
Health of Electronic Communities: An Evolutionary Game Approach

XIANJUN GENG, ANDREW B. WHINSTON, AND HAN ZHANG

XIANJUN GENG is an Assistant Professor in Information Systems at the University of Washington. He received his Ph.D. in information systems from the University of Texas at Austin. His current research interests include pricing and competition issues in IT-enabled new markets, trust in electronic communities, and corporate governance in new technology firms.

ANDREW B. WHINSTON is the Hugh Roy Cullen Centennial Chair Professor in Information Systems at the McCombs School of Business at the University of Texas at Austin. He is a professor in the Departments of Economics and Computer Science as well, and Fellow of the IC² Institute, Austin, and the director of the Center for Research in Electronic Commerce, for several years a pioneering research facility in Electronic Commerce. He received his Ph.D. in management from Carnegie Mellon University. His current research spans various realms of electronic commerce, its effect on business protocols and processes, organizational structure and corporate networks, electronic publishing, electronic education, complementarity of convergent computational paradigms, and business value of IT. He is the coauthor of several books on electronic commerce, prominent among them being the *Frontiers of Electronic Commerce*, *Economics of Electronic Commerce*, and *Electronic Commerce—A Manager's Guide*. He has also published over 300 papers in leading academic journals in economics, business, and computer science.

HAN ZHANG is an Assistant Professor of Information Technology Management at the College of Management, Georgia Institute of Technology. He received his Ph.D. in information systems from the University of Texas at Austin. His research focuses on online trust and reputation related issues, trusted third parties or intermediaries, and the evolution of electronic markets. He currently serves on the editorial review board of *Journal of Database Management*.

ABSTRACT: Creating electronic communities is a critical venture in the digital economy. However, fraud and misrepresentation have led to widespread skepticism and distrust of electronic communities. We develop an evolutionary model to explore the issue of trust within an electronic community from a dynamic process perspective. This model emphasizes large populations, continuous change in community memberships, and imperfect information and memory. As the term *trust* is often used in the context of individual interaction, at a group level we propose using the term *health* to measure the sustained competitive advantages of honest members over cheaters throughout the evolution of a community. We find conditions under which an electronic community is healthy and attracts outside population. We find that many factors, such as information dissemination speed, honest players' payoffs and possible losses, new community members' initial trust status, and the replacement rate of community members, all affect the health of an electronic community, and that some of them also

affect a community's size. We then discuss the implications of our research for e-community practices.

KEY WORDS AND PHRASES: dynamic process, electronic community, evolutionary game theory, trust in e-commerce, trust status.

IN A MERE DECADE, COMMERCIAL ACTIVITIES on the World Wide Web have led to the proliferation of electronic communities (e-communities). E-communities have increasingly intrigued both the popular press and academic researchers. Building on the research of Chang et al. [6], we define an e-community as a social aggregation of people whose members interact principally via electronic communication channels. E-communities formed by those bidding on eBay, ordering books from Amazon.com, or chatting on various online forums have become a part of daily life for an increasing number of Internet users.

Increasingly, fraudulent activities have encroached on e-community life, and thus have threatened the bonds of trust among community members [14]. For example, according to the Internet Fraud Watch, operated by the National Consumers League, fraudulent online auction sales remained the number one type of Internet fraud in 2002 [10]. As a result, promoting trust in electronic marketplaces has become an important goal for e-commerce [4, 11, 14, 18].

Understandably, the notion that trust is essential to social interaction and social order spans several research disciplines. The current literature related to the study of information systems tends to focus on trust in long-term relationships among individuals within a small virtual group or an organization [11] as well as on consumers' trust in online storefronts based on laboratory experiments [12, 20]. Among research efforts concentrating on e-communities, Kollock [14] conceptually explores the emergence of endogenous solutions to the problems of risky trade in e-communities, and Resnick et al. [18] review the current challenges of Internet-based reputation systems.

Prior research, however, has largely ignored three crucial aspects of trust stemming from the unique dynamics of e-communities. First, due to the global availability and pervasive nature of the Internet, e-communities have attracted large numbers of participants. Consequently, economic incentives, rather than personal relations, play a larger role in determining individual behavior in an e-community because of the small probability of rematch under random matching [3]. Second, e-communities are characterized by the continuous inflow and outflow of members. Such freedom for individuals to enter or exit the community facilitates the perpetration of fraud, because cheaters can easily leave the community to avoid punishment. Third, in order to build trust among members, e-communities should have a mechanism to record members' identities and later punish members who cheat [1, 4, 13, 17]. However, in large-scale e-communities, it is impossible for a member to recall the histories of all other mem-

bers. This inevitably leads to imperfect memory of past cheating behavior, and as a result facilitates deceit.

These three dynamic aspects of an e-community—*large populations, continuous change in community memberships, and imperfect information and memory*—call for a new research framework that is based on an evolutionary, rather than a static, perspective. To build such a research framework, we adopt an evolutionary game approach. As the term *trust* emphasizes individual interactions, in this paper we introduce the term *health*, which emphasizes the evolution of an entire e-community. In this dynamic model, a healthy e-community has the following two characteristics. First, from a static perspective, as the result of the evolutionary process, an e-community is composed of honest members who trust each other. Second, from a dynamic perspective, in the evolutionary process, honest members have sustained competitive advantages over cheating members in the e-community.

Evolutionary game theory, as pioneered by Maynard Smith and Price [16], proposes the concept of an evolutionarily stable strategy (ESS). It has been used extensively by biologists (e.g., [8, 15, 21]), economists (e.g., [23]), and social scientists (e.g., [2]). Young [23], for example, studies individual strategy and social structure and proposes an evolutionary theory of social and economic institutions. Axelrod [2] explores the evolution of cooperation by providing computerized tournaments of the repeated prisoner's dilemma, and proposes a valuable framework to analyze the conditions promoting cooperative play. Generally, evolutionary game theory offers a novel approach to the classical game theory, and can be applied to the study of social and economic evolution. The current research is one of the first efforts to introduce an evolutionary model to understand trust in e-communities.

The Evolutionary Model

OUR EVOLUTIONARY MODEL OF E-COMMUNITIES consists of three major components: a *payoff matrix*, which describes payoff functions in a single period; *trust status*, which indicates the extent to which e-community members trust each other in a single period; and *dynamic processes of evolution*, which define how new players replace old ones and how new players choose their strategies.

The Basic Setup and Payoff Matrix

We begin with an e-community containing two subsets of players: subset *H*, consisting of honest players, and subset *C*, consisting of cheating players. There is also a subset *Q*, consisting of players outside the e-community who pursue other activities. The possibility for players inside the e-community (in subset *H* or *C*) to move out and pursue other activities (in subset *Q*) is called an outside option. Our model permits outside options, which is in agreement with the fact that Internet-based e-communities allow the free flow of members.

For simplicity (and for tractability of results), we treat the population of each subset as a continuous variable. A fundamental assumption of this research is that each player's

strategy, either to be honest or to cheat, is exogenously given. That is to say that a player never changes her strategy throughout her lifetime. The assumption that players' actions are exogenously given was first used by Maynard Smith and Price [16], and later followed by many scholars such as Binmore [5], Hofbauer and Sigmund [9], Taylor and Jonker [21], and Van Damme [22].

The behavioral foundation of this basic evolutionary game assumption derives from the observation that players are limited to bounded rationality and a predilection toward inertia and sticking to certain habits [16]. Furthermore, this assumption of bounded rationality on the individual level is not very restrictive, since, for the entire population, shifting strategies is still possible given another fundamental assumption: limited lifetime. Our model assumes that each player has a limited lifetime, and on her death will be replaced by a newborn player, with the size of the whole population remaining unchanged. This newborn player will not necessarily inherit the dead player's strategy, as we will discuss shortly. Consequently, the percentage of players in the whole population who use a certain strategy can change. If this percentage shrinks to almost zero, it reflects that almost all members of this e-community reject this specific strategy.

Furthermore, this bounded rationality assumption can still be used to model players who can change their strategies, albeit with limitations. A player who switches her strategy from Strategy 1 to Strategy 2 can be viewed, analytically, as two sequential players: the first, using Strategy 1, dies and a second player, using Strategy 2, is born.¹

In each time period, each player in this e-community randomly pairs up with another player to play the game, which can be, for example, a business transaction or information exchange. This symmetric model emulates e-communities such as online auctions, in which a player could be a buyer in one period and a seller in another period. Consider two players, 1 and 2, who are randomly drawn at period k to play the game. The payoff matrix is given in Table 1.

In the matrix of Table 1, a , b , and d are all positive. a is the normal payoff to each player in one match given no cheating behavior. d is the extra payoff to any player who cheats an honest player. b is an honest player's loss if she is cheated. Examples of cheating in an online auction could be providing damaged or misrepresented goods or refusing to pay for goods received. If both players cheat, we normalize their payoffs to zero, since the evolution depends only on the relative payoffs. Here we assume that cheating harms social welfare—that is, $2a > a + d - b$, or equivalently, $a > d - b$.

Trust Status and Payoff Functions

Generally, a player may participate in an e-community for multiple periods. Our model accommodates this by assigning age t as the number of periods, except the current period, that a player has participated in the e-community, or "lived." By excluding the current period, we know that each newborn player will have an age of 0 when she is born.

If a cheating player is not punished in the e-community, cheating behavior will sustain a competitive advantage over honest behavior—in Table 1, the cheating strat-

Table 1. The Payoff Matrix

		Player 2	
		Honest	Cheat
Player 1	Honest	a, a	$-b, a + d$
	Cheat	$a + d, -b$	$0, 0$

egy dominates honest behavior, and "cheat, cheat" is the unique Nash equilibrium for a single-period model. This situation will lead to the collapse of an e-community as cheaters predominate. Therefore, a mechanism for recording, disseminating, and punishing cheating behaviors must exist in any healthy e-community to counteract the natural advantage of cheating strategies over honest strategies. We define the following punishment mechanism: for a player (e.g., "Alice") with trust status α , at any period k , her counterpart will attempt to do business with her with probability α . In other words, with expected probability $1 - \alpha$, her counterpart will refuse to do business with her. This refusal is her punishment. Since trust status is a probability concept, we assume that it always takes values in $(0, 1)$.

Central to this punishment mechanism is the player's trust status. There are two ways to explain trust status α . First, trust status can be explained as the number of players in the e-community who know Alice's type, whereby such knowledge evolves over time. If Alice is honest, more and more players will know this from her history, and her trust status will improve over time.

Trust status is also influenced by reputation; that is, although other players may not know exactly whether Alice is honest or dishonest, normally they will feel more comfortable doing business with her if she has a positive reputation. However, we are hesitant to use the word reputation, as it often implies that the whole e-community is in a consensus regarding the extent to which Alice should be trusted. In cases where there is strong heterogeneity in information dissemination in the e-community, this is often not true.

If a player cheats in a period, information about her fraud is recorded (either by the system or by word of mouth) and disseminated throughout the e-community, which eventually leads to a reduction in her trust status. We model the information dissemination process in the following way. Suppose a cheating player enters the e-community at period k_0 , and her initial trust status is α_0 . An initial trust status reflects how members of the e-community treat newcomers. The larger the initial trust status, the more receptive the e-community is to strangers. After period k_0 , her cheating behavior will be recorded in the e-community and disseminated, and her trust status will decrease. At period k , $k \geq k_0$, her age is $t = k - k_0$. We define

$$\alpha_c(k, t) = \alpha_0 r_c^t, \quad t = 0, 1, 2, \dots, \quad (1)$$

which is this cheater's trust status at period k when her age is t . r_c is the status decreasing factor, $r_c \in (0, 1)$. We can also interpret $1 - r_c$ as "the information dissemination




Figure 1. The dynamics of $\alpha_c(k, t)$ and $\alpha_h(k, t)$

speed for cheating behavior," since, for larger values of $1 - r_c$, players in the e-community come to know more quickly the true character of a new player. Figure 1 shows the dynamics of $\alpha_c(k, t)$ (and of $\alpha_h(k, t)$, which we will define shortly). Notice that $\alpha_c(k, t)$ does not depend on k ; that is, the e-community has a status evolving rule that does not change over time.

The changing dynamics of the trust status of $\alpha_c(k, t)$ is an exponential process, which has the steepest slope at the beginning of the player's life. This means that her status is more sensitive to her cheating behavior in the early rounds. The behavioral explanation is that players treat information from different periods equally. At the start of her life, since there is little or no history about her, other players will pay more attention to her current behavior. Later, however, when she has a longer track record in the e-community, the weight of her current behavior relative to her whole history becomes smaller. Therefore, the magnitude of the change in her trust status decreases with time.²

Similarly, an honest player's honest behavior will also be recorded and disseminated. Thus, if the honest player enters the e-community at period k_0 , then at period k

$$\alpha_h(k, t) = 1 - (1 - \alpha_0) r_h^t, \quad t = 0, 1, 2, \dots, \quad (2)$$

where $t = k - k_0$. r_h is the trust status increasing factor, $r_h \in (0, 1)$.

At period k , let the proportion of honest players of age t ($t \leq k$) in the whole population be $x_h(k, t)$, and the proportion of cheating players of age t ($t \leq k$) be $x_c(k, t)$. We assume a large enough population N so that we can view $x_h(k, t)$ and $x_c(k, t)$ as continuous variables. Let $x_h(k) = \sum_{t=0}^k x_h(k, t)$, $x_c(k) = \sum_{t=0}^k x_c(k, t)$, which are the proportions of honest players and cheating players of all ages, respectively. Notice that $x_h(k) + x_c(k) < 1$ occurs if there are options outside of the e-community; that is, players can choose to leave the e-community.

In period k , an honest player of age t_0 has a chance of $(x_h(k, t))/(x_h(k) + x_c(k))$ to meet another honest player of age t . This honest player also has a chance of $(x_c(k, t)) /$

$(x_h(k) + x_c(k))$ to encounter a cheating player, therefore, this honest player's payoff in period k is

$$u_h(k, t_0) = \sum_{t=0}^k \frac{x_h(k, t)}{x_h(k) + x_c(k)} \alpha_h(k, t) \alpha_h(k, t_0) a - \sum_{t=0}^k \frac{x_c(k, t)}{x_h(k) + x_c(k)} \alpha_c(k, t) \alpha_h(k, t_0) b. \quad (3)$$

Similarly, in period k , a cheating player of age t_0 has a payoff of

$$u_c(k, t_0) = \sum_{t=0}^k \frac{x_h(k, t)}{x_h(k) + x_c(k)} \alpha_h(k, t) \alpha_c(k, t_0) (a + d). \quad (4)$$

The Dynamic Processes of Evolution

Suppose that on entering any period k , all players in the population have equal probability, $1/T$, to die. A newborn player immediately replaces each dead player; thus, the total population size remains constant. Given a large population, this implies that on entering at each time k , $1/T$ of the whole population dies and is replaced by $1/T$ new players. Note that since all players have the same probability to die no matter their age, the life span of a player is a random integer variable.

On one hand, when a player in or out of the e-community dies, the newborn player that replaces her will not necessarily inherit her strategy. This leads the e-community to change between periods. On the other hand, there should be appropriate constraints to a newborn player's choice to reflect that an e-community develops according to an evolutionary process instead of simple chaos. These considerations lead to three guidelines for modeling the evolution process, as stated below. Note that these guidelines are widely accepted in evolutionary game theory and are applied in multiple academic disciplines including biology, social science, and economics [2, 5, 8, 16, 21, 23]. First, although a newborn player may switch to strategies other than the one used by her predecessor, she should have a certain degree of inertia in switching away. This is equivalent to saying that the evolutionary process should be gradual instead of radical. Second, the tenets of bounded rationality and incomplete memory imply that a newborn player cannot choose her strategy based on the complete history of the e-community and the outside population. For simplicity, we consider the extreme case where, at period $k + 1$, a newborn player makes her decision based only on the information available to her about period k —that is, newborn players are myopic. This assumption is not very restrictive, as the trust status at period k for any existing player incorporates the history of this player. Third, if a newborn player switches, it should be assumed that the more profitable a strategy is in period k , the more likely she will choose that strategy in period $k + 1$.

Given the above constraints, we use a modified version of replicator dynamics of overlapping generations [5, 21] to model the evolutionary dynamics of the

e-community, as shown in Equations (5) and (6). Replicator dynamics is a central component of evolutionary game theory that provides mathematical formulations for modeling evolution in an economy.

$$x_h(k+1, 0) = \sum_{t=0}^k \left(\frac{1}{T} \right) x_h(k, t) (1 + u_h(k, t) - \underline{u}(k)) \quad (5)$$

$$x_c(k+1, 0) = \sum_{t=0}^k \left(\frac{1}{T} \right) x_c(k, t) (1 + u_c(k, t) - \underline{u}(k)). \quad (6)$$

Here $\underline{u}(k)$ is the average payoff of the whole population in period k , which includes players both in and out of the e-community. In the following analysis, we will omit the dynamics for players outside of the e-community, since this aspect is easy to derive given the constant total population. Notice that all the payoff functions are constrained by the fact that in Equations (5) and (6), $x_h(k+1, 0)$ and $x_c(k+1, 0)$ must always be nonnegative. This can be satisfied, for example, if a , b , d , and the outside option's payoff are all positive numbers that are smaller than 1.

The evolutionary processes we describe in Equations (5) and (6) are the simplest ones that can incorporate the above three constraints. In Equation (5) (and similarly in Equation (6)), only information about period k affects the dynamics from period k to $k+1$. The term $(1/T)x_h(k, t)$ shows that newborn players have a tendency to stick to their predecessors' strategy, and the term $(1 + u_h(k, t) - \underline{u}(k))$ shows that newborn players favor high-payoff strategies from the last period.

It is worthwhile to note that Equations (5) and (6) could have a broad behavioral foundation. We have used the phrase "limited memory" to refer to the case where players have little memory of prior periods other than the last one, as we just argued. We can also argue that some players do not have complete memory—or in other words they have a biased memory—of the outcome of the last period. As long as the whole population, as a group, has unbiased, aggregated memory, Equations (5) and (6) still follow. Intuitively, one newborn player may overfavor the honest strategy if she overestimates the value of $u_h(k, t)$, whereas another newborn may underfavor the honest strategy if she underestimates this value. If their biases are of similar magnitude, they cancel out each other in Equation (5), since the right side of Equation (5) is a linear function of $u_h(k, t)$.

Players who have survived after period k simply enter period $k+1$ with their age and trust status changed:

$$x_h(k+1, t+1) = \left(\frac{T-1}{T} \right) x_h(k, t), \quad 0 \leq t \leq k \quad (7)$$

$$x_c(k+1, t+1) = \left(\frac{T-1}{T} \right) x_c(k, t), \quad 0 \leq t \leq k. \quad (8)$$

Equations (1) through (8) jointly define how the e-community evolves.

Health and Size Issues in the e-Community

THE MODEL PROPOSED IN THE LAST SECTION incurs considerable computational complexities. Specifically, the $(2k+2)$ -tuple $(x_h(k, 0), x_h(k, 1), \dots, x_h(k, k), x_c(k, 0), x_c(k, 1), \dots, x_c(k, k))$ is the population composition of the e-community at time k . Analyzing the evolution of the e-community directly using this $(2k+2)$ -tuple is difficult, however, because its size increases as k increases. In this section, we first reduce this tuple to a tractable size by focusing on average trust status and average payoffs. It turns out that these are the sufficient statistics for describing the evolution of the e-community as a whole, as we show in Lemma 1.

In this section, we also propose *health* as a measure of the level of trust among members in e-communities, which naturally reconciles two different objectives of research on trust for communities—(1) give honest players a sustained competitive advantage over cheating players, and (2) to realize an e-community where almost all are honest players. Finally, we examine the option of remaining outside of the e-community.

Model Reduction

Central to model reduction is the definition of *average trust status* and *average payoff*, which omit individuals' information but still capture the characteristics of the community as a whole. Let $\underline{\alpha}_h(k) = \sum_{t=0}^k ((x_h(k, t))/(x_h(k)))\alpha_h(k, t)$, which is the *average trust status* of the honest player population in the e-community at period k . Similarly, $\underline{\alpha}_c(k) = \sum_{t=0}^k ((x_c(k, t))/(x_c(k)))\alpha_c(k, t)$ is the average trust status of the cheating player population in the e-community. Let $\underline{u}_h(k) = \sum_{t=0}^k ((x_h(k, t))/(x_h(k)))u_h(k, t)$, which is the *average payoff* for the honest player population in the e-community at period k . Similarly, let $\underline{u}_c(k) = \sum_{t=0}^k ((x_c(k, t))/(x_c(k)))u_c(k, t)$, which is the average payoff for the cheating player population in the e-community. Based on these average values, we present the evolutionary processes using six discrete time equations.

Lemma 1: Given any starting period k_0 and $x_h(k_0)$, $x_c(k_0)$, $\underline{\alpha}_h(k_0)$, $\underline{\alpha}_c(k_0)$, and $\underline{u}(k)$ for any $k \geq k_0$, the evolution of the e-community is uniquely determined by the following set of six discrete time equations.

$$x_h(k+1) = \left(\frac{1}{T}\right) x_h(k) (T + \underline{u}_h(k) - \underline{u}(k)) \quad (9)$$

$$x_c(k+1) = \left(\frac{1}{T}\right) x_c(k) (T + \underline{u}_c(k) - \underline{u}(k)) \quad (10)$$

$$\underline{u}_h(k) = \underline{\alpha}_h(k) \underline{\alpha}_h(k) \frac{x_h(k)}{x_h(k) + x_c(k)} a - \underline{\alpha}_h(k) \underline{\alpha}_c(k) \frac{x_c(k)}{x_h(k) + x_c(k)} b \quad (11)$$

$$\underline{u}_c(k) = \underline{\alpha}_h(k) \underline{\alpha}_c(k) \frac{x_h(k)}{x_h(k) + x_c(k)} (a + d) \quad (12)$$

$$\underline{\alpha}_h(k+1) = \frac{1 + \underline{u}_h(k) - \underline{u}(k)}{T + \underline{u}_h(k) - \underline{u}(k)} \alpha_0 + \frac{T-1}{T + \underline{u}_h(k) - \underline{u}(k)} (1 - r_h + r_h \underline{\alpha}_h(k)) \quad (13)$$

$$\underline{\alpha}_c(k+1) = \frac{1 + \underline{u}_c(k) - \underline{u}(k)}{T + \underline{u}_c(k) - \underline{u}(k)} \alpha_0 + \frac{T-1}{T + \underline{u}_c(k) - \underline{u}(k)} r_c \underline{\alpha}_c(k). \quad (14)$$

In other words, given any starting period k_0 , the dynamics of the evolution of an *e*-community are uniquely determined by the proportion and the average trust status of the honest player population and the cheating player population, respectively, at period k_0 and the average payoff of the outside population at any period.

All proofs are in the Appendix. Based on Lemma 1, it is sufficient for us to study the reduced model composed of Equations (9) through (14). It is useful to point out that in this reduced model, $x_h(k)$, $x_c(k)$, $\underline{\alpha}_h(k)$, and $\underline{\alpha}_c(k)$ are endogenous characteristics of the *e*-community, whereas $\underline{u}(k)$ includes information about outside options.

The Definition of Health: Sustained Competitive Advantage Versus Population Dominance

To study the general case, at period k_0 let both $x_h(k_0)$ and $x_c(k_0)$ be nonzero—that is, both honest players and cheating players exist in the *e*-community at the start period. Notice that the choice of the start period is arbitrary—it can be any period in the evolution.

We define a healthy *e*-community to be an *e*-community in which honest players have sustained competitive advantages over cheating players. More precisely,

Definition 1: An e-community is healthy at period k_1 if and only if for any $k \geq k_1$, $\underline{u}_h(k) > \underline{u}_c(k)$, and $\lim_{k \rightarrow \infty} (\underline{u}_h(k) - \underline{u}_c(k)) > 0$ if this limit exists.

From this we immediately know that if an *e*-community is healthy at period k_1 , it is healthy at any period $k \geq k_1$. Thus our definition emphasizes that the dynamics of *e*-community evolution are healthy, rather than that the result of the evolution is healthy.

Another plausible way to define health is from a population dominance perspective—that is, honest players are the dominating majority in the *e*-community. Before characterizing this second perspective, we first give a definition of the Evolutionarily Stable Strategy (ESS) from the asymptotic perspective. Simply speaking, an ESS is a Nash equilibrium satisfying an additional stability property; that is, an ESS should be able to withstand the pressures of mutation and selection once it becomes established in a population [15, 16].

Definition 2: Given any constant nonnegative population proportion x_{h0} and x_{c0} , $x_{h0} + x_{c0} \leq 1$. (x_{h0}, x_{c0}) is an ESS if and only if given any $\varepsilon > 0$, there exists $k_1 > k_0$ such that for any $k \geq k_1$, $|x_h(k) - x_{h0}| < \varepsilon$ and $|x_c(k) - x_{c0}| < \varepsilon$.

Then from the population dominance perspective, we can also define a healthy e-community to be an e-community in which $(x_{h0}, 0)$ is the realized ESS, where $x_{h0} \geq 0$. Notice that we need the asymptotic property of ESS because from Equation (8) we know that, even if the number of cheating players keeps decreasing, at any period $k > k_0$ it will still be positive.

It is natural to argue that in any reasonable model, both definitions of health should be consistent with each other. Formally, the following proposition shows the relationship between these two perspectives of a healthy e-community.

Proposition 1: If the e-community is healthy according to Definition 1, then $(x_{h0}, 0)$ is the realized ESS, where $x_{h0} \geq 0$.

From Proposition 1, we know that Definition 1 is stronger than the ESS version. Therefore, in this paper, we adopt Definition 1 to define the health of an e-community. We call $(x_c(k))/(x_h(k))$ the *unhealthy index* of the e-community at period k . A healthy e-community is characterized by a decreasing unhealthy index with a limit of 0.

Size of the e-Community

In addition to health, the size of the entire population inside the e-community is also important in analyzing the evolution of an e-community. Given $(x_c(k))/(x_h(k)) \rightarrow 0$, honest players will eventually dominate the e-community. Nevertheless, it is possible that at the same time $\lim_{k \rightarrow \infty} x_h(k) = 0$ —that is, both the honest player population and the cheating player population are diminishing, which makes “healthy” valueless. Intuitively, this means that although, inside the e-community, behaving honestly is better than cheating, it is still worse than outside options, and thus players will go away. We will come to this point again when we analyze outside options in the section “Size of the e-Community and Outside Options.”

A Health Condition for the e-Community

THE EMPHASIS ON SUSTAINED COMPETITIVE ADVANTAGES for honest players in the definition of health fits well into our view that the trust issue of e-communities should be studied in an evolutionary perspective. Nevertheless, this definition of health creates measurability issues: at any given period, it is impossible to directly measure the health of an e-community, since health depends on the unknown future.

In this section, we derive a condition that enables us to measure the health of an e-community after the e-community evolves for only limited periods, thus solving the measurability issue. We then discuss the implications of this health condition.

The Health Condition

The process for assessing the health condition of our e-community begins by determining the fixed health condition—that is, the condition under which $(x_c(k))/(x_h(k))$ does not change between two consecutive periods. For convenience, let $e_h(k) = 1 + \underline{u}_h(k) - \underline{u}(k)$, and $e_c(k) = 1 + \underline{u}_c(k) - \underline{u}(k)$.

$$\frac{x_c(k+1)}{x_h(k+1)} = \frac{x_c(k)}{x_h(k)} \Leftrightarrow e_c(k) = e_h(k) \Leftrightarrow \frac{x_c(k)}{x_h(k)} = \frac{1}{b} \left[a \frac{\underline{\alpha}_h(k)}{\underline{\alpha}_c(k)} - a - d \right]. \quad (15)$$

The notation “ \Leftrightarrow ” means the two equations are mathematically equivalent. Equation (15) is the fixed health condition. We should note that Equation (15) does not always have solutions. If $[a(\underline{\alpha}_h(k))/(\underline{\alpha}_c(k)) - a - d] < 0$, it has no solution. Furthermore, under such conditions the unhealthy index is increasing. The intuitive explanation is that if the extra payoff to the cheating player, d , is too large, then evolution tends to favor cheating behavior more than honest behavior.

The fixed health condition implies that

$$\frac{x_c(k+1)}{x_h(k+1)} < \frac{x_c(k)}{x_h(k)} \Leftrightarrow \frac{x_c(k)}{x_h(k)} < \frac{1}{b} \left[a \frac{\underline{\alpha}_h(k)}{\underline{\alpha}_c(k)} - a - d \right]. \quad (16)$$

Thus, given any k_1 , if inequality (16) holds for all $k \geq k_1$, then at period k_1 we can say that the e-community is healthy.

However, at any arbitrary period k_1 , whether inequality (16) will be satisfied for all future periods is hard to determine. This is because the satisfaction of inequality (16) at period k_1 does not guarantee that at any period beyond k_1 inequality (16) is still satisfied. This is illustrated in the following example.

Consider the simulation results in Figure 2. They are based on settings of $T = 4$, $r_h = 0.8$, $r_c = 0.8$, $a = 0.5$, $d = 0.15$, $b = 0.2$, $\alpha_0 = 0.4$, $\underline{x}_h(0) = 0.1$, and $\underline{x}_c(0) = 0.2$. Also in this example, we assume that the outside option gives a constant payoff of 0.15 each period. The simulation is from period 0 to period 500, as shown in Figure 2b. Figure 2a amplifies the first 30 periods to better display the start of the evolution. In each figure, the solid line stands for the evolution of unhealthy index $(x_c(k))/(x_h(k))$, and the dashed line stands for $(1/b)[a(\underline{\alpha}_h(k))/(\underline{\alpha}_c(k)) - a - d]$. Figure 2a shows that the unhealthy index has a short-term surge, and Figure 2b shows that it finally drops down to zero. Figure 2b also shows that $(1/b)[a(\underline{\alpha}_h(k))/(\underline{\alpha}_c(k)) - a - d]$, and therefore $(\underline{\alpha}_h(k))/(\underline{\alpha}_c(k))$, quickly surges and then slowly drops to a stable level.

One implication of this example is that an e-community may deteriorate in the short run, as the unhealthy index increases when $0 \leq k \leq 5$ (as shown in Figure 2a). However, in the long run the unhealthy index may still decrease to almost zero when $k > 5$ (as shown in Figures 2a and 2b). This observation again strengthens our view that the health of an e-community should be defined from an evolutionary perspective rather than a stationary or a short-term perspective.




Figure 2. An Example of Evolution

Remark: An increase in the proportion of cheating players in the e-community does not necessarily imply that the health of the e-community is deteriorating: if at the same time the difference between honest and cheating players' average trust status keeps increasing, honest players may actually be in the process of gaining a competitive advantage over cheating players.

The complex evolution process makes it impossible to identify a *static* condition that is both a necessary and sufficient condition for a *dynamically* defined healthy e-community. Nevertheless, it is possible to give a sufficient condition for a healthy e-community (which we shall soon present). This comes from the finding that, no matter how $(x_c(k))/x_h(k)$ evolves, $(\underline{\alpha}_h(k))/(\underline{\alpha}_c(k))$ will converge into a small region in a limited number of periods, as we will prove shortly. For example, in Figure 2, $(1/b)[a(\underline{\alpha}_h(k))/(\underline{\alpha}_c(k)) - a - d]$ converges into $[3, 3.5]$ when $k = 10$. As the ratio $(\underline{\alpha}_h(k))/(\underline{\alpha}_c(k))$ implies the effect of punishment on cheating behaviors, the convergence of $(\underline{\alpha}_h(k))/(\underline{\alpha}_c(k))$ implies that the punishment will converge to a relatively stable level. Therefore, the e-community is healthy if this relatively stable punishment is large enough to establish a competitive advantage for honest players.

Thus, to analyze the evolution of the e-community, we divide it into two eras. In era one (period $k_0 \leq k < k_1$), $\underline{\alpha}_h(k)$ and $\underline{\alpha}_c(k)$ converge to smaller ranges. In era two (period $k \geq k_1$), the boundaries of these small ranges constitute a lower bound of $(\underline{\alpha}_h(k))/(\underline{\alpha}_c(k))$, which provide a sufficient condition for the e-community to be healthy.

Era One: The Convergence of Average Trust Status (From Period k_0 to Period k_1)

From Equations (A.3) and (A.4) in the proof of Lemma 1, we know both $e_c(k)$ and $e_h(k)$ are nonnegative. Since a , b , and d are all bounded from above, $e_c(k)$ and $e_h(k)$ are

also bounded from above. Select two numbers e_{c0} and e_{h0} that are slightly larger than the upper bounds of $e_c(k)$ and $e_h(k)$, respectively; thus, $0 \leq e_c(k) < e_{c0}$, $0 \leq e_h(k) < e_{h0}$. Intuitively, e_{c0} and e_{h0} are the upper bounds of how much better off cheating and honest players can become in a single period. For convenience, we define

$$\gamma_h = \frac{1}{(T-1)(1-r_h)/e_{h0} + 1}, \quad \gamma_c = \frac{1}{(T-1)(1-r_c)/e_{c0} + 1}.$$

Suppose era one starts at period k_0 , and ends at period k_1 .

Lemma 2:

- i. If $\underline{\alpha}_c(k_0) < \gamma_c \alpha_0$, then for all $k \geq k_0$, $\underline{\alpha}_c(k) < \gamma_c \alpha_0$. If $\underline{\alpha}_c(k_0) \geq \gamma_c \alpha_0$, then there exists $k_1 > k_0$ such that for all $k \geq k_1$, $\underline{\alpha}_c(k) < \gamma_c \alpha_0$.
- ii. If $\underline{\alpha}_h(k_0) > 1 - \gamma_h(1 - \alpha_0)$, then for all $k \geq k_0$, $\underline{\alpha}_h(k) > 1 - \gamma_h(1 - \alpha_0)$. If $\underline{\alpha}_h(k_0) \leq 1 - \gamma_h(1 - \alpha_0)$, then there exists $k_1 > k_0$ such that for all $k \geq k_1$, $\underline{\alpha}_h(k) > 1 - \gamma_h(1 - \alpha_0)$.

Lemma 2 establishes that after long enough periods, $\underline{\alpha}_h(k)$ converges to a small range of relatively large average trust status values, whereas $\underline{\alpha}_c(k)$ converges to a small range of relatively small average trust status values. Explicitly, we define k_1 to be the period when the convergence is finished. Such convergence differentiates the honest player population from the cheating player population, showing that the e-community's punishment mechanism was effective in punishing the cheating players.

Era Two: A Sufficient Condition for a Healthy e-Community (From Period k_1 On)

In era two, we propose a sufficient condition in Proposition 2 under which the punishment is strong enough to maintain a healthy e-community, given the convergence of average trust statuses.

Proposition 2: At period k_1 , if

$$\frac{x_c(k_1)}{x_h(k_1)} \leq \frac{1}{b} \left[a \frac{1 - \gamma_h(1 - \alpha_0)}{\gamma_c \alpha_0} - a - d \right], \quad (17)$$

then the e-community is healthy.

Although Proposition 2 provides a guideline for measuring the health of the e-community, it is not a necessary condition. We have pointed out that a closed form solution of the necessary and sufficient condition is not available due to the complex evolution. Given this, Proposition 2 gives us the best available criterion for a healthy e-community. Equation (17) is also a robust condition—that is, small disturbances will not disrupt healthy outcomes. Hereafter in this paper, we refer to inequality (17) as the *health condition*.

Implications of the Health Condition

The left side of the health condition is the unhealthy index at period k , which evolves as k increases. The right side of the health condition is a static value that is independent of the evolution. Clearly, an e-community with a large $(1/b)[a(1 - \gamma_h(1 - \alpha_0))/(\gamma_c\alpha_0) - a - d]$ is more favorable in the sense that the e-community is healthy even if $(x_c(k_1))/(x_h(k_1))$ is large. If an e-community has a large $(x_c(k_1))/(x_h(k_1))$, while still maintaining health, we say that this e-community is of *robust health*. We define *robust health* as having the ability to endure the existence of a large proportion of cheating players inside the e-community while still maintaining a competitive advantage for honest players, so that in the long run the e-community can steadily eliminate the population of cheating players. Again, we limit this analysis to era two, so that we may assume that average trust statuses have converged.

From Proposition 2, conditional on $(1/b)[a(1 - \gamma_h(1 - \alpha_0))/(\gamma_c\alpha_0) - a - d]$ being positive, we have the following:

Corollary 1: The robustness of the health of an e-community is positively related to a , the payoff of an honest player when she plays with another honest player, negatively related to b , a cheated honest player's loss, and negatively related to d , a cheating player's extra profit gained by cheating an honest player.

Corollary 1 shows that an e-community's health is closely related to its endogenous payoff structure. For instance, in e-communities where products of high cost (and thus large b) but low profit (thus low a) are exchanged, it is difficult to maintain health. As another example, in mailing lists where people exchange field-specific information, cheating usually does not benefit the cheater, as the information sent out seldom triggers any direct benefit for the sender. As a result, mailing lists can often successfully establish themselves as an effective way for people with common interests to help each other.

This corollary can also explain the effect of transaction fees in an e-community. Suppose that the e-community charges each player a fixed fee for each transaction. This effectively decreases a and increases b , which results in a reduced robustness of e-community health. Intuitively, transaction fees affect honest players no matter whether they are doing normal business or get cheated, whereas they do not affect the extra payoff, d , that a cheater can get by cheating, although a cheater's overall payoff is reduced.

Corollary 2: The robustness of the health of an e-community is positively related to $1 - r_h$ and $1 - r_c$, the information dissemination speed inside the e-community.

If a player cheats in an e-community, the ideal situation is that all other community members immediately know who the cheater is. That is, the speed of information dissemination is so rapid that the members of the community are well informed and thus know whom to punish for cheating. This condition may be satisfied in small-scale virtual communities. However, it is more problematic in large communities, such as a global e-community that crosses geographical or national borders. On eBay,

for example, most transactions take place in the e-community where players change trading partners often and may never face the same partner again. Unfortunately, not every cheated player leaves negative comments about her trading partner in eBay's Feedback System—people may be afraid of retaliatory negative feedback, or they merely want to avoid further unpleasant interactions [18]. Moreover, not every player checks her trading partner's feedback rating before making a transaction. In such an e-community, information is incomplete. Nevertheless, Corollary 2 points out that the more information players can get, the better it helps honest players, thus partial information is also valuable. Thus, in e-communities, gathering better information on a trading partner's past behavior is very important. This result is also related to Shapiro's finding [19] that the quicker consumers gather information, the better is the product that a firm will offer. However, in Shapiro's model, the firm produces better products because it correctly *expects* retaliation from consumers and thus rationally *avoids* it, whereas in our model players are not that smart: the e-community gets more and more honest players simply because cheating behaviors are *already* severely punished, which makes them unpopular.

Corollary 3: The robustness of the health of an e-community is negatively related to α_0 , a new member's initial trust status.

Corollary 3 gives one a sense that the more conservative an e-community is, the healthier it is. For example, on eBay, some players may come directly into the e-community from the outside, whereas other players who get negative feedback ratings may reregister. In essence, they are throwing away their bad reputations and playing as a newcomer. To combat this, eBay attaches a "changed ID" icon next to new user IDs and informs community members that a user has changed her user ID within the past 30 days.

This result is consistent with Friedman and Resnick's finding [7] that it is easy for someone to obtain a new identity on the Internet, which introduces opportunities to misbehave without incurring reputation costs. Their game theory analysis demonstrates that there are inherent limitations to the effectiveness of reputation systems when people can change their identities. Newcomers without any feedback rating will be distrusted until they have somehow paid their dues, either through an entry fee or by accepting more risks or inferior prices while building up a reputation.

Corollary 4: The robustness of the health of an e-community is negatively related to $1/T$, the population change rate between periods.

Let us assume that there are two e-communities, one with smaller $1/T$, and the other with larger $1/T$. Intuitively, the former means that players stay longer in the community (on average) and have more opportunities to meet trading partners, which affects their trust status (e.g., reputation rating) more. The latter means that there are more newcomers without any reputation entering the e-community, and often players may not know whether they should trust the newcomers. Therefore, in an e-community where there is a rapid turnover of members, the robust health of the e-community is difficult to maintain. For example, as a business operation, eBay wants to see more

new users entering the market. However, eBay also has to pay attention to $1/T$, the population change rate between periods, and keep it within an acceptable range. In other words, an unreasonably large percentage of new users entering the market is not good for eBay because its reputation system loses its force, and people may choose not to trade in the market.

On one hand, $1/T$ is often an endogenous characteristic of an e-community. For example, online fan groups devoted to a specific pop musician often have a very high turnover rate, as popularity changes fast, whereas mailing lists for professional groups usually have a small turnover rate, as it is harder to change professions. On the other hand, e-community management may use various mechanisms, such as registration, to help keep $1/T$ within an acceptable range.

Corollary 5: The robustness of the health of an e-community is negatively related to e_{h0} and e_{c0} , the upper bound of $e_h(k)$ and $e_c(k)$, respectively.

Recall that $e_h(k) = 1 + \underline{u}_h(k) - \underline{u}(k)$, and $e_c(k) = 1 + \underline{u}_c(k) - \underline{u}(k)$. If the upper bounds of $e_h(k)$ and $e_c(k)$ —that is, e_{h0} and e_{c0} —are small, that means both the average payoff $\underline{u}_h(k)$ for honest players and the average payoff $\underline{u}_c(k)$ for cheating players are not much better than $\underline{u}(k)$, which is the average payoff of the whole population for players both within and outside the e-community. In other words, the temptation for outside players to enter the e-community is not very strong. As a result, population dynamics will not be too dramatic, and the difference between the trust status for honest players and cheating players in the e-community will remain relatively stable, which helps to sustain the competitive advantages that honest players enjoy. We further discuss various types of outside options in the next section.

Corollaries 4 and 5 are similar in that they both slow down the evolution of the e-community: the former works by restricting the turnover rate of the whole population, the latter by restricting newcomers' possibilities of switching away from their ancestors' strategies.

Although it is clear that smaller e_{h0} and e_{c0} helps smooth the evolutionary process, thus increasing the robustness of the health, one may expect that larger e_{h0} may also benefit the e-community, since honest players have a chance to get high payoffs. Nevertheless, the effect of a larger e_{h0} is actually unclear: any period during which honest players enjoy a high payoff may trigger a large inflow of new honest players in the next period, which reduces the average trust status of honest players. Such a reduction could be large enough that the competitive advantage that honest players enjoyed no longer exists. Therefore, increasing the upper bound of $e_h(k)$ could be a double-edged sword to the robust health of the e-community. Given the limitations in the current model, it is not clear which effect will be stronger.

Size of the e-Community and Outside Options

IN THE SUBSECTION "SIZE OF THE E-COMMUNITY," we mentioned that whenever both the honest player population and the cheating player population are diminishing inside a community, the community's health makes little difference for business

purposes. Whereas in the previous section we focus on the *earlier periods* of the evolution, in which the central issue is to find a condition as early as possible that guarantees the evolutionary process is healthy, in this section we focus on the *later periods* of the evolution, in which the central issue is whether the e-community or the outside option will eventually attract the most players when the evolutionary process stabilizes.

In this section, we first find out the average payoff for honest players when the evolution eventually stabilizes. We then discuss its implications and, along with the outside option, how it affects the competition between the e-community and the outside world.

Size of the e-Community

As we discussed earlier, if condition (17) is satisfied at period k_1 , the e-community is healthy. Then, from Proposition 1 we know that the e-community will evolve toward dominance by honest players. Nevertheless, even when it reaches a period where almost all players are honest, the players still will not fully trust each other and therefore will forgo a nonnegligible proportion of profitable transactions, as shown in the following proposition:

Proposition 3: Given a healthy e-community, suppose the outside option yields a payoff per period that converges to a constant $u_s(\infty)$, and suppose the evolution eventually stabilizes:

- if $[1 - (1 - \alpha_0)/(T - (T - 1)r_h)]^2 a > u_s(\infty)$, an honest player's per-period payoff converges to $[1 - (1 - \alpha_0)/(T - (T - 1)r_h)]^2 a$. Furthermore, almost the entire population will eventually enter the healthy e-community and behave honestly;
- if $[1 - (1 - \alpha_0)/(T - (T - 1)r_h)]^2 a < u_s(\infty)$, almost all members in the healthy e-community will eventually leave.

From the proof of Proposition 3, we know that $\alpha_h(k)$ converges to $1 - (1 - \alpha_0)/(T - (T - 1)r_h)$ if the healthy e-community dominates the outside option. This may look counterintuitive: if players know that the e-community consists of almost all honest players in all future periods, why do they not still fully trust each other? And does not this contradict the assumption that $\alpha_0 < 1$; that is, the e-community is always suspicious of incoming new members?

It is true that the assertion that "the e-community is always suspicious of incoming new members" cannot be explained *statically* in any given period. Nevertheless, the existence of such a suspicion on the part of members exerts an indispensable punishment mechanism on any potential future cheaters, and thus is justified from a *dynamic* process perspective. Corollary 5 provides additional support: an increase in α_0 decreases the robustness of the health of an e-community, and therefore may overturn condition (17), and thus possibly lead to an unhealthy e-community.

On the other hand, we want the punishment to fit the crime, since excessive punishment may also end up harming honest players, as shown in the following corollary:

Corollary 6: Given a healthy e-community, suppose the outside option yields a payoff per period that converges to $u_s(\infty)$, and suppose $[1 - (1 - \alpha_0)/(T - (T - 1)r_h)]^2 a > u_s(\infty)$, when the evolution stabilizes:

- *an honest player's per-period payoff is positively related to α_0 , a new member's initial trust status;*
- *an honest player's per-period payoff is positively related to $1 - r_h$, the information dissemination speed for honest players.*

Corollary 6 follows directly from Proposition 3. α_0 and $1 - r_h$ together determine how the trust status evolves for an honest player. Corollaries 3 and 6 together provide a balanced view of how an e-community should treat new incoming members: do not trust them too much, since some of them are cheaters, but do not trust them too little, either, since honest players may also get hurt in the long run.

From Proposition 3, we also know that eventually an honest player's per-period payoff will be negatively correlated with $1/T$, the population change rate between periods. Intuitively, a slowly changing e-community gives an honest player a longer expected life, and thus makes it possible for more community members to learn about her honesty, which in turn increases her trust status. Taken together with Corollary 4, we know that a smaller $1/T$ both helps to achieve a healthy e-community and increases honest players' payoffs in the long run.

Outside Options and Competing e-Communities

Proposition 3 also sheds light on how outside options may affect the size of an e-community. In this subsection, we consider two specific types of outside options. The first is a constant reserve price p , which can be interpreted as the average payoff to a player if she chooses to stay outside of the e-community and use the time to pursue other options. A constant reserve price implies that the whole society is relatively stable. The second option is more specific: another e-community. Thus, in this second type we are effectively considering two competing e-communities.

Although the evolutionary rules for players outside the e-community are not explicitly written, they are implicit from Equations (11) and (12), since the size of the total population is always one at any period. Notice that in order to have players outside of the e-community, we need to have $x_h(k) + x_c(k)$ strictly smaller than 1 at any period, which is true if $x_h(k_0) + x_c(k_0) < 1$. Therefore, in this subsection we assume that $x_h(k_0) + x_c(k_0) < 1$.

A Constant Reserve Price as the Outside Option

Given a constant reserve price, Proposition 3 directly implies

Corollary 7: Given a healthy e-community, suppose the evolution eventually stabilizes:

- *if $[1 - (1 - \alpha_0)/(T - (T - 1)r_h)]^2 a > p$, the outside population will be gradually eliminated;*
- *if $[1 - (1 - \alpha_0)/(T - (T - 1)r_h)]^2 a < p$, the e-community will gradually diminish.*

Following Corollary 7, if the environment undergoes a relatively disruptive change—for example, if a new online firm with attractive properties opens for business—the reserve price could experience an upward shift. In that case, the ordering between $[1 - (1 - \alpha_0)/(T - (T - 1)r_h)]^2 a$ and p could be reversed. As a result, a formerly large e-community could shrink.

It is worthwhile to also consider size changes between consecutive periods. If more players move into the e-community than move out between consecutive periods, we have

$$x_c(k+1) + x_h(k+1) > x_c(k) + x_h(k),$$

or

$$\frac{x_h(k)}{x_h(k) + x_c(k)} \underline{u}_h(k) + \frac{x_c(k)}{x_h(k) + x_c(k)} \underline{u}_c(k) > \underline{u}(k). \quad (18)$$

Equation (18) says that the e-community's size increases if and only if the average payoff inside the e-community is larger than the average payoff to the whole population, or equivalently, if and only if the average payoff inside the e-community is larger than that outside the e-community. Nevertheless, simulation shows that when the conditions in Corollary 7 hold, Equation (17) may not hold.

Remark: The size of a healthy e-community that satisfies $[1 - (1 - \alpha_0)/(T - (T - 1)r_h)]^2 a > p$ may temporarily decrease, although in the long run it will increase to almost 1.

A Competing e-Community as the Outside Option

Now we consider the case where the whole population exists in two competing e-communities where there are no other outside options. In this case, one can be viewed as the outside option for the other. We label them e-community 1 and e-community 2, and put superscripts on parameters to distinguish them.

Corollary 8: Suppose the outside option is another e-community. Suppose both e-communities are healthy and the evolution eventually stabilizes. If $[1 - (1 - \alpha_0^i)/(T - (T - 1)r_h^i)]^2 a^i > [1 - (1 - \alpha_0^j)/(T - (T - 1)r_h^j)]^2 a^j$, where $i, j \in \{1, 2\}$, $i \neq j$, then eventually all players will choose e-community i .

It is worthwhile to discuss Corollary 8 along with Proposition 2, since together they provide a picture of how competing e-communities will evolve. Specifically, during early periods of the competition, e-communities focus on being healthy, whereby controlling the unhealthy index is vital. Once the objective of health is achieved, the result of competition eventually depends on the endogenous and fixed properties of both e-communities. This provides managerial implications for e-communities, which we will discuss in the next section.

Implications and Conclusions

Implications for e-Communities

IN SEPARATE SECTIONS ABOVE WE HAVE DISCUSSED the sustained competitive advantages for honest players in an e-community. Specifically, we addressed the health issue in Proposition 2, and the size issue in Proposition 3. Together these two propositions provide several implications for the success of an e-community. On one hand, factors including high information dissemination speed, limiting population turnover rate, increasing payoff to honest players, and limiting the loss honest players suffer all benefit both the health and growth of an e-community. On the other hand, factors such as initial trust status could affect health and growth in opposite directions, thus their overall effect on an e-community needs to be carefully balanced.

A high information dissemination speed helps the e-community to quickly distinguish between honest and cheating players, thus effectively punishing cheating players by refusing to trade with them, and benefiting honest players by reducing their chances of being cheated. Consequently, a high information dissemination speed benefits an e-community by both maintaining health and attracting outside players. One way to achieve a higher speed is through the use of ratings or peer review systems. When the sheer volume of rating information overwhelms a person's ability to digest it all, an effective rating aggregation mechanism, such as counting the number of negative or positive reviews, also helps raise the information dissemination speed.

A low population turnover rate gives players long expected lives in the e-community, which helps the e-community better know each player, thus giving honest players a higher expected payoff. While a low turnover rate may limit the expansion speed the e-community can achieve, it actually prevents dramatic population changes inside the e-community in the long run, which in turn protects the established trust among players. The organizer of e-communities may use various levels of registration (e.g., from registration using just e-mail addresses to registration using e-mail addresses and credit card numbers or even charging a registration fee) to keep the proportion of newcomers within a preferable range, therefore keeping the community turnover rate within an acceptable range. Similarly, placing strict limitations on player payoffs relative to the average payoff of the entire population of the e-community (e_{h0} and e_{c0}) also restricts dramatic population changes, thus improving the health and size of the e-community. One example of limiting player payoffs is to discourage players from trading exceptionally high-value goods in the e-community.

The e-community's trust over new members—that is, initial trust status—serves as a double-edged sword. On one hand, relative distrust of new members makes it hard for new cheaters to benefit from cheating, which is beneficial to the health of the e-community. On the other hand, as the e-community cannot distinguish new honest players from new cheating players, skepticism also hurts honest players, which may propel them to choose outside options, which negatively affects the size of the e-community. An appropriate choice of initial trust status needs to balance these two affected sides. The initial trust status is often largely formulated endogenously by players, depending

on the specific context and purpose of an e-community. For instance, online professional discussion groups often have a high initial trust status—we usually believe the truthfulness of information posted on ISWorld (www.isworld.org) mailing list by strangers—because of the high integrity of these groups. Comparatively, online auction sites have low initial trust status because of the high possibility of fraud. Note that managerial mechanisms do have effects on initial trust status. For example, in the early ages of e-commerce, people often assume unknown partners are trustworthy. eBay now offers educational information on the possibility of fraud to anyone who becomes a member, which makes traders on eBay now more cautious in dealing with unknown partners. Alternatively, eBay may enforce even stricter registration procedures than those currently adopted to ensure that cheaters cannot flee easily, which in turn generally raises the initial trust status of all the newcomers.

If the e-community provides high payoffs to players who interact with each other honestly, it both encourages honest behavior and gives honest players an advantage over outside players. Payoffs related to cheating—that is, the extra profit a cheater enjoys and the loss an honest player suffers—do not affect a healthy e-community in the long run, since cheating behaviors will be gradually eliminated. To protect health, an e-community should reduce both a cheater's extra profit and an honest player's loss when cheating behavior happens. Examples of possible mechanisms include escrow services and trusted third parties.

Conclusions

In this paper, we have developed an evolutionary model to study the trust issue in e-communities, with emphasis placed on large populations, continuous changes in community members, and imperfect information and memory. We proposed the health of an e-community as the central notion in our model, which is defined from an evolutionary perspective. This research has revealed that the health of an e-community is positively related to payoffs when honest traders meet, and negatively related to honest players' losses and cheating players' extra profits. Increasing the information dissemination speed in an e-community plays a crucial role in maintaining the health of the e-community. Further, our findings indicate that both a new member's initial trust status and the rate of replacement in an e-community negatively affect the health of an e-community. Moreover, payoffs when honest traders meet, the information dissemination speed, and a new member's initial trust status all positively affect an honest player's equilibrium payoff, which is also negatively affected by the rate of replacement.

The major contribution of this research is to propose an evolutionary framework to understand trust in e-communities, define what it means for an e-community to be healthy, and derive conditions under which one is healthy. This paper serves as a theoretical foundation for future research on e-communities using the evolutionary approach. For example, as business entities pursue profit maximization, e-communities have an interest in implementing fee schemes that can yield high profits for them without losing honest members to other e-communities. Moreover, a profit-seeking

business entity may not prefer very high levels of trust if it costs too much to maintain such levels. How we can expand the evolutionary e-community model introduced in this paper to address the profitability issue for e-community owners remains an open and intriguing question.

Acknowledgments: The authors thank the three anonymous reviewers and the Editor-in-Chief for their careful reviews and detailed suggestions on previous versions of the paper. Their comments and suggestions helped strengthen the analysis and greatly improved the quality of the research.

NOTES

1. Nevertheless, this analytical equivalence has its limitations. When an old player dies and is replaced by a new player, all information about the old player is lost, although this is not the case for a player who merely switches her strategy. As a result, the current model cannot explain complex individual behaviors, such as when a player first builds a reputation for being honest and then exploits her reputation.

2. This exponential process is just one example of processes that have decreasing slopes. There are three reasons we pick this specific form. First, it is a good approximation of numerous other processes that have smooth, decreasing slopes. Second, it offers the simplest way to model the dynamics of the evolution. Third, and as shown in Proposition 2, if Equation (17) holds, the healthy e-community is robust in that small disturbances to the evolution will not change the result. Therefore the analysis still holds true, even if the dynamic of trust status is slightly different from a strict exponential process.

REFERENCES

1. Abreu, D. On the theory of infinitely repeated games with discounting. *Econometrica*, 56, 2 (1988), 383–396.
2. Axelrod, R. *The Evolution of Cooperation*. New York: Basic Books, 1984.
3. Ba, S.; Whinston, A.B.; and Zhang, H. The dynamics of the electronic market: An evolutionary game approach. *Information Systems Frontiers*, 2, 1 (2000), 31–40.
4. Ba, S.; Whinston, A.B.; and Zhang, H. Building trust in the electronic market through an economic incentive mechanism. *Decision Support Systems*, 35, 3 (June 2003), 273–286.
5. Binmore, K. *Fun and Games*. Lexington, MA: D.C. Heath, 1992.
6. Chang, A.; Kannan, P.K.; and Whinston, A.B. Electronic communities as intermediaries: The issues and economics. In R.H. Sprague Jr. (ed.), *Proceedings of the Thirty-Second Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Computer Society Press, 1999 (available at csdl.computer.org/comp/proceedings/hicss/1999/0001/05/00015041.PDF).
7. Friedman, E., and Resnick, P. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10, 2 (2001), 173–199.
8. Hammerstein, P., and Riechert, S.E. Payoffs and strategies in territorial contests: ESS analyses of two ecotypes of the spider, *Agelenopsis aperta*. *Evolutionary Ecology*, 2, 2 (1988), 115–138.
9. Hofbauer, J., and Sigmund, K. *The Theory of Evolution and Dynamical Systems*. Cambridge: Cambridge University Press, 1988.
10. Internet Fraud Watch. Internet fraud statistics. National Consumers League, Washington, DC, 2002 (available at www.fraud.org/2002intstats.htm).
11. Jarvenpaa, S.L.; Knoll, K.; and Leidner, D. Is anybody out there? Antecedents of trust in global virtual teams. *Journal of Management Information Systems*, 14, 4 (Spring 1998), 29–64.

12. Jarvenpaa, S.L.; Tractinsky, N.; and Vitale, M. Consumer trust in an Internet store. *Information Technology Management*, 1, 1-2 (2000), 45-71.
13. Klein, B., and Leffler, K.B. The role of market forces in assuring contractual performance. *Journal of Political Economy*, 89, 4 (1981), 615-641.
14. Kollock, P. The production of trust in online markets. In E.J. Lawler, M. Macy, S. Thyne, and H.A. Walker (eds.), *Advances in Group Processes*, vol. 16. Greenwich, CT: JAI Press, 1999, pp. 99-123.
15. Maynard Smith, J. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press, 1982.
16. Maynard Smith, J., and Price, G.R. The logic of animal conflict. *Nature*, 246, 5427 (1973), 15-18.
17. Milgrom, P.R., North, D.C., and Weingast, B.R. The role of institutions in the revival of trade: The law merchant, private judges, and the champagne fairs. *Economics and Politics*, 2, 1 (March 1990), 1-23.
18. Resnick, P.; Zeckhauser, R.; Friedman, E.; and Kuwabara, K. Reputation systems: Facilitating trust in Internet interactions. *Communications of the ACM*, 43, 12 (2000), 45-48.
19. Shapiro, C. Consumer information, product quality, and seller reputation. *Bell Journal of Economics*, 13, 1 (1982) 20-35.
20. Stewart, K.J. Trust transfer on the World Wide Web. *Organization Science*, 14, 1 (January-February 2003), 5-17.
21. Taylor, P.D., and Jonker, L.B. Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40, 1-2 (July 1978), 145-156.
22. Van Damme, E. *Stability and Perfection of Nash Equilibria*. Berlin: Springer, 1991.
23. Young, H.P. *Individual Strategy and Social Structure*. Princeton: Princeton University Press, 1998.

Appendix

Proof of Lemma 1

EQUATIONS (3) AND (4) CAN BE REWRITTEN AS

$$u_h(k, t_0) = \underline{\alpha}_h(k) \alpha_h(k, t_0) \frac{x_h(k)}{x_h(k) + x_c(k)} a - \underline{\alpha}_c(k) \alpha_h(k, t_0) \frac{x_c(k)}{x_h(k) + x_c(k)} b \quad (\text{A.1})$$

$$u_c(k, t_0) = \underline{\alpha}_h(k) \alpha_c(k, t_0) \frac{x_h(k)}{x_h(k) + x_c(k)} (a + d). \quad (\text{A.2})$$

Given the expressions about average payoff, Equations (5) and (6) can be rewritten as

$$x_h(k+1, 0) = \left(\frac{1}{T} \right) x_h(k) (1 + \underline{u}_h(k) - \underline{u}(k)) \quad (\text{A.3})$$

$$x_c(k+1, 0) = \left(\frac{1}{T} \right) x_c(k) (1 + \underline{u}_c(k) - \underline{u}(k)). \quad (\text{A.4})$$

Combining Equations (7), (8), (A.3), and (A.4), we get Equations (9) and (10). From Equations (A.1) and (A.2), we have Equations (11) and (12). From the definition of $\underline{\alpha}_h(k)$ and Equation (A.3), we get Equation (13). Finally, from the definition of $\underline{\alpha}_c(k)$ and Equation (A.4), we get Equation (14).

Note that Equations (9) through (14) uniquely determine the evolution of the e-community given any starting period k_0 and $x_h(k_0)$, $x_c(k_0)$, $\underline{\alpha}_h(k_0)$, $\underline{\alpha}_c(k_0)$, and $\underline{u}(k)$ for any $k \geq k_0$.

Proof of Proposition 1

From Equations (9) and (10), we have $[(x_c(k+1))/(x_h(k+1))]/[(x_c(k))/(x_h(k))] = (T + \underline{u}_c(k) - \underline{u}(k))/(T + \underline{u}_h(k) - \underline{u}(k))$. Since for any $k \geq k_1$, $\underline{u}_h(k) > \underline{u}_c(k)$ and $\lim_{k \rightarrow \infty} (\underline{u}_h(k) - \underline{u}_c(k)) > 0$, if this limit exists, it must be true that $[(x_c(k+1))/(x_h(k+1))]/[(x_c(k))/(x_h(k))] < 1$ for any $k \geq k_1$, and $\lim_{k \rightarrow \infty} [(x_c(k+1))/(x_h(k+1))]/[(x_c(k))/(x_h(k))] < 1$. Thus $\lim_{k \rightarrow \infty} (x_c(k+1))/(x_h(k+1)) = 0$, $(x_{h0}, 0)$ is an ESS.

Proof of Lemma 2

Part i: If $\underline{\alpha}_c(k_0) < \gamma_c \alpha_0$, then for $k = k_0$, we want to show $\underline{\alpha}_c(k+1) < \gamma_c \alpha_0$; that is,

$$\alpha_0 - \frac{T-1}{T-1+e_c(k)} (\alpha_0 - r_c \underline{\alpha}_c(k)) < \alpha_0 - \frac{(T-1)(1-r_c)}{(T-1)(1-r_c)+e_{c0}} \alpha_0,$$

that is,

$$(1-r_c)e_c(k) < e_{c0} - r_c \frac{\alpha_c(k)}{\alpha_0} [(T-1)(1-r_c) + e_{c0}].$$

This is true since

$$\begin{aligned} e_{c0} - r_c \frac{\alpha_c(k)}{\alpha_0} [(T-1)(1-r_c) + e_{c0}] &> e_{c0} - r_c \gamma_c [(T-1)(1-r_c) + e_{c0}] \\ &= (1-r_c)e_{c0} > (1-r_c)e_c(k). \end{aligned}$$

Iteratively, we can show that this result holds for $k = k_0 + 1, k = k_0 + 2, \dots$. Thus, we know that whenever $\alpha_c(k_0) < \gamma_c \alpha_0$ happens, it will be true for all future periods.

If $\alpha_c(k_0) \geq \gamma_c \alpha_0$, we can rewrite Equation (14) as

$$\alpha_0 - \alpha_c(k+1) = \frac{T-1}{T-1+e_c(k)} (\alpha_0 - r_c \alpha_c(k)).$$

Then from $0 \leq e_c(k) < e_{c0}$, we know that

$$\alpha_0 - \alpha_c(k+1) > \frac{T-1}{T-1+e_{c0}} (\alpha_0 - r_c \alpha_c(k)).$$

We also have

$$\alpha_c(k_0) \geq \gamma_c \alpha_0 \Rightarrow \frac{T-1}{T-1+e_{c0}} (\alpha_0 - r_c \alpha_c(k_0)) \geq \alpha_0 - \alpha_c(k_0).$$

Therefore, $\alpha_c(k_0+1) < \alpha_c(k_0)$.

To show that there exists $k_1 > k_0$ such that for all $k \geq k_1$, $\alpha_c(k) < \gamma_c \alpha_0$, we only need to show that there exists $k_1 > k_0$ such that for k_1 , $\alpha_c(k_1) < \gamma_c \alpha_0$. Assume, to the contrary, that for all $k > k_0$, $\alpha_c(k) \geq \gamma_c \alpha_0$. Then $\alpha_c(k)$ is a decreasing sequence with a lower bound of $\gamma_c \alpha_0$, thus it has a limit of no less than $\gamma_c \alpha_0$. Suppose this limit is A ; then $\alpha_c(k) - \alpha_c(k+1) \rightarrow 0$ when $\alpha_c(k) \rightarrow A$. However,

$$\begin{aligned} \alpha_c(k) - \alpha_c(k+1) &\rightarrow 0 \Rightarrow (\alpha_0 - \alpha_c(k+1)) - \frac{T-1}{T-1+e_{c0}} (\alpha_0 - r_c \alpha_c(k)) \rightarrow 0 \\ &\Rightarrow e_{c0} - e_c(k) \rightarrow 0, \end{aligned}$$

and this is a contradiction. (Recall that e_{c0} is slightly larger than the upper bound of $e_c(k)$.)

Part ii. We prove Part ii in the same way as Part i, noticing the symmetry between $\alpha_c(k)$ and $(1 - \alpha_h(k))$.

Proof of Proposition 2

From Lemma 2, we know

$$\frac{1}{b} \left[a \frac{\underline{\alpha}_h(k)}{\underline{\alpha}_c(k)} - a - d \right] > \frac{1}{b} \left[a \frac{1 - \gamma_h(1 - \alpha_0)}{\gamma_c \alpha_0} - a - d \right]$$

for all $k \geq k_1$.

At period $k = k_1$,

$$\frac{x_c(k)}{x_h(k)} \leq \frac{1}{b} \left[a \frac{1 - \gamma_h(1 - \alpha_0)}{\gamma_c \alpha_0} - a - d \right] < \frac{1}{b} \left[a \frac{\underline{\alpha}_h(k)}{\underline{\alpha}_c(k)} - a - d \right] \Rightarrow \frac{x_c(k+1)}{x_h(k+1)} < \frac{x_c(k)}{x_h(k)}.$$

Thus

$$\frac{x_c(k)}{x_h(k)} \leq \frac{1}{b} \left[a \frac{1 - \gamma_h(1 - \alpha_0)}{\gamma_c \alpha_0} - a - d \right]$$

still holds for $k = k_1 + 1$. By iteration we know $(x_c(k))/(x_h(k))$ is a decreasing sequence. $(x_c(k))/(x_h(k))$ is lower bounded by 0, thus it has a limit. Given $\varepsilon > 0$, notice that $(x_c(k))/(x_h(k)) - (x_c(k+1))/(x_h(k+1)) \rightarrow 0$ when $(x_c(k))/(x_h(k)) > \varepsilon$. Therefore the limit is 0.

Proof of Proposition 3

Let the payoff per period from the outside option be denoted as $u_s(k)$. Since the e-community stabilizes, the ESS exists and average trust statuses converge. Let the ESS be $(x_{h0}, 0)$, where $x_{h0} \geq 0$. Let $\underline{\alpha}_h(k) \rightarrow \underline{\alpha}_h(\infty)$. Then $\underline{u}_h(k) \rightarrow \alpha_h^2(\infty)a$, and $\underline{u}(k) \rightarrow x_{h0}\alpha_h^2(\infty)a + (1 - x_{h0})u_s(\infty)$. From Equation (2), we know

$$\begin{aligned} \underline{\alpha}_h(\infty) &= \frac{1}{T} [1 - (1 - \alpha_0)] + \frac{1}{T} \frac{T-1}{T} [1 - (1 - \alpha_0)r_h] \\ &+ \frac{1}{T} \left(\frac{T-1}{T} \right)^2 [1 - (1 - \alpha_0)r_h^2] + \dots = 1 - \frac{1 - \alpha_0}{T - (T-1)r_h}. \end{aligned}$$

Then

$$\underline{u}_h(k) \rightarrow \left(1 - \frac{1 - \alpha_0}{T - (T-1)r_h} \right)^2 a.$$

However, if

$$\left(1 - \frac{1 - \alpha_0}{T - (T-1)r_h}\right)^2 a > u_s(\infty),$$

in each period more outside players join the e-community than inside players move out, so it is an ESS only if $x_{h0} = 1$; if

$$\left(1 - \frac{1 - \alpha_0}{T - (T-1)r_h}\right)^2 a < u_s(\infty),$$

in each period more inside players leave the e-community than outside players join in, so it is an ESS only if $x_{h0} = 0$.

Proof of Corollary 8

Suppose the ESS is $(x_{h0}^i, 0)$ for e-community i and $(x_{h0}^j, 0)$ for e-community j .

1. If both x_{h0}^i and x_{h0}^j are strictly positive. Then for e-community i and j , honest players' average payoffs converge to $[1 - (1 - \alpha_0^i)/(T - (T-1)r_h^i)]^2 a^i$ and $[1 - (1 - \alpha_0^j)/(T - (T-1)r_h^j)]^2 a^j$, respectively. But then, given the condition in this corollary, $x_h^j(k)$ cannot converge, which is a contradiction.
2. If $x_{h0}^i = 0$ and $x_{h0}^j = 1$. Then for e-community j , honest players' average payoff converges to $[1 - (1 - \alpha_0^j)/(T - (T-1)r_h^j)]^2 a^j$. According to the first half of Proposition 3, this is impossible. Therefore the only feasible case is $x_{h0}^i = 1$ and $x_{h0}^j = 0$.

Copyright of *Journal of Management Information Systems* is the property of M.E. Sharpe Inc. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.