

Dynamics of Rule Revision and Strategy Revision in Legislative Games

Moshe Looks^{a,1}, Ronald P. Loui^a and Barry Z. Cynamon^b

^a *Dept. of Computer Science, Washington University in St. Louis, USA*

^b *Graduate School of Business, University of Chicago, USA*

Abstract. Many legislative games of interest defy classical assumptions and techniques; they tend to be open-ended, with weakly defined objectives, and either non-competitive or pseudo-competitive. We introduce a conceptual and mathematical framework for grappling with such systems. Simulation results are presented for basic specifications of the framework that exhibit a number of qualitative phenomena overlapping with real-world dynamics across a broad spectrum of settings, including aspects of financial regulation and academic decision procedures, that as we demonstrate, may be viewed through the lens of our framework.

Keywords. Multiagent Systems, Social Simulation, Legislative Revision

1. Introduction and Overview

It is customary for mathematical economics and political economy to assume that self-interested social agents can be effectively incentivized in the presence of the right economic controls. Almost every institution seeks to achieve its goals by publishing criteria for rewarding and penalizing the behavior of agents. In this way, legislators seek a meritocratic society that provides more room for free choice than an authoritarian society.

There are many problems with the assumption of effective incentivization. For one thing, agents may place value on unanticipated substantive dimensions, unforeseen by the legislators. Worse, agents are known to attach value to procedural, process, and contextual components of their experience, such as the valuing of individuality, the devaluing of conformity or compliance, the mechanization of their behavior, or the fairness of the outcomes, whether the outcomes result mainly from individual contribution to collective action or not. In many cases, the non-ideal incentive drives the behavior of the system more than the non-ideal behavior of the agents, or the non-ideal responsiveness of the process.

A couple of issues are appropriate for investigation through agent modeling and computational simulation of the resultant society. Here, we are concerned with the legislator's problem of *abridgement*, which results from the legislator's inability to publish incentives that are precisely in line with the social objectives. In jurisprudence there is a tradition of assuming that the rules are imperfect mainly for linguistic reasons, the legis-

¹Correspondence to: Moshe Looks, Dept. of Computer Science, Washington University, St. Louis, MO 63130, USA Tel.: +1 202 641 2157; mlooks@cse.wustl.edu. Also loui@cse.wustl.edu.

lature’s inability to express explicitly and precisely the behavior that is desired. There are also pragmatic reasons why abridgement occurs: the process of producing controls has many imperfections, such as a limit of agreement among conflicting legislative views, a limit to revision, a limit to the practicable complexity of rules, and a limit to the legislators’ understanding of the agents it seeks to incentivize.

Most importantly, we consider the race conditions that occur when the legislature has bounds on how quickly it can revise its edicts, and the agents have bounds on how quickly they can revise their positions.

All of these issues are within the purview of the proposed model. In addition, in a nod to Herbert Simon, agents are restricted in their knowledge of admissible strategies, both in terms of their legality and their potential as a revised position.

The aim of this paper is to produce a simple and flexible mathematical model that captures some of the interesting phenomena we claim to observe in practice. In this way, we seek a better understanding of the process of controlling agents through the revision of non-ideal incentives.

2. Related Work

Broadly speaking this paper is founded on the use of multi-agent system models within computational social science, as advocated/demonstrated in [7,9,5]. There is also prior work focusing in particular on incentivizing agents; e.g., with respect to violation of norms (*cf.* [4,3]), and specific scenarios such as rent control [2] and water demand [11]. More generally, our work is based on fundamental concepts of bounded rationality [13], legal positivism [10], and the study of incentives for individual advancement [14].

3. Veridical Utility and Public Incentive

Consider a legislature incentivizing a population of agents toward some goals. Real goals are complex and may be implicit; they must be approximated via a legislative *abridgement*, \mathbb{A} . The global utility function or *truth* (\mathbb{V}), embodying the actual goals, is over some very high dimensional space where agents live.¹ Over time, agents move through this space by changing their *positions*. A position corresponds to a strategy, action, or mode of behavior in the given domain (e.g., a particular kind of stock sale, managerial style, research focus, etc.).

Assume a veridical utility function mapping from \mathbb{R}^n to \mathbb{R} ; the legislative abridgement used for public incentive, as a function of agent position $\vec{x} \in \mathbb{R}^n$, might be

$$\mathbb{A}(\vec{x}) = \mathbb{V}(\vec{x}A), \quad A \in \mathbb{R}^{n,n}, \quad \text{rank}(A) = k \ll n. \quad (1)$$

Here, \mathbb{A} downsamples \mathbb{V} by projecting n -dimensional points onto a k -dimensional hyperplane passing through the n -dimensional space. In contrast to \mathbb{V} , \mathbb{A} is assumed to be publicly available, explicit, and relatively inexpensive to compute.

¹Throughout this paper, we assume that \mathbb{V} may be additively summed over agents; this need not be the case, but is often a useful simplification.

The social counterparts of veridical and abridged utilities are simply the summations over the entire population; we will assume a population of N agents, with positions $\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N$. Agents are not omniscient. Each agent has a set of known positions, $\Upsilon_i \subseteq \mathbb{R}^n$, to choose from. Furthermore, many positions are legislatively inadmissible, represented by a set of banned positions, $\Omega \subset \mathbb{R}^n$, which is determined by the legislature.

4. Change

At discrete time steps, all agents may update their positions. Similarly, the legislature may update \mathbb{A} and Ω . This may be denoted by subscripting them with a time-step (e.g., \mathbb{A}_t and $\vec{x}_{i,t}$ specify the legislative abridgement and i th agent's position at time t , respectively). The spaces over which these are defined are assumed to be static, however.

If \mathbb{A} perfectly mirrors \mathbb{V} , the legislature has no need for rule revision. When this is not the case, inaccuracies in the abridgement can lead to *legislative misdirection*. That is, agents end up in the wrong regions of the space (from a \mathbb{V} -maximizing perspective), and the public incentives must be revised to compensate.

4.1. Bounds

There are often temporal constraints on the rates of legislative *change* and *growth* over time. Legislation is ρ - δ -bounded if

$$\begin{aligned} \forall t \geq 0, \quad |A_t, A_{t+1}| \leq \rho & \quad \forall t \geq 0, \quad |\Omega_t, \Omega_{t+1}| \leq \rho \\ \forall t \geq 0, |A_{t+1}| - |A_t| \leq \delta & \quad \forall t \geq 0, |\Omega_{t+1}| - |\Omega_t| \leq \delta, \end{aligned} \quad (2)$$

where $|\cdot, \cdot|$ and $|\cdot|$ are normalized metrics of distance and size (or complexity), respectively, over $\mathbb{R}^n \rightarrow \mathbb{R}$ on the left, and $2^{\mathbb{R}^n}$ on the right.

Constraints on an agent's change (in position) and growth (in the set of known positions) are analogous. The i th agent ($1 \leq i \leq N$) is λ - μ -bounded if

$$\begin{aligned} \forall t \geq 0, \quad |\vec{x}_{i,t}, \vec{x}_{i,t+1}| \leq \lambda & \quad \forall t \geq 0, \quad |\Upsilon_{i,t}, \Upsilon_{i,t+1}| \leq \lambda \\ \forall t \geq 0, |\vec{x}_{i,t+1}| - |\vec{x}_{i,t}| \leq \mu & \quad \forall t \geq 0, |\Upsilon_{i,t+1}| - |\Upsilon_{i,t}| \leq \mu, \end{aligned} \quad (3)$$

with normalized metrics $|\cdot, \cdot|$ and $|\cdot|$ over \mathbb{R}^n (left) and $2^{\mathbb{R}^n}$ (right). A population of agents may be said to be λ - μ -bounded if all of its members are.²

5. Dynamics

In this section, we will outline in more detail various agent and legislative strategies which can give rise to the emergent dynamics we are interested in studying.

²Note that if an agent's position is *extinguished* (i.e., added to Ω) λ - μ -bounds may need to be violated to move to a new position. Similarly, if $\Upsilon_{i,t} \cap \Omega_t = \emptyset$, an agent may be unable to field *any* position.

5.1. The Character of Simple Agents

In the simplest case, all agents are interchangeable, but begin play with different initial positions. The set of known positions may simply be a spherical neighborhood of the player’s position defined by a constant radius (d). More powerful agents might be specified by increasing d . At any given time-step, agents simply pick whichever known position maximizes the legislative abridgement (i.e., hill-climbing). Putting it all together, we have

$$\begin{aligned} \forall 1 \leq i \leq N, & \quad \vec{x}_{i,0} = c_i \\ & \quad \Upsilon_{i,0} = \{\vec{x} \in \mathbb{R}^n : |\vec{x}, \vec{x}_{i,0}| \leq d\} \\ \forall 1 \leq i \leq N, t > 0, & \quad \vec{x}_{i,t} = \arg \max_{\vec{x} \in \Upsilon_{i,t-1}} A_t(\vec{x}) \\ & \quad \Upsilon_{i,t} = \{\vec{x} \in \mathbb{R}^n : |\vec{x}, \vec{x}_{i,t-1}| \leq d\}, \end{aligned} \quad (4)$$

where c_i , the initial agent locations, are drawn from some prior distribution. Such agent behavior is clearly λ - μ -bounded, with $\lambda = d, \mu = 0$, using straightforward metrics (in this context, behavioral complexity is assumed to be constant). This protocol for agent behavior (or, *agent strategy*) will be referred to as *greedy* $_{A,d}$.³

However, many of the effects we are interested in studying will only emerge from a substrate containing inter-agent effects. The simplest of these is *imitation*. Let \bar{x}_t be the average of all agent positions at time t . We may define a new agent strategy, *imitate* $_{A,d}$, as moving distance d from one’s previous position toward \bar{x}_t . This behavior is also clearly λ - μ -bounded ($\lambda = d, \mu = 0$). We may also of course consider composite agents that weight greediness versus imitation, or alternate between greedy and imitative steps.

5.2. Simple Agents with Idealized Public Incentive

Before proceeding to more complex agents and dynamic legislation, let us demonstrate a few basic effects with the simple agents described above and idealized public incentive (hence a static legislature). Here we will have an unchanging $\mathbb{V} = \mathbb{A}$, reducing the model to a swarm of optimizers with a fixed utility function. Consider $n = 2$ and the extremely simple one-peak function $A(s) = -|\vec{x}, (0, 0)|$, where $|\cdot, \cdot|$ is normalized Euclidean distance.⁴

The initial player positions are uniformly randomly sampled from $[0, 1] \times [0, 1]$. Position selection is implemented by sampling from the current position, along with 100 uniformly spaced points on a circle with radius $d = 0.01$ centered at the current position. Simulations are run for 100 time-steps, with 20 agents.

Figure 1 shows the results for three different agent populations. The first (upper-left) consists entirely of greedy agents that converge quickly to the global optimum. The second (upper-right) consists of imitative agents that fail to reach the global optimum. Since none of the agents are novelty-seekers, the global utility will remain roughly con-

³As a technical note, agents here are exploring what might be called the \mathbb{V} -neighborhood of their positions, rather than the \mathbb{A} -neighborhood. To clarify, \mathbb{A} will typically map agent position to a subspace (i.e., reducing the intrinsic dimensionality of the space from n to k). However, this does not circumscribe agent explorations; they search through the full n -dimensional space, albeit in a very incremental and limited fashion.

⁴It is important to note that we are not interested in function optimization per se; rather, these examples are used to illustrate the emergence of interesting dynamics in our model that mirror some real-world phenomena.

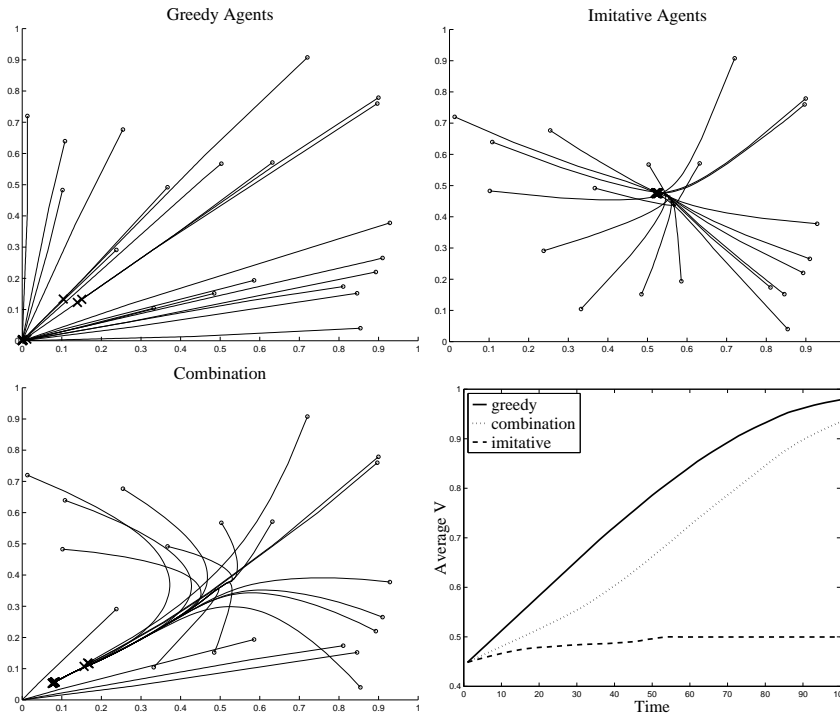


Figure 1. Trajectories for three populations of agents over time: greedy, imitative, and a combination of the two. The global optimum is the lower-left corner, $(0, 0)$. Initial agent positions are the same for all populations, and indicated by small dots. Final agents positions are denoted by \times es.

stant (see lower-right). The third case (lower-left) consists primarily of imitative agents, combined with a small number (4) of greedy agents that act as bellwethers for the population as a whole. Since these agents are imitated, but do not imitate, they will have a disproportionate effect on the final outcome of the simulation.

5.3. The Character of Simple Legislatures

In the simplest case, the legislature is merely reactive, modifying A (used to compute \mathbb{A} as described above) and Ω in response to trends in agent movement through the space. The set of inadmissible positions may be implemented as an initially empty queue that is slowly filled with agent positions which are undesirable from the legislature's perspective as a \mathbb{V} -maximizer. If the legislature is unable to compute this directly, it may be approximated via a background function (described below). Similarly, \mathbb{A} may be adjusted (e.g., in terms of which dimensions are supported) to maximize the expected gain in \mathbb{V} with respect to current agent positions.

This legislature will be ρ - δ -bounded, with $\delta = 0$. The exact value of ρ will depend on the particular updating rule and distance metric used. For example, if at most m dimensions are modified, and the distance between legislations is the number of dimensions by which they differ, we will get $\rho = m$.

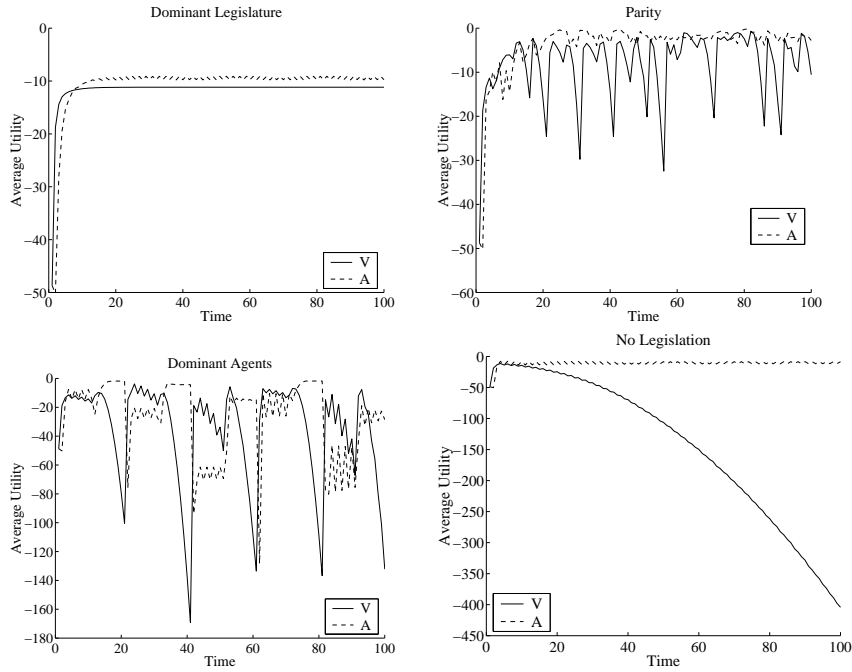


Figure 2. Veridical and abridged utility over 100 time steps for four legislative configurations.

5.4. Simple Agents with Circumscribed Public Incentive

This section will explore simulations where public incentive is circumscribed, leading to an abridgement that changes over time. Again we will consider $n = 2$, but this time with $k = 1$. The legislator is limited to mapping the agents' positions onto a line. The function to be optimized here is again quite simple, $-(1-x)^2 - (1-y)^2$, which has a global optimum at $(1, 1)$. Initial agent positions are uniformly drawn from $[0, 10] \times [0, 10]$.

At each *legislative update*, a new \mathbb{A} may be computed (Ω is not used here). That is, the legislature must choose a line to project the agent positions onto. Ideally, each agent would be mapped onto a line passing through its own position and the global optimum. Since the legislation must be identical for all agents however, this is generally not possible. Instead, a line passing through the optimum and the average of all agent positions, \bar{x}_t , is used.

What dynamics will arise if all of the agents follow the *greedy*_{A,d} strategy? It depends on how often legislative updates (i.e., rule revisions), occur. Figure 2 shows average \mathbb{V} and \mathbb{A} values over 100 time steps for four regimes; dominant legislature (updates at every time step), parity (updates every other time step), dominant agents (updates every ten time steps), and no legislative updates.

What is going on here? Let's first consider the case where there are no legislative updates (lower-right). To begin with, agents are directed toward the optimum, and both \mathbb{V} and \mathbb{A} rise. Quickly however the legislation becomes out of date, and while \mathbb{A} remains roughly constant, it no longer mirrors veridical utility, which falls dramatically as agents move further and further away from the high- \mathbb{V} regions of the space.

Next, consider a dominant legislature (upper-left). Both \mathbb{A} and \mathbb{V} rise precipitously, as expected. But why is the global optimum (0) never reached? The problem is that we are dealing with many agents, and there is no mechanism for instituting personalized legislation. That is, the simulation reaches a local optimum where it is simply not possible for the legislature to nudge any agents closer to the global optimum without nudging other agents correspondingly further away.

Finally, some words on the most dramatic scenarios, dominant agents (lower-left) and parity (upper-right). In the former, legislature has some power to mitigate disastrous crashes of \mathbb{V} , but only at the cost of introducing similar chaotic fluctuations in \mathbb{A} , as infrequent but massive legislation wracks the fitness landscape. The latter still experiences periodic crashes of \mathbb{V} as the legislature falls temporarily behind, but at the same time, reaches the highest levels of \mathbb{V} more often than any of the other scenarios shown.

5.5. Agent Behavior in Complex Domains

The agent behavioral schema outlined above neglected the issue of strategy complexity by assuming that all strategies were of equal complexity. To understand how this may not be the case, consider a space \mathbb{R}^n with very large n , where the starting positions of agents are zero along the vast majority of the dimensions. Non-zero dimensions are said to *support* a particular strategy; $\text{support}(s) = \{1 \leq i \leq k : \vec{x}[i] > 0\}$, where $\vec{x}[i]$ denotes the i th component of strategy \vec{x} . We may now define the complexity of a given position, $|s|$, as $|\text{support}(s)|$.

Now, assume a λ - μ -bounded population; movement along known dimensions and discovery along new (previously unsupported) dimensions may both be limited, to differing extents. This leads to a familiar dichotomy of exploration (seeking new dimensions worth optimizing on) versus exploitation (optimizing along currently supported dimensions).

5.6. Legislative Behavior in Complex Domains

Legislation operating in a complex domain may induce diverse and nuanced constraints. For example, an agent may discover a previously unknown strategy (position) along a new dimension which spreads to other agents via imitation and plays havoc with \mathbb{V} (possibly leading to a crash). The legislature may retaliate via modifying \mathbb{A} and/or extending Ω (leading to *strategy extinction*), but is limited here by ρ - δ -bounds.

5.7. Background Functions

To improve the model, consider a *background* function (\mathbb{B}), providing additional insights into \mathbb{V} , which is not determined by the legislature. Agents that integrate this background knowledge into their choice of strategies may have a better chance of improving \mathbb{V} . Furthermore, such agents might be able to avoid strategy extinction, as a region of the space with high \mathbb{A} and low \mathbb{B} would be a prime target for future legislative action.

Background knowledge can also play an important role in overcoming local optima and/or deceptive regions of the search space with respect to \mathbb{A} . It may alternatively be seen as injecting new strategies into play. In a population where some agents are imitative, these innovations can then spread to others.

With respect to some background knowledge \mathbb{B} , one might taxonomize agents into four classes as follows:

- **positive-novel** agents try to maximize \mathbb{B} , without regard for \mathbb{A} .
- **nonpositive-novel** agents try to maximize \mathbb{B} , while avoiding decreasing \mathbb{A} .
- **positive-compliant** agents try to maximize \mathbb{A} , while avoiding decreasing \mathbb{B} .
- **nonpositive-compliant** agents try to maximize \mathbb{A} , without regard for \mathbb{B} .

According to this taxonomy, *greedy_A* agents, for example, would be classified as *nonpositive-compliant*.

6. Scenarios

There is perhaps no better real-world example of the type of legislative games we address in this paper, and hence inspiration for future work, than the complex financial markets that allocate capital in the modern global economy. These scenarios exhibit great complexity in terms of the interactions between regulator and regulated (strategy extinction vs. strategy discovery).

The legislator's motivation for intervening in financial markets is unquestionable. Given populations' exposure to finance, and the tremendous power of the financial institutions that control the flow of capital, the finance industry is seen as being too important to be left to its own devices. Every time there is a bank failure or an investment scandal, the resulting hysteria either gives members of the legislature an opportunity to push through a bill, or forces them to in order to appease their frightened or angry constituents.

Thus the legislator bans certain strategies while encouraging others. It uses rule revision to incentivize such behavior as it expects to reduce vulnerability, over-exposure, and over-borrowing, in order to promote long-term cautious investment. The agents respond with an impressive rate of strategy discovery. As Pixley [12] writes, "the [financial] sector has diced and packaged social relations of debts and mortgages, of future promises into future securities, futures, and derivatives of futures, into reinsuring assurances of insurance."

For example, commercial banks are required by law to hold a certain percentage of all funds deposited in checking accounts aside. This is known as a reserve requirement, and is intended to decrease volatility. Furthermore, it allows the Federal Reserve system to be self-funded, as it accrues more than enough to meet its operating expenses from interest on the treasury bills that it purchases with banks' reserves. Thus, congress has a clear incentive to maintain the reserve requirement—without it, the banks would be free of that implicit tax and the money would go to them instead of the Treasury. Contrariwise, banks (the agents) have a clear motivation to avoid maintaining unnecessarily large reserves.

A new strategy, the retail sweep program, was first initiated in 1994 by a North Carolina bank. Under this scheme, banks move forecasted excess funds out of checkable deposits into money market deposit accounts that have no reserve requirements. It could not have been clear at the time how the bank's regulator, the Federal Reserve, would respond. To date, the Federal Reserve has neither sanctioned nor banned such schemes. Anderson and Rasche [1] state that retail sweep programs have led to an extraordinary unwinding of statutory reserve requirements in the US, reducing required reserves in December 1999 by an estimated \$34.1 billion.

Another illustration in the world of finance is the practice of short selling, the sale of stock that one does not yet own; this is profitable if the price of the stock later falls. In

comparison with “long” selling, where one makes a profit if the price of the stock later rises, selling short is particularly open to manipulative and deceptive techniques. This can give rise to cycles of loophole discovery and closure between brokers (the agents) and regulators (the legislature) – *cf.* Finnerty [8].

Ideally, the legislature has the goal of eliminating manipulative behavior, in this context the “intentional interference with the free forces of supply and demand” [8], leading to artificially depressed share prices. In lieu of this inaccessible \mathbb{V} , various abridgements are used. Historically, the British parliament banned short selling altogether between 1734 and 1860, while the U.S. Securities Exchange Commission’s “uptick rule” (1934) permits short selling only following an uptick (trade where the price of a stock rises).

Unlike these financial scenarios, academic tenure, our final example, involves a *decision procedure*. Conceptually, agents are continually being removed from the space (a fixed interval after their creation) and replaced with new agents at “random” initial positions. Tenure is presumed to be granted if an agent’s \mathbb{A} -value is above a fixed constant at termination.

The legislator wishes to make a decision based on the long-term research potential and teaching ability of the assistant professor. Typical proxies are teaching evaluations, letters of recommendation, and candidate-compiled research dossiers. It has been argued, somewhat dubiously, that “the minimum publication requirement of the tenure contract induces the optimal level of research with less variation in expected income, avoiding inefficient behavioral responses to the greater riskiness of a contract rewarding only realized publications” [6]. However, it seems unavoidable that abridgements based on proxies such as number of publications may cause legislative misdirection away from the hard-to-measure cost and risk of innovation.

7. Conclusion

In this paper we have developed a simple conceptual and mathematical framework for describing a legislature incentivizing a population of agents toward complex goals. We describe some qualitative phenomena of interest exhibited in by the implementation of a simple version of the framework. The relationship to real-world phenomena is also explored.

A taxonomy of agents is given, distinguishing between agents optimizing in socially productive and socially unproductive ways (as seen from the legislator’s point of view). The results of simulation are given, in which the rapidity of legislative revision and the rapidity of agents’ strategy revision are in competition. The model is shown to produce societies in which agents behave pathologically, as well as societies in which agents can be effectively controlled, exactly as one would hope from a simulation model with the right structural relationships.

References

- [1] Richard G. Anderson and Robert H. Rasche. Retail sweep programs and bank reserves, 1994-1999. Technical Report 2000-023A, Federal Reserve Bank of St. Louis, St. Louis, MO, 2000.

- [2] R.N. Bernard. Using adaptive agent-based simulation models to assist planners in policy development: the case of rent control. Technical Report Working Paper 99-07-052,, Santa Fe Institute, Santa Fe, NM, 1999.
- [3] G. Boella and L. van der Torre. Local policies for the control of virtual communities. In *Procs. of IEEE/WIC WI-03*, Halifax, Canada, 2003. IEEE Press.
- [4] G. Boella and L. van der Torre. Norm governed multiagent systems: The delegation of control to autonomous agents. In *Procs. of IEEE/WIC IAT-03*, Halifax, Canada, 2003. IEEE Press.
- [5] C. Castelfranchi. The theory of social functions: challenges for computational social science and multi-agent learning. *Cognitive Systems Research*, 2, 2001.
- [6] Zhiqi Chen and Stephen Ferris. A theory of tenure for the teaching university. *Australian Economic Papers*, 38:9–25, 1999.
- [7] Joshua M. Epstein and Robert Axtell. *Growing artificial societies: social science from the bottom up*. The Brookings Institution, Washington, DC, USA, 1996.
- [8] John D. Finnerty. Short selling, death spiral convertibles, and the profitability of stock manipulation. *SEC Regulatory Action*, 2005.
- [9] G.N. Gilbert and J. Doran. *Simulating societies: the computer simulation of social phenomena*. UCL Press, London, UK, 1994.
- [10] H.L.A. Hart. *The Concept of Law*. Clarendon Law Series. Oxford University Press (Clarendon), London, 1961.
- [11] Scott Moss, Thomas E. Downing, and Juliette Rouchier. Demonstrating the role of stakeholder participation: An agent based social simulation model of water demand policy and response. Technical Report CPM Report 01-79,, Centre for Policy Modeling, Manchester Metropolitan University, 2000.
- [12] Jocelyn Pixley. *Emotions in Finance: Distrust and Uncertainty in Global Markets*. Cambridge University Press, Cambridge, UK, 2004.
- [13] H.A. Simon. *Models of Bounded Rationality*. MIT Press, Cambridge, MA, 1982.
- [14] Thorsten Veblen. *The Theory of the Leisure Class*. Number 1899 in History of Economic Thought Books. McMaster University Archive for the History of Economic Thought, 1899. available at <http://ideas.repec.org/b/hay/hetboo/1899.html>.