# Understanding Taxpayer Behavior and Assessing Potential IRS Interventions Using Multiagent Dynamic-Network Simulation

*Kathleen M. Carley, Carnegie Mellon University, and*
*Daniel T. Maxwell, Innovative Decisions, Inc.*

The IRS Strategic Plan points to improving taxpayer compliance as an important goal of the Service. The traditional response to improving compliance in the area of taxes and other regulated social behaviors (e.g., traffic control, drug use) is to increase enforcement activities. Often, these efforts are targeted on specific segments of the population in an effort to achieve maximum effect as efficiently as possible. The need for these efficiencies is usually a direct result of the need to allocate finite resources across a very large, heterogeneous population. The heterogeneity of the population further complicates the Service challenge because different types of people have different reasons and intentions that cause them to be noncompliant; they communicate those motivations and knowledge differently; and they respond differently to interactions with the Government, such as enforcement or other types of interventions.

This paper describes how emerging research in the computational social sciences, specifically the combination of multiagent simulations and dynamic social network analysis could assist the Service in better understanding taxpayer behavior, as well as how taxpayer behavior changes in response to their interactions with others in their social and family circles, perceived tax experts, and various types of potential IRS interventions.

Achieving maximum voluntary taxpayer compliance is an important goal of the IRS. This compliance can be organized into three multifaceted components: filing, reporting, and payment (Brown and Mazur, 2003). This research effort is focusing on improving Service understanding and effectiveness in the area of reporting compliance. Early analysis of the 2001 National Research Program (NRP) indicates that this segment of the compliance challenge is responsible for approximately $80 billion annually in underreported income and $5 billion annually in overreported income (Bennett, 2005). The revenue implications of the underreported income are significant. Additionally, both types of compliance shortcomings could potentially serve to undermine taxpayer confidence in the voluntary tax system with devastating long-term consequences.

From the perspective of the taxpayer, tax laws are often seen as complicated and difficult to interpret. Because of this perception, many people turn to others to determine what taxes they should pay or they might make mistakes in filling out tax forms. In some cases, people consult with reputable tax professionals; sometimes, they consult with friends or family; and sometimes, information is obtained from unethical people who prey on people's ignorance of tax law. As a result, tax opinions and decisions are based on a collection of information and misinformation. This could cause even well-intentioned individuals to sometimes over- or underpay taxes by taking, or failing to take, appropriate tax liability adjustments. Another motivation for some people might be that taxes appear to be an unfair and unacceptable burden or depreciation of personal wealth. In such cases, individuals may knowingly take part in inappropriate or even illegal tax avoidance schemes.

To meet these challenges, the IRS provides a number of education services that are in many service channels, ranging from Web sites and special tax preparation classes or seminars to Walk-in Assistance Centers. Interventions such as these educational products and tax payer assistance services, as well as enforcement measures such as audits and arrests, are intended to increase awareness of who should pay what taxes and to increase voluntary compliance across the taxpaying population.

These interventions vary in both cost and effectiveness. Their success depends, in part, on reaching the right group of potential taxpayers at the right time and mobilizing the taxpayers' own social networks so that the broader population becomes self-educating in appropriate tax behavior. The success of the different interventions also depends on the particular inappropriate behavior the Service is trying to curtail. For example, there are four basic cases of tax-compliant behavior that need to be accounted for: those who pay appropriately, those who underpay while thinking they are entitled to an adjustment to tax liability but are not, those who overpay while not taking a legitimate adjustment, and those who underpay intentionally. These groups have different motivations and tend to have different social networks, as well as different communication habits. Consequently, different types of interventions will have different levels of effectiveness in improving (or reducing) compliance across these four groups. Finally, the success of the different interventions depends on the way in which general tax-related information spreads through the taxpaying population, independent of Service interventions. That is, the nature of the social networks used to encourage fraud or to propagate misinformation about tax preparation will in turn impact what type of interventions are most likely to stem the tide of misinformation.

What is needed is a way of pre-assessing these interventions prior to their use, in terms of their likelihood of success. However, assessing the impact of these interventions is difficult. Even after an intervention has occurred, there is often little available data on how it impacted the diffusion of information through the social network about a service, a tax credit, the illegality of a particular scheme, etc. Further, such data are rarely captured at a fine enough level of detail and in sufficient quantity to enable the systematic evaluation of future interventions prior to those interventions being used. As such, there is a need for a systematic approach to thinking through intervention strategies.

Dynamic-network simulations can be used to effectively and systematically evaluate the relative efficacy of different intervention strategies. Dynamic-network simulations (Carley, 2003) are multiagent simulation systems in which the agents are enabled and constrained by their positions in dynamic metanetworks that include both social and knowledge networks. Such simulation systems provide a framework for characterizing differences in populations, tax credits, fraudulent schemes, and interventions and then assessing how these differences play out over time in affecting both knowledge about tax law and the level of compliance. The strengths of such an approach include the ability to: characterize the dynamic behavior of large heterogeneous populations, rapidly and systematically assess novel types or timings of interventions on the population as a whole as well as on targeted subgroups, and engage in proactive planning.

The remainder of this paper demonstrates how dynamic-network simulations can be used to evaluate intervention strategies. First, background on networks and dynamic-network simulations are described. Second, a high-level description of a specific dynamic-network model is described. Then, a virtual experiment for assessing the impact of an intervention is defined and the results presented.

## What Are Networks?

A network is a set of nodes and relations; graphically, this looks like a set of dots and the lines connecting them. Networks of many types are a ubiquitous feature of human life (Carley, 2002). Herein, we are primarily concerned with three types of networks: the social network, the knowledge network, and the beliefs network.

Consider first the social network; i.e., who talks to whom. For example, humans are connected through family ties, work relationships, and friendships into a vast social network that impacts all aspects of life from who has access to what information when, to who will watch each other's children, to who will infect whom with what disease. Individuals are more likely to be connected with others if they are related; share the same race, gender, or age; have gone to the same school; or work in the same area. Within this social network, some individuals play more key roles; e.g., salespeople, teachers, and ministers are often connected to more people than is the average citizen. Such people critically influence the flow of information.

With respect to taxpaying behavior, these networks influence the likelihood that people will learn of and engage in various fraud schemes or learn of and take various tax credits. Both promoters of abusive schemes and the IRS use knowledge of the network to design interventions, locate opinion leaders, and tailor activities to increase the number of people who could potentially be "reached" by their messages. While knowing the details of a specific network, exactly who talks to whom may not be feasible, general features of networks, and how they vary by cities can be assessed from the way in which people in that city are distributed across high-level sociodemographic information such as gender, race, age, economics, and occupation. Such indicators give a first approximation of the underlying social network, as there is a general human trait to, ceteris paribus, engage in homophilous interactions.

Another critical network, particularly when considering the diffusion of information and innovation, is the knowledge network (Carley and Hill, 2001). The knowledge network is a network connecting people and ideas. That is, the knowledge network specifies who knows what. An interesting feature of this network is that it evolves as people learn. People of course learn by talking to each other (learning by being told) and by engaging in tasks (learning by doing).

The last critical network with which we will be concerned is the belief network, i.e., who believes what. Like the knowledge network, the belief network changes as people interact, the main difference being that beliefs describe people values rather than their knowledge. In general, people's beliefs are a function of many things including their expertise, their prior beliefs, and the beliefs of those with whom they interact (social influence).

From a network perspective, there are two types of IRS interventions. The first aims at altering the underlying social and knowledge networks simultaneously. Examples of this might be when the availability of services is adjusted or when an enforcement action is taken. The second aims at altering the knowledge network by "educating" people about tax law. In both cases, these changes may alter not just the knowledge network (who knows what) but also the belief network (who believes what).

## Dynamic-Network Simulations

Some research exists on the use of multiagent simulations to explore the effects of enforcement on compliance. Davis, Hecht, and Perkins (2003) use a multiagent simulation to explore the movement of populations between compliant and noncompliant states. Similarly, Bloomquist (2004) uses a multiagent simulation to explore the impact varying audit rates have on the compliance level of the population. In both cases, information and belief-related concepts are imputed somehow in the characteristics of the agents. The agents move about on a grid with little attention to social networks in which real human actors would have been embedded. This was consistent with the state of the art and practice for the simulation environments that were employed. This research extends that work in three ways. First, we represent explicitly the concepts of knowledge and beliefs, thereby allowing these factors and the consequent behavior to co-evolve over the simulated period. Second, the agents we model do not move on grids but are enabled and constrained by their network positions in networks that dynamically adapt in response to agent behavior. Additionally, this work (compliments of Moore's Law and improving software) significantly increases the fidelity of the population representation.

Dynamic-network simulation systems can be used to examine how networks evolve and change over time and the repercussions of those changes for individual behavior. A dynamic-network simulation is an agent-based simulation in which the agents exist in a multidimensional or "metanetwork" space that changes as they interact. Note, this is in contrast to the traditional multiagent simulations in which the agents populate points in the grid and interact with neighbors or physically move through "squares" on the grid surface. Agent-based simulation systems are valuable for studying complex socio-cultural systems as they admit reasoning about the behaviors of large populations of heterogeneous agents.

To anchor and validate these systems, real world data are used to initialize the model and tune internal processes. The result is a highly constrained system that enables the analyst to explore a wide range of behaviors in a virtual environment. This virtual environment has been narrowed through the initialization and tuning process, such that the range of emergent behaviors is within the range of possibility.

An analyst can use such a model to assess various changes in the environment or systematically evaluate alternative interventions. This is typically done by setting up a series of virtual experiments and then analyzing the resultant response surface. Note, a virtual experiment is an experiment conducted using a computer simulation. For example, an experiment might examine the relative impact of no intervention versus an IRS-generated Web page containing general information versus a targeted ad campaign in a newspaper.

Dynamic-network simulation systems have three key uses. First, the development of the model helps the participants understand the relationships which come together to effect complex behavior, such as failing to take a proper adjustment to tax liability. The process of simply building the model lays bare relationships that may not have been evident before. Second, the model itself supports detailed analysis and enables more systematic evaluation of effects in a way that supports both explanation and forecasting. Finally, because such models can be used to examine a broad range of interventions under diverse sociodemographic conditions, the model can be used to engage in a series of "what-if" analysis sessions and thereby support planning.

## The Model

For this work, we use a multiagent dynamic-network simulation system called CONSTRUCT (Carley, 1991; Schreiber and Carley, 2004). As we are concerned here with the diffusion of information and change in beliefs, we use the Construct simulation as our baseline. (CMU: http://www.casos.cs.cmu.edu/projects/construct/) In Construct, each agent is an information processor who interacts with others, communicates information, learns, and uses their information to make decisions. Construct has been used to examine information diffusion, cultural change, and the evolution of social networks at the small group, corporate, and community level. The basic tenets for interaction are based on well-documented logics for social interaction, specifically, homophily-based and expertise-based interaction. Construct has been used at the societal level to study integration of subcultures, and at varying levels from team to nation-state to understand the diffusion of information and the resultant impact of that diffusion on cultural norms (Breiger and Carley, 2003).

Within Construct, the basic elements are as follows:

- Agents (different types of agents are distinguished by their information-processing characteristics and their knowledge).
- Knowledge (the set of facts that agents either know or do not and that can be categorized into areas such as knowledge that the scheme exists, knowledge about how to take part in a scheme, and knowledge that the scheme is legal or illegal).
- Beliefs (a set of opinions or beliefs that agents hold and that can impact their behavior, such as whether a scheme is legal).

Agents in Construct are sophisticated socially-realistic information-processing agents subject to structural and cognitive limitations on their behaviors, and differentiated from each other in terms of sociodemographic factors. Table 1 illustrates the sociodemographic characteristics that are currently represented in the data.

**Table 1. Demographic Characteristics of Agents**

| Characteristic | Number of Categories |
|---|---|
| Gender | 2 |
| Age | 5 |
| Education | 3 |
| Income | 4 |
| Race | 5 |
| Parent (Dependents?) | 2 |

This population decomposition allows one to represent over 1200 unique types of agents in the simulation. In addition to the taxpaying agents, the simulation currently has three other special types of agents. The first we will call a "Promoter." This agent spreads misinformation throughout the taxpaying population through a series of one-on-one interactions. The second is a "Seminar." Seminars can attract multiple agents at each time step and can be used to serve as a misinformation threat to compliance or a treatment that spreads positive information. The third and final type of agent is called a Web site. This allows taxpaying agents to seek out information either from the IRS or potentially from agents proliferating misinformation.

For Construct, at each time step, agents are selected to initiate communication with other agent(s). This communication is done as follows: An agent is selected, and, depending on that agent's capabilities, that agent might initiate an interaction with one or more others and then communicate one or more facts and or beliefs. For example, a Web site or ad campaign as an agent can send facts to other agents, but it cannot have its facts modified, i.e., the information is read-only, and it cannot initiate an interaction (i.e., it sits passively waiting for others to interact with it). On the other hand, a promoter can initiate an interaction and then communicate beliefs and facts. The likelihood that two agents interact is a function of whether they are available for interaction (i.e., not interacting with others) and their relative similarity/expertise when compared with others. For example, when agents are not actively seeking expertise, they interact with those to whom they are relatively similar (homophily-based interaction). Homophily-based interaction is a function of similarity both in terms of knowledge and in terms of characteristics.

Finally, as a function of their knowledge and beliefs, agents make decisions. The core decision is, of course, whom to interact with when. However, there can be other decisions that dictate different behaviors. For each decision, there is a decisionmaking logic. In the case of interaction, this logic is a function of homophily and expertise. The agents spend roughly 60 percent of their time engaged in homophilous interactions where they interact with those with whom they have much in common. On the other hand, about 40 percent of the time, they actively seek out those whom they believe to have specialized knowledge. Other logics can be instituted for other decisions.

For this project, we operationalized Construct by making the following identifications.

- We set the number of agents in the simulation to be proportional to the number of people in a canonical Midwestern city (Figure 1).
- We defined the agent's characteristics in terms of sociodemographic characteristics that are relatively available to the IRS and that are liable to impact taxpaying behavior, e.g., age, education, race, income, gender, and parental status. The distribution of these categories across the population of agents was proportionate to the real census data for the city in question.
- We segregated knowledge into four categories: knowledge of the scheme, knowledge of how to engage in the scheme, knowledge about the legality/illegality of the scheme, and general social knowledge.
- We identified two core beliefs: belief that the scheme is legal and belief that they should engage in the scheme.
- We identified a single decision other than with whom to interact. This decision is whether to engage in a scheme.
- We instituted a specialized logic for choosing to engage in a scheme.

**Figure 1. Sample Population Distribution**

# Canonical City Demographics

Population Percentages in pums data and number of agents used in a 1,000 agent sample

**white** 93% by pums data (940 agents)

**black** 2.3% by pums data (25 agents)

**hispanic** 0.9% by pums data (17 agents)

**over 30 years old** 88% by pums data (867 agents)

-- lowest age bracket was 0-29 so its unclear what the over 18 percentage was

**married** this was not extracted

**parent** 32% by pums data (319 agents)

The logic for participation is as follows:

1. The agent must know of the scheme.
2. The agent must know enough about how to participate in the scheme to do so, i.e., have 50 percent or more of the "how to" facts.
3. The agent must have the "resources" to pursue the scheme. We assumed that the agent had the resources if the agent was a good match to the sociodemographic group being targeted by the scheme promoter.
4. The agent is essentially a risk taker or has other psychosocial behavioral patterns that lead him or her to participate. We operationalized this as simply a random tendency to participate.
5. The agent must believe that the scheme is legal.
6. The agent must believe that he or she should engage in the scheme.

Then, given these six factors, an agent will participate if the first four conditions hold and either of the last two. This results in agents who can participate and do not; agents who can, believe it is legal, and do participate; those who can, believe it is legal, and do not participate; and those who can, believe it is illegal, but participate anyway.

## Illustrative Results

A key feature of dynamic-network simulation systems is that they can be used to generate a large number of virtual experiments whose results enable a large number of issues to be addressed. Rarely are all ramifications of such models identified and described. Rather, a model, once tuned to fit the known input data, is then used to evaluate a few select issues and in those realms shed some light.

Here, we use the Construct model to examine the relative impact of ad campaigns. We contrast the expected response of the public, as simulated in our canonical Midwestern city, to no intervention, a short ad campaign by the IRS, and a lengthier ad campaign. We chose ad campaigns for this example because they are a common mode of intervention used by the IRS that can vary in intensity based purely on the length of time they are run. Ad campaigns are simple in the sense that they are not directed to a specific subpopulation but are rather directed to the general reading public.

The results illustrated in Figure 2 show that ad campaigns as an intervention are not particularly effective. Rather than reducing participation in a scheme, ad campaigns can actually increase participation. On the surface, this result seems counterintuitive. However, this effect makes sense when we consider how ad campaigns are structured. Ad campaigns tend to be fairly general. They contain information about schemes, discuss their illegality, and also discuss related legal ways of engaging in tax reduction. In other words, they contain information about the scheme and information on alternatives. In contrast, information provided by promoters tend to spread information exclusively on the scheme and how to engage in it. Ad campaigns are not focused on a particular group, whereas, scheme promoters tend to focus on, seek out, and work to engage those for whom the scheme is most relevant.

**Figure 2. Sample Participation Results**

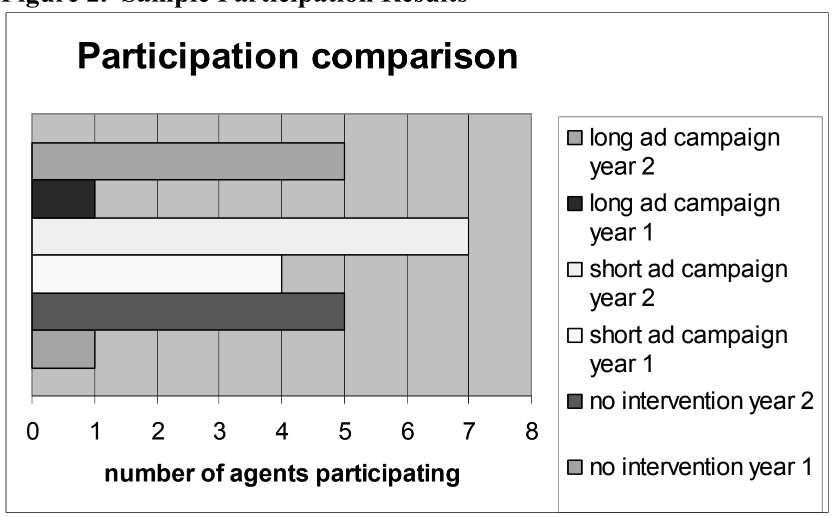


Ad campaigns contain a broader range of information, which is communicated to more of the taxpaying public than in campaigns waged by promoters. Consequently, individuals who know nothing of the scheme can learn of it and choose to engage, simply by reading the IRS ad. This effect is more pronounced in the short term or for shorter adds as there is less time and opportunity for the taxpayer to learn from the add, not just that the scheme exists but that it is illegal.

## Conclusion

This paper demonstrates that dynamic-network simulations can be used to conduct analysis and develop insights into the types of taxpayers most likely to enter into abusive tax deduction schemes, and to assess the relative impact of alternative interventions. Dynamic-network simulations have significant potential for representing the decisions and behaviors of the taxpaying population. Using the model described herein, we have observed a differential spread of information and participation across the target subpopulations, and examined the impact of a wide variety of IRS interventions. We found instances where such interventions had the potential to change taxpayer behaviors--sometimes counterintuitively--as in increasing the likelihood that a taxpayer would engage in a scheme, possibly unknowingly. The promise of such models in general, and Construct in particular, is great.

Models such as Construct can be progressively refined to provide relevant and focused exploration of the interactions of taxpayers, response strategies, and other relevant variables. For example, for this study, refinements of Construct included the tailoring of the population to match a canonical city, identification and characterization of interventions and promoters, and the addition of logic of participation on the part of taxpayers. Refinements such as these increase the relevance and utility of the results and enable more reasoned policy setting.

Additional refinements are of course feasible given sufficient programming time. For example, we could augment the agents to include information on occupation or number of children. We could augment the city level description to include physical locations of convention centers, churches, universities (i.e., locations where seminars might be held), and locations of IRS assistance centers. Large numbers of refinements are of course possible. However, we find that, in general, if the goal is to support policy, it is better to add such refinements slowly and only if the following two conditions are met: 1) the refinement can be supported by empirical data, and 2) the refinement enables using the simulation to reason about an important outcome or behavior that could not otherwise be reasoned about and for which there is some empirical data against which to tune the results. Note, it is relatively easy to build simulation models that are highly complex and have so many features that the results are as difficult to analyze as the real world.

It is also worth noting that there are two key tradeoffs: feature-speed and feature-analysis. Every refinement brings with it one or more new features to the model. As these features are added, on average, the speed of processing slows. As such, the model itself takes longer to run or requires more powerful computers. While it is true that simulations with millions of agents can be run on laptops, such simulations tend to have very simplistic and unrealistic social and cognitive agents. The higher the social and cognitive accuracy of the agents, on average, the longer the simulations take to run. The current model can run a small city in about 30 minutes; however, increasing the size of the city, adding occupation, adding decisions about multiple tax credits, linking populations into family units, and so on, will increase execution time--though how much is unknown. The problem here is simply that the longer the execution time, the longer it takes to generate virtual experiments to test the impacts of different interventions.

The second tradeoff has to do with analysis. As more features are added, more potential analyses are possible. In general, it is easy to add so many features and generate so much data that no existing statistics package can handle all the generated data and that all disk-space on a normal desktop is filled up. The key here is to grow the model in such a way that you get increased veridicality at the same time as you ensure that the results can be analyzed. Further, since an increase in features also tends to decrease speed, you cannot trade speed for analysis and save less output but have more runs.

In part, these dilemmas speak to the state of the art in large-scale computing. Clearly, as we move to grid-based computing, distributed data storage, simulation feeds to databases, and more service-oriented analysis techniques, these tradeoffs will be less pronounced. However, even with such technological advances, we need a reasoned approach to adding features that are empirically driven, particularly when the results are used to inform policy. The need for empirics is driven by the fact that it is easy to add a feature, but adding features in ways that are legitimate means linking them to some form of data whether qualitative or quantitative. For example, it is relatively easy to add occupation and to differentiate groups on the basis of whether the occupation is white or blue collar. However, from a taxpaying perspective, the issue is not white or blue, but more specifically the type of occupation and its relation to income. As such, in information about the relative range of salaries, wages are as important as information about the distribution of occupations across sociodemographic groups. As with most models that are relevant for policy setting, "the devil is in the details," and getting the details to be reasonable requires working hand-in-hand with empirical data.

With these caveats in mind, there are of course clear next steps for Construct. Key features would be the addition of occupation and family groups as this would facilitate examining a variety of taxpaying behaviors, such as those related to credits as well as alternative deductions. Additional interventions, such as IRS service centers and TV commercials, should be examined. The tool as a whole should be linked directly to a database to ease analysis. Technologies for multithreading should be investigated and so on. The key here will be to refine the simulation infrastructure and agent representations in a buildingblock fashion, ensuring that each addition augments results from the one before.

Dynamic-network simulations are very promising tools for examining tax-related issues as they enable refined reasoning about both a set of heterogenous agents and the socio-cultural context, i.e., the networks, they inhabit. Part of this promise lies in their ability to be used for both policy setting and education. As such, it is important that such models be developed carefully and with full attention to the needs of the users and the uses to which they will be put. The value of these models derives both from their results and from the process of development which brings to light the constraints and relations among the various factors influencing taxpaying behavior. As we move to the future, our goal should be the development of a set of simulation tools that provide a flexible and easy-to-use system that can sit on the analyst's desk and enable the analyst through a series of "what-if" analyses to preassess alternative interventions and scenarios relative to specific possible taxpaying behaviors so as to pre-evaluate their efficacy and so reduce the cost of these interventions to the taxpayer.

## Acknowledgments

The authors gratefully acknowledge the support and assistance of the members of the Service's Office of Policy Evaluation and Risk Analysis (OPERA), the members of Carnegie Mellon's Center for the Analysis of Social and Organizational Systems (CASOS), and researchers at Evidence Based Research. Without their collective contributions, this research would not have been possible.

## References

Bennett, C. (2005), Preliminary Results of the National Research Program's Reporting Compliance Study of Tax Year 2001 Individual Returns, Proceedings of the 2005 IRS Research Conference, June 7-9 2005, Washington D.C.

Bloomquist, Kim M. (2004), Multiagent-Based Simulation of the Deterrent Effects of Taxpayer Audits, paper presented at the 97th Annual Conference of the National Tax Association, November 2004, available online at http://www.irs.gov/pub/irs-soi/04blonta.pdf.

Breiger, R. and Carley, Kathleen (2003), Dynamic Social Network Modeling and Analysis. Workshop Summary, The National Academies Press, *National Research Council of the National Academies*, 2003, The National Academies Press, Washington, D.C.

Brown, R.E. and Mazur, M.J. (2003), IRS's Comprehensive Approach to Compliance Measurement, IRS White Paper, Washington D.C.

Carley, Kathleen M. (2002), Smart Agents and Organizations of the Future, The Handbook of New Media, edited by Leah Lievrouw and Sonia Livingstone, Sage, Thousand Oaks, CA, Chapter 12, pp. 206-220.

Carley, Kathleen M. and Vanessa Hill (2001), Structural Change and Learning Within Organizations, in Dynamics of Organizations: Computational Modeling and Organizational Theories, edited by Alessandro Lomi and Erik R. Larsen, MIT Press/AAAI Press/Live Oak, Chapter 2, pp 63-92.

Carley, Kathleen M. (2003), Dynamic Network Analysis, in Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers, Ronald Breiger, Kathleen Carley, and Philippa Pattison, editors, Committee on Human Factors, National Research Council, National Research Council, pp. 133-145.

Carley, Kathleen M.(1991), A Theory of Group Stability, *American Sociological Review*, Volume 56, Number 3, pp. 331-354.

Davis, J.; Hecht, G.; and Perkins, J. (2003), Social Behaviors, Enforcement, and Tax Compliance Dynamics, *The Accounting Review*, Volume 78, Number 1, pp. 39-69.

IRS Strategic Plan 2005-2009, Washington D.C.

Schreiber, Craig and Carley, Kathleen M. (2004), Construct--A Multiagent Network Model for the Coevolution of Agents and Sociocultural Environments, Carnegie Mellon University, School of Computer Science, Institute for Software Research International, Technical Report CMU-ISRI-04-109.