The Game of Scale: Decision Making with Economies of Scale

Christopher J. Hazard
North Carolina State University
Raleigh, NC 27695
cjhazard@ncsu.edu

Peter R. Wurman
North Carolina State University
Raleigh, NC 27695
wurman@ncsu.edu

ABSTRACT

While diffusion of innovation topics in economics and majority games in game theory have been widely studied, the impact of economy-of-scale effects in aggregated decision making has received little attention. In this paper, we present a basic model, the Game of Scale, to study the effects of economy-of-scale in decision making among a large pool of self-interested agents. We solve the model's static equilibria and present two dynamic decision models, one myopic and one trend-following. Most of the parameter space converges quickly; however, the behaviors exhibited near critical input values show drastic changes. We demonstrate how trend-following can improve global outcomes over myopic decision making. Finally, we describe how the game can be externally controlled.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Science—*Economics*; K.4.4 [**Computing Milieux**]: Computers and Society—*Electronic Commerce*

General Terms

economics, theory

Keywords

economy of scale, innovation diffusion, decision model, game theory

1. INTRODUCTION

Understanding dynamics of agent cooperation in decisions exhibiting economy-of-scale utility is of critical importance to society. In technology adoption, we frequently hear about the economy's dependence on oil and its concomitant effects on environment and global politics. By the U.S. government's estimates,[1] price manipulations by OPEC have cost the U.S. economy $7 trillion in the past 20 years. Despite availability of cleaner alternative sources of energy, the world is still chiefly driven by oil. Adoption of cleaner fuels is driven by economics; when it becomes in a firm or individual's best interest, they will adopt the new technology.

While some technologies are immediately cheaper than their predecessor, others are characterized by an economy of scale effect. This means that the technologies do not become generally cost effective until a large enough population adopts them, either because the increased market for the product drives innovation and pushes costs down (e.g., the recent advances in affordable wind power), or because the technology is characterized by network externalities (e.g., hydrogen cars depend on a hydrogen distribution system, which does not become cost effective until enough people own hydrogen cars).

Surprisingly, though innovation diffusion and network externalities have been widely studied, the vast majority of this work does not explicitly model economy-of-scale effects. To the best of our knowledge, no current work in innovation diffusion and similar topics meets the following criteria: simplicity of the model in terms of low dimensionality, explicit modeling of switching costs, and utility functions that capture economy-of-scale effects. In this paper, we propose an iterative, game theoretic model of technology adoption that we call the *Game of Scale*, in which agents make choices about whether or not to adopt a new technology. The cost of the technology is a function of the number of agents who have adopted it.

Real world examples of the game of scale abound in both decision support systems and economics. For example, adoption of biodiesel may be limited due to high initial costs, despite renewability and lower amounts of harmful pollution. Another example is that the market success of more efficient hybrid and diesel vehicles depends on manufacturers' decisions to build and market them. With larger scale production, the reduced marginal cost might allow these technologies to succeed in the market, but initial adoption cost and uncertainty slow adoption.

From a consumer's perspective, many substitutable yet incompatible products exhibit switching costs and network externalities. Some examples are cell phone networks that offer free in-network calls, social networking websites that allow individuals to advertise their friendships, and portable music players with services that only offer music in a non-transferable proprietary format. In such products, switching costs can be in terms of terms of money, time, and effort. Similarly, economies of scale are manifested in forms such as cost, availability of aftermarket components, and compatibility with products owned by friends, family, and coworkers. Our model also sheds light on the adoption—and obsoletion—of social trends facilitated by advances in communication, such as the Internet. In particular, our results show that trend-following leads to more rapid adoption than myopic decisions. Mass media provides

[1] Source: http://www.fueleconomy.gov/feg/oildep.shtml

ICEC'07, August 19–22, 2007, Minneapolis, Minnesota, USA.

a feedback loop for such trends, though the feedback may be real or exaggerated.

In this paper, we generally refer to adoption of innovation rather than diffusion of innovation or joining a coalition. We make this distinction as we are focusing on cost-benefit mechanics instead of dissemination of knowledge and risks of an innovation. Although we focus primarily on economic decisions, the model captures aspects of multi-agent negotiation (as when robots decide to clear a highway through rugged terrain) or the adoption of new Internet technologies (like social networking applications).

We propose the Game of Scale as a means of exploring dynamics of economy-of-scale effects on self-interested decisions. In this iterative model, agents are repeatedly given the choice of switching to a new innovation (or switching back if they are currently employing a new innovation), but only make the switch if it incurs less cost. Each agent is characterized by its switching cost, drawn from a probability distribution. The game has relatively low dimensionality which keeps the parameter space manageable, while maintaining the basic features for modeling economies of scale. Our findings show the innovation adoption characteristics of the system, including adoption time and penetration. We also show how the system can get stuck in non-optimal equilibria, how slow transitions can be more costly, and how trend-following can produce better solutions than myopic decisions. This framework offers a way to synthesize and analyze multi-agent systems incorporating economy-of-scale decisions.

2. RELATED WORK

Modeling innovation diffusion has a rich history. Baptisa's survey [1] is a useful synopsis of well-studied economic theories of innovation diffusion. Taking a different approach, Sarkar [15] compares and contrasts differing ideologies of many of the main modeling approaches. Though the field has advanced since Fichman [9] surveyed empirical results of information diffusion models, Fichman's work illustrates complexities of applying these models to the real world.

Reinganum [14] found the Nash equilibria in a game theoretic approach to innovation adoption timing in a duopoly setting. This work differs from ours in two primary ways. First, it assumes that the first firm to switch to the new innovation has explicit extra benefit, implying a first-mover advantage. Second, while Reinganum's model generalizes to n players, she admits that it would be "extremely messy". In contrast, the focus of our work is on the aggregate behavior.

In the Minority Game, agents choose either 1 or 0 at each iteration, and those in the minority are rewarded as a function of the size of the minority [6]. Conversely, in the Majority Game, those in the majority are rewarded accordingly. Dindo [7] extends these models with asymmetric payoffs, making one side, either 1 or 0, more favorable than the other. Intuitively, Dindo's asymmetric majority game is very similar to our model. However, because Dindo's model does not use a switching cost, each successive iteration is independent of its predecessor. This independence sharply changes the analysis and results, yielding chaotic behavior for many regions of Dindo's model.

Loch and Huberman [13] approach innovation adoption in a manner similar to our work, with three major differences. First, they address uncertainty explicitly, whereas we address it indirectly as expected values through marginal and switching costs to reduce system dimensionality. Second, Loch and Huberman only briefly explore a small model with switching costs, whereas switching costs are a primary feature of our work, and are based on a probit model. Third, they use a linear model of utility from additional firms using an innovation to compare performance (profit) of innovations. Our approach is from the other direction; costs may be non-linear, and are reduced asymptotically as more firms adopt the innovation.

Caselli and Coleman [5] present a model of ethnic conflict which exhibits similarities to our model. While their ethnic conflict model encompasses switching costs, it focuses on diseconomies of scale, where as members of one ethnicity switch to another, and resources are shared among a larger group. Their model additionally features a method of expending the resources of one group to take over ownership of a common resource.

The model Stoneman presents [17] is a major inspiration to our work, in that a firm switches to a new innovation if the current benefits of adopting the innovation, based on the current number of subscribing firms, outweigh the cost of switching. Ireland and Stoneman model innovation diffusion in terms of the supply and demand for the given innovation [11]. This is in contrast to our model where the innovation directly lowers a particular cost. Stoneman and David [18] investigated subsidy policy of innovation. Like our model, this model separates a probit characteristic from the added benefits of diffusion. It differs from our model in that it only measures the effect of subsidy between two specific times and assumes an epidemic learning model. Battisti and Stoneman [2, 3] investigate a large case study of inter- and intra-firm innovation diffusion. They find that epidemic learning models do not seem to be significant in intra-firm diffusion, but that rank/probit effects do seem significant [3]. As epidemic learning is not central to our model, this work suggests that our model may be most applicable to systems where knowledge of the innovation is widespread.

Silverberg et al. [16] construct an explicitly evolutionary approach that incorporates factors such as mark-up price, equipment costs, and investments. In contrast, we chose to focus on the dynamics of more simple mechanics.

Farzin et al. [8] apply dynamic programming to model innovation adoption which can model an environment of many simultaneous innovations. While their model is capable of determining the optimal time for firms to adopt a new innovation, they assume that improvements and maturation of the innovation are completely independent of the number of firms using it. Alternatively, our model provides positive feedback for an innovation based on the number of firms that have adopted it.

Congestion games are resource utilization games where each player's payoff is a function of the number of players utilizing the same resource. While the game we present is a congestion game by this definition, congestion games are typically studied with negative externalities and without switching costs. Blumrosen and Dobzinski [4] find the computational complexity of finding optimal welfare for different classes of congestion games, including those with positive externalities.

3. THE GAME OF SCALE

3.1 Description

The Game of Scale has N agents and I distinct competing innovations. At time, t, $n_i(t)$ agents are using innovation i. For convenience we will sometimes denote $n_i(t)$ as n_i. Each agent can only use one innovation at a time, such that

$$N = \sum_{i=0}^{I-1} n_i. \tag{1}$$

We limit most of the discussion in this paper to cover the case when $I = 2$. For the purposes of discussion, we will consider innovation 0 to be the legacy technology, process, or strategy used, and innovation 1 to be the new or invading innovation. An agent in the game can represent a subset of a firm, or an entire firm, thus combining inter- and intra-firm diffusion. To keep the game simple, we assume that each agent uses a single unit of one of the two technologies. We denote the initial number of agents using an innovation as $n_i(0)$. The $n_1(0)$ agents may be considered the entrepreneurs and proponents of innovation 1, who have taken an initial risk (before the game begins) to implement innovation 1.

The agent's objective is to minimize cost. Each innovation has a cost per unit time, c_i, which is composed of a base cost, b_i, and a cost that decreases with the number of agents using it. In this paper, we focus on a cost model with constant overhead cost, where some resource is needed to adopt the innovation, but may be spread without further cost over many adopters. The constant a is divided equally among the number of agents with the innovation, n_i, in proportion to the total number of agents, N, yielding a cost of

$$c_i(n_i) = b_i + \frac{a_i}{n_i/N}. \tag{2}$$

In situations where further learning, experience, and refinement of an innovation yield better results, an exponential cost model may be more appropriate. Such a cost function may be represented as

$$c_i(n_i) = b_i + a_i e^{-d_i n_i/N}. \tag{3}$$

While any decreasing function of n_i/N would suffice, we focus on equation (2). Note that we do not account for risk in terms of the innovation's cost; we assume risk neutral agents and that risk has been accounted for in the innovation cost functions.

Each agent has a cost of switching from one innovation to another. Switching cost is effectively a probit/rank model, where we assume that some characteristic of the agents defines their switching cost. This switching cost is constant for each agent throughout the game. While an agent's switching cost is private information, all agents know the switching cost distribution, making the Game of Scale a Bayesian game. To keep the game simple, the switching cost, S, is the same regardless of the innovation being adopted.

In this paper we examine two switching cost distributions, normal and uniform, with mean μ_S. For the normal distribution, the standard deviation is σ_S, yielding

$$S \sim N(\mu_S, \sigma_S). \tag{4}$$

The uniform distribution has a range about the mean, w_S, corresponding to the standard deviation as $\sigma_S = \frac{w_S}{\sqrt{12}}$,[2] is expressed as

$$S \sim U(\mu_S - w_S/2, \mu_S + w_S/2). \tag{5}$$

The game of scale is iterated, with each agent choosing whether to switch innovations at each step. The game concludes when an iteration occurs where no agents switch.

3.2 Nash Equilibria

While congestion games that include positive reinforcement and player-specific payoffs are not guaranteed to have pure-strategy Nash equilibria [12], the switching costs, iterative nature, and common payoff scaling allow the game of scale to have pure-strategy equilibria for many parameterizations. We define an innovation schedule as a strategy consisting of a chronological list of innovations and times, such that an agent will switch to a specified innovation at a given time (e.g. switch to innovation 1 at time 7, switch to innovation 2 at time 12, etc.).

[2]The variance of a uniform distribution on $[a,b]$ is $\frac{(b-a)^2}{12}$.

Any innovation schedule is a pure-strategy Nash equilibrium provided that the game's parameters meet two criteria. The first criterion is that no innovation, i, when used by a single agent, is cheaper for an agent than any innovation in the schedule, j, when j is used by all other agents. This criterion can be expressed as,

$$c_i(1) > c_{j \neq i}(N). \tag{6}$$

If the first criterion is not met, then for all pure-strategy equilibria, a number of agents will use this cheaper strategy permanently. Intuitively, an innovation that is strictly dominated by another innovation cannot be part of a pure-strategy Nash equilibrium.

The second criterion for an innovation schedule to be a pure-strategy Nash equilibrium for all agents is that no agent's switching cost can be greater than the loss it will incur by skipping the first switch when all other agents are performing two innovation switches. Consider an innovation schedule where all agents start with innovation h, switch to innovation i at time T_i and innovation j at time T_j. If the agent with the largest switching cost will not profit during the interval from switching twice, first to i and then to j, then the agent's optimal strategy in this case would be to only switch once to innovation j at time T_j. Using S_{max} to represent the largest switching cost, this criterion can be expressed as

$$(T_j - T_i)c_h(1) > (T_j - T_i)c_i(N) + S_{max}. \tag{7}$$

The Pareto frontier in the game of scale consists of all allocations except for those with $n_i = 1$, where the single agent has a low enough switching cost such that the agent can switch to a different innovation and still profit. All other allocations are Pareto efficient, because each agent is using the same innovation of at least one other. Because innovation cost is decreasing with the number of agents, each agent would reduce at least one other agent's outcome by switching innovations.

Among innovation schedules, each switch from one innovation to the next incurs additional cost to all agents. Therefore, the socially optimal solution is for all agents to initially switch to the innovation that scales the best (smallest $c_i(N)$), and make no further switches. If the game is being played for only a small number of iterations, then the globally optimal strategy may be for some or all agents to stay with their initial innovation.

Despite being Nash equilibria, the socially optimal equilibrium is payoff dominant and not necessarily risk dominant. Risk-neutral agents' optimal strategy, given that all other agents are also risk-neutral, is the socially optimal strategy in many parameterizations. Similarly, risk-averse agents' optimal strategy may yield a socially costly outcome. In parameterizations and game states where private switching cost is highly relevant in agents' decisions with respect to their risk tolerance, the best solution concept that can be achieved is a Bayesian-Nash equilibrium. Because agents' optimal strategies involve simultaneous switching (and as soon as possible), quick communication is important. For this reason, we look toward decision models that allow agents to signal and commit to their intentions, with the goal of reaching close to the socially optimal solution with risk-averse agents.

4. DECISION MODELS

Because each agent is self-interested, the optimal decision for any agent may not be the optimal decision globally. Here we describe three decision models based on different perceptions of cost.

When using two innovations, such as in our first two decision models, we can reduce the number of variables by combining base costs into a single variable. Consider an agent using innovation 0 in a constant overhead game with two innovations. The cost savings of switching to innovation 1 for the next unit of time is

$$c_0(N-n_1)-c_1(n_1+1)-S = b_0+\frac{a_0}{1-n_1/N}-b_1-\frac{a_1}{(n_1+1)/N}-S. \quad (8)$$

It is convenient to let $b=b_0-b_1$. We can express the profit for switching for one time step, $\pi_{0\to 1}$, by incorporating the expected number of agents using innovation 1 at time t, $E(n_1(t))$. The cost of switching to innovation 1 is

$$\pi_{0\to 1} = -S+b+\frac{a_0}{1-E(n_1(t))/N}-\frac{a_1}{E((n_1(t)+1))/N}. \quad (9)$$

4.1 Equilibrium Decision Model

Since an agent knows the switching cost distribution, but not the other agents' actual switching costs, the agent must determine the expected equilibrium of costs between innovations to make its decision whether to switch innovations. Here we assume that switching costs are significant and that agents know that other agents' switching costs will also be significant. As no agent stands to benefit by switching innovations at the equilibrium, it is static with respect to time. The equilibrium can be expressed as $n_1(t+1)=n_1(t)=E(n_1)$, allowing us to substitute n_1 for $E(n_1(t))$.

For this section, we assume that $n_1(0)=0$. By ranking the agents according to switching cost, we can find the equilibrium points at which the innovation will either succeed or fail based on which side of the point the number of agents using innovation 1 resides. Here we will use a uniform distribution of switching costs, but note that this may be solved for other switching costs as well. Based on agent rank, p, the expected switching cost for agent, S_p, can be found as

$$S_p = \frac{p+\frac{1}{2}}{N}w_S+\mu_S-\frac{w_S}{2}. \quad (10)$$

To find the equilibrium points, we find the number of agents at which the profit, $\pi_{0\to 1}$, from equation (9) is 0. The equilibrium point will only be reached if is profitable for at least one agent to switch from the beginning; this initial switch will cause a cascade until no agent has any further incentive to switch. We substitute the agent switching cost as S_{n_1} from equation (10), and arrange for n_1, yielding

$$\begin{aligned} 0 &= 2w_S n_1^3-((2b+2\mu_S+w_S)N-3w_S)n_1^2 \\ &+((2(a_0+b+\mu_S+a_1)-w_S)N^2-2(b+\mu_S+w_S)N+w_S)n_1 \\ &-(2N^2a_1-(2(a_0+b+\mu_S)-w_S)N+w_S)N. \quad (11)\end{aligned}$$

The roots of n_1 in equation (11) represent the number of adopters of innovation 1 at which the optimal global behavior diverges for adopting innovation 1. For all other values of n_1 other than the aforementioned roots, the agents will do one of three behaviors: all agents will adopt innovation 1, all agents will revert back to innovation 0, or, in the case of multiple real roots, the agents will converge to one of the equilibrium points. The sign of the first derivative of $\pi_{0\to 1}$ represents the direction of innovation adoption on each side of the cost equilibrium point. Direction, $D(n_1)$, is negative when innovation 1 is cheaper. The derivative of $\pi_{1\to 0}$ is similar, in that it is negative when innovation 0 is cheaper. If both derivatives are positive, the system is in equilibrium. Adoption direction from 0 to 1 is represented as

$$D_{0\to 1}(n_1) = \frac{Na_0}{(N-n_1)^2}+\frac{Na_1}{(n_1+1)^2}-\frac{w_S}{N}. \quad (12)$$

4.2 Myopic Decision Model

We define myopic decision making as each individual agent deciding to switch innovations if and only if it yields a decrease in cost for the next time step. When all agents in the game of scale employ myopic decision making, the game simplifies in a number of ways. Further, with myopic decision making, agents exhibit the same behavior whether they decide simultaneously at each discrete time, decide individually in sequence, or decide based on a randomized process.

Using generic cost savings functions, c_i with equation (9), we find the best strategy with all myopic decision making agents is to switch innovations when

$$c_0(N-n_1)-c_1(n_1+1)-S>0, \text{ if using innovation 0}, \quad (13)$$

and

$$c_1(n_1)-c_0(N-n_1+1)-S>0, \text{ if using innovation 1}. \quad (14)$$

Using equations (13) and (14), and the distribution of switching costs, we can find the expected number of agents using innovation 1, n_1. Given a cumulative distribution function of switching cost, $F(x)$, the quantile function gives the expected value for the qth quantile, expressed as $F^{-1}(q)$. We find the ranked expected values of switching costs by evenly dividing the quantile space, $q\in[0,1]$, by the corresponding number of agents. Each agent's expected switching cost is represented by the center of its corresponding quantile range, so we must offset each by $1/2$ agent. We assume the $n_1(0)$ agents initially starting with innovation 1 are chosen independent of their switching cost. The agents with the lowest switching cost using either innovation are the most likely to switch. From this, we can break agents into the two innovations and rank them by their switching costs, with the lowest switching cost at $n_1(0)$, and highest switching cost at $n_1=0$ and $n_1=N$. We can express each agent's rank, $k(n_1)$, as

$$k(n_1) = \begin{cases} 1-\frac{n_1+1/2}{n_1(0)} & \text{if } n_1<n_1(0) \\ \frac{n_1-n_1(0)+1/2}{N-n_1(0)} & \text{if } n_1\geq n_1(0) \end{cases} \quad (15)$$

We can now express the expected switching cost as,

$$E(S(n_1)) = F^{-1}(k(n_1)). \quad (16)$$

We can find the equilibrium by solving for the number of agents using innovation i at time t, $n_i(t)$ when the game is in equilibrium, meaning switching has no benefit or loss. This equilibrium is expressed as,

$$0 = c_0(N-n_1(t))-c_1(n_1(t)\mp 1)\pm F^{-1}(k(n_1(t+1)), \quad (17)$$

where the sign of $F^{-1}(\dots)$ is negative if adopting innovation 0 and positive if adopting innovation 1, opposite the sign of the additional agent to the function c_1. To find the adoption direction, we use equations (13) and (14). If equation (13) is true, innovation 1 is adopted, if equation (14) is true, innovation 0 is adopted, and if both are false, then no agents will switch. Both equations (13) and (14) cannot be true at the same time because a difference between values cannot be simultaneously positive and negative. If we substitute into equation (17) the uniform switching distribution (equation (5)) and the rank formula (equation (15)) with direction chosen by equations (13) and (14), we can find the difference equation for n_1 by solving for $n_1(t+1)$ as if $c_0(N-n_1)-c_1(n_1+1)-S>0$, then

$$n_1(t+1)=n_1(0)\frac{c_0(N-n_1(t))-c_1(n_1(t)+1)}{w_S}+\frac{n_1(0)-1}{2} \quad (18)$$

if $c_1(n_1)-c_0(N-n_1+1)-S>0$, then

$$n_1(t+1)=(N-n_1(0))\frac{c_1(n_1(t))-c_0(N-n_1(t)+1)}{w_S}+\frac{N+n_1(0)-1}{2}. \quad (19)$$

Equations (18) and (19) iterate until the game concludes with n_1 converging when $n_1(t+1)=n_1(t)$.

4.3 Trend-Following Decision Model

When the only information known by each agent is the number of agents in each technology, myopic decision making is rational. Due to the game of scale's iterations, however, each agent could maintain knowledge of the previous states (the number of agents in each innovation at each time in the past). In this section, we present a decision strategy that utilizes these historical trends as signals. In this model, each agent approximates rational expectations with a discount factor. We assume a non-trivial number of agents are playing, as this gives trends statistical significance.

For our trend-following decision model, we combine a Taylor series approximation with a discount factor. Taylor series approximations are widely used in finance as they strongly reflect local trends to predict short-term future results. The number of expansion terms is typically two to three, with two often being an excellent approximation for financial applications [10]. We argue that the Taylor series approximation is well-suited for the game of scale as a short-term trend prediction model. Agents' decisions are based on feedback and trends can rapidly change in the game of scale, as prices in a market often do. Taylor series expansions offer good local approximations and depend only as much history as is needed for a given order of expansion (to make sure there is ample data to approximate derivatives). Taylor series expansion is also a simple trend model which can be varied in accuracy by the number of terms. Further, such trend-following can work even when agents do not know the switching cost distribution.

The first three terms of the Taylor series are represented as

$$\tilde{n}_i(t+\Delta t)\approx n_i(t)+n_i'(t)\Delta t+n_i''(t)\frac{\Delta t^2}{2}+n_i'''(t)\frac{\Delta t^3}{6}, \quad (20)$$

where Δt is the change in discrete time, and $n_i(t)$ is the number of agents using innovation i at time t. The derivatives of $n_i(t)$ are approximated by finding the difference since the previous time. This yields

$$n_i'(t)\approx n_i(t)-n_i(t-1), \quad (21)$$

$$n_i''(t)\approx n_i'(t)-n_i'(t-1), \text{ and} \quad (22)$$

$$n_i'''(t)\approx n_i''(t)-n_i''(t-1). \quad (23)$$

Because the Taylor series is an approximation and loses accuracy further into the future, we employ a discount factor, γ. We can combine the discount factor with each approximated time-step to produce an exponentially weighted expected cost value, $E(n_i)$. To follow a trend, an agent needs to decide if it is better to switch innovations at the current time, t, or to wait some additional time, k, at which to reevaluate. The agent must also apply the discount factor to the switching cost so that the switching cost scales properly with the corresponding cost term when the switch occurs. Using equation (20) as an approximation of the number of agents using an innovation at a specified time, the predicted cost, $\tilde{c}_{i\to j}(k)$, to transition from innovation i to innovation j, at time $t+k$ is

$$\tilde{c}_{i\to j}(k)=(1-\delta_{ij})\gamma^k S+\sum_{l=0}^{k-1}\gamma^l c_i(\tilde{n}_i(t+l))+\sum_{l=k}^{\infty}\gamma^l c_j(\tilde{n}_j(t+l)). \quad (24)$$

In equation (24), the delta function, δ_{ij},[3] prevents adding a switching cost when the innovation is not changing. Using the predicted cost for a given switching time, an agent may find the optimal time, $k_{i\to j}$, to switch from its current innovation, i, to innovation j as

$$k_{i\to j}=\underset{k\geq 0}{\operatorname{argmin}}\,\tilde{c}_{i\to j}(k). \quad (25)$$

We can find the optimal innovation to adopt, ι, by evaluating the predicted cost at each optimal switching time, $k_{i\to j}$, as

$$\iota=\underset{j\in I}{\operatorname{argmin}}\,\tilde{c}_{i\to j}(k_{i\to j}). \quad (26)$$

From equation (26), an agent should switch to innovation ι only if the optimal switching time is the current time, that is, $k_{i\to\iota}=0$. If the optimal switching time is in the future, then it is in the agent's best interest to wait until the next time step and reevaluate its situation. Both the infinite series in equation (24) and the unbounded k in equation (25) can be effectively approximated by truncating the terms. $O(\gamma^l)$ will dominate many practical cost model compositions of a finite Taylor series expansion, such as our fixed overhead cost model, making each successive term less significant.

While equation (24) can predict fast-moving trends regardless of current innovation adoption, large increases and decreases are not as meaningful near the boundaries where an innovation is nearly ubiquitous or unused. Further, the cost functions may be undefined for $n_i<1$ and $n_i>N$. For both of these reasons, the number of agents predicted for any innovation by equation (20) should be clamped to within the range $[1,N]$ by the cost function. Because each agent is evaluating the utility of adopting each innovation, it must consider itself when determining the costs.

Our method of approximation for the derivatives of each innovation yields global predispositions as inputs. The three additional initial conditions with the trend-following dynamic analysis are $n_i'(0)$, $n_i''(0)$, and $n_i'''(0)$. If these values are known by all agents to be 0 (as well as $n_i''(1)=0$ and $n_i'''(1)=n_i'''(2)=0$, because insufficient data will have been gathered up to that point), then the system starts without any anticipation or preconceived beliefs about how the particular game of scale will unfold. These values may be nonzero due to "cheap talk" in a game theoretic or agent perspective, or due to marketing and rumors in a more traditional economic framework. Further, agents may collectively have a distribution of initial derivatives, based on the information each obtained before the start of the game. However, in our investigation of the impact of these initial derivatives, we found that they only had significance when the parameters were near a critical transition region.

[3] The Kronecker delta function yields 1 if $i=j$, 0 otherwise.

5. ANALYSIS

5.1 Behavior Characteristics of the Myopic Decision Model

In this section we examine effects of parameters in the game of scale with constant overhead cost and myopic decisions. We implemented a simple simulation to generate switching costs and iterate equations (13) and (14) to convergence, as an exact analytic solution is infeasible due to the reciprocals in the difference equations.

While parameter choice has a strong impact on the game dynamics, many choices of parameters yield trivial results. For example, if switching cost is low and overhead cost, a_i, is insignificant compared to base cost, b, all agents immediately switch to the better innovation. Alternatively, if switching cost is comparatively high, no agent will switch.

Despite the wide ranges of parameters yielding trivial results, the game of scale has interesting critical ranges where the outcome is very sensitive to small changes in the parameters. To find how each of the parameters impact the results, we initially sampled and explored the parameter space, and found two types of behaviors. Most parameterizations yielded quick convergence, where either a new innovation had no benefit, or completely dominated the prior innovation. The other parameterizations, with less drastic innovation disparity, showed a slower convergence. To demonstrate this range of behavior, we used a typical set of inputs that took many iterations for agents to converge to using innovation 1, and explored the dynamics around this critical region. The set of parameters used is $N = 1000$, $n_1(0) = 74$, $\mu_S = 5.5$, $b = 6$, $a_0 = 0.180$, $a_1 = 0.282$, and $w_S = 2.2$ for a uniform distribution and $\sigma_S = 1.7$ for a Gaussian distribution.

Figure 1 shows the profit for each agent as computed by equation (13) for $\pi_{0\to1}$ and (14) for $\pi_{1\to0}$, given that all agents with lower switching costs (or agents initially using innovation 0) are using innovation 1. This graph uses the aforementioned parameters and the Gaussian distribution of S. The small peak around $n_1(0)$ is caused by the low switching costs for agents on both sides of $n_1(0)$, since the innovation that each agent begins with is independent of switching cost. Note that $\pi_{1\to0}$ is only positive for very low values of $n_1(0)$, where the cost is too high for the few agents to absorb. The profit $\pi_{0\to1}$ reaches 0 after $n_1 = 668$, which is the equilibrium according to the myopic decision making model.

Each graph in Figure 2 shows the results of simulations over a range of one parameter, while holding all of the others constant. The variable μ_S is not depicted since it has the same effect as b when using myopic decision making. The outer graphs have a uniform switching cost distribution, whereas the inset graphs have a Gaussian switching cost distribution. The solid line with square data points is the percentage of agents that have chosen innovation 1, as given by n_1/N. The dashed line represents time until convergence.

Common across all graphs is increased convergence time near critical points, where behavior undergoes a sudden, drastic change. The uniform distribution has sharp linear regions, with the exception of high w_S values, while the Gaussian distribution shows more curved regions. The Gaussian distribution runs showed greater variance, which is expected since the distribution's variance was 2.7 times larger than the uniform distribution. Though not shown in the figure, when the Gaussian distribution's variance is decreased to the depicted uniform distribution, the graphs are nearly identical. The only significant difference between the two plots is that the convergence time was scaled to slightly longer times for the Gaussian distribution. Likewise, but to a lesser extent, when uniform distribution was run with a correspondingly high variance, its graphs resembled that of the Gaussian distribution. The Gaussian distribution runs took longer to converge, independent of input variance. This is due to the increase in spread between samples further from the mean, causing fewer agents to switch innovations at the start. The longer convergence time can be seen in the graphs; while the parameter value and percent in innovation 1 scales are the same between the two distribution graphs, the scale of convergence time for the Gaussian distribution graphs is 4 times larger.

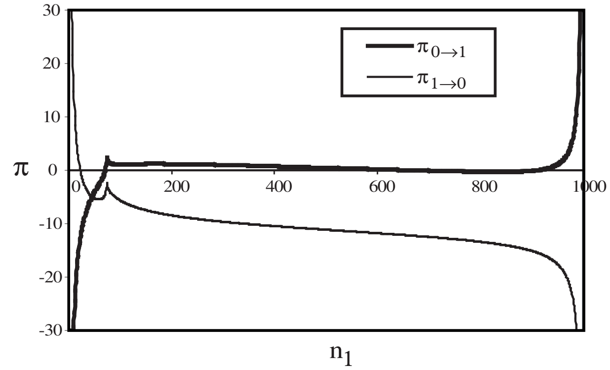

Figure 1: Profit from equations (13) and (14) for each ranked n_1 assuming that $n_1 - 1$ agents are using innovation 1. The profit of switching to innovation 1 is indicated by the dark line, while the profit of switching to innovation 0 is indicated by the light line.

The variable b has 3 distinct regions. The first region covers negative values, where innovation 1 is more expensive than innovation 0. In the second region, even though one innovation costs less than the other, the switching cost dominates the agents' decisions. The third region is where the cost of innovation 0 pushes agents to choose innovation 1.

The overhead cost, a_1, has similar behavior to that of b, only reversed. The only major difference is in the Gaussian distribution in the transition before the full switch to innovation 1. Here, b has a steep transition region, whereas a_1 has a large jump and a shallow transition region. The variable a_0 closely resembles the region of a_1 with lower values. As b is fixed at a significant positive value, we only observe a_0 speeding up the transition to innovation 1.

The initial number of agents starting with innovation 1, $n_1(0)$, also has 3 regions. While the high and low regions follow the same explanation of b, the middle region increases linearly with $n_1(0)$. This linear increase is due to switching costs, which prevent agents from switching back to innovation 0, even though it is less costly. The linear region is much smaller for the Gaussian, which is caused by the higher variance and distribution itself. Though the higher variance has a greater impact in shortening this linear region, the linear region in the Gaussian distribution is still shorter than that in the uniform distribution with corresponding variance. However, the high variance adds a second high-variance linear region in both Gaussian and uniform, with occasional peaks where all agents adopt innovation 1.

When the variance (w_S and σ_S) is low, switching costs are too high for agents to switch with this parameter set. Once the variance is raised enough, the lowest switching values reach the cost difference threshold and a number of agents switch to innovation 1. As the variance continues to increase, the upper bound of switching costs continues to rise, putting more agents' switching costs too high to switch. High variance can also cause agents to continu-

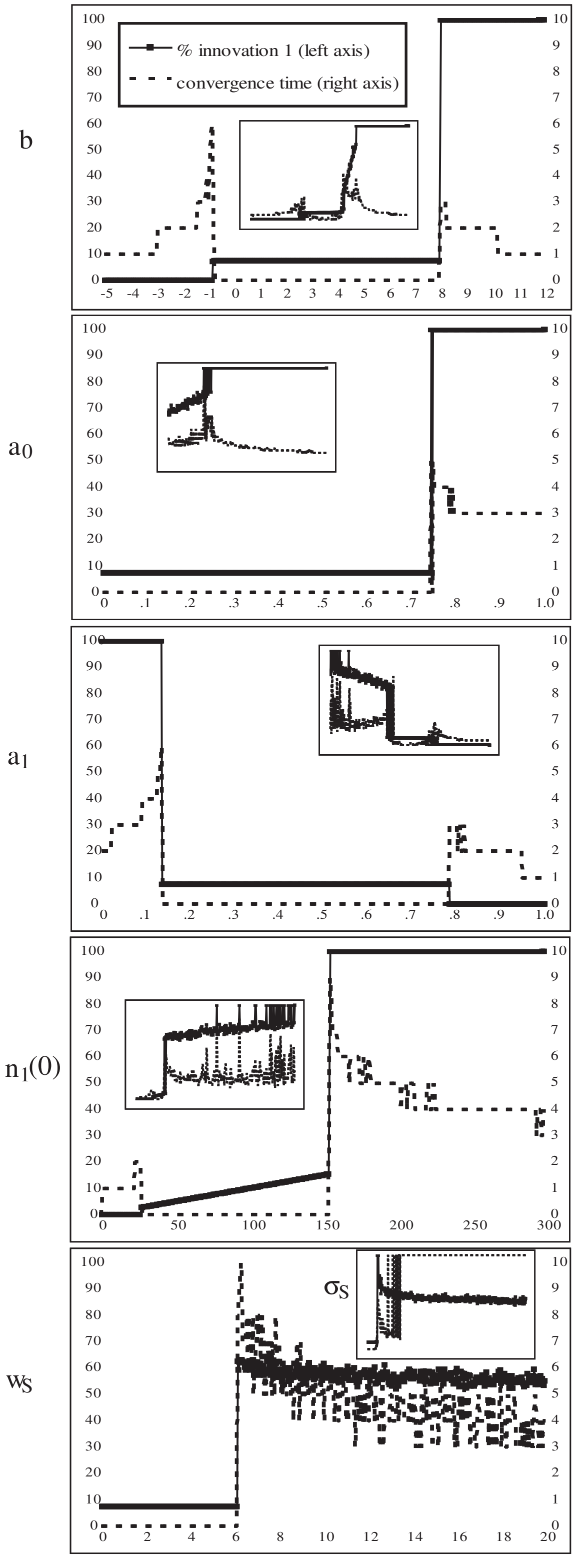

Figure 2: Convergence time and final fraction of agents using innovation 1 with the myopic decision model, each row plotted for the input variable shown as the horizontal axis, with uniform (outer graphs) and Gaussian (inset graphs) switching costs, constant overhead cost, and myopic decision model.

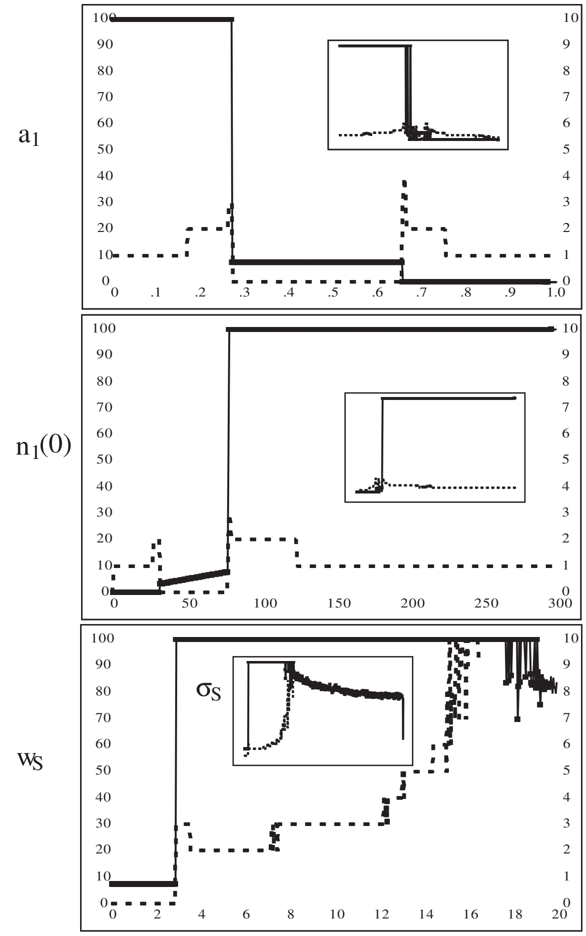

Figure 3: Results of the trend-following decision model, depicted in the same manner as Figure 2.

ally flip when their switching costs are less than the innovation cost difference. This behavior can be observed in the σ_S graph, where convergence time is clamped to 40 iterations.

5.2 Behavior Characteristics of the Trend-Following Decision Model

Trend-following decision making performed better than myopic decision making in virtually all cases in terms of both global efficiency and time until convergence. Figure 3 shows the results of three of the variables with graph scales and representation matching that of Figure 2. For the depicted results, we used a discount factor of $\gamma = 0.4$ and $\tilde{n}_i$ approximated by three derivatives. Variables b and a_0 are omitted for brevity, as the changes in their behavior due to the decision model closely resemble the changes of a_1.

The graph of a_1, which shows the varying overhead cost for the trend-following decision model, closely resembles the corresponding graph of the myopic decision model for the uniform switching cost distribution. The two benefits from the trend-following are that the range of non-ubiquitous adoption of either innovation is smaller, and the convergence time for smaller a_1 values is reduced. The Gaussian distribution shows how trend-following allows the agents to quickly push through costly transition regions to reach the global optimum of full innovation 1 adoption. Spending less time in these transition regions decreases the global cost.

The variables $n_1(0)$ and w_S show similar behavior to that of a_1.

For higher values of w_S and σ_S, trend-following yields a more efficient global optimum in terms of cost, but does not converge as quickly as myopic decision making. The convergence times sometimes exceeded the graph, oscillating indefinitely.

With agents considering their future costs beyond the next time-step, the switching cost becomes amortized. Switching cost amortization explains why the graphs are scaled on the horizontal axis; a decreased switching cost allows innovation adoptions for larger regions of the parameters. We found that the varying the discount factor, γ, scales the graph along the horizontal axis.

Amortizing switching costs in the myopic decision model often produces results that beat the trend-following model in terms of global efficiency, but take significantly longer to converge. While amortization effectively lowers the relative switching cost, trend-following accelerates the adoption of an innovation. However, the trend-following decision model may also determine that waiting is least costly and will reevaluate whether an innovation should be adopted.

We also found that varying the degree of the Taylor series expansion had little effect on the results. This was to be expected in most cases, as a third derivative approximation requires at least three time steps, so higher derivatives would tend to be unutilized. The benefit of using the approximation with one term offered significant benefit over a static model. While the second and third terms only changed some behavior in higher values of variance, w_S and σ_S, we did not find conclusive evidence that one performed better than the other.

6. PERTURBING THE GAME

An influencer overseeing the game may wish to have more or all agents adopt innovation 1. If the switching costs are too high, agents may not switch over to innovation 1 without subsidy, even if the innovation is cheaper when widely adopted. Such behavior can also exist if not enough agents have initially adopted innovation 1, as a critical mass is required to make the adoption profitable. Even if it is possible for the influencer to force all agents to adopt innovation 1, the agents may be unwilling to incur all the costs themselves.

The quickest solution is to subsidize all of the differences between the costs of innovation 0 and innovation 1 for all agents in innovation 0 at the same time. However, this incurs the largest expense. Given each agent p's switching cost, S_p, the subsidy required for all agents to switch to innovation 1 within 1 unit of time, $\eta_{t=1}$, is

$$\eta_{t=1} = \sum_{p=n_1}^{N} c_1(n_1+1) - c_0(N-n_1) + S_p. \qquad (27)$$

If an influencer wishes to spend as little as possible to switch all agents to innovation 1, but does not care how long it takes, subsidy cost may be minimized. With this method, only the agent closest to switching must be subsidized, and only if the agent needs a subsidy. The influencer thus leverages the cost reduction of the game mechanics. If all S_p for agents using innovation 0 at the start of subsidy are ranked ascending with value (with agents already using innovation 1 at the low end), the minimum subsidy cost to switch all agents to innovation 1, η_{min}, can be written as

$$\eta_{min} = \sum_{p=n_1}^{N} \max\{c_1(p+1) - c_0(N-p) + S_p, 0\}. \qquad (28)$$

An influencer may wish to compromise between minimal time and minimal subsidy. These compromises may be met by choosing to increment p by more than one for each iteration in equation (28).

An influencer will likely base the subsidy value from the societal cost of using a mix of innovations instead of purely using innovation 1. Societal cost per unit time, Ω, is expressed as

$$\Omega = n_1 \cdot c_1(n_1) + (N-n_1) \cdot c_0(N-n_1) - N \cdot c_1(N). \qquad (29)$$

From societal cost, we can see that cost will often be maximized at some point between full adoption of either innovation. This is important in terms of policy, since slow transitions can be more costly than fast ones, and because agents may be stuck in a non-optimal equilibrium.

If the game consists of agents using the trend-following decision model, an influencer can also change the course of the game by manipulating the perceived trend derivatives. This method is applicable in environments where the derivatives are estimated or computed in a distributed manner.

7. CONCLUSIONS

The game of scale presented in this paper is a model of dynamics of economy-of-scale decisions, such as innovation adoption. The model is of low input dimensionality while preserving the core behavior, and is easily extendible. We explore four different basic decision models. The most interesting behavior occurs around inputs where innovation adoption occurs slowly. Critical input values cause sharp spikes in the convergence times, and in other ways resemble phase changes. When too few agents adopt an innovation, the new innovation may either be discarded, or the agents will remain in equilibrium. Trend-following decision making can push through these local optima equilibria to obtain a better global solution. Once a critical mass of an innovation has been reached, the innovation may be widely and rapidly adopted. Adopting policies that expedite transition to new cost effective innovations is beneficial to a system in terms of societal cost.

Obtaining better models for economy-of-scale decisions is important in several regards. Applications include forecasting a technology's market penetration for planning or investment purposes, predicting the rises and falls of consumer trends, and engineering a multi-agent system with self-interested agents to make effective global agreements. Governments may wish to encourage adoption of new technologies despite myopic views of the affected industries, just as the manager of a deployed multi-agent system may wish to reward agents for incurring switching costs when switching to a better global optimum. Effectively applying these tools requires understanding their effects on adoption. The mechanism we present allows some global control over systems with decentralized decisions based on self-interested agents' individual utilities.

8. REFERENCES

[1] R. Baptista. The diffusion of process innovations: A selective review. *International Journal of the Economics of Business*, 6(1):107–129, 1999.

[2] G. Battisti and P. L. Stoneman. Inter- and intra-firm effects in the diffusion of new process technology. *Research Policy*, 32(9):1641–1655, 2003.

[3] G. Battisti and P. L. Stoneman. The intra-firm diffusion of new process technologies. *International Journal of Industrial Organization*, 23(1-2):1–22, 2005.

[4] L. Blumrosen and S. Dobzinski. Welfare maximization in congestion games. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 52 – 61, Ann Arbor, Michigan, 2006.

[5] F. Caselli and W. J. Coleman. On the theory of ethnic conflict. The London School of Economics and Political Science, March 2006.

[6] D. Challet, M. Marsili, and Y.-C. Zhang. *Minority Games: Interacting Agents in Financial Markets*. Oxford University Press, 2005.

[7] P. Dindo. A tractable evolutionary model for the minority game with asymmetric payoffs. *Physica A*, 355(1):110–118, 2005.

[8] Y. Farzin, K. Huisman, and P. Kort. Optimal timing of technology adoption. *Journal of Economic Dynamics and Control*, 22:779–799, 1998.

[9] R. G. Fichman. Information technology diffusion: A review of empirical research. In *Proceedings of the Thirteenth International Conference on Information Systems*, pages 195 – 206. University of Minnesota, 1992.

[10] W. Hlawitschka. The empirical nature of taylor-series approximations to expected utility. *American Economic Review*, 84(3):713–719, June 1994.

[11] N. Ireland and P. L. Stoneman. Technological diffusion, expectations and welfare. *Oxford Economic Papers*, 38(2):283–304, 1986.

[12] H. Konishi, M. L. Breton, and S. Weber. Pure strategy nash equilibrium in a group formation game with positive externalities. *Games and Economic Behavior*, 21(1-2):161–182, October 1997.

[13] C. H. Loch and B. A. Huberman. A punctuated-equilibrium model of technology diffusion. *Management Science*, 45(2):160–177, February 1999.

[14] J. F. Reinganum. On the diffusion of new technology: A game theoretic approach. *The Review of Economic Studies*, 48(3):395–405, 1981.

[15] J. Sarkar. Technological diffusion: Alternative theories and historical evidence. *Journal of Economic Surveys*, 12(2):131–176, 1998.

[16] G. Silverberg, G. Dosi, and L. Orsenigo. Innovation, diversity and diffusion: A self-organisation model. *The Economic Journal*, 98(393):1032–1054, 1988.

[17] P. L. Stoneman. Technological diffusion: The viewpoint of economic theory. *Richerche Economiche*, 40:585–606, 1986.

[18] P. L. Stoneman and P. A. David. Adoption subsidies vs information provision as instruments of technology policy. *The Economic Journal*, 96(380a):142–150, 1986.