The Dynamics of Trust and Trustworthiness on EBay. An Evolutionary Analysis of Buyer Insurance and Seller Reputation*

Werner Güth†

Max Planck Institute of Economics Jena

Friederike Mengel‡

University of Alicante

Axel Ockenfels§

University of Cologne

February 11, 2006

Abstract

Applying an evolutionary framework, we investigate how a reputation mechanism and a buyer insurance (as used on Internet market platforms such as eBay) interact to promote trustworthiness and trust. Our analysis suggests that the costs involved in giving reliable feedback determine the gains from trade that can be obtained in equilibrium. Buyer insurance, on the other hand, can affect the trading dynamics and equilibrium selection. We find that, under reasonable conditions, buyer insurance crowds out trust and trustworthiness.

*Ockenfels gratefully acknowledges the support of the Deutsche Forschungsgemeinschaft. Mengel's research was supported by the Spanish Ministry of Education and Science (grant SEJ 2004-02172).

† *Max Planck Institute of Economics*, Kahlaische Straße 10, D-07745 Jena, Germany; *E-mail:* gueth@econ.mpg.de (www.econ.mpg.de).

‡ *Universidad de Alicante,* Departamento de Fundamentos del Análisis Economico, 03080 Alacant, Spain. *E-mail*: friederike@merlin.fae.ua.es.

§ *University of Cologne*, Department of Economics, Albertus Magnus Platz, D-50923 Cologne, Germany; *E-mail:* ockenfels@uni-koeln.de (http://ockenfels.uni-koeln.de). Ockenfels advised several online market provider on auction design, including eBay. The views expressed are his own.

1 Introduction

Trading on Internet platforms such as eBay typically takes place between anonymous and geographically dispersed parties in one-shot interactions (e.g., Resnick and Zeckhauser, 2002). This creates moral hazard on the seller side, hampering trade efficiency. EBay responds to this challenge with a mixture of policies and mechanisms that are devised to promote trustworthy sellers and trusting buyers.[1] The most prominent mechanism to promote trustworthiness is eBay's so-called Feedback Forum, where traders can publicly post comments on the transaction partner (Dellarocas, forthcoming a). The most prominent policy to promote trust is eBay's so-called Purchase Protection Program, an insurance which reimburses buyers "where an item was purchased on eBay and either not received or was received but significantly not as described."[2] While, to our knowledge, there is no study dealing with the effect of eBay's buyer insurance on trading patterns, there is quite some empirical and theoretical work on the impact of the Feedback Forum. Overall, the literature seems to agree that the Feedback Forum contributes to the trust and trustworthiness levels observed on eBay.[3]

However, there are also signs that there are problems with seller trustworthiness. A report by the research group GartnerG2 (2002) concludes that "Internet transaction fraud is 12 times higher than in-store fraud." More recently, controlled field experiments by Jin and Kato (2004) led to similar conclusions: The online fraud rate was significantly higher than the fraud rate observed in corresponding offline transactions. Furthermore, problems with seller trustworthiness seem to accelerate over time. Total reported losses from Internet auction fraud tripled from $17 million

[1] http://pages.ebay.com/help/confidence/index.html (December 2005).

[2] http://pages.ebay.com/help/tp/esppp-coverage- eligibility.html (December 2005). A buyer can only be reimbursed up to a maximum of $200 (less $25 to cover processing costs).

[3] See the survey articles by Bajari and Hortacsu (2004) and by Ockenfels and Reiley (forthcoming) on economic research dealing with eBay, the paper by Dellarocas (forthcoming b) on empirical and theoretical studies investigating the impact of eBay's reputation system, and the work by Bolton et al. (2004) for an experimental analysis of online reputation mechanisms.

in 2002 to $54 million in 2003. Alarmed by these dynamics, the Identity Theft Resource Center warned that, "Online fraud is becoming as big an issue for ... eBay and AOL as security is for Microsoft."[4] More recently, in 2005, the Federal Trade Commission (FTC) reported that "auction fraud is on the rise, with an increasing number of consumers complaining about sellers who deliver their advertised goods late or not at all, or deliver something far less valuable than promised."[5] Even eBay's director of trust and safety in the U.K. admitted, according to the BBC, "extreme growth" in the number of fraud incidents during 2005.[6]

In this paper, we simultaneously investigate the effectiveness of buyer insurance and seller reputation to reverse the adverse dynamics of cooperative trading patterns on eBay (and similar Internet market platforms). We apply an evolutionary approach, which - in our context - has several advantages over standard game theoretic approaches. For one, evolutionary analysis better captures out-of-equilibrium dynamics of trade behaviors, as we seemingly observe them on eBay. Also, when it comes to matters of trust and trustworthiness in trade relationships, there are often multiple equilibria, both in the standard and the evolutionary analysis. The evolutionary approach reveals the 'attractiveness' of each equilibrium and how it relates to the path of play (as captured by the initial population of traders).[7] Finally, the evolutionary approach does not require the usual assumptions like (commonly known) rationality of traders but rather allows trading patterns to emerge from simple learning and imitating behaviors. This appears to be a reasonable starting point for analyzing a C2C-platform like eBay, which brings together many million traders.

---

[4] http://www.usatoday.com/tech/news/2003-10- 23-fraud_x.htm (December 2005).

[5] See the FTC's "Top Ten Dot Cons" on http://www.ftc.gov/bcp/conline/edcams/dotcon /auction.htm (July 2005). Similarly, in 2004, the Internet Crime Complaint Center (IC3) reported that "auction fraud was by far the most reported offense, comprising 71,2% of [207,449] referred complaints" (see http://www1.ifccfbi.gov/strategy/statistics.asp July 2005).

[6] For example, http://news.designtechnica.com/article 9083.html (December 2005).

[7] Trade insurance may have a quite different impact if introduced in a trustworthy rather than in a fraudulent community. For example, a cross-cultural study of trust by Vishwanath (2004) conducted on eBay suggests that, given the same market institutions, the dynamics of social behavior systematically differ with the social capital measured in the respective communities.

2 A simple trade model

Figure 1 shows how sellers and buyers interact in a basic trade game. The seller can either be trustworthy and deliver properly ($\delta = 0$) or not ($\delta = 1$). Then, in case of bad delivery, the buyer must either leave a positive (or neutral) feedback ($\rho = 0$) or a negative feedback ($\rho = 1$).[8] Assume for the moment that the payoff parameters Π, E, and d are all zero. Then, a successful trade creates a surplus of 1, shared equally between seller and buyer. Untrustworthiness, on the other hand, yields $q \in (1/2, 1)$ to the seller and zero to the buyer. Clearly, in this simple scenario, inefficiency will arise because of the seller's incentive not to deliver properly.

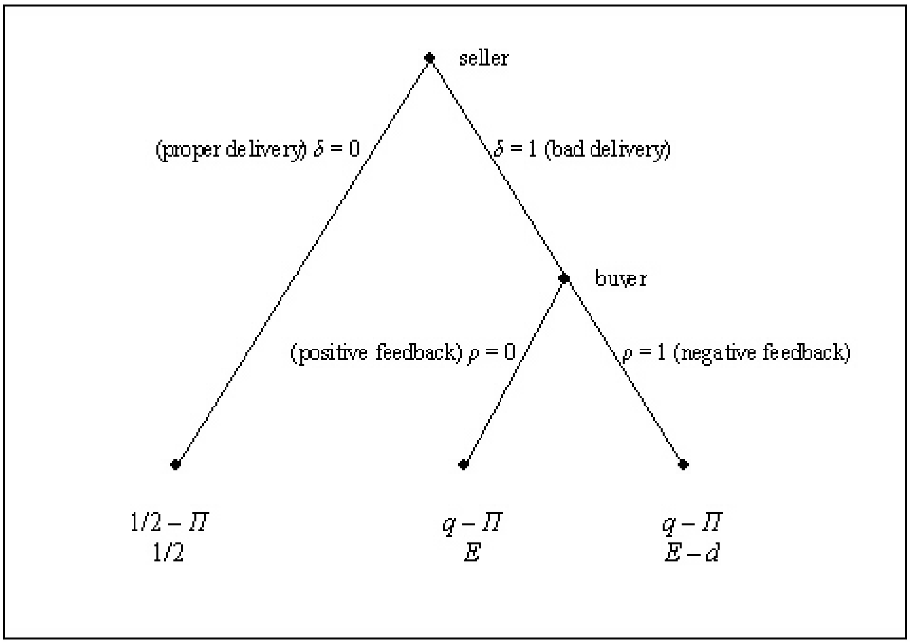

Figure 1: Basic trade game without outside option. $E \leq 1/2$ and $E - d$

Let us now assume that the basic trade game is played repeatedly, such that in each period τ, each of an infinite number of traders can assume both the role of a buyer or that of a seller. In each

[8]As we will explain in a moment, the buyer's decision whether to trust a seller will be implicitly captured by the payoffs of the repeated version of the game and is thus not discussed here. Note also that, for simplicity, buyers are not allowed to leave no feedback, as is permitted on eBay. However, no feedback can be (re)interpreted either as positive, neutral, or negative feedback.

period, buyers and sellers are randomly matched to play the basic trade game. Over time, the base games are linked via reputation information $t \in \{g, b\}$ of each seller, where b stands for "bad" and g for "good" reputation. In particular, the reputation t in period τ is given by the following rule:

$$t_{\tau+1} = \begin{cases} b & \text{if } t_\tau = b \ \vee \ \delta\rho = 1 \\ g & \text{otherwise} \end{cases} \tag{1}$$

That is, bad reputation occurs after bad delivery ($\delta = 1$) plus negative feedback ($\rho = 1$), and then stays forever. Without explicitly modeling the decision of a buyer whether to trust the seller, we implicitly assume that a seller with a bad reputation is never trusted and thus receives zero payoffs in the current and all future periods.[9] Sellers with a good reputation, on the other hand, can always be certain to be trusted. This implies that the feedback mechanism offered by the trading platform is, in principle, an effective punishment device: if exerted properly, it can promote trustworthy sellers and trusting buyers. However, as we will see below, giving proper feedback may come at a cost hampering trade efficiency. Even if there are no costs, there might be multiple equilibria with different degrees of cooperation.[10]

According to Figure 1, giving negative feedback in the basic trade game implies a payoff d to the buyer. d can be negative, zero or positive. $d = 0$ implies that the choice of feedback, bad or good, is not payoff-relevant to the buyer. In fact, there are no charges and no rewards on eBay for giving feedback, and we would argue that the transaction costs of giving feedback on eBay is approximately zero. Thus, $d = 0$ would be a natural choice. However, giving negative feedback on eBay may involve an indirect cost because negative feedback runs the risk of being

[9] Another assumption here is that sellers cannot change their trading identity. See Friedman and Resnick (2001) and Ockenfels (2003) for an analysis of the economic implications of free identity change.

[10] Our assumptions regarding the effect of the reputation mechanism are strong but qualitatively in line with Resnick and Zeckhauser (2002), among many others, who found that sellers with relatively bad reputation have more difficulties to find buyers than sellers with relatively good reputation. Note that our evolutionary approach can be applied even if all sellers have a bad reputation when we assume small trembles in the sense of Selten (1983, 1988). The trembles make sure that the probability of a seller being trusted is always positive.

retaliated by 'reciprocal negative feedback' by the seller to the buyer (see Dellarocas et al., 2004, Klein et al., 2005, and the references cited therein for empirics and some theory). The expected cost of retaliation can be captured by $d > 0$ without explicitly modeling eBay's two-sided feedback mechanism.

On the other hand, $d < 0$ might be interpreted in terms of a motive to "altruistically punish" defectors. Fehr and Gächter (2002), among others, have shown experimentally that altruistic punishment of this sort is one of the key drivers of cooperation. To the extent the benefit from punishing defectors is a material, 'fitness-relevant' payoff component, it is relevant to our analysis. If there is no (fitness-relevant) altruistic punishment, or if the problem of retaliatory feedback is negligible or can be considered solved (see our conclusions), we are back to the case $d = 0$.

Finally, to complete the description of our model environment, buyer insurance is described by the payoff parameters E and Π. E is paid to all buyers who become the victims of bad delivery, and Π is the insurance premium that, in our formulation in Figure 1, has to be paid by all sellers. We note that the analysis below would go equally through if the premium is paid by the buyers, or if the premium is a function of the probability of bad delivery in the seller population. We also note that we will assume that the buyers' outside options are sufficiently low and the insurance E is sufficiently large so that, with buyer insurance, participation in the market is always profitable to buyers. This explains why we do not explicitly model the buyers' decision to trust; our focus is on the trade behavior of those who already decided to buy and sell on eBay.

3 Evolutionary analysis: equilibria and dynamics

We assume that in each of infinitely many trading periods τ, each individual plays with two anonymous partners, once in the seller and once in the buyer role. Within this one-population perspective, each trader can be described by his mode of behavior in the seller as well as in the buyer role. That is, after each period τ, each trader is described by his dispositions $\delta \in \{0,1\}$ and $\rho \in \{0,1\}$.

Evolutionary success is measured by the payoffs. We assume monotone evolutionary dynamics; strategies that yield higher payoffs increase their population share over time at the cost of the less successful strategies. Obviously, this approach is in line with every reasonable concept of evolutionary dynamics such as Replicator Dynamics. The following matrix shows the payoffs of the row-trader with profile (δ, ρ) when facing the corresponding column-trader. Payoffs are based on the basic trade game in Figure 1, and are averaged over the infinite time horizon and over both trader roles. We do not include the insurance premium Π here, because it must be paid by all sellers and is thus evolutionarily neutral. Note that, as bad reputation leads to being forever excluded from trade, payoffs for any trader with a bad reputation approach zero when the time horizon goes to infinity.

(δ, ρ)	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
$(0,0)$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1+2E}{4}$	$\frac{1+2E}{4}$
$(0,1)$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1+2E-2d}{4}$	$\frac{1+2E-2d}{4}$
$(1,0)$	$\frac{1+2q}{4}$	0	$\frac{q+E}{2}$	0
$(1,1)$	$\frac{1+2q}{4}$	0	$\frac{q+E-d}{2}$	0

(2)

Table 1: Evolutionary success

Let us define a *population state* (or strategy distribution) $\overrightarrow{x} = (x_1, x_2, x_3, x_4)$ as the description of the shares of the four pure strategies in the population. x_1 denotes the share of traders using strategy $(0,0)$, x_2 denotes the share of traders using strategy $(0,1)$, x_3 belongs to $(1,0)$, and x_4 to $(1,1)$. The set of population states is given by the three-dimensional simplex $X = \{ \overrightarrow{x} = (x_1, x_2, x_3, x_4) \in \mathbb{R}^4 : x_i \geq 0 \ (i = 1, ..4), \sum_{i=1}^{4} x_i = 1\}$.

3.1 Multiplicity of equilibria and the effect of feedback costs (d)

The following proposition shows that the gains from trade in any evolutionarily stable rest point are determined by the feedback costs d - independent of buyer insurance.

Proposition 1 *(i) If* $d < 0$*, there is a continuous set of neutrally evolutionarily stable states (NSS) in game (2) given by* $\{\overrightarrow{x} \in X : x_1 < \frac{2}{2q+1}, x_2 > \frac{2q-1}{2(q+E)}, x_3 = x_4 = 0\} =: \Theta$*. Note that, regardless of* E*,* Θ *only includes trustworthy sellers.*

(ii) If $d = 0$*, there are two sets of NSS:* Θ *and the continuous set* $\Phi := \{\overrightarrow{x} \in X : x_1 = x_2 = 0, x_3 > \frac{1+2E}{2(q+E)}, x_4 < \frac{2q-1}{2(q+E)}\}$*. Note that, regardless of* E*,* Φ *only includes untrustworthy sellers.*

(iii) If $d > 0$*, there is a unique evolutionarily stable state (ESS) given by* $(1, 0)$*. In addition, all population states included in the set* Θ *are neutrally stable (NSS).*

Proof. Recall that a strategy $x^* \in X$ is an evolutionarily stable strategy (ESS) if for every alternative strategy $x \neq x^*$ the following two conditions hold: (i) $u(x, x^*) \leq u(x^*, x^*) \forall x \in X$ and (ii) $u(x, x^*) = u(x^*, x^*) \Rightarrow u(x, x) < u(x^*, x) \ \forall x \neq x^*$. $u(x, x^*)$ is the payoff (fitness) of a player using strategy x against x^*. NSS is a weaker requirement in that the latter condition has to hold only with weak inequality (Maynard-Smith 1974). Consider first the case $d < 0$. Here all strategies contained in Θ are NSS, as the only strategies (except for those contained in Θ itself) satisfying condition (i) involve proper delivery and thus yield the same payoff as the strategies in Θ when used in a population of cooperative traders (condition (ii)). These strategies are not NSS, though, as they can be invaded by either $x_3 = 1$ or $x_4 = 1$. On the other hand, strategies with $x_2 > 0$ (i.e., those placing positive probability on $(0, 1)$) can be invaded by $x_4 = 1$, and strategies involving $x_4 > 0$ can be invaded by $x_1 = 1$ or $x_2 = 1$ (check condition (i)). In case $d = 0$ all strategies have a duplicate strategy.[11] Consequently, no ESS exists. All elements of Θ and Φ are NSS, as

[11] Strategies $\widehat{s}$ and $\widetilde{s} \neq \widehat{s}$ are duplicate strategies iff $u(\widehat{s}, s) = u(\widetilde{s}, s) \forall s \in S$.

can be proved analogously to the case $d < 0$. Finally, if $d > 0$, there is a unique ESS given by $x_3 = 1$, because $x_3 = 1$ yields strictly higher payoffs against itself than any of the other strategies yield against $x_3 = 1$. Condition (i) implies that there is no other ESS. The strategies contained in Θ are neutrally stable, as both strategies $x_3 = 1$ and $x_4 = 1$ yield strictly lower payoffs than the equilibrium strategy (from the set Θ) in a population, where all players play this equilibrium strategy (condition (i)). Analogously to case $d < 0$, there are no other NSS. ∎

Proposition 1 implies that buyer insurance cannot affect trade efficiency in equilibrium, which is rather shaped by the cost d of giving negative feedback to untrustworthy sellers. If $d < 0$, there is a unique set of NSS, which involves efficient trade. Efficiency is stabilized by a sufficiently strong threat of punishing untrustworthy sellers. That is, a motive for punishing bad behavior, as measured by d, yields a robust cooperative equilibrium - as long as d is fitness-relevant.[12] Buyer insurance (E) may somewhat affect the buyers' average willingness to punish untrustworthiness (as it affects the share of x_2 in the equilibrium set Θ). However, the effect is small in the sense that buyer insurance cannot influence seller behavior or trade efficiency.

If $d = 0$, there are two sets of equilibria. In one set (Θ), all sellers deliver properly and a sufficient number of buyers are willing to punish untrustworthiness. In the other one (Φ), no seller is trustworthy and only few buyers are willing to leave negative feedback. As before, buyer insurance may affect the share of buyers that is willing to leave negative feedback, but in each equilibrium set it leaves the efficiency of trade unaffected. However, for $d = 0$, buyer insurance plays an additional role, as it shapes the basins of attraction of each of the two equilibrium sets. This important role will be described along with the implications for equilibrium selection in the next subsection.

[12]Within the so-called "indirect evolutionary analysis," the result could also be obtained even if $d < 0$ would be merely a preference parameter, not associated with any material payoffs; see Güth (1995) for the indirect evolutionary approach and Güth and Ockenfels (2000, 2003 and forthcoming) for applications in trust games.

If $d > 0$, there is a unique ESS, in which cooperation of both sellers (trustworthiness) and buyers (punishment of untrustworthiness) completely collapses. There are also NSS, in which proper delivery can be stabilized by the threat of punishment. However, as we show in the next section, all equilibria but the (maximally inefficient) ESS fail to be robust.[13] So, as in case $d = 0$, buyer insurance cannot influence trade efficiency.

3.2 Basins of attraction and the effect of buyer insurance (E)

While buyer insurance does not affect trade efficiency in equilibrium, it may affect the path of play and thus the selection of an equilibrium in case of multiple possible outcomes of evolutionary selection. According to Proposition 1, there are multiple equilibria both for $d = 0$ and for $d > 0$.

In this subsection, we rely on the standard dynamic model of evolutionary game theory, the Replicator Dynamics (Taylor and Jonker 1978), in order to analyze the "attractiveness" of the different equilibria in a dynamic setting. The two propositions that we prove demonstrate that only in case of $d = 0$ can buyer insurance influence the dynamics in a way that affects trade efficiency.

Proposition 2 *For* $d > 0$ *the only asymptotically evolutionarily stable equilibrium of the Replicator Dynamics is the ESS* $(1, 0)$*. The other equilibria fail to be attractive.*

Proof. After an affine transformation of payoffs, the dynamic system representing the Replicator

[13] Take, e.g., the momorphism (0,1) with full cooperation. Whenever there is a positive probability that delivery is improper, those buyers who do *not* punish fare better (because of $d > 0$), so that cooperation is driven out in the presence of small trembles.

Dynamics is given by

$$\begin{aligned}
\dot{x}_1 &= [2(x_1+x_2)+(1+2E)(x_3+x_4)-\overrightarrow{x}^T A\overrightarrow{x}]x_1 \\
\dot{x}_2 &= [2(x_1+x_2)+(1+2E-d)(x_3+x_4)-\overrightarrow{x}^T A\overrightarrow{x}]x_2 \\
\dot{x}_3 &= [x_1(1+2q)+2x_3(q+E)-\overrightarrow{x}^T A\overrightarrow{x}]x_3 \\
\dot{x}_4 &= [x_1(1+2q)+2x_3(q+E-d)-\overrightarrow{x}^T A\overrightarrow{x}]x_4
\end{aligned}$$

where A denotes the payoff matrix (2). Asymptotic stability of $x_3 = 1$ (i.e. the monomorphism, where all players play $(1,0)$) in the Replicator Dynamics follows directly from the fact that it is evolutionarily stable (Hofbauer, Sigmund and Schuster 1979). The only additional rest points (apart from $x_3 = 1$) of the Replicator Dynamics that are Nash-equilibria of the underlying 4×4 game are given by the set $\Theta = \{\overrightarrow{x} \in X : x_1 < \frac{2}{2q+1}, x_2 > \frac{2q-1}{2(q+E)}, x_3 = x_4 = 0\}$. These are the only candidates for asymptotic stability. While the Jacobian matrix corresponding to the mixed strategy equilibrium $\overrightarrow{x} = (2/(2q+1), (2q-1)/(2q+1), 0, 0)$ is indefinite - so that the equilibrium is unstable - it is negative semi-definite for all other equilibria in this set. Consequently, while these equilibria can be Lyapunov stable, they are not attractive and therefore not asymptotically stable.

■

Proposition 2 implies that, with or without buyer insurance, there is no hope that players cooperate given a positive cost d of leaving negative feedback. For $d = 0$, on the other hand, buyer insurance may assume a decisive role. Note first that in this case the profiles $(\delta, \rho) = (0,0)$ and $(0,1)$ as well as $(1,0)$ and $(1,1)$, respectively, are duplicate strategies which always yield the same payoffs. Because evolutionary drift cannot select between duplicate strategies, we fix the proportions of these types and conduct the analysis conditional on these proportions.[14]

[14] In principle, the actual proportion of agents who are reliably and truthfully willing to punish improper conduct, in the absence of any material incentive to do so, can be empirically or experimentally determined, and might differ

Denote agents with $\delta = 0$ as $0-$types, and agents with $\delta = 1$ as $1-$types. Let $x \in [0,1]$ be the proportion of $0-$types that is *not* giving negative feedback ($\rho = 1$), and $y \in [0,1]$ be the corresponding proportion of $1-$types. Then, average payoffs are described by the following 2×2 matrix:

	$0 - type$	$1 - type$
$0 - type$	$\frac{1}{2}$	$\frac{1+2E}{4}$
$1 - type$	$x\frac{1+2q}{4}$	$y\frac{q+E}{2}$

Table 2: Evolutionary success for fixed proportions x and y

Let $z \in [0,1]$ be the proportion of $0-$types in the population, and denote fitness of the two types by $R_0(z)$ and $R_1(z)$, respectively. Then a necessary and sufficient condition for 0 to be an absorbing type is[15]

$$R_0(1) > R_1(1) \Leftrightarrow x < \frac{2}{1+2q} \tag{3}$$

Analogously for 1 to be an absorbing type, we need

$$R_1(0) > R_0(0) \Leftrightarrow \frac{1+2E}{2q+2E} < y \tag{4}$$

Let

$$z^* = \frac{1+2E-2y(E+q)}{2(E+q)(1-y)-(1-x)}$$

be the share that solves $R_1(z) = R_0(z)$. Then we can state the following proposition:

across communities and cultures.

[15] Note that these treshholds just coincide with the shares of "punishers" in the mixed strategy states that delineate the sets Θ and Φ from Proposition 1.

Proposition 3 *Let $d = 0$.*

(i) If both (3) and (4) hold, then the basin of attraction of $z = 0$ is given by $[0, z^*]$, and the basin of attraction of $z = 1$ is given by $[z^*, 1]$.

(ii) If (3) holds but not (4), $z = 1$ is globally stable (with basin of attraction $[0, 1]$).

(iii) If (4) holds but not (3), $z = 0$ is globally stable.

(iv) If neither (3) nor (4) hold, then $z = z^*$ is globally stable.

Proof. Assuming a payoff-monotonic evolutionary dynamics, like the Replicator Dynamics, we know that $z = 1$ is locally stable iff (3) holds, and $z = 0$ is stable iff (4) holds. In addition, as any interior zero of the Replicator Dynamics has to be a Nash-equilibrium of the underlying game, there can be at most one interior stable equilibrium. As the Replicator Dynamics is continuous on the interior of the strategy space, Proposition 3 follows. ■

Using

$$\frac{\partial z^*}{\partial E} = -\frac{2(-2 + 2q + x)(-1 + y)}{(-1 + 2E + 2q + x - 2(E + q)y^2)^2} \begin{cases} < 0 & \text{iff } x < 2(1 - q) \\ \geq 0 & \text{otherwise} \end{cases},$$

Figure 2 illustrates the impact of buyer insurance on the basin of attraction of the efficient equilibrium at $z = 1$. $\widehat{z}$ denotes the situation without buyer insurance ($E = 0$) and z^* the situation with a buyer insurance ($E > 0$). The thick arrows indicate the direction of change of z when introducing a buyer insurance (according to the derivative above). The thin arrows indicate the respective evolutionary drifts (according to Proposition 3).

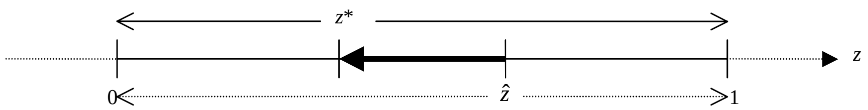

Fig. 2a: $z = 0$ and $z = 1$ locally stable.

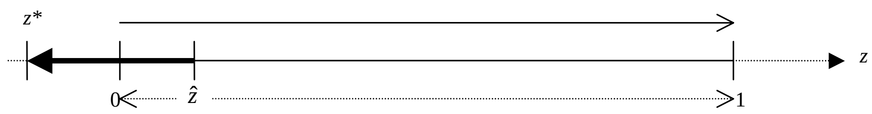

Fig.2b: $z = 1$ globally stable.

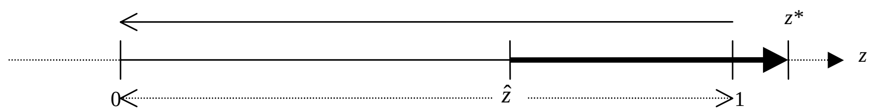

Fig.2c: $z = 0$ globally stable.

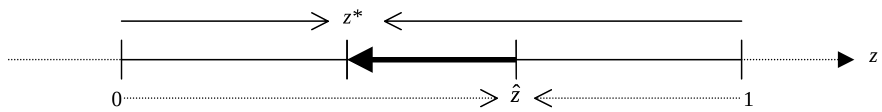

Fig.2d: Interior Equilibrium.

Figure 2 shows that in cases (i) and (ii) of Proposition 3, in which $z = 1$ (efficient trade) is an absorbing state, there is a crowding in by insurance as measured by the basin of attraction for the 0-types. In the other cases, (iii) and (iv), in which $z = 1$ is not an absorbing state, there is crowding out. In particular, in the unique mixed equilibrium population of case (iv), efficiency is smaller with than without buyer insurance.

In other words, buyer insurance can only help in those cases, where efficient trade is a stable state *regardless* of whether there is a buyer insurance or not.[16] In those cases, depending on the population composition z, the introduction of buyer insurance may 'push' the traders into the basin of attraction of the efficient equilibrium. However, this situation does not seem to reflect eBay's situation. Starting from a rather cooperative population (see Cohen 2002 for the beginnings of eBay), fraud is becoming an increasing concern. This suggests that efficient trade is *not* an absorbing state. Rather, eBay appears to be either in case (iii) or (iv), in which z close to one

[16]According to condition 3, this is the case when there are sufficiently many traders who are cooperative in both roles.

is *pulled down* to less efficient states. In particular, case (iv), the stable bimorphism in which trustworthy and untrustworthy traders coexist, appears to be a natural description of eBay. But introducing a buyer insurance in this situation does not only do no good, but rather countervails efficiency and cooperative behavior in the long run.

4 Conclusions

In this study, we ask how buyer insurance and seller reputation interact to reverse a decline of cooperation and proper conduct in online markets such as eBay. Our evolutionary analysis suggests that a reputation mechanism that elicits truthful feedback would be very effective in promoting efficient trade - independent of the existence of buyer insurance. Assuming that bad reputation has sufficiently strong negative payoff consequences, even a small reward from punishing defectors yields a unique and fully efficient equilibrium. Buyer insurance, on the other hand, has no impact on trade efficiency in equilibrium. This implies that if there is no efficient equilibrium without buyer insurance, there will not be any with buyer insurance. Rather, buyer insurance can affect the basins of attraction of equilibria that differ with respect to trade efficiency. Yet in those important cases, in which efficient trade is *not* an absorbing state, or in which trustworthy and untrustworthy sellers coexist in a stable bimorphism, buyer insurance can make things only worse. This is because, in the long run, insurance renders improper seller behavior more 'attractive.'

One more general implication of our study is that certain design and policy choices that are meant to promote trust, such as buyer insurance, may backfire. There are a number of policies that may potentially fall into this category. In particular, eBay sometimes makes it unecessarily difficult to get a realistic sense of the risks involved in buying on its platform. For one, eBay does not publish aggregate data on fraud conducted on its platform.[17] Second, as mentioned earlier, by

[17]Quite to the opposite, eBay removed the link to fraud.org from its site in 2003, so that the number of auc-

allowing retaliatory feedback, reputation information looks better than it actually is.[18] Finally, by not allowing buyers to search for negative feedbacks regarding a particular seller, eBay makes it difficult to asses the trustworthiness of a seller.[19] Policies and design choices like this may, in the short run, promote trust. In the long run, however, they may well aggravate the trustworthiness problem and thus undermine mutually profitable trade.

References

[1] Bajari, P. and A. Hortaçsu, (2004). Economic Insights from Internet Auctions. Journal of Economic Literature, XLII, 457-486.

[2] Bolton, Gary, Elena Katok, and Axel Ockenfels (2004). How Effective are Online Reputation Mechanisms? An Experimental Investigation. Management Science, 50(11), 1587-1602.

[3] Cohen, Adam (2002). The Perfect Store: Inside EBay. Little, Brown, New York.

[4] Dellarocas, C. (forthcoming a). The Digitization of Word-of-Mouth: Promise and Challenges of Online Reputation Mechanisms. Management Science.

[5] Dellarocas, C. (forthcoming b). Reputation Mechanisms. Handbook of Information Systems and Economics.

[6] Dellarocas, C., M. Fan, and C. A. Wood (2004). Self-Interest, Reciprocity, and Participation in Online Reputation Systems. Working Paper.

tion complaints reported to NCL has dropped to one-sixth the previous level after the removal of eBay's link (http://www.fraud.org/internet/internet_scams_report-jan-sept_2005.pdf, December 2005).

[18] This is reflected in our model which involves positive costs of giving negative feedback. Of course, there are various ways to avoid the costs associated with retaliatory feedback. One is not to allow sellers to assess buyers. Another one is a "blind period," in which traders give feedback on each other, which is only revealed after the blind period has ended (see, e.g., www.rent-a-coder.com). Rewards for reliable feedback may come from scoring systems that induce honest reporting of feedback (Miller et al. forthcoming).

[19] Recently, eBay has implementend this search option in selected countries, but not in the U.S.

[7] Fehr, Ernst, and Simon Gächter (2002). Altruistic Punishment in Humans. Nature 415, 137-140.

[8] Friedman, E. and P. Resnick (2001). The Social Cost of Cheap Pseudonyms. Journal of Economics and Management Strategy, 10(2), 173-199.

[9] Güth, Werner (1995). An Evolutionary Approach to Explaining Cooperative Behavior by Reciprocal Incentives. International Journal of Game Theory, 24, 323-344.

[10] Güth, Werner, and Axel Ockenfels (2000). Evolutionary Norm Enforcement. Journal of Institutional and Theoretical Economics, 156(2), 335-347.

[11] Güth, Werner, and Axel Ockenfels (2003). The Coevolution of Trust and Institutions in Anonymous and Non-anonymous Communities. In: M.J. Holler, H. Kliemt, D. Schmidtchen and M. Streit (Eds.), Jahrbuch für Neue Politische Ökonomie, 20, 157-174, Tübingen: Mohr Siebeck.

[12] Hofbauer,J., P.Schuster, and K.Sigmund (1979). A Note on Evolutionary Stable Strategies and Game Dynamics. Journal of Theoretical Biology, 81, 609-612.

[13] Jin, Ginger Z. and Andrew Kato (2004). Consumer Frauds and the Uninformed: Evidence From an Online Field Experiment. University of Maryland Working Paper.

[14] Klein, Tobias J., Christian Lambertz, Giancarlo Spagnolo, Konrad O. Stahl (2005). Last Minute Feedback, Consip Working Paper VI/2005.

[15] Resnick, P. and Zeckhauser, R. (2002). Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. In Michael R. Baye (Ed.), The Economics of the Internet and E-Commerce, Advances in Applied Microeconomics, 11, Amsterdam, Elsevier Science.

[16] Maynard - Smith, J. (1974). The Theory of Games and the Evolution of Animal Conflicts. Journal of Theoretical Biology, 47, 209-221.

[17] Miller, Nolan, Resnick, Paul, and Richard Zeckhauser (forthcoming). Eliciting Honest Feedback: The Peer Prediction Method. Management Science.

[18] Ockenfels, Axel, and David Reiley (forthcoming). Online Auctions. Handbook of Information Systems and Economics.

[19] Ockenfels, Axel (2003). Reputationsmechanismen auf Internet-Marktplattformen: Theorie und Empirie. Zeitschrift für Betriebswirtschaft, 73(3), 295-315.

[20] Selten, Reinhard (1983). Evolutionary Stability in Extensive Two-person Games. Mathematical Social Sciences, 5, 269-363.

[21] Selten, Reinhard (1988). Evolutionary Stability in Extensive Two-person Games: Corrections and Further Developments. Mathematical Social Sciences, 16, 223-266.

[22] Taylor, P. and L. Jonker (1978). Evolutionary Stable Strategies and Game Dynamics. Mathematical Biosciences, 40, 145-156.

[23] Vishwanath, Arun (2004). Manifestations of Interpersonal Trust in Online Interaction. A Cross-cultural Study Comparing the Differential Utilization of Seller Ratings by eBay Participants in Canada, France, and Germany. News Media and Society, 6(2), 219-234.