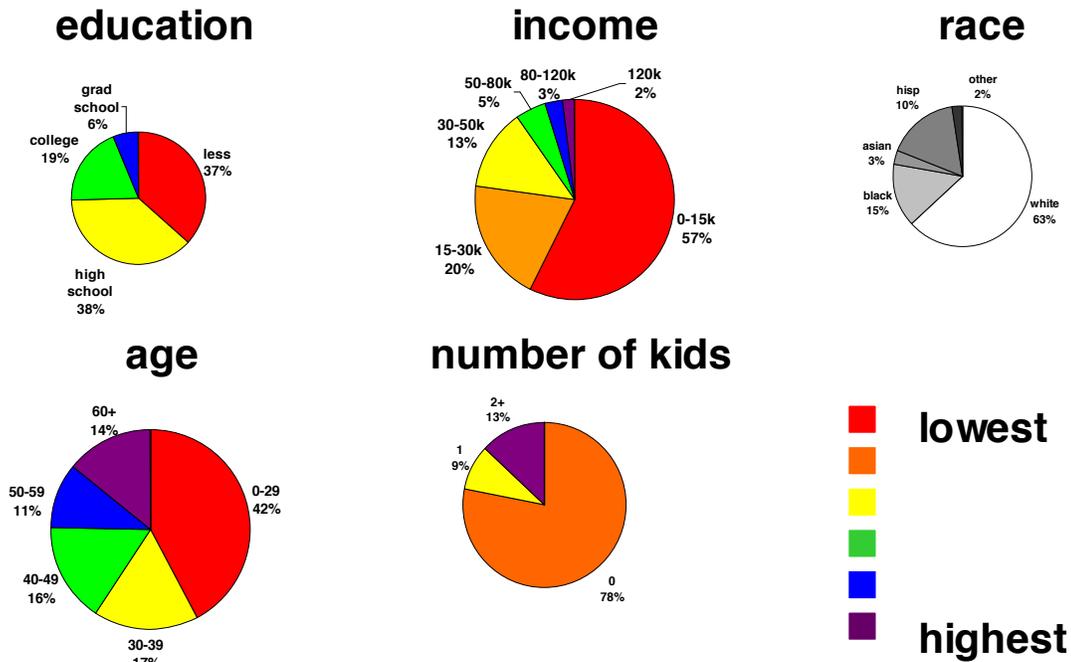Figure 4: Socio-demographic breakdown for Orlando, FL

Figure 5: Socio-demographic breakdown for Kansas City, MO

Unlike the experiments discussed in Section 5 the cities were not directly over-sampled to include low-income and low-education individuals. Still, there were a large number of such individuals in each of the cities. However, in these cities it was not possible to compute a matching function for the number of agents who performed the action due to the fact that the fact that different attributes were used for these agents when compared to the agents in the general society. For this reason, the statistics presented in Section 7.2 use the total number of agents who knew of the activity, held the belief, or performed the action for the entire society.

7.2   Experiment Design

The experiment results were first generated as a 12x5x4 matrix, with twelve interventions crossed by five different access mechanisms crossed by four real cities. The script creating the input deck did create the full factorial design, even cases (such as Internet access with the web site intervention) that would be no different than another case (here, the Internet access without the access intervention). A total of about one hundred and fifty experiment input decks were created and run on machines local to Carnegie Mellon University. A total of fifty iterations were performed for each input deck for a total of 7,500 Construct runs.

Computation time took about two weeks on three servers available at the CASOS center at Carnegie Mellon. In comparison to other experiments which were run on the TeraGrid super computer cluster at Purdue University, this experiment performed much more slowly since the runs could only be replicated across a handful of systems, rather than drawing from a pool of about a thousand, machines [13]. The local machines used were three four-processor machines with 64GB of RAM each.

The same cases run using the stylized city were again run for each of the real cities; Table 5 summarizes the modeled cases which contained the most realistic results examined – the results

Table 7: The most and least effective intervention combinations for real cities

Behavior Only	Hartford	San Diego	Kansas City	Orlando
Most effective	all	mail + web	web alone	all but ad
mean	0.543%	0.546%	0.527%	0.557%
std dev	± 0.144%	± 0.149%	± 0.151%	± 0.165%
Second most	none	ad alone	all but ad	ad alone
mean	0.543%	0.560%	0.545%	0.563%
std dev	± 0.134%	± 0.146%	± 0.150%	± 0.182%
Third most	mail + web	all	radio + web	ad + web
mean	0.550%	0.561%	0.548%	0.564%
std dev	± 0.116%	± 0.146%	± 0.146%	± 0.179%
Least effective	radio + web	radio	ad alone	web
mean	0.587%	0.614%	0.612%	0.603%
std dev	± 0.145%	± 0.149%	± 0.138%	± 0.155%

with the literacy, internet access, and newspaper readership constraints enabled. Any case with the advertisement, web page, or mail message had the literacy mechanism active; any case with the advertisement had the newspaper readership mechanism active, and any case with the web site had the Internet access mechanism active. In some intervention combinations, all three of the mechanisms could be active simultaneously. However, only relevant interventions were affected each individual mechanism even when multiple mechanisms were active, as was done in the stylized city case.

8   Results for Real Cities

8.1   General Effects

The experiment results were examined for trends and to determine the best intervention combination for the society. While there was considerable variation in the effectiveness of an intervention over the course of the two hundred and fifty replication experiment, some general trends were apparent. The most effective interventions examined, as well as the number of agents who knew of the action, had sufficient belief to complete the action, and actually followed through with the action are summarized in Table 7.

In reading Table 7, it is important to recognize three important factors. First, the baseline cases for each of the cities are different, as will be discussed in detail shortly. Thus, while the all-intervention case was the most effective in Hartford, it was essentially identical to the no-intervention case there. On the other hand, about half of the interventions in the other cities lead to decreases in activity. Second, the individual cities examined were each slightly different in terms of socio-demographic composition. While they differed at the aggregate level, as highlighted in figures Figure 2 through Figure 5, they also differed at the individual level. This could lead to slight differences in the effectiveness of different interventions in different cities. Last, only fifty iterations of each virtual city were run for this set of experiments in the most realistic case. Thus, there is more room for noise effects and individual differences to manifest themselves in this experiment. The standard deviations observed in these cities are in some cases smaller than the variances observed in the stylized city of Table 6, although the experimental power of the results obtained may not be as great.

Nevertheless, Table 7 illustrates a number of important points. First, the interventions that are effective in one city may not always be effective in another city. For example, the web

intervention in Kansas City lead to a nearly 10% drop in activity, resulting in the lowest mean activity level of any of the interventions examined in any of the virtual experiments examined. On the other hand, the same intervention in Orlando performed the worst, with the mean percentage of agents performing the activity rising by more than 5%. The same intervention setup was used in both cities: the intervention had the same percentage of action-related facts, the same percentage of belief-related facts, and the same percentage of social facts. While the aggregate socio-demographic breakdowns of each of the cities differed slightly, the underlying socio-demographic composition at the agent level differed substantially more. This lead to striking differences in the results, as the tables and figures demonstrate.

Secondly, in comparing the effective interventions in this table to the stylized city discussed in Section 6, it is clear that some of the most effective interventions for the stylized city are not present in the table while other interventions which were ineffective for the stylized city. For instance, the advertisement was the least effective intervention for the stylized city, yet the second most effective intervention in both San Diego and Orlando. Though minor fluctuations in knowledge and belief may have played a part in making the advertisement so effective in these real cities and so ineffective in the stylized city, it is important to note the socio-demographic distinctions between the cities examined. In the stylized city, there was no correlation between any of the attributes, meaning that the small amount of information coming from the advertisement may have circulated in some social circles and the knowledge from the advertisement circled. This was due to the assumption that agent attributes were uncorrelated, an assumption that increased socio-demographic similarity between agents of different categories due to the increased chance of overlap in other socio-demographic factors. When socio-demographic attributes became more tightly coupled, there were more pronounced differences in literacy rates, differences in preferred socio-demographic partner selection, and inter-group knowledge – factors which could constrain information and knowledge flow to stay within the same socio-demographic group. Something like an advertisement intervention, which provided minimal knowledge and belief information, could decrease an activity mainly by providing a minimal amount both to the same group of agents. Provided that the agents in the group learned both the knowledge and the belief from the advertisement or via second-hand information, a decrease could still be observed. In a stylized city in which knowledge could go to one place and belief to another, the chances of creating agents who learn only the knowledge but not the belief could be greatly increased.

When the virtual experiments were performed, the background cases for the individual cities differed quite substantially by the baseline number of agents performing the activity. For instance, while the baseline cases for San Diego and Kansas City were close to equal – 0.581% and 0.579% respectively – the baseline activity for Orlando was slightly less at 0.573% and that for Hartford was even less at 0.543%. Some of this difference is due to individual fluctuations in each of the runs, as none of these differences are statistically significant for the number of runs performed. However, some of this difference may have been due to minor fluctuations within the cities themselves, as some populations may have been more easily able to spread information coming from the promoter as compared to others. In Hartford, for instance, the agents with which the promoter initially communicated may have been of minority races or may have been moderate-income agents, factors which might have inhibited the flow of information to the majority of the low-income agents in the society. Alternatively, the Hartford agents may have been less likely to pass on information internally within the society. A moderately low standard

deviation in the Hartford case (±0.134 percentage points) may suggest that the low mean is due to the lack of high outliers, a factor which may have had some influence on the observed results.

Figure 9 through Figure 12 at the end of this document present reliability versus effectiveness plots for each of the four cities run. Figure 9 presents the reliability versus effectiveness plot for activity in Hartford. As can be seen in the plot, most of the interventions lead to an increase in the percentage of agents who performed the activity and most of them lead to decreased reliability of the intervention. Some of this may be due to the fact that the no-intervention case was much lower than the other base cases, as discussed before. However, other factors may be due to changes in knowledge and belief in the city. While the changes in knowledge were not statistically significant for the city, the decrease in belief was; the all intervention combination and all of the combined intervention cases lead to a statistically significant reduction in the total number of agents who believed that the action was right. Nevertheless, while there were overall decreases in the number of agents who held the belief, there was no substantial decrease in the number of agents who acted upon the belief. The all intervention combination lead to a total number of agents participating that was about as low as the total number of agents participating when no intervention was enacted. The mail + web intervention combination lead to a 1% overall increase in activity, while the ad alone, the ad + web combo, the mail + radio + web combo, and the all but ad combos lead to increases in the 2-3% range. Generally speaking, the standard deviations on these results were greater than the results of the no-intervention case, suggesting that the interventions were less reliable than letting the promoter run its course in the city. While some interventions such as the mail + web and the all but ad were able to decrease the variability of results, they usually did so by providing a guarantee that the higher outcomes were more likely, not by making the low-activity outcomes occur more often.

The effects of the interventions on San Diego, as presented in Figure 10, suggest that interventions lead to larger percentage decreases than observed in Hartford. The most effective intervention, the mail + web intervention, lead to a decrease of nearly 6% as compared to the no intervention case; most other intervention combinations contributed to similar decreases. Other interventions which lead to a mean decrease in belief were the ad, all but ad, all, web, ad + web, and mail + radio + web interventions, most of which decreased the activity slightly. As in the Hartford case, the effects of the interventions on knowledge were not statistically significant but the effects on belief were; the all and all but ad interventions were able to demonstrate a statistically significant decrease in the number of overall agents who held the necessary belief to perform the activity. However, these decreases in belief did not lead to overall decreases in the amount of activity observed – while decreasing the number of agents who held the belief was important, it was neither a necessary nor a sufficient condition for decreasing the number of agents who were both knowledgeable and believed the action was right (though it often could be correlated with such changes). In San Diego, the interventions lead to situations in which the overall effects of the interventions were less reliable than the effects of not performing any interventions. While the ad + web case lead to only a slight increase in the standard deviation, most cases lead to about a 10% increase in the standard deviation, meaning that the mean estimates for the interventions became slightly less reliable and more prone to fluctuation. As the standard deviation in the no-intervention case in San Diego was similar to the no-intervention cases in the other cities (three of the four had deviations of about ±0.135 percentage points, the same as San Diego), this suggests that interventions in San Diego were more likely to have large fluctuations. While some initialization scenarios might lead to large decreases in the percentage

who perform the action, other scenarios might be likely to lead to no observed changes or even to increases.

Of all the cities examined in this second set of virtual experiments, Kansas City was the most like the stylized city examined earlier. Four of the top five intervention combinations for Kansas City were the same as those observed for the stylized city, though the exact ordering and relative strengths of these interventions and intervention combinations differed.

Figure 11 presents the reliability versus effectiveness plot for the number of agents who performed the activity in Kansas City, suggests that the web site intervention was the most effective, followed by the all but ad intervention and the radio + web intervention. All three of these interventions lead to 5-9% decreases. The differences in belief for the web and the all but ad interventions were statistically significant, factors which may have been associated with the decrease in activity observed since the decreased belief may have lead to situations in which fewer knowledgeable agents were able to take part in the activity. No statistically significant decreases in the number of knowledgeable agents were observed, though most of the most successful interventions were associated with small decreases in the total number of agents who knew how to perform the activity. Yet another similarity with the stylized city – and a difference from both the Hartford and San Diego cities – was the fact that many of the more effective interventions lead to increased reliability of the mean estimates. The means on the web, the all but ad combo, and the radio + web combo all were 5-10% more reliable than those on the no-intervention case, a result which was partially caused by the ability of these interventions to decrease the number of high outliers.

Figure 12 presents the reliability versus effectiveness graph for the last city modeled, Orlando Florida. In this city, only four of the interventions decreased the number of agents who performed the activity: the all but ad intervention, the ad alone, the ad + web combo, and the radio + web combo. While the advertisement intervention alone and the all but ad intervention combination both lead to a net decrease in the mean percentage of agents performing the behavior, the joint effect of the all intervention combination was a very slight increase in the number of agents who performed the behavior, though as before all values examined were not statistically significant. The only statistically significant decreases in belief were due to the all intervention combination, the all but ad intervention, the mail + web intervention, and the mail + radio + web intervention. Nevertheless, only one of these four was effective at reducing the overall amount of activity in the society; the remainder actually led to a slight increase in behavior. The reliability estimates for the behavior decreased relative to the baseline case; the standard deviations increased relative to the no-intervention case. The most effective interventions had results which were nearly 30% less reliable than the baseline cases, due to the fact that some of the decreases were due to several cases in which there were large decreases in the number of agents who performed the action. Most of the interventions had slight increases in the standard deviations, but these interventions in particular lead to large increases.

8.2 Comparison of Stylized and Real Cities

While the real cities and stylized cities vary somewhat in their compositions, it is worthwhile to compare the effectiveness of interventions among the different cases model summarize overall conclusions. Figure 8 attempts to do just this. It brings the top five interventions together and compares them for each of the five cities modeled, presenting the ordinal ranking of each of the interventions.

Table 8: Intervention effectiveness at deterring activity

Intervention Effectiveness	Ad + Web	Web	All	All But Ad	Mail + Web
Stylized	1st	2nd	3rd	5th	10th
Hartford	5th	8th	1st	7th	3rd
San Diego	5th	10th	3rd	4th	1st
Kansas City	3rd	1st	6th	2nd	7th
Orlando	3rd	12th	11th	1st	4th

Two caveats are necessary when examining Table 8. First, it is important to note that the table does not account for the relative effectiveness of each of the interventions. For instance, the web intervention is significantly better than the all but the ad intervention in Kansas City, while the ad + web and the web alone interventions are effectively identical for the stylized city. As such information is not clearly presentable in this table format, the reader should keep in mind that the differences between each of the places within each of the interventions may not be identical, and the differences between any two places when comparing between any of the interventions may not be identical either. Second, it is also important to note that the effectiveness scales being compared in Table 8 are slightly different. In the stylized city, the percentages and relative percentages are computed using the entire society. This means that there may be some agents affected by the interventions who are actually not low income and would not have matched the activity had the criteria been used for the real cities virtual experiment. For instance, while the all-intervention combination was the most effective at deterring the activity in Hartford, it is possible that many of the agents deterred were agents who might not have participated because their income levels were too high. Nevertheless, the vast majority of the agents who participated in each of these cities were low income agents who were likely to have been participants had a matching criteria been used. Thus, while it is probably inaccurate to directly compare the percentages observed in Table 6 and Table 7, the relative ordering of the interventions was less likely to greatly changed by the use of matching or non-matching agents.

Keeping these warnings in mind, it is still possible to analyze Table 8 and extract a variety of meaningful interpretations. As an example, to reiterate what was mentioned before in Section 8, interventions that are effective in one environment may not be effective in another. For instance, the web site was effective in the stylized city as well as Kansas City, but actually increased the number of agents performing the activity in Orlando and other cities. While some of this effect

is due to the micro-variations between runs, it is a testament to the fact that the cities each differ slightly in terms of ethnic, economic, educational, age, and other socio-demographic elements, and this can lead to slightly different interaction patterns in each. The runs performed in this set of experiments do not lead to an overall dominant winner that is effective in every environment; while some interventions may be more effective than others in different places, the interventions are certainly subject to the nuances of the underlying society.

Some general patterns can also be drawn comparing the different cities modeled. First, results observed for Orlando are strikingly different than the results observed for either the stylized city or for any of the other real cities examined in this work. The web and all interventions, which were effective in other cities, were among the worst performers in this environment. The fact that these interventions performed so poorly suggests that there may be something unique about Orlando which may be worth investigating in follow-up experiments to understand if the web intervention is leading to too large of an increase in knowledge which in turn is leading to additional interaction with the promoter. Second, there are several similarities between Hartford and San Diego even while they have relatively different socio-demographic patterns. Both exhibit decreases in activity when either the all intervention or the mail + web intervention are active, do reasonably well with the ad + web and the all but ad interventions, and exhibit relatively large increases in behavior when the web site is modeled. Such results suggest follow-up experiments that seek to try and understand why the two exhibit similar behavior, and to try and understand whether the similarities are spurious or perhaps the result of underlying socio-demographic similarities uncovered in the process of simulation.

While no overall conclusions can be drawn about the overall effect of interventions, several patterns can be observed from the table which – while not true in every case – may provide a reasonable summary of the overall results observed. The web site intervention was a double edged sword: it could have the positive effect of reducing belief, or it could be counter-productive and increase the amount of knowledge available. In conjunction with other interventions – specifically, with either the advertisement alone or with all interventions except the advertisement (the all but ad case) – it had the possibility of being one of the best interventions. The ad + web case was slightly superior to the all but ad case for most of the cities examined, though due to the different reliabilities observed for different cities it is impossible to say exactly which one might be more effective in a general environment. Other interventions not included in Table 7 are generally less effective or less consistently effective than the ones portrayed; nevertheless, it is possible that there exist some specific cities for which such combinations may be worthwhile or at least worth investigating further.

9 Conclusion

This technical report described a series of virtual experiments conducted to try to isolate the effects of particular educational interventions and information access mechanisms on agents with specific socio-demographic attributes. It suggested that different interventions may be more effective for agents of certain socio-demographic characteristics, including age, income, and education. By modeling three different information access mechanisms – literacy, Internet access, and newspaper readership – it was possible to better model some of the nuances that restrict the flow of information in a society. By breaking this down by agent traits, it was possible to examine the effect of the interventions on different groups in the society. While

members of these different groups interact, and the effect of an access constraint on one group may be slightly modified due to interaction with members of another group that does not suffer from that access constraint as strongly, it is still possible to get an approximate idea of what kinds of behavior are possible in a society when different access constraints are enabled.

Additionally, this report sought to describe the parameters used in conducting these virtual experiments in order to document the input decks used in this experiment. As many of the experimental parameters used in this deck have been used for some time, this document also serves as retroactive documentation for some of the input parameters used in past work. It is also likely that many of these parameters may be in use in future simulations in the CASOS center.

It has been suggested that the interventions analyzed here be used in the creation of a stylized city generator in which it would be possible to examine the socio-demographic constraints of an arbitrary stylized city, then use the data gathered from cities like those analyzed in Sections 6 and 8 to try and predict the relative effectiveness of the intervention outcomes. The expectation is that, if successful, such a methodology could be used to help predict the types of intervention combinations which might be useful on real cities of an arbitrary socio-demographic composition. These ideas, and others, provide excellent stepping stones for future extensions to this work.

10 Acknowledgements

11 References

[1] Brian Hirshman and Kathleen Carley. Interactions in educational interventions. *Submitted to IEEE SMC*, pages 1–20, 2008.

[2] Brian Hirshman and Kathleen Carley. Modeling information access in construct. Technical report, Carnegie Mellon University School of Computer Science, May 2008.

[3] Kathleen Carley. Communication technologies and their effect on cultural homogeneity, consensus, and the diffusion of new ideas. *Sociological Perspectives*, 38(4):547–571, 1995.

[4] Kathleen Carley. A theory of group stability. *American Sociology Review*, 56(3):331–354, June 1991.

[5] Brian Hirshman and Kathleen Carley. Specifying agents in construct. Technical report, Carnegie Mellon University School of Computer Science, April 2007.

[6] Internal Revenue Service. *EITC for individuals*. March 2008.

[7] National Assessment of Adult Literacy. A first look at the literacy of america's adults in the twenty-first century. *Institute of Education Sciences*, pages 1–28, 2006.

[8] Pew Internet. Pew internet and american life project. Web site, viewed 12/12/2007., 2007.

[9] Mediamark Research Inc. The daily and sunday newspaper audience report 2007. *Interactive Market Systems, Inc*, pages 1–31, March 2006.

[10] Paul Erdos and Alfred Renyi. On random graphs i. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.

[11]   David Wegner. Transactive memory: A contemporary analysis of the group mind. In B. Mullen and G.R.Goethals, editors, *Theories of group behavior*, New York, 1987. Springer-Verlag.

[12]   Gretchen Livingston. *Latinos Online*, pages 1–23. Pew Internet & American Life Project, 2007.

[13]   Charlie Catlett et al. Teragrid: Analysis of organization, system architecture, and middleware enabling new types of applications. *IOS Press*, 2007.

[14]   United States Bureau of the Census. "US Census Data for the Year 2000: 5% PUMS data". August 2003.

Figure 6: Reliability versus effectiveness plot for knowledge, stylized city
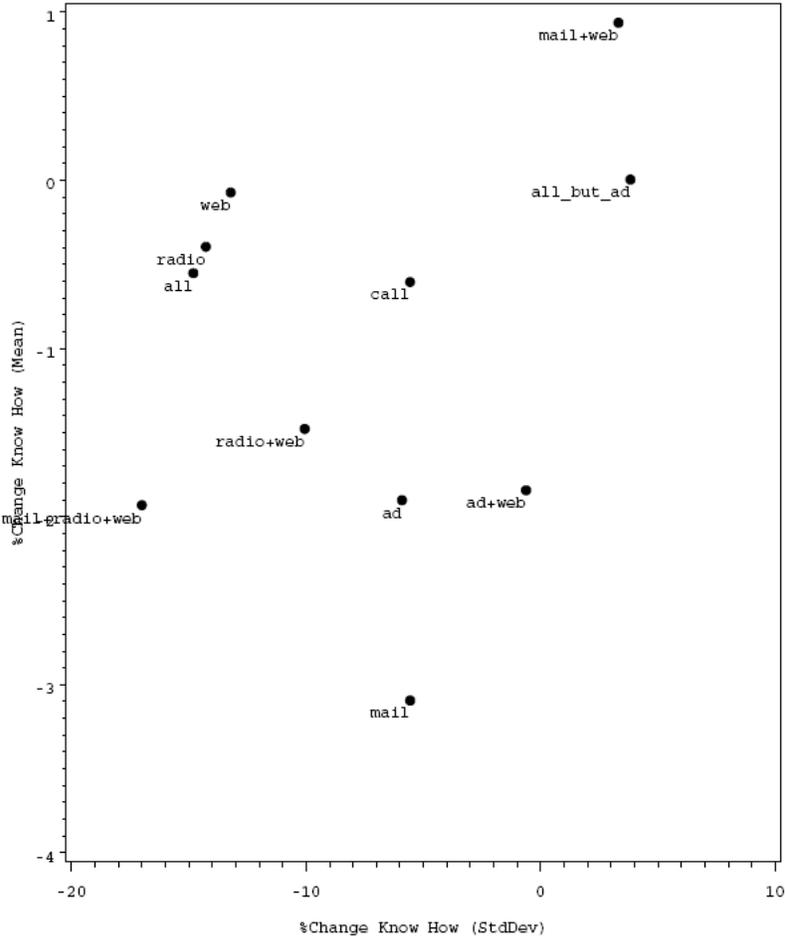
Figure 7: Reliability versus effectiveness plot for belief, stylized city
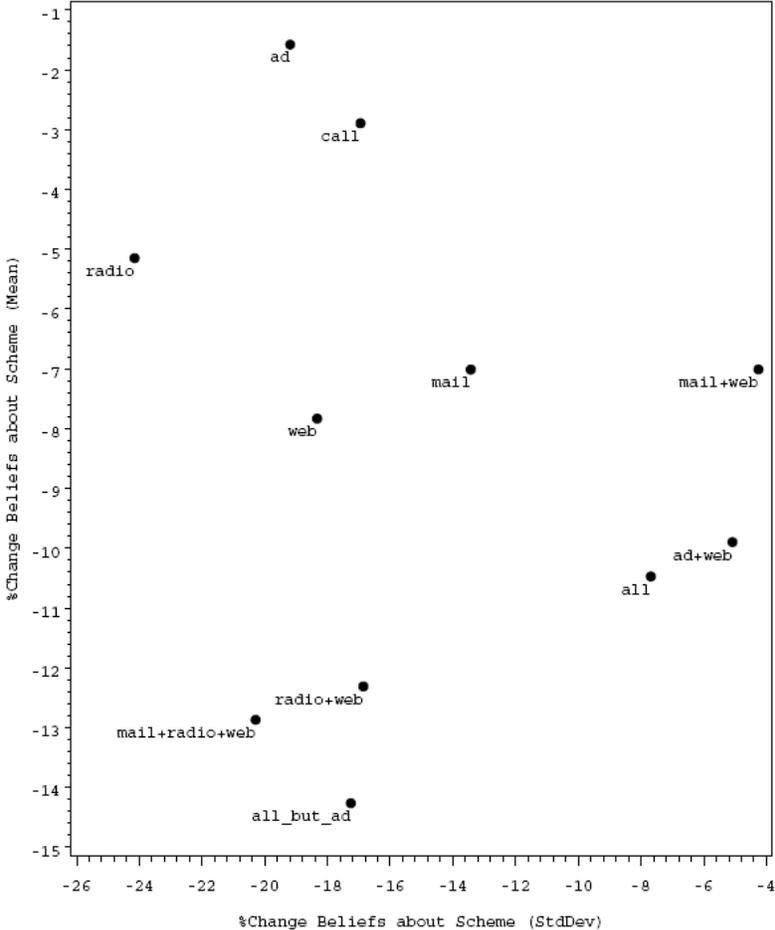
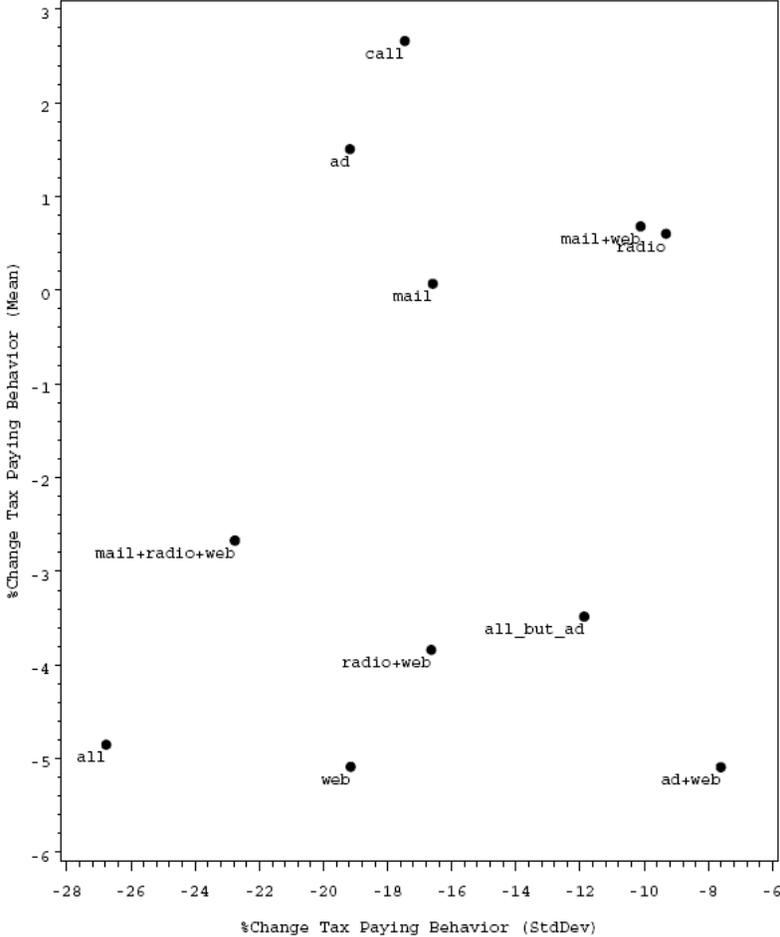Figure 8: Reliability versus effectiveness plot for behavior, stylized city

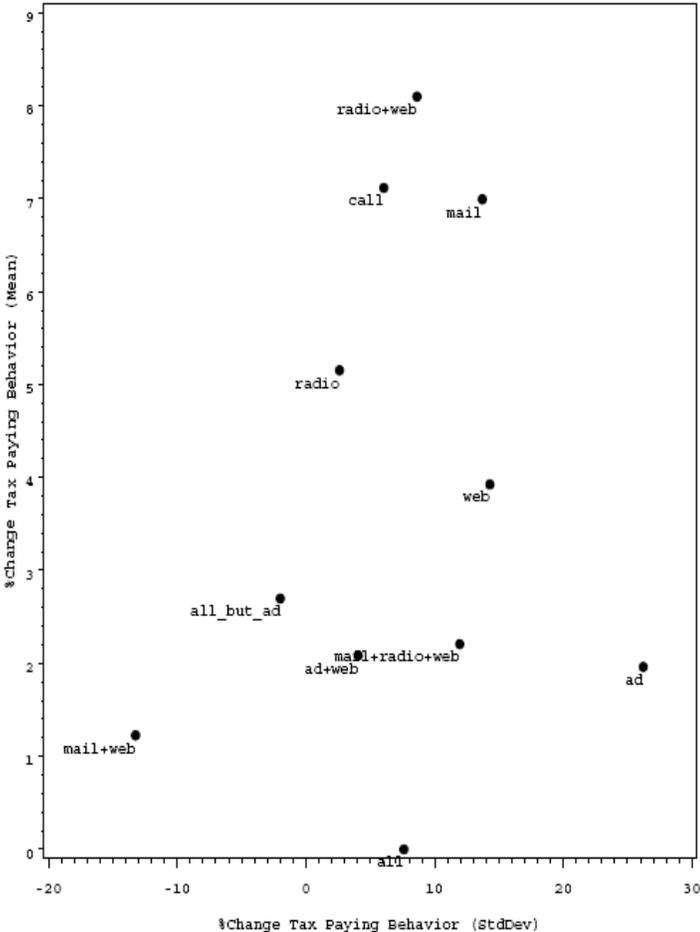Figure 9: Reliability versus effectiveness plot for behavior, Hartford

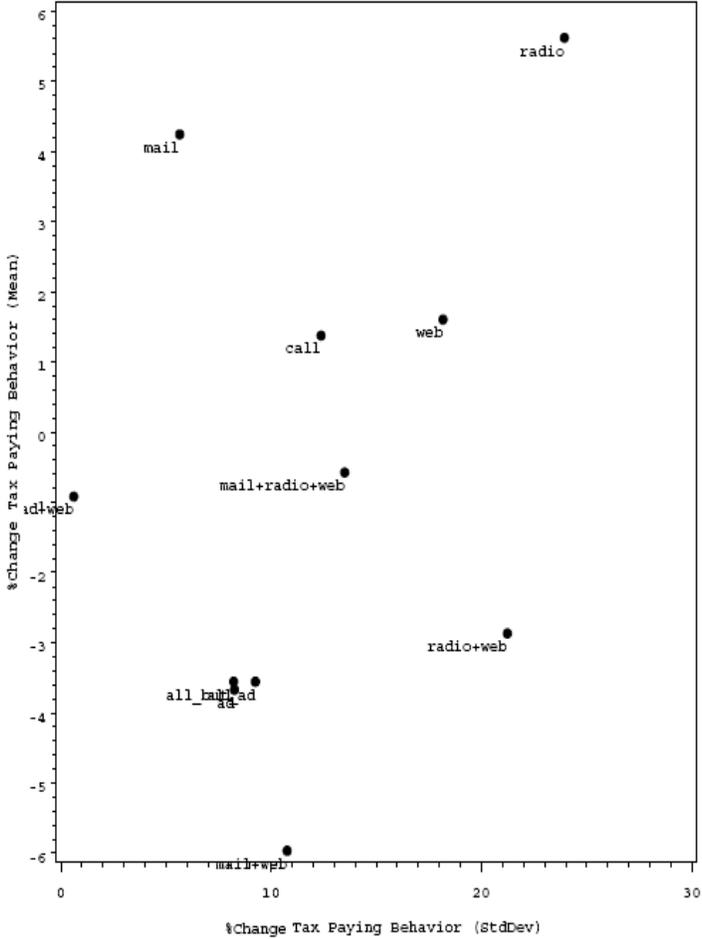Figure 10: Reliability versus effectiveness plot for behavior, San Diego

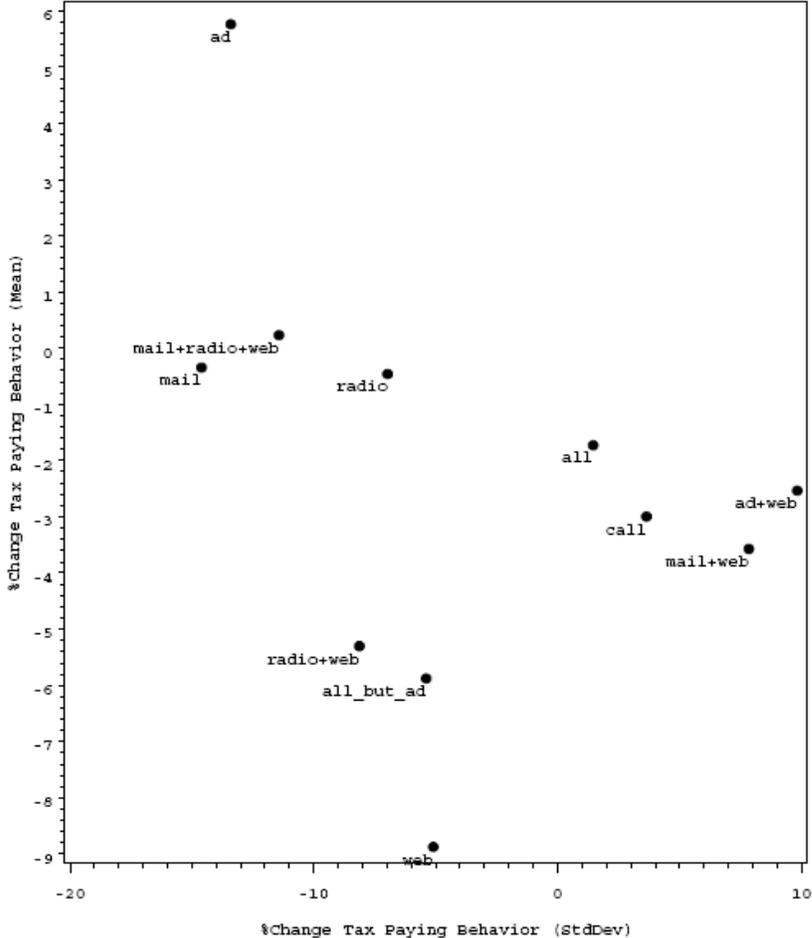Figure 11: Reliability versus effectiveness plot for behavior, Kansas City

Figure 12: Reliability versus effectiveness plot for behavior, Orlando