



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers Collection

2014-04-15

Evolutionary Agents and Dynamic Analysis of Accident Law

Seagren, Chad W.

<http://hdl.handle.net/10945/48653>

Downloaded from NPS Archive: Calhoun



Calhoun is a project of the Dudley Knox Library at NPS, furthering the precepts and goals of open government and government transparency. All information contained herein has been approved for release by the NPS Public Affairs Officer.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

Evolutionary Agents and Dynamic Analysis of Accident Law

Chad W. Seagren
Assistant Professor
Operations Research Dept
Naval Postgraduate School
1411 Cunningham Road
Monterey, CA 93943
cwseagre@nps.edu

15 April 2014

Abstract:

Keywords: law & economics, accident law, agent-based modeling, evolutionary economics

We employ an agent-based computational model to gain insights regarding relationships in accident law that are often overlooked in neoclassical theorizing. We consider the situation faced by the inhabitants of Eggtopia, a fictitious society whose members search for and collect precious eggs. In the course of their productive efforts they occasionally experience destructive accidents with other individuals. The inhabitants face trade-offs in that engaging in behavior that increases their productivity may also increase the possibility of an accident. We demonstrate how the agent-based model developed in this paper can be used to answer the questions commonly addressed in the mainstream literature, while also opening up new avenues of investigation. While certain outcomes are consistent with neoclassical theory, we find that agents may elect to be careful even when the neoclassical theory predicts otherwise, and particular negligence rules tend to differ significantly in their ability to rid society of negligent behavior.

1 Introduction

The purpose of this paper is to demonstrate the efficacy with which evolutionary economics can contribute to economic analysis of the law. We employ an agent-based computational model with which we examine the effects of different tort regimes under a variety of circumstances. We find that macro-level stability does not necessarily imply micro-level equilibrium, and we find that agents occasionally elect to engage in negligent behavior when neoclassical theory predicts otherwise.

We present the situation faced by the inhabitants of Eggtopia, a fictitious society whose members search for and collect precious eggs. In the course of their productive efforts they occasionally experience destructive accidents with other individuals. The inhabitants face trade-offs in that engaging in behavior that increases their productivity may also increase the possibility of an accident. We construct an agent-based model of Eggtopian society in REPAST¹, a Java-based simulation package, and employ it as a laboratory with which to evaluate different legal liability rules as well as other questions. The model enables us to examine how boundedly-rational and heterogeneous agents interact with each other outside of an equilibrium framework.

Analysis within the neoclassical framework typically revolves around the economic efficiency of various rules given certain extensions and complications. In the models most neoclassical economists employ to analyze accident law, rational agents tend not elect to engage in negligent behavior as the strategy to employ due care typically dominates all others. Yet, in reality, legal findings of negligence occur rather frequently. To explain the puzzle, economists typically extend the models to include unintentional lapses in the exercise of care; uncertainty on behalf of agents regarding the legal level of due care; and improperly set legal standards

¹ For more on REPAST, see <http://repast.sourceforge.net/>.

(Shavell, 1987). As an alternative, we argue examining the law through the lens of evolutionary economics provides insight into this puzzle, and illuminates other important characteristics inaccessible to the mainstream neoclassical theorizing.

We place this project within two threads of the literature. The first is Rizzo's (1980) effort to outline some of the weaknesses of applying static general equilibrium concepts to analyze a dynamic system that is constantly in a state of flux. From such a perspective, the present project is an attempt to imbue the otherwise static neoclassical model with heterogeneous agents and heterogeneous time, in order to allow for out of equilibrium behavior.

We also draw from Smith's (2003) notions of ecological rationality. We explore how the institution of tort law provides otherwise boundedly-rational individuals with the appropriate information and incentives to ultimately reduce the social costs of accidents. Our approach enables us to examine the extent to which individuals rely upon the institution for guidance as opposed to their own rational faculties to pursue goals that are privately and socially beneficial.

1.1 Static and Dynamic Analyses of Accident Law in the Literature

Law and economics scholars in the neoclassical tradition, such as Diamond (1974), or Landes and Posner (1981), typically develop an explicit mathematical equation that describes the social welfare function for actors related to a particular type of accident (see Posner (2007), Shavell (1987), and Cooter and Ulen (2007) for textbook treatments). Landes and Posner consider a representative injurer and a representative victim to simplify the analysis and delineate the roles of the disputants. Diamond models a single individual's rational choice in response to $n-1$ other agents, all of whom employ the same strategies in equilibrium.

These approaches are consistent with the standard mode of analysis for neoclassical economists that upon selecting a phenomenon for investigation one assumes a system inhabited by rational utility maximizers who have stable preferences and some form of rational expectations regarding their world. Agents are then imagined to relate in a linear manner, which enables the researcher to choose one or perhaps several representative agents that select strategies to maximize their utilities in response to other agents maximizing their own. Agent behavior in equilibrium is deduced from the basis of these relationships. These assumptions enable the theorist to obtain an explicit closed form mathematical expression of the system which can then be analyzed to determine global or local optima subject to the theorized constraints.

In this project, we relax the rationality assumption in order to explore the role that the institution of tort law provides individuals with the information and incentives to ultimately conduct themselves appropriately. In keeping with Axtell's (2007) notion of the neoclassical sweet spot, relaxing one assumption such as agent rationality necessarily requires one to jettison other assumptions such as agent homogeneity, equilibrium, and non-interaction. While neoclassical economics does not have much to say about disequilibrium behavior, the agent-based model we employ in the remainder of the article will enable us to maintain tractability of the problem.

Relaxing the neoclassical sweet spot in this application serves two purposes. The first is that in reality, agents are not intimately familiar with the cost functions of their behaviors, let alone the respective functions for all other individuals engaged in similar activity, yet the neoclassical model requires an individual to take these factors into consideration and determine a response that is jointly optimal. In fact, individuals must learn these relationships through their

experience. Such experience may involve trial and error, observation, imitation, and any number of other modes of learning that humans employ to make sense of their world. The present effort is an attempt to examine the role of tort law in guiding individuals in the process of discovery. In other words, we know actual people are not hyper-rational utility maximizers and according to Smith (2003), some of the rationality we see is actually because of institutions that have emerged to provide individuals guidance in making correct decisions. This approach seeks to examine the extent to which the institution of tort law provides individuals guidance in these situations.

The other purpose is to gain a proper understanding of the extent to which the neoclassical model is an adequate description of reality. As George Box famously quipped, “all models are wrong, but some are useful.” The analytical power of the neoclassical paradigm is one reason why it is the mainstream model for scientific thought among economists. It has the strengths of generality, syntactic clarity, and tractability. However, these qualities come at a price of relying on unrealistic assumptions. The use of an agent-based model is an attempt to move towards the more realistic end of the spectrum.

While it is true that the artificial societies depicted in agent-based models lack the complexity and richness of human society, and the agents that populate these virtual worlds lack the intelligence of human beings, it is also true that the agents’ relative ability to act within their society might be comparable (Lavoie, 1994, pg. 554). Virtual agents certainly are not as creative or innovative as the individuals they mean to portray, but relative to their world, they could be considered creative as they are capable of learning from experience and adopting courses of actions as a result of trial and error (Lavoie, 1994, pg 554). Thus, artificial agents are subject to

the criticism that they are not and perhaps never will be capable of achieving the intelligence and creativity of human beings, but they are capable of innovation relative to the worlds they inhabit.

A search of the literature to find an agent-based simulation model applied to this aspect of economic analysis of tort law has proven unsuccessful. Diianni (2007) applies such techniques to model disputants and the evolution of precedent. See also Yee (2005), and Picker (1997; 2002) for treatments on similar topics. If simulation as a tool is used at all, such as in Katz (1987), Hylton (2002), Parisi (2002) and Hylton and Miceli (2005), it typically employed as a computational technique to numerically examine models that otherwise defy a tractable analytical approach, rather than as the focus of effort. This project is an attempt to remedy this gap in the literature as it pertains to tort law.

1.2 The Fictional Society of Eggtopia

Imagine a society known as Eggtopia. The inhabitants are human beings, and are just like any other human beings in their ability to use their senses to collect information about their environment, as well as their ability to take action on the basis of that information in conjunction with their subjective valuation of the relative benefits of means and ends. In addition, these individuals possess the same physical attributes as any other human being such as visual acuity, strength, ability to move, etc. In many respects, one can consider the mental construct of Eggtopian society as indistinguishable from nearly any other society in the Western world. A significant source of income in the Eggtopian economy is based on the sale of eggs and egg products. These eggs are scarce, extremely fragile, and incredibly valuable, thus a large proportion of Eggtopians are employed in their collection. The eggs are found in the vast Egg

Fields, where they lie just below the surface, relatively easy to spot by the trained eye and easily extracted. Individuals search the fields with varying levels of intensity and fill their baskets with the eggs they find. Upon filling their basket to capacity, they return home as expeditiously as possible to store them for later use.

Unfortunately, for all their skill and talent the process of searching for eggs, extracting them from the ground, and storing them securely is fraught with danger. In their haste to collect as many eggs as possible, accidents between Eggtopians are relatively commonplace. When two individuals collide, the force of their collision completely destroys the eggs the victim was carrying. Eggtopians face a challenge similar to that of many members of nearly any society. That is, to engage in productive activity also brings with it distinct risk of accidental damage to self or property. Changing behavior along margins that improve productivity, such as speed of travel, may also increase the possibility an accident occurs.

When an accident occurs, the parties to the accident decide how to proceed pursuant to the relevant accident law and the facts of the case at hand. The effect of accident law as an institution is to provide guidance to individuals regarding “appropriate” behavior as it pertains to accidents and duties owed to others. Participants in market exchanges have the luxury of performing economic calculation to determine profit and loss as a means of assisting *ex ante* in decisions regarding considered courses of action and *ex post* in evaluating the success of those actions. However, price alone may not be a reliable guide for the decision to engage in potentially dangerous activity, due to the fact damages that occur as a result of accidents are not necessarily implied in the market price. A typical analysis of such a decision parallels that of Becker’s description of criminal behavior and punishment (1976). It is reasonable to believe that

individuals subjectively weigh the payoff of risky behavior against the concomitant damage discounted by their estimate of the probability that the destructive event occurs.

The fictional Eggtopian society is rather stylized concerning the circumstances surrounding the egg collection and production, but is still recognizable enough to reality to glean insights common to individuals' behaviors regarding accidents. The general keys to analysis of accidents in the imaginary world of the model are no different than analysis of accidents in the real world. In reality, accidents typically occur while both parties are engaged in otherwise productive behavior. In the model, agents are constantly in the process of collecting and storing valuable eggs. In the course of evaluating their behavior, individuals may take measures that reduce the likelihood of an accident occurring, however, these activities may simultaneously inhibit productive behavior, i.e. as when the driver of a delivery truck opts to maintain a relatively slower speed, it may reduce the probability of an accident, but it also increases the time it takes to transport the goods she is hauling. There is a clear parallel between this line of thinking one's decision to engage in almost any productive, yet risky, endeavor.

The intent of this mental construct is to serve as the target of an economic analysis of accident law. Since Eggtopia is a figment of the author's imagination (and now the reader's), no empirical data, case law, or historical record exists that describe in detail the activities of its inhabitants. The agent-based computational model provides the mechanism with which we generate data to test the effectiveness of various analytical approaches. Part of the reason for the relative dearth of empirical law and economics studies is the difficulty of obtaining data conducive to analysis and testing. In this case, agent-based modeling helps to overcome this challenge.

2 The Model: An Artificial Implementation of Eggtopia

The agent-based model of Eggtopia is implemented in REPAST, a simulation toolkit that uses the Java programming language. The environment is comprised of a two-dimensional torus grid that is populated with heterogeneous agents and eggs. The agents possess various attributes (instance variables), to include egg carrying capacity, visual range, and speed of movement. An agent's vision is the radius of the circle (centered on the agent), in grid cells within which the agent is capable of viewing and locating an egg. An agent's capacity is the maximum number of eggs an agent may carry before he must return home and unload his collection, and an agent's speed is the maximum number of grid cells an agent moves in a single timestep. Agents take their maximum step size in most cases.

Agents are always engaged in one of four modes of activity. An agent in *search* mode actively explores the environment in search of eggs. While searching, the agent randomly selects a heading within ninety degrees of his current heading and moves in that direction. The number of cells the agent moves is equal to his speed. Upon locating an egg, the agent enters *collect* mode and selects the most direct route towards the target egg. When the agent reaches the cell containing the egg, it picks it up and adds it to its collection. The agent returns to *search* mode if there remains excess capacity in his basket. However, if the number of eggs in its collection meets his capacity, the agent enters *return* mode in which case the agent takes the most direct route to his home cell and drops off his collection of eggs. The eggs are then deposited into a virtual savings account, rendering them invulnerable to damage.

Regardless of the agent's mode, during the *Pre-Step* stage the agent selects a candidate location on the grid to occupy in that timestep. If another agent lies anywhere on the path to the

target cell, an accident occurs between the two agents. The agent attempting to enter the occupied cell is declared the injurer and the eggs the victim was carrying are destroyed. If more than one agent lies on the path of the agent in question, only the closest agent is selected for involvement in an accident. A tort rule is immediately employed to adjudicate the disputes that arise surrounding the aftermath of these accidents. Figure 1 is a graphical depiction of the simulation schedule of the model.

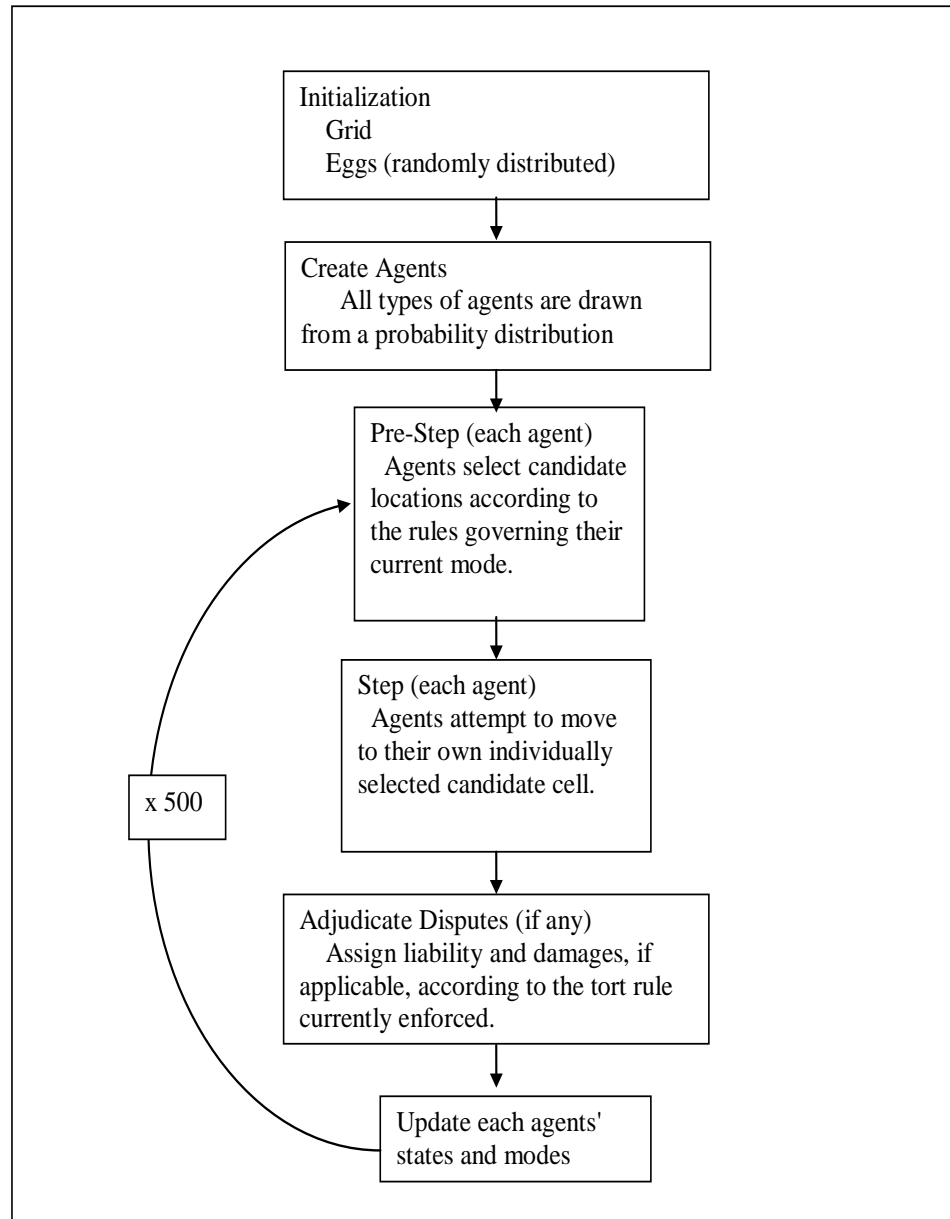


Figure 1 Graphical Depiction of Simulation Schedule

The process of collecting eggs and bringing them home is a productive endeavor and the efficiency with which agents collect the eggs distributed in the environment is a function of the individual agents' attributes of vision, capacity, and speed. All else equal, agents with greater vision and speed will tend to collect eggs more quickly than their slower counterparts. An agent's

higher capacity allows it to spend more time actively searching for eggs, since agents with lower capacities spend relatively more time returning home delivering their eggs. Like wealth, the number of accidents and the magnitude of accident losses are also functions of agent attributes. An agent with greater speed travels longer paths with each step which increases the probability of interacting with other agents, while agents with a higher capacity are more likely to suffer high losses when involved in an accident.

It is important to note that there are no exogenous “cost” or fixed “profit” functions related to agent behavior. The consequences to any agent’s behavior are always endogenously determined through their interactions with other agents and the environment. Agents employ a “strategy,” which consists of a choice of vision, capacity and speed. One can easily imagine that a solitary agent’s optimal strategy would be the maximum allowable levels of each parameter. However, such a strategy would likely result in too many destructive accidents if other agents were present. In a social setting with other agents included in the model, causing accidents, suffering from accidents, and ultimately competing with other agents in pursuit of collecting eggs, it is clear that the success an agent enjoys is highly dependent upon the agent’s relationship to (most) other agents in the model. An individual could employ the same strategy and experience highly variable results due entirely to chance, or due to the strategic behavior of other agents.

The agents that populate the model single-mindedly pursue the goal of collecting eggs. They possess no explicit choice algorithm nor do they form expectations regarding the future. One might say that if they possess a utility function at all, it is a lexicographic preference for eggs to the exclusion of all other goods, i.e. leisure, safety, etc. Modeling such zero intelligence

agents has precedence in the computational and behavioral finance literature. In an influential article, Gode and Sunder (1993) utilize what they term “zero-intelligence traders” to examine the institutional effects of particular auction rules. See Duffy (2006) for a comprehensive survey and assessment of this literature. Ultimately, we employ zero-intelligence agents to highlight the institution’s role in guiding behavior, as opposed to relying on notions of rationality or a particular level of intelligence.

While Gode and Sunder’s agents select their bids randomly, all agents in the present model must pursue egg collection. However, they purposively select their strategy on the basis of their local information. While such an algorithm is perhaps an unambitious description of human action, it emphasizes the institution’s role in guiding behavior, as opposed to relying on notions of rationality or a particular level of intelligence.

2.1 Accidents and Liability Regimes

The three classes of liability regimes we consider are no liability, strict liability, and simple negligence. Under a regime of no liability, the parties to the accident simply have no recourse to adjudicate their dispute and the losses fall where they may. A strict liability rule requires the injurer to fully compensate the victim of the accident for all of his lost eggs, even if the injurer takes on a negative balance of eggs. Finally, the negligence rule requires the injurer to compensate the victim if the injurer was not exercising due care, and relieves the injurer from liability for losses if he was exercising due care. Numerous definitions of due care are examined in the subsequent analysis to include various “reasonable” levels of speed, vision, and egg capacity.

In reality, if the (perceived) cost of preventive actions is lower than the expected cost of the accident, a judge may prescribe them in the course of determining a proper definition of "due care." Likewise, various behaviors may contribute to the frequency and severity of accidents in different ways. Due care may be defined along numerous margins, some of which may not effectively reduce the possibility of an accident. In the model, the speed an agent travels is a significant factor in determining an agent's productivity. Under certain circumstances, higher speed also contributes to higher accident rates. Thus, the speed at which an agent travels is a potentially effective margin on which to analyze the intricacies of due care. Consider also, that while the attributes that an agent possesses are obvious and available to the researcher, such attributes as visual acuity (attentiveness, intelligence, reflexes) are not necessarily as accessible to the court charged with adjudicating the dispute.

2.2 Evolutionary Agents

Let the term, strategy, refer to a particular combination of *vision*, *capacity*, and *speed* levels. Parameter values may range from 1 to 33. The parameters *vision* and *speed* are measured in cells, while *capacity* is measured in eggs. Thus, to say that an agent's *speed* = 9 is to say that such an agent takes "steps" 9 cells in length each timestep. Agents take their maximum step size in all cases except when it is necessary to take a smaller step to pick up an egg or to occupy their home cell.

The task placed before each agent is daunting. Individuals must select the strategy that, given all other agents strategies and behavior, will improve or at least maintain their current level of income over the course of a generation. However, they have exactly 33^3 or 35,937 strategies

from which to choose. Their choice is based on the perceived effectiveness of their current strategy and their most recent reasonably successful strategy. In contrast to the elegant continuous, twice differentiable field upon which neoclassical agents are assumed to maximize their utility, these agents are confronted with a more realistic and complicated combinatoric problem.

Potts (2000) provides an analysis of the shortcomings inherent in basing a model of individual decision-making on a mathematical field and illustrates how such a model ultimately assumes an impossible level of knowledge concerning the state of the world on the part of the individual. He argues that the economic space upon which individuals operate is better perceived as a less than fully connected set of elements. The individual's task is to explore this space by experimenting with various combinations of elements in order to discover the means which best serve her ends. This combinatoric problem goes beyond the largely artificial necessity of only choosing integer values for strategy parameters, rather it implies massive amounts of uncertainty in the agent's choice due to a paucity of information concerning the relative values of various strategic choices. Figure 2 is a graph theoretic representation of the strategy component for any given agent. The connected elements represent the strategy the agent currently employs. The elements outlined with dotted lines are outside of the agent's present strategy neighborhood, and are thus inaccessible at the present generation.




Figure 2 Agent Strategy Choice as Combinatoric Problem

An agent may only select a candidate strategy from its neighborhood, that is, those levels that are within two elements from the agent's incumbent strategy. Those elements that are omitted from the figure or otherwise outlined with dotted lines reside outside the neighborhood

and are only reachable after multiple iterations of the evolutionary algorithm. Because the agent is faced with such ignorance and uncertainty, this framework should be a good test to see how well the tort law as an institution and the rules therein help to teach agents the appropriate behavior.

2.3 The Satisficing Algorithm

Agents are allowed to change their strategies upon completion of each generation, which lasts 500 timesteps. The evolutionary algorithm the agents employ contains both satisficing and hill-climbing aspects. Suppose an agent employs a particular strategy. If her wealth at the end of the current generation exceeds that gained from the previous generation, meaning that the current (incumbent) strategy has succeeded in bettering her condition, it remains her incumbent strategy and she employs it in the next generation. This is the satisficing characteristic, in that the agent is satisfied with a relatively well performing strategy and does not seek to “fix what is not broken.” See Brenner (2006) for a description of satisficing strategies, and see Nelson, Winter, & Schuette (1976) and Nelson & Winter (1982) for well-known and successful implementations in terms profit seeking firms. Figure 3 is a flow chart for the algorithm.

If the agent’s current wealth is less than that observed in her previous generation, that is, if the incumbent strategy does not succeed in improving her condition, she selects a candidate strategy to employ in the next generation. Candidate strategies are selected with equal probability from the set of neighbors of the current strategy. In the current implementation, neighbors are all strategies whose parameters are no more than +/- 2 levels from the current one. This means that most strategies have 5^3 , or 125, neighbors including itself.

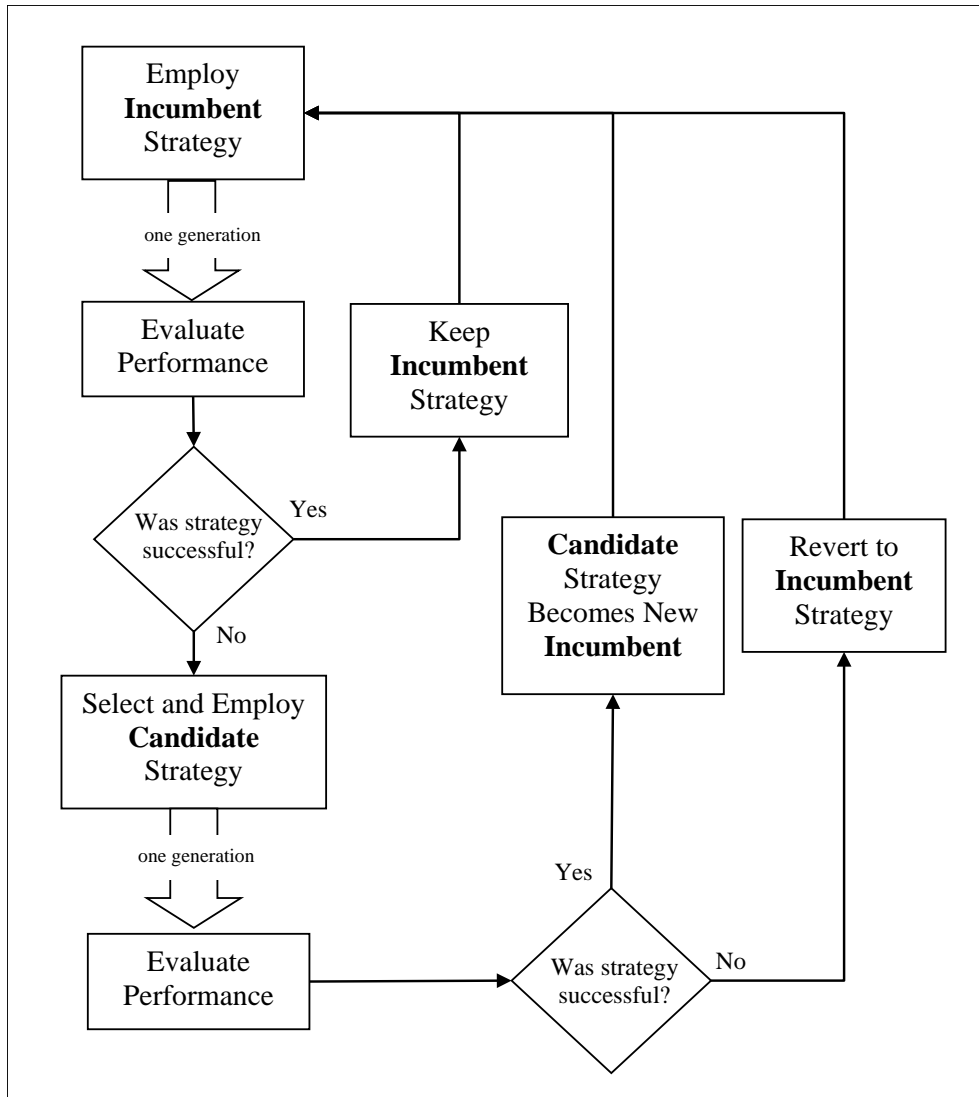


Figure 3 Flow Chart Depicting Agent Level Evolutionary Algorithm

At the end of the next generation, the candidate strategy's success is measured against the level of wealth the incumbent strategy garnered the last time it was employed. If the candidate strategy yields greater wealth, it is deemed the new incumbent strategy and employed in the next generation. If the candidate strategy's performance is less than that of the incumbent, then the incumbent remains as such and is employed in the next generation. Thus, the upper level of the

model is a classic evolutionary process in which the agents adapt their strategies or attributes according to the strategies and attributes that have exhibited success in that particular generation, or stage, of the model.

From a practicality perspective, the complexity of the model precludes a straightforward hill-climbing or similar optimization algorithm, due to the fact that it is impossible to “test” the expected success of the candidate strategy prior to employing it. For example, in the classic hill-climbing algorithm, the agent selects a candidate strategy in the neighborhood of his current strategy. Given the current strategies of the other agents, and a *known* function that relates individual agent strategies to the response variable of interest, the agent tests the candidate strategy to determine an expected value of the strategy. The agent then employs the strategy with the greatest expected value among a short list of candidates.

There are two aspects of the current model that render such a straight-forward choice algorithm impossible. The first is that since an agent’s performance is dependent upon the strategies employed by the other agents, it is impossible for an agent to practically and effectively predict the success of a particular candidate strategy, because a nontrivial number of agents are likewise changing their strategies in preparation for the next generation. The second is that the strategy space itself is so broad and sparsely populated, it is impractical to even learn about successful strategies from other agents. The domain is so large, that many strategies are never even employed by any agent through the course of entire replications.

3 Analysis Scheme

In this section, we outline the experimental design and the response variables we employ in the following analysis.

3.1 Experimental Design

We employ a full-factorial experimental design with three factors. Systematically varying the most important factors in the model enables us to gain a better understanding of the relationships between these factors and the various response variables. Table 1 summarizes these experimental factors.

Table 1 Full-Factorial Experimental Design

Factor	Description	Levels
numAgents	number of agents	[100, 500]
numEggs	number of eggs	[100, 1000]
liabilityRule	liability rule in effect	20 levels*

We choose the levels of *numAgents* and *numEggs* in order to achieve a range of agent-to-egg ratios. Intuitively, the agent-to-egg ratio is a measure of the scarcity of eggs relative to agents. These combinations enable us to examine ratios that range from 0.2 (100 eggs:500 agents) to 10 (1000 eggs:100 agents).

We consider three classes of liability rules: no liability, strict liability, and negligence. For the negligence rules, we examine rules in which due care is defined in terms of a particular level of vision, capacity, or speed. Table 2 summarizes these rules.

Table 2 Summary of Liability Rules

liabilityRule	Description	Due care levels
no_liable	No liability regime	n/a
strict_liable	Strict liability regime	n/a
negligence_vis_nmt_c	Due care defined as <i>vision</i> no more than c	9, 17, 25
negligence_vis_nlt_c	Due care defined as <i>vision</i> no less than c	9, 17, 25
negligence_cap_nmt_c	Due care defined as <i>capacity</i> no more than c	9, 17, 25
negligence_cap_nlt_c	Due care defined as <i>capacity</i> no less than c	9, 17, 25
negligence_spd_nmt_c	Due care defined as <i>speed</i> no more than c	9, 17, 25
negligence_spd_nlt_c	Due care defined as <i>speed</i> no less than c	9, 17, 25

For example, under the rule *negligence_cap_nlt_17*, agents with *capacity* < 17 are held liable for the accidents they cause.

Thus, our experimental design consists of 80 design points, with ten replications of each. In an effort to reduce variance between the outputs at each design point, we employ common random numbers to the greatest extent possible. Every i th replication of each design point is initialized with the same seeds in the random number generators. For more on Common Random Numbers, see Law and Kelton (2000, pp. 582-598).

Finally, table 3 outlines those model parameters that do not vary during the experiment.

Table 3 Static Parameters

Parameter	Description	Level
generation_duration	Timesteps per generation	500
numGenerations	Number of generations per run	800
worldSize	surface area of torus	100,000
agentPicker	Distribution from which initial agent strategy is chosen	Uniform
agentDirector	Distribution from which agent search direction is chosen	Uniform

3.2 Response Variables

One of the most important contributions of economic analysis of accident law is the notion that the common law of torts evolved so as to maximize efficiency (Landes and Posner, 1981). Thus, the primary measure of effectiveness (MOE) we employ in the subsequent analysis is per-capita wealth. We define per-capita wealth as the sum of all eggs in agents' possession (at the end of a particular generation) divided by the number of agents. Likewise, per-capita accidents and losses are the sum of all accidents (or losses) that agents experience during a generation, divided by the number agents.

The number of times each agent changes strategies is tallied and used to estimate the probability an agent changes its strategy at the end of each generation. To estimate this probability we sum the number of agents who change their strategy at the end a particular generation and divide by the total number of agents. Finally, agent strategies (*vision*, *capacity*, and *speed*) are also examined over time. One such aspect of agent strategies we examine is the range of each strategy element an agent explores over time. For a given duration of time, the range of a strategy element is given by the maximum level of that element an agent chooses during that time minus the minimum level the agent chooses during that time. So, an agent whose *capacity* exhibits a maximum of 23 and a minimum of 10 during a particular duration has a capacity range of 13.

Populating the model with heterogeneous agents enables us to examine how certain MOEs are distributed across the population. Towards this end, we also calculate (in post-processing) an agent's wealth quantile. For each run and time combination, we sort the agents by descending order of wealth and assign each agent into corresponding quantiles according to their

rank in the wealth distribution, with the first quantile as the wealthiest. Using quantiles enables us to analyze the characteristics of the most and least successful agents during each generation.

The agents in the model form a complex adaptive system with many characteristics that maintain a state of flux. The level of aggregate wealth, that is, the sum of the wealth of all agents, exhibits transient behavior before settling into a steady-state after approximately 700 generations. Each replication consists of 800 generations, or 400,000 timesteps, but only data after 350,000 timesteps is considered for analysis of steady-state variables². The entire run is considered for the analysis of transient variables, such as the trajectory of agent's strategy choices through time. In order to economize on time and computer space, only the results of every 10th generation are output and thus available for analysis. Unless indicated otherwise, we average the response variables for each design point over the ten replications. Note also that 40 simulation runs, (i.e. four design points with ten replications each) requires approximately 24 hours of computer time on an Intel Core i5-3470 3.20 GHz processor. Thus, to run the entire experiment takes approximately three weeks on a single desktop computer.

4 Results

Consistent with typical neoclassical analysis, this more dynamic evolutionary approach is capable of differentiating among any of several liability rules and identifying the superior performing rule. The framework allows for the researcher to actually employ the various liability rules in artificial societies to examine the rules' effects directly, as opposed to only analyzing the effects of a rule of no-liability and deducing the effects of other rules. A number of other

² We employ Welch's method to confirm that the model achieves steady state after 700 generations (Law and Kelton, 2000, pp. 520-525)

questions concerning the distribution of wealth, the distribution of strategies, individuals' steady state behavior, and transient behavior are examined as well.

4.1 Identify the Wealth-Maximizing Rule

Table 4 provides an overall comparison of the best performing tort rules for each of the four combinations of agents and eggs. The values shown are the per-capita averages of the generations that occur during steady-state (after $t = 350,000$). For each type of negligence rule, one that defines due care in terms of *vision*, *capacity*, or *speed*, we only include the one that obtains the highest per-capita wealth relative to the other rules of that type.

Table 4 Assessment of Liability Rule Performance

Eggs	Agents	liabilityRule	Net Wealth Per-Capita**		Accidents Per-Capita	Losses Per-Capita	Losses / Accident
100	500	negligence_cap_nmt_9	A	11.0	106.2	53.2	0.50
		no_liable	A B	6.9	120.0	61.5	0.51
		negligence_spd_nmt_25	B	6.7	104.6	57.8	0.55
		negligence_vis_nmt_25	B	6.6	99.4	57.4	0.58
		strict_liable	C	-312.8	45.9	346.7	7.55
	100	negligence_cap_nmt_17	A	66.2	66.5	86.0	1.29
		negligence_vis_nmt_9	B	57.6	75.1	106.0	1.41
		no_liable	B	56.9	68.0	99.7	1.47
		negligence_spd_nmt_25	C	46.8	56.1	94.6	1.69
		strict_liable	D	-78.2	17.4	154.0	8.86
1000	500	negligence_cap_nmt_25	A	121.0	121.7	149.2	1.23
		no_liable	A B	113.9	117.8	161.0	1.37
		negligence_vis_nmt_25	B	110.8	114.0	164.6	1.44
		negligence_spd_nmt_25	B	108.4	114.5	160.9	1.41
		strict_liable	C	-785.8	102.3	951.3	9.30
	100	negligence_cap_nmt_25	A	224.4	42.7	121.9	2.85
		no_liable	A	223.9	40.2	128.9	3.20
		negligence_vis_nmt_25	A	222.6	36.8	129.6	3.52
		negligence_spd_nmt_25	B	211.7	38.7	133.5	3.45
		strict_liable	C	58.7	25.0	234.8	9.41

**Levels not connected by the same letter are significantly different ($\alpha < 0.1$).

In each case, the negligence rule where due care is defined in terms of *capacity* is optimal, though *no_liable* is not statistically different than *negligence_cap* for three of the four cases. Perhaps surprisingly, *strict_liable* performs profoundly worse in terms of wealth relative to any of the other rules under consideration.

A relationship exists between losses and accidents, as well. While *strict_liability* achieves the lowest levels of wealth, it also achieves the fewest number of accidents. But the accidents that it does allow are substantially more destructive according to the high loss per accident ratios. Alternatively, the *negligence_cap* rules, which achieve the highest levels of wealth for each design point also obtain the highest accident rates. But most importantly, the negligence rule that maximizes wealth also seems to minimize losses per accident. So, while there may be more accidents under the *negligence_cap* rule, they tend to be less damaging than under the other tort regimes. Thus, the more successful tort rules appear to be those that provide agents with incentives that not only reduce the total damage from all accidents, but also to reduce the destructiveness of the accidents they do experience. Models with endogenous damage levels are difficult to construct in the neoclassical framework, but the evolutionary framework renders such aspects tractable.

The evolutionary perspective provides a rich description of the complexities of Eggtopian society. We show that unraveling the sweet spot to enable heterogeneous agents to pursue their own interests on the basis of their local knowledge is sufficient to achieve a steady state condition for each of the liability rules under investigation. Agent interaction tends to drive social wealth asymptotically to the vicinity of the maximum achievable for a given liability rule and is an unintended, though seemingly beneficial, consequence of agent behavior.

These results, namely that strict liability performs poorly relative to other negligence rules and even no liability rule would seem to militate against Rizzo's argument that in a world of flux, the rule of strict liability is superior. However, Rizzo's thesis is that given an uncertain world of flux, the rule of strict liability provides an institutionally efficient rule by reducing uncertainty as to how the courts will handle disputes. Given the relative *ex ante* certainty of the rule of strict liability, individuals are better able to plan their activities and assess the consequences of risky behavior. This effect is magnified in a world of technological change that affects relationships in unimaginable ways and where people make subjective predictions on the basis of local knowledge. Since the present version of the model includes neither agent expectations nor a mechanism for technological change, these results should serve to simply inform this debate rather than provide weight to either side. In this stage of its development, the model may be considered a useful foil. Simply adding time and process to the analysis is not sufficient to conclude that strict liability is superior to negligence rule in a dynamic society.

4.2 Macro-Steady State, Micro-Turbulence³

The previous section outlined the relative performance of various liability regimes in terms of per-capita wealth achieved during steady state as a measure of effectiveness. Steady-state is determined to have arrived when aggregate wealth ceases to vary significantly with time. One of the benefits of the agent-based modeling approach to the analysis of the current problem is the ability to examine the behavior of the entire distribution of agents. While the macro-level outcome of time-invariant aggregate wealth is the result of the interaction of heterogeneous

³ The graphs presented in this section are from the design point for which there are 100 agents and 100 eggs. We find this qualitative behavior is robust over all agent-egg combinations considered.

agents and their environment, it is not obvious what individual behavior is necessary to achieve it. Agent based modeling enables us to answer this question.

If it were the case that individual equilibrium is a necessary condition for system level steady-state, we would expect that agents would decrease the frequency with which they change their strategies as the system achieves equilibrium, at or around 350,000 timesteps. Figure 4 is a graph of how the probability an agent changes its strategy varies over time, with separate lines depicting behavior under different liability regimes. As the graph indicates, there is no perceptible reduction in the frequency with which agents change their strategies upon entering steady state, towards the right-hand side of the chart. That is, agents do not seem to settle on a particular strategy that is a robust response to the strategies employed by all of the other agents. Agents appear to continue to grope for strategies as a means to improve upon the outcomes they currently experience.




Figure 4 Probability Agent Changes Strategy Through Time

As Figure 4 illustrates, the rate at which agents alter their strategies does not appear to change as the system enters steady-state. While agents may continue to furiously change their strategies, perhaps they are merely jumping back and forth within a narrow band of the “optimal” area of the strategy domain.

Figure 5 shows the distributions of the ranges of agents’ strategy elements during the last 40 generations in steady-state. The box-plots depict the 25th, 50th, and 75th percentiles, while the whiskers depict the 10th and 90th percentiles of each distribution. For example, the left-most box-plot depicts the distribution of agent ranges for *vision* when agents are under the *negligence_cap* rule. The box plot shows that half of all agents (those included in the inter-quartile range) experienced *vision* ranges between 19 and 10 during this time-period and that 75% of all agents experienced a *vision* range of 10 or more. If agents were merely toggling between relatively close locations of the strategy space, we would expect the ranges agents experience to be much smaller (i.e. 3-5). The figure belies the fact that agents continue to search wide swaths of the strategy domain even while the system is in steady-state.

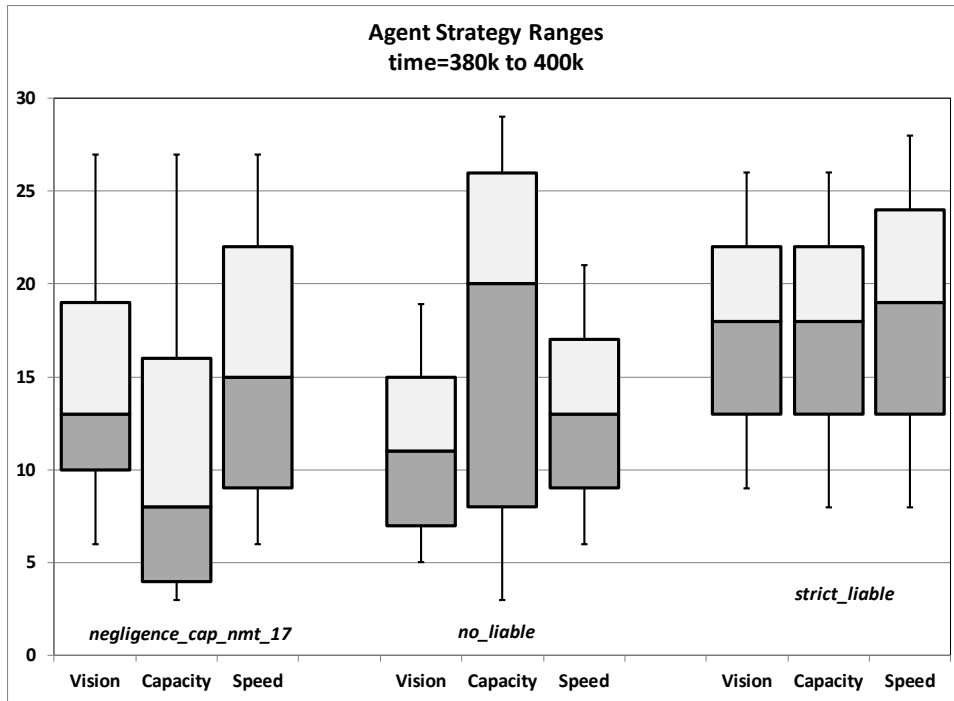


Figure 5 Agent Strategy Range Distributions in Steady-State

While we establish that agents continue to change their strategies in light of their attempt to respond to an ever changing world, it is not clear whether this flux also affects the agents' outcomes that ultimately obtain. Figure 6 outlines the how widely agents' outcomes in terms of relative wealth vary during steady state.




Figure 6 Comparison of Average Agent Mobility Among Wealth Quantiles in Steady State

If agents' outcomes achieved in steady-state were commensurate with that of the macro-level steady state, we would expect that agents that appear in a given quantile at the onset of steady state would remain there throughout. Indeed, a conservative expectation would be for most agents to appear in no more than two different quantiles. However, as the figure indicates, between approximately 74% and 87% of agents appear in three or more quantiles throughout steady state, depending upon the liability rule in effect. The most variable outcomes result under the rule of strict liability, where approximately 87% of agents experience significant fluctuations in their success relative to other agents even while the system is in steady state.

It is fortunate that the model achieves a steady state in terms of an important aggregate variable such as total wealth and enables a straight-forward method to compare different liability regimes. Whether this system we have created would behave in this manner was a question that

could only be answered empirically. Furthermore, while the aggregate system may achieve a pattern of steady state behavior, this section demonstrates that the agents do not exhibit behavior recognizable as individual equilibrium. This result is counterintuitive in that one would expect that individual stability or “equilibrium” would be a necessary condition for system level stability.

4.3 Out of Equilibrium Dynamics⁴

Another feature of the evolutionary approach is that it provides the researcher with the ability to analyze out of equilibrium behavior. Rather than assert the existence of equilibrium and then deduce the conditions that must be present in order to sustain it, this approach begins from out of equilibrium conditions and enables analysis of the process through which agents achieve equilibrium – or a steady-state, if at all.

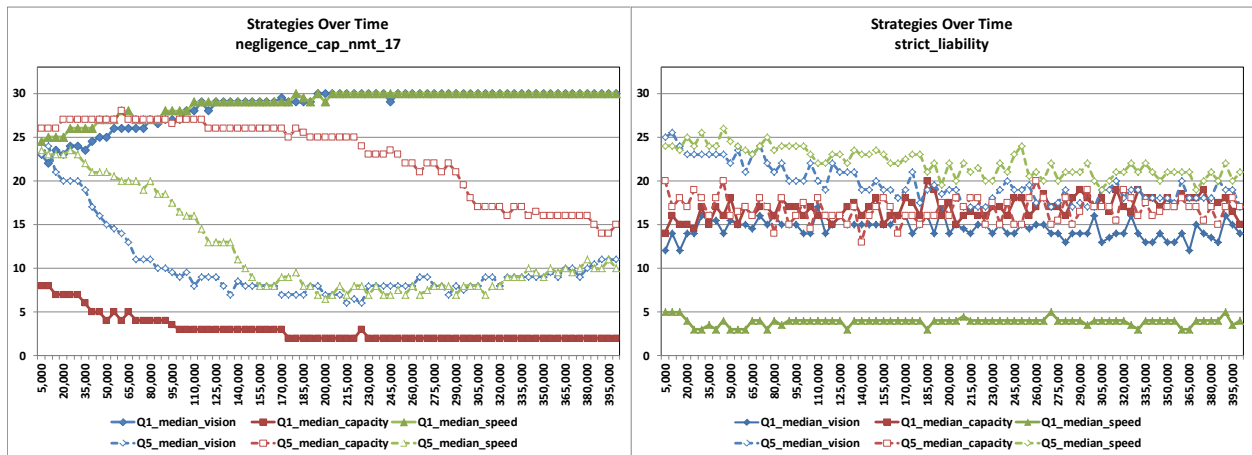


Figure 7 Transient Strategy Behavior for Negligence and Strict Liability Rules

⁴ As in the preceding section, the graphs presented in this section are from the design point for which there are 100 agents and 100 eggs. We find this qualitative behavior is robust over all agent-egg combinations considered.

The left-hand panel of Figure 7 is a depiction of how the strategies of agents under the negligence rule (due care defined as *capacity* no more than 17) change over time. The median *vision*, *capacity*, and *speed*, of those agents in the first and fifth wealth quantiles are shown. The behavior of those agents who fall into the top wealth quantiles achieve consistency at or around $t = 200,000$. The median strategy of those agents in the bottom wealth quantile seems to respond more slowly. In fact, the median agent in the fifth quantile is negligent (*capacity* > 17) until $t = 350,000$.

In contrast, the strategies for agents under the strict liability rule are shown in the right-hand panel of Figure 7. The trajectories of the strategies appear substantially more volatile, than those under the negligence rule. In fact, the strategies of both the first and fifth quantiles seem centered around the middle of the strategy-space, which is consistent with the notion that the agents are unable to reliably ascertain the correct strategy selection from their experiences. However, it does appear that agents that discover traveling slowly is wealth-enhancing are rewarded with positions in the top of the wealth distribution, due to the difference in *speed* choices employed by agents in the first and fifth quantiles.

In the neoclassical framework, agent rationality rules out negligent behavior by definition under most negligence rules. Agents know that if they fail to exercise care they will bear the full costs of accidents they cause, so it is rational to exercise the level of care that meets the legal standard and nothing more. The persistent presence of negligent injurers is indicative of an institutional failure to provide appropriate incentives to agents to engage in non-negligent behavior. However, it could also mean that individuals select negligent strategies because it is profitable for them to do so despite bearing liability for accidents.

Figure 8 is a graph of the average number of negligent agents through time for several negligence rules where due care is defined in terms of capacity. The negligence rules under which due care is defined as $capacity < 25$ and $capacity > 9$ both succeed in ridding society of negligent behavior. However, even the optimal rule (for the 100 agents and 100 eggs design point) in terms of wealth production, namely due care defined as $capacity < 17$, fails to completely rid agents of negligent behavior. The other rules considered fail to rid society of negligent behavior, as well. Whether some number of negligent agents are present in the wealth maximizing scenario is an empirical matter, but one that Figure 9 begins to address.

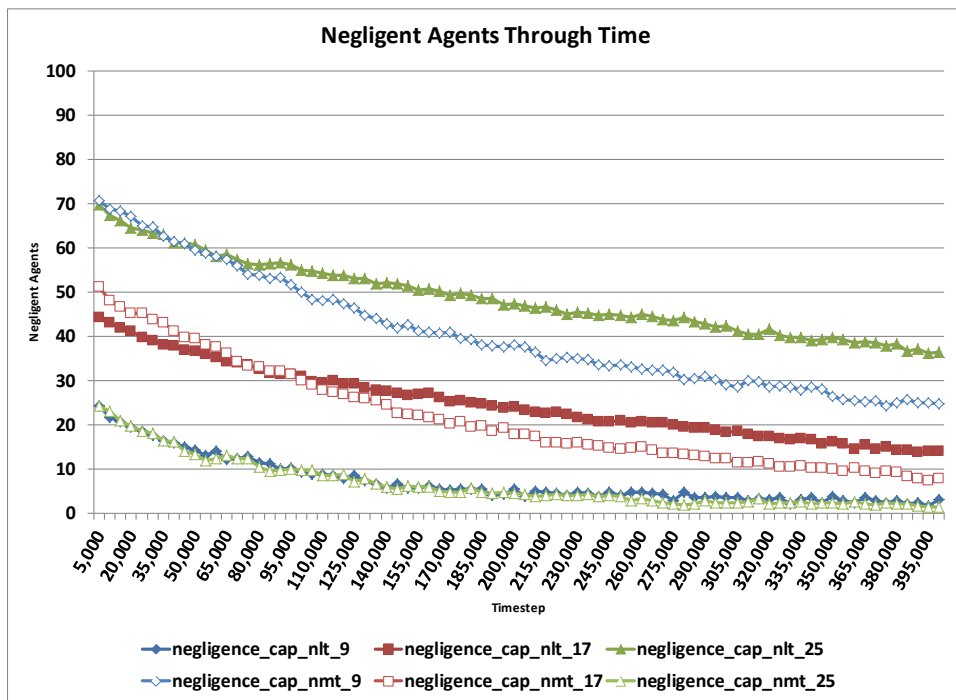


Figure 8 Persistence of Negligent Agents Through Time

It appears the suboptimal negligence rules are possibly so due to their inability to provide agents with sufficient incentives to behave non-negligently. In cases where due care is improperly set,

the incentives seem to be such that some agents are not guided to engage in non-negligent behavior.

5 Conclusion

In this paper we demonstrate that the agent-based modeling approach is able to adjudicate between numerous liability rules and determine the rule or set of rules that achieve a particular performance standard in regards to any number of effectiveness measures. The evolutionary choice algorithm that agents employ to select their strategies drives the complex adaptive system of this artificial society to eventually achieve an institutionally contingent social wealth maximizing steady state. Agents, in diligent pursuit of ever more eggs, search for strategies that tend to result in higher egg production, and gradually push the system to achieve a steady state level that is in the neighborhood of the highest achievable under that liability rule.

The power of the evolutionary approach is not simply that it provides a dynamic method for differentiating between various liability rules, it is that the approach enables the exploration of population dynamics and the close examination of out of equilibrium behavior. Among the realizations this framework yields is the notion that while the aggregate system appears to achieve a relatively stable level, the agents who populate the system continue to furiously grope around the domain in search of wealth enhancing strategies. In addition, agents may elect to be careful even when the neoclassical theory predicts otherwise, and particular negligence rules tend to differ significantly in their ability to rid society of negligent behavior.

Acknowledgements

For guidance, discussion, and useful suggestions I would like to thank Richard Wagner, Charles Rowley, Robert Axtell, Virgil Storr, David Prychitko, Steven Horwitz, and Valentina Dimitrova-Grajzl for comments on earlier versions of this paper. Any errors remain mine alone. The opinions contained herein do not necessarily reflect the opinions of the Department of Defense or those of the Federal Government.

References

- Axtell, R., 2007. What Economic Agents Do: How Cognition and Interaction Lead to Emergence and Complexity. *The Review of Austrian Economics*, 20, 105-122.
- Becker, G.S., 1976. *The Economic Approach to Human Behavior*. Chicago: University of Chicago Press.
- Brenner, T., 2006. Agent Learning Representation: Advice on Modeling Economic Learning. In: L. Testfatsion and K. Judd (Eds), *The Handbook of Computational Economics: Agent-Based Computational Economics, Volume 2*. Amsterdam: North-Holland.
- Cooter, R., Ulen, T., 2007. *Law and Economics*. Boston: Pearson Addison Wesley.
- Diamond, P.A., 1974. Single Activity Accidents. *Journal of Legal Studies*, 3, 107-164.
- Diianni, I., 2007. *An Economic Analysis of Judicial Doctrine*. Ph.D. dissertation. Fairfax: George Mason University.
- Duffy, J., 2006. Agent-Based Models and Human Subject Experiments. In: L. Testfatsion and K. Judd (Eds), *The Handbook of Computational Economics: Agent-Based Computational Economics, Volume 2*. Amsterdam: North-Holland.
- Gode, D.K., Sunder, S., 1993. Allocative Efficiency of Markets with Zero Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, 101, 119-137.
- Hylton, K. 2002. An asymmetric-information model of litigation. *International Review of Law and Economics*, 22, 153 – 175.
- Hylton, K.N., Miceli, T.J., 2005. Should Tort Damages be Multiplied? *Journal of Law, Economics and Organization*, 21, 388-416.
- Katz, A., 1987. Measuring the Demand for Litigation: Is the English Rule Really Cheaper? *Journal of Law, Economics, and Organization*, 3, 143-176.
- Landes, W.M., Posner, R.A., 1981. The Positive Economic Theory of Tort Law. *Georgia Law Review*, 15, 851-924.
- Lavoie, D. 1994. The interpretive turn. In P. Boettke (Ed.) *The Elgar companion to Austrian economics*. Cheltenham, UK: Edward Elgar.

- Law, A.M., Kelton, D.W., 2000. *Simulation Modeling and Analysis*. 3rd Ed. Boston: McGraw Hill.
- Nelson, R.R., Winter, S.G., 1982. *An Evolutionary Theory of Economic Change*. Cambridge: Harvard University Press.
- Nelson, Richard R., Sidney G. Winter, and Scott Schuette. 1976. Technical Change in an Evolutionary Model. *The Quarterly Journal of Economics*, 90, 90-118.
- Picker, R.C., 1997. Simple Games in a Complex World: A Generative Approach to the Adoption of Norms. *University of Chicago Law Review*, 64, 1225-1288.
- Picker, R.C., 2002. SimLaw 2001. *University of Illinois Law Review*, 2002, 1019-1032.
- Posner, R.A., 2007. *Economic Analysis of Law*. Austin: Wolters Kluwer.
- Potts, J., 2000. *The New Evolutionary Microeconomics: Complexity, Competence, and Adaptive Behavior*. Cheltenham, UK: Edward Elgar.
- Rizzo, M.J., 1980. Law Amid Flux: The Economics of Negligence and Strict Liability in Tort. *The Journal of Legal Studies*, 9, 291-318.
- Smith, V., 2003. Constructivist and Ecological Rationality in Economics. *American Economic Review*, 93, 465-508.
- Shavell, S., 1987. *Economic Analysis of Accident Law*. Cambridge: Harvard University Press.
- Yee, K.K. 2005. Common Law Efficiency under Haphazard Adjudication. *Available at SSRN: <http://ssrn.com/abstract=270593>, last accessed 9 April 2014.*