Communication, Risk and Efficiency in Games*

Andreas Blume
Department of Economics
University of Iowa

April, 1996

Abstract

This paper studies the evolution of effective pre-play communication in games where a single communication round precedes a simultaneous-move, complete-information game. The paper identifies stable outcomes under population learning dynamics in which individuals with some probability replace their current strategy with a best reply against beliefs supported on a sample of currently used strategies. It is shown that under these conditions the effectiveness of one-sided pre-play communication is inversely related to risk in the underlying game, and to the size of the message space. Multi-sided communication can be shown to be more effective than one-sided communication; i.e., risk and the size of the message space play no role. This requires that all players communicate, have the same preferred equilibrium and messages have some small a priori information content that identifies message profiles that signal agreement on a strict equilibrium in the underlying game.

*I am grateful to Eddie Dekel, Roger Lagunoff, Joel Sobel, Eric van Damme and two referees for helpful comments and suggestions. I have benefitted from comments by seminar participants at Northwestern University, the Midwest Mathematical Economics Meetings (University of Michigan) and the Econometric Society Winter Meetings (Washington D.C.). Support by the National Science Foundation under award number SBR-9410588 is gratefully acknowledged.

1 Introduction

This paper studies pre-play communication in games where a single communication round precedes a simultaneous-move, complete-information game. In the communication round, a subset of players send a message from a finite message space. Messages are costless and have no (or limited) a priori meaning. The main concern of the paper is whether meaningful communication can evolve endogenously under a large class of dynamic rules. We show that this is the case, even if the dynamics are not restricted to be gradual. In addition, the departure from gradual dynamics reveals two issues that are novel in the evolutionary literature on pre-play communication; these are a role for risk in the underlying game, and for the size of the message space. Both factors potentially hinder the evolution of effective pre-play communication.

The result, that Pareto efficiency of an equilibrium alone is no guarantee for it being reached via pre-play communication has an interesting precedent outside of an evolutionary formulation. Aumann [1990] has argued that pre-play communication may not lead to Nash equilibrium, even if the underlying game has a unique strict and efficient equilibrium.[1] He illustrates his argument with a version of the familiar *Stag Hunt* game, in which there is a tension between Pareto dominance and Harsanyi and Selten's [1988] risk dominance. It is interesting that the concern about the effectiveness of pre-play communication in this game can be made operational in a setting where messages have no a priori meaning.

The evolutionary literature on pre-play communication, as exemplified by Bhaskar [1992], Fudenberg and Maskin [1991], Kim and Sobel [1993], Matsui [1991], Sobel [1993] and Wärneryd [1991], has shown that meaningful communication can evolve endogenously. Roughly, this approach envisions a large population of players who are repeatedly and randomly matched to play a given communication game. Players gradually adjust their strategies in the direction of successful ones. If the players' interests are sufficiently closely aligned, they will learn over time to communicate successfully.

The evolutionary approach to pre-play communication thus far does not distinguish among equilibria according to their risk. In games with common interests, the common interest outcome cannot be destabilized through evolutionary forces. Intuitively this is so because the evolutionary process moves the population via

[1] Farrell [1988] argues that pre-play communication need not lead to a Nash equilibrium even in a game with a unique Nash equilibrium.

mutations that affect only a small fraction of the entire population. Since the mutant population is small, the strategic problem faced by the mutants in the post-entry population reduces essentially to an optimization problem. The mutants face no significant strategic uncertainty. This point is expressed most clearly in Sobel [1993] where one player at a time gets to adjust his strategy while the strategies used by the rest of the population remain unchanged.

The present paper refers to the same scenario, a large population of players who are repeatedly and randomly matched. The paper departs from the above mentioned evolutionary studies in not postulating that the population dynamics are gradual. It does not rule out the possibility that large fractions of the population adjust their strategies simultaneously or that some players try to anticipate the population dynamics.

We propose a class of population learning dynamics where each individual with some probability replaces its current strategy with a best reply to beliefs that are supported on a sample of currently used strategies. The stable outcomes of such dynamics can be conveniently characterized in terms of subsets of the strategy space. *Curb (closed under rational behavior) retracts* (Basu and Weibull [1991]) are minimal sets of strategies closed under inclusion of best replies; they are convex and spanned by pure strategies. Curb retracts exist in every game and always contain a *(curb) equilibrium.*[2] We will show that if the dynamics reflect multiple levels of depth of reasoning and cautious behavior among members of the population, then in games with one-sided communication, curb equilibria will be observed with high frequency. With multi-sided communication curb retracts have no predictive power, unless we permit some prior differentiation of messages. With such differentiation, curb retracts consist entirely of equilibria.

The paper contains two major results. The first deals with two-player games in which only one of the players can send a message. Assume that in the underlying game there is a unique strategy combination that maximizes the communicating player's payoff. Let this strategy combination be a strict equilibrium, i.e. its own unique best reply. Fix the size of the message space and the number of strategies in the underlying game. Then the payoffs associated with the communicating player's favorite equilibrium will be the only curb equilibrium payoffs in the communication game, provided the communicating player's *risk* at that equilibrium is sufficiently low; with a message space of size two, the appropriate risk measure

[2]Hurkens [1995] examines an alternative dynamic whose stable outcomes can also be characterized via curb set. His dynamic is inspired by Young's [1993] learning dynamic.

is related to, and in the case of a symmetric 2×2-game, coincides with Harsanyi and Selten's [1988] definition of risk-dominance for 2×2-games.

The second major result concerns multi-sided communication in N-player games in which all players can talk and messages have *limited information content (LIC). LIC* is modelled via an (arbitrarily) small variation in payoffs in the communication game that links message profiles to equilibria in the underlying game.[3] We suggest a way to amend our dynamics that would generate such message space differentiation endogenously.

Without *LIC,* the curb condition has no predictive power in games with multi-sided communication. With *LIC,* the curb concept distinguishes two-sided from one-sided communication. *LIC* does not affect the results with one-sided communication. However, with two-sided communication, and a unique strict common interest equilibrium in the underlying game, only the payoffs of that equilibrium are curb equilibrium payoffs in the communication game.

The paper is organized as follows. The next section discusses two examples to motivate the results on one-sided communication. Section 3 introduces the dynamics and characterizes their stable outcomes. Sections 4 and 5 deal with one- and multi-sided communication, respectively. Section 6 concludes with a discussion of the literature and some thoughts about the appearance of risk dominance here and elsewhere in the literature.

2 Examples: Dodo and Stag Hunt

This section examines the evolution of effective pre-play communication for two underlying games. The examples show that risk in the underlying game matters under a simple adaptive rule.

For future reference it is useful to recall Harsanyi and Selten's definition of risk dominance. We will confine ourselves to symmetric games to simplify the exposition. Consider the game $\mathbf{G_0}$ below and assume that $a > c$ and $b < d$ such that (U,L) and (D,R) are two strict Nash equilibria. According to Harsanyi and Selten (U,L) risk dominates (D,R) if $a - c > d - b$ and vice versa, where $a - c$ $(d - b)$ is the deviation loss associated with the risk dominant (risk dominated) equilibrium.[4]

[3]Examples in the literature where communication games are analyzed after adding such small payoff variations include Blume, Kim and Sobel [1993] and Hurkens [1993].

[4]It is also true that (U,L) is risk dominant exactly when the probability of the column player

	L	R
U	a,a	b,c
D	c,b	d,d

$\mathbf{G_0}$

An example of such a game is $\mathbf{G_1}$, which Binmore calls *Dodo.* Assume for the moment that players can rely on a commonly understood language. Then, it is commonly accepted (see for example Farrell and Rabin [1996] and the references therein) that we would expect the row player's announcement "I will play U" to be believed by the column player. After all, the message is *self-signaling*; i.e., the row player wants it to be believed if and only if it is true. It is also *self-committing,* because conditional on it being believed the row player prefers to conform to his announcement.

	L	R
U	3,3	0,0
D	0,0	1,1

$\mathbf{G_1}$

	L	R
U	9,9	0,8
D	8,0	7,7

$\mathbf{G_2}$

Game $\mathbf{G_2}$ is a version of Rousseau's *Stag Hunt.* Just like in game $\mathbf{G_1}$ players have common interests; i.e., there is a unique efficient payoff vector. Nevertheless Aumann [1990] has used the *Stag Hunt* game to argue that pre-play communication need not lead to Nash equilibrium. Suppose that without pre-play communication players are (say by convention) coordinated on the inefficient (D,R) equilibrium. This could be the result of payoffs having changed over time; e.g., the Pareto dominant equilibrium could correspond to switching to a technology that has recently improved. Would it help if a player were able to make a pre-play announcement? Not necessarily. Presumably the row player wants to convince the column player to play L and thus might announce to play U herself. The problem is that he wants to induce the column player to play L no matter what he intends to play herself. If he is not completely convinced that his announcement is successful, he may well play D. If that possibility is given sufficient weight by the column player, he will play R himself. According this argument pre-play communication is not successful in the *Stag Hunt* game because a message does not

playing L which makes the row player indifferent between U and L is less than $1/2$.

reveal anything about the intentions of the announcer; the message "I will play U" is not self-signaling.

Aumann's analysis of his example suggests one reason for why pre-play communication need not be effective in games with multiple Pareto-ranked equilibria. It is less clear what conclusions to draw from it about the role of pre-play communication in general. There are at least four reasons for reexamining the difficulty with pre-play communication in Aumann's example. Two of these can be illustrated by varying the stag hunt game, G_3 and G_4 below.

	L	R
U	9,9	6.9,8
D	8,6.9	7,7

G_3

	L	R
U	9,9	0,6
D	6,0	7,7

G_4

First, Aumann's argument does apply to game G_3 and yet its force seems to be diminished in this game. U is optimal for player 1 for a wide range of beliefs, which lessens the burden on communication to discredit the (D,R) equilibrium.[5] Second, there may be reason's for communication to be ineffective, even if Aumann's argument does not apply; this is illustrated by game G_4. Having made an utterance such as "I intend to play U," in game G_4, player one must have a high degree of confidence in its effectiveness before following his own recommendation because D is a best reply against a large set of beliefs. Third, Farrell and Rabin [1996] insist even in Aumann's original example that the *self-committing* property of the message "I will play U" is enough to ensure its effectiveness. Fourth and related to the last point, the distinction made between the roles of pre-play communication in *Dodo* vs *Stag Hunt* is informal and relies on an interpretation of the status of a commonly understood language.

This paper attempts to formalize the concern about the effectiveness of pre-play communication in Aumann's example without assuming that messages have an *a priori* meaning; rather, like the evolutionary literature, we ask under what conditions meaningful communication can evolve endogenously, and show that these conditions are related to risk in the underlying game. The above examples show that there is no containment relationship between risk considerations and

[5] According to Aumann [1990] David Kreps has a raised a similar point. Aumann argues, the fact that the efficient outcome appears likely in this game should not be attributed to a role for communication.

Aumann's argument. And yet, the two are related; in G_0, let $a > d > b$, and $a > c$. By raising c, we both increase the risk of U and make it more likely that Aumann's argument applies; loosely speaking both Aumann's argument and the risk argument are monotonic in c.

Consider two large finite equal-sized populations of communicating and silent players. They repeatedly play the following communication game. In each period each of the communicating players is randomly matched with a silent player. The communicating player sends one of two messages, m_1 or m_2, and then both players play the underlying game, which is either $\mathbf{G_1}$ or $\mathbf{G_2}$. The reduced normal forms of the communication games corresponding to *Dodo* and *Stag Hunt,* denoted by $\blacksquare_1$ and $\blacksquare_2$ are shown below.

	LL	*LR*	*RL*	*RR*
(m_1,U)	3,3	3,3	0,0	0,0
(m_1,D)	0,0	0,0	1,1	1,1
(m_2,U)	3,3	0,0	3,3	0,0
(m_2,D)	0,0	1,1	0,0	1,1

$\blacksquare_1$

	LL	*LR*	*RL*	*RR*
(m_1,U)	9,9	9,9	0,8	0,8
(m_1,D)	8,0	8,0	7,7	7,7
(m_2,U)	9,9	0,8	9,9	0,8
(m_2,D)	8,0	7,7	8,0	7,7

$\blacksquare_2$

Suppose play evolves according to a simple adaptive rule. In each period each player uses last period's strategy with probability one-half or else moves to a best reply against last period's population play; if there are multiple pure best replies, he randomizes uniformly.

What can be said about the long run properties of this dynamic process in either game? Consider first Γ_1. One can show that irrespective of the initial distribution of population play, the postulated dynamic process will almost surely converge to the set Q of strategies consisting of all mixtures of (m_1,U) and (m_2,U) for the row player and of all mixtures of *LL*, *LR* and *RL* for the column player. To see this note that from any initial state, there is positive probability, bounded away from zero, that in two steps the population moves to a state where all players of a

given kind use the same strategy. From any such state, there exists a sequence of best-reply iterations which takes the population into Q in no more than two steps; these sequences all have probability bounded away from zero. Thus, from any initial state, the probability that after four periods Q is not reached is bounded away from one. This implies that from any initial state, Q is reached almost surely. Furthermore, once the process has entered Q it can never leave Q because strategies outside of Q are never best replies against beliefs concentrated on Q.

Next consider the set Q in the game Γ_2. The same argument as above shows that from any initial state, Q is reached almost surely. However, for Γ_2, the set Q is not invariant under the dynamic process. To understand this observe that there is a bounded number of steps, such that from any initial state in Q the population moves to a state where half of the silent players use the strategy *LR* and the other half uses the strategy *RL*. At that point none of the strategies in Q are a best reply for the communicating players, and the process exits Q with positive probability.

In the next section we characterize the stable outcomes of a class of dynamics in the spirit of this example. We will use this characterization in the subsequent sections to analyze the roles of message space size, risk and message differentiation in communication games

3 Dynamics

This section proposes a class of population learning dynamics for finite games and characterizes its stable outcomes in terms of the curb retracts of those games.

For a finite strategic form game G with player set $\mathcal{P}$ and $P = \#(\mathcal{P})$ let S_p be player p's set of pure strategies, with typical element s_p and let Σ_p be his set of mixed strategies, with typical element σ_p. S and Σ are the sets of pure and mixed strategy profiles. σ is a strategy profile and σ_{-p} a partial profile that excludes player p's mixed strategy. Denote the convex hull of any set Z by $\mathrm{co}(Z)$. $\mathrm{MBR}_p(\cdot)$ is player p's (mixed) best reply correspondence and $\mathrm{MBR}(\sigma) = \times_{p=1}^{P}\mathrm{MBR}_p(\sigma_{-p})$. $Q \subseteq \Sigma$ is a retract if $Q = \times_{p=1}^{P} Q_p$ where $Q_p \subseteq \Sigma_p$ is nonempty, closed and convex. A retract Q is a *curb retract* if for all $\sigma \in Q$, $\mathrm{MBR}(\sigma) \subseteq Q$, and if Q is a minimal retract with that property. A Nash equilibrium contained in a *curb retract* is called a *curb equilibrium.*

In every finite game there is at least one *curb retract* and a *curb equilibrium. Curb retracts* are spanned by pure strategies; i.e., with any mixed strategy $\sigma_p \in Q_p$, Q_p contains all pure strategies in the carrier of σ_p. Therefore we can (and will)

identify any curb retract with the set of pure strategies that span the retract. The intersection of two retracts closed under inclusion of best replies is itself closed under inclusion of best replies.[6]

Curb retracts are attractive because they conveniently characterize the stable outcomes for a large class of learning dynamics. We propose a class of *population learning dynamics* for finite games played between multiple populations of players. We show that every dynamic in this class converges almost surely to a set of states that is supported on a curb retract and that from states supported on a curb retract every other state supported on the same curb retract is reached with positive probability in finite time. Those dynamics in this class in which players are cautious and reason at sufficient depth have an additional property: If a player has a dominant strategy in a curb retract, then conditional on reaching that set individuals representing that player will use the dominant strategy with high probability.

Consider any finite game G with player set $\mathcal{P}$ whose cardinality we denote by P. $\mathcal{P}$ also denotes the set of populations from which individuals are drawn at random in any given period to play the game G. Let S_p denote the finite set of strategies available to a player from population $p \in \mathcal{P}$, and $S := \times_{p\in\mathcal{P}} S_p$. Assume that $\#(p) > \#(S_p)$ for every population p. For every set of strategies $X_p \subseteq S_p$ let $\Delta(X_p)$ stand for the set of probability distributions with support X_p, and for $X \subseteq S$, let $\Delta(X) := \times_{p\in\mathcal{P}}\Delta(X_p)$ indicate the set of uncorrelated distributions supported on X, and let $\Delta^0(X)$ denote all those uncorrelated distributions on X that have full support. Let $BR_p(\Delta(X))$ denote the set of player p's pure best replies to uncorrelated beliefs supported on X and define $BR(\Delta(X)) := \times_{p\in\mathcal{P}} BR_p(\Delta(X))$.

At each time t the state of population $p \in \mathcal{P}$ is the vector $\omega_{pt} = \{s_{it}\}_{i\in p}$ of strategies adopted by each member i of population p. The state $\omega_t = \{\omega_{pt}\}_{p\in\mathcal{P}}$ at time t is the concatenation of all time t population states. Write Ω for the set of all possible states.

The dynamic process on Ω follows a transition rule $\phi(\cdot|\cdot)$ where $\phi(\omega_{t+1}|\omega_t)$ is the probability that the state in period $t+1$ will be ω_{t+1} given that it is ω_t in period t. The transition rule is a function of individual agent behavior. Agents enter period t with their adopted strategy $s_{i,t}$ and use that strategy to play the game G with various random selections of players from other populations. Through these random matches agents gain information about the strategies currently in use by members of the different populations. With probability λ, $1 > \lambda > 0$, an individual

[6]For more on curb sets see Basu and Weibull [1991], Balkenborg [1992] and Hurkens [1993].

is active and uses his information to update his strategy, with probability $1-\lambda$ he is inactive and adopts last period's strategy, in which case $s_{i,t+1} = s_{i,t}$.

Denote by $X_{itp'}$ the set of pure strategies that an active individual $i \in p$ observes at time t among individuals in population p'. This will be a subset of the strategies currently in use in that population. Call $X_{it} := \times_{p' \in \mathcal{P}} X_{itp'}$ individual i's sample at time t. Active individuals adopt strategies that are best replies to some belief based on their sample.

Define the operator B by $B(X) := X \cup BR(\Delta(X))$ and let B^K be its K-fold iteration. For each active individual $i \in p$ and for all $p \in \mathcal{P}$ define

$$S_K(X_{it}) := \{s_p \in S_p | s_p \in BR_p(\Delta(B^K(X_{it})))\},$$

and

$$S_K^0(X_{it}) := \{s_p \in S_p | s_p \in BR_p(\Delta^0(B^K(X_{it})))\}.$$

We complete the description of individual agent behavior by imposing the following restrictions on the rules by which active agents adopt new strategies. For an active agent i with period t sample X_{it}, who *reasons at depth K, $K \geq 0$*, assume that $\text{Prob}\{s_{i,t+1} = s_p | X_{it}\} > 0$ if and only if $s_p \in S_K(X_{it})$. If in addition, the agent is ε-*cautious* then $\text{Prob}\{s_{i,t+1} \in S_K^0(X_{it})\} \geq 1-\varepsilon$, where $\varepsilon \in (0,1)$. These transition probabilities are time-invariant. Note that we allow that $K = 0$ and agents need not be ε-cautious, unless explicitly assumed otherwise.

For any $X \subseteq S$, there exists a T such that for all $t > T$, $B^{t+1}(X) = B^t(X)$, since $B(\cdot)$ is monotone and S is finite. Let $\text{supp}(\omega) := \{s \in S | \forall p \in \mathcal{P}, \exists i \in p, \omega_i = s_p\}$ and define $t(\omega) := \min\{t \in \mathcal{N} | B^{t+1}(\text{supp}(\omega)) = B^t(\text{supp}(\omega))\}$. Then $W(\omega) := B^{t(\omega)}(\text{supp}(\omega))$ is closed under the inclusion of best replies and thus contains a curb retract for all $\omega \in \Omega$. For any curb retract Q define $K(Q) := \max\{t(\omega) | \text{supp}(\omega) \in Q\}$.

For any curb retract Q call the set of states $\omega \in \Omega$ in which every individual $i \in p$ has adopted a strategy $s_{it} \in Q_p$ the *curb-state set* supported on Q. The class of dynamic processes described above consists of Markov chains with stationary transition probabilities $\phi(\cdot|\cdot)$ on the state space Ω. The stable set of states for such a process are the so called *recurrent communication classes* of the process. The recurrent communication classes are subsets of Ω such that (i) from every state there is a finite length sequence of positive probability transitions to at least one of these classes, (ii) within each class every state can be reached from every other state via a finite length sequence of positive probability transitions, and (iii) no state outside one of the classes can be reached from a state inside through a

positive probability transition. The following result identifies the curb-state sets as the recurrent communication classes for any population learning dynamic ϕ.

Proposition 1 *(1) The curb-state sets are the recurrent communication classes of every population learning dynamic ϕ. (2) From any initial state, the population learning dynamic converges almost surely to a curb-state set. (3) If $Q \subset S$ is a curb retract in which a player $p \in \mathcal{P}$ has a dominant strategy (with respect to Q) $s_p \in S_p$, and players reason at least at depth $K(Q)$ and are ε-cautious, then for all active $i \in p$ and $t > 0$, $Prob(s_{i,t+1} = s_p | supp(\omega_t) \in Q) \geq 1 - \varepsilon$.*

Proof: For any $\omega \in \Omega$ and any $X \subseteq S$, ϕ induces a probability $P^{\tau}(X|\omega)$ that the support of the state τ periods from the present will be X. There exists a number of periods T such that $P^T(W(\omega)|\omega) > 0 \quad \forall \omega \in \Omega$. To see this note that each $s \in BR(\Delta(\text{supp}(\omega_t)))$ has positive probability of being in the support of ω_{t+1}. Since $1 - \lambda > 0$, any $s \in \text{supp}(\omega_t)$ also has positive probability of being in the support of ω_{t+1}. Since $\#(p) > \#(S_p)$, there is positive probability that all strategies in $B(\text{supp}(\omega_t))$ are present in the population state ω_{t+1} in period $t+1$. Let $T = \max\{t(\omega)|\omega \in \Omega\}$.

Once $\text{supp}(\omega_t) = W(\omega)$ for some $\omega \in \Omega$, there is positive probability of all agents being active and drawing a sample whose support is contained in the same curb retract. Thus, from all initial states the dynamic ends up in a curb-state set with positive probability after no more than $T' = T + 1$ periods. Once the dynamic has entered a curb-state set $\hat{\Omega}$, it will not leave it and since $P^T(W(\omega)|\omega) > 0 \ \forall \omega \in \hat{\Omega} \subset \Omega$, all states in $\hat{\Omega}$ are reached with positive probability from any other state in $\hat{\Omega}$.

Let $\tilde{\Omega} \subset \Omega$ be the subset of states that belong to some curb-state set. Since there are only finitely many states, $\exists\, \pi > 0 : \ \forall \omega \in \Omega$,

$$\text{Prob}(\omega_{t+T'} \in \tilde{\Omega} | \omega_t = \omega) \geq \pi \Rightarrow$$

$$\text{Prob}(\omega_{t+kT'} \notin \tilde{\Omega} | \omega_t = \omega) \leq (1-\pi)^k.$$

Therefore the probability that the dynamic process will never reach $\tilde{\Omega}$ equals

$$\lim_{k \to \infty} (1-\pi)^k = 0.$$

If $\text{supp}(\omega_t) \in Q$, where Q is a curb set and players reason at least at depth $K(Q)$, then the support of their sample X_{it} will be contained in Q for all $i \in p, \ p \in$

$\mathcal{P}$ and $S(X_{it}) = Q$. Thus, with probability of at least $1 - \varepsilon$, active players will use a best reply against a belief in $\Delta^0(Q)$, which must be the dominant strategy with respect to Q for players with such a strategy. □

Thus, in the long-run we expect the dynamic to end up in a curb-state set; once a curb-state set is reached, the dynamic never leaves it and returns to every one of its member states with positive probability, regardless of the initial state. This much requires neither multiple levels of depth of reasoning nor caution on part of the individuals. With caution and depth of reasoning, strategies that are dominated with respect to a curb set will be used with low probability. The communication games considered in the next section have curb sets in which every strategy profile that does not use a strategy dominated with respect to that curb set is an equilibrium. This characteristic, shared with many normal form games derived from an extensive form, lends credence to the focus on curb equilibria in our analysis.

4 One-Sided Communication

In this section we generalize our observation about *Dodo* versus *Stag Hunt.* First we derive a necessary condition on message space size for one-sided communication to be effective. Next we provide a complementary sufficient condition. Then we specialize the sufficient condition for the case of only two messages, and finally we specialize it further for the case where the underlying game is a symmetric 2×2-game. In each of these cases the result is that effective communication is a stable outcome for our population learning dynamics, if the *risk* of the preferred equilibrium is low relative to a *standard of comparison* that depends on the size of the underlying game and the size of the message space.

In the general setting the standard of comparison is inversely related to both the size of the message space and the size of the underlying game. Both of these factors increase the number of possible strategies, which can be interpreted as increasing strategic uncertainty. With only two message the standard of comparison becomes independent of the size of the underlying game (and trivially the message space). If the underlying game is in addition a symmetric 2×2-game, then the appropriate condition reduces to Harsanyi and Selten's [1988] *risk dominance* criterion.

Denote the strategies of the row (column) player in the underlying game G by $i \in I$ ($j \in J$). Let the row player be player one and the column player be player two with payoffs $u_k(i, j)$, $k = 1, 2$. Assume that the underlying game has multiple

strict Nash equilibria; this is the interesting case. Assume also that there is a unique strategy combination $(\hat{i},\hat{j})$ that maximizes the row player's payoff, i.e.

$$(\hat{i},\hat{j}) = \arg\max_{(i,j)} u_1(i,j),$$

and that this strategy combination is a strict Nash equilibrium. Examples of such games are *Dodo, Stag Hunt* and the *Battle of the Sexes.* Let M be the message space available to player one in the communication game $\Gamma_1(G,M)$. In this game player one first sends a message from the set M, then both players play the game G. Player one's strategies in the reduced normal form of the communication game are of the form (m,i) with $m \in M$ and $i \in I$. Player two's strategies are functions f that map messages $m \in M$ into actions $j \in J$. Let $F(G,M)$ be the set of pure strategies of player two in the communication game induced by G and M, and denote the cardinality of any finite set X by $\#(X)$. We will require $\#(M) \geq 2$ to ensure that the communication game does not become degenerate.

Our first result demonstrates that a restriction on the size of the message space is a necessary condition for one-sided communication to be effective.

Proposition 2 *For every game G, and every strict equilibrium $(\tilde{i},\tilde{j}) \neq (\hat{i},\hat{j})$ in G, if the message space M is sufficiently large, i.e.,*

$$\#(M) - 1 > \max_{i \neq \tilde{i}} \frac{u_1(i,\hat{j}) - u_1(\tilde{i},\hat{j})}{u_1(\tilde{i},\tilde{j}) - u_1(i,\tilde{j})},$$

then there exists a curb equilibrium in the communication game with payoff $(u_1(\tilde{i},\tilde{j}), u_2(\tilde{i},\tilde{j}))$.

Thus, for any game G we can find a message space large enough to ensure that communication does not discriminate among the game's strict equilibria. Note that if the underlying game G is a 2×2-game, then the bound on message space size is the ratio of the deviation losses that enter Harsanyi and Selten's definition of risk dominance.

The idea of the proof is simple. First, one shows that every curb set of the communication game contains every strategy of the silent player that responds to one message with $\hat{j}$ and to all other messages with $\tilde{j}$. This follows from iterating best replies and the fact that any best reply against a strategy of the sender can be altered arbitrarily after unused messages, and still be a best reply. Second, one considers beliefs of the communicating player that are concentrated on this set of strategies and assign equal probability to all strategies in the set. Clearly, the more

messages, the more attractive is $\tilde{i}$ against such beliefs. In contrast taking action $\hat{i}$ becomes increasingly risky as the number of messages increases. How attractive $\tilde{i}$ will be depends on how well $\tilde{i}$ does against mixtures of $\hat{j}$ and $\tilde{j}$ versus how well other actions i do against such mixtures. This creates the trade-off between message space size and payoffs that appears in the statement of the proposition. Since $(\tilde{i},\tilde{j})$ is a strict equilibrium, a sufficiently large message space guarantees that $\tilde{i}$ will be the unique best reply against the postulated beliefs.

Proof: First, we show that player one's preferred equilibrium $(\hat{i},\hat{j})$ appears as a continuation equilibrium in any curb retract of the communication game, regardless of the size of the message space. Let (m,i) be any strategy of the sender that is part of a curb retract Q. With at least two messages, there exists a best reply, f, for player two such that $f(m') = \hat{j},\ m' \neq m$. Therefore, $(m',\hat{i}) \in Q_1$, and the profile $((m',\hat{i}),f)$ is a curb equilibrium.

Next we show that if the size of M satisfies the condition stated in the proposition, then other strict equilibria will also appear in the curb set. Let $(\tilde{i},\tilde{j})$ be any other strict equilibrium of G. Consider the set of strategies $\tilde{F}$ of the form

$$\tilde{F} := \left\{ f \in F(G,M) | f_l(m) = \begin{cases} \hat{j} & \text{if } m = m_l \\ \tilde{j} & \text{otherwise} \end{cases},\ \ l = 1,2,\ldots \right\}$$

Note that each of the strategies in $\tilde{F}$ is in Q_2, where Q contains the curb equilibrium $((m',\hat{i}),f)$. To see this observe that the strategy f' that responds to all messages with $\hat{j}$ is a best reply to $(m',\hat{i})$, and in turn all strategies of the form $(m,\hat{i})$ are in Q_1. Each of the strategies in $\tilde{F}$ is a best reply against one of the latter strategies.

The set $\tilde{F}$ is constructed such that if player one believes that player two uses a strategy in $\tilde{F}$, but is not certain which of these strategies is used, then regardless of which message he considers, action $\tilde{i}$ is very attractive. Suppose player one has beliefs corresponding to uniform randomization over the set of strategies $\tilde{F}$ of player two. Against such beliefs, any strategy $(m,\tilde{i})$ gives player one the payoff

$$\frac{1}{\#(M)} u_1(\tilde{i},\hat{j}) + \frac{\#(M)-1}{\#(M)} u_1(\tilde{i},\tilde{j}).$$

Any strategy of the form (m,i), $i \neq \tilde{i}$, yields the payoff

$$\frac{1}{\#(M)} u_1(i,\hat{j}) + \frac{\#(M)-1}{\#(M)} u_1(i,\tilde{j}).$$

The former payoff is larger than the latter exactly when

$$\#(M) - 1 > \frac{u_1(i,\hat{j}) - u_1(\tilde{i},\hat{j})}{u_1(\tilde{i},\tilde{j}) - u_1(i,\tilde{j})}.$$

□

Having demonstrated that in our environment a restriction on message space size is necessary for effective communication to emerge, we now turn to developing a sufficient condition. In doing so, we let Proposition 2 guide us. From that proposition we learn that message space size matters because it increases the number of strategies in which player two responds to *some* message with (player one's) preferred equilibrium action but responds with other actions after different messages. We learn also that deviation losses matter, because they provide a measure of the payoff impact of these other actions.

Define the set of strategies of player two which sometimes respond with the preferred action $\hat{j}$ as

$$\hat{F}(G,M) := \{f \in F(G,M) | (\exists m \in M : f(m) = \hat{j})\}.$$

We will develop a sufficient condition, with an analog to the ratio of deviation losses from Proposition 2. This will be a measure of risk. Harsanyi and Selten's deviation losses only involve equilibrium actions. Here however, player one has to consider *all* of his alternative actions if he is uncertain about which messages induce the preferred response *and* which responses are induced otherwise. If effective communication is to be stable, it will be necessary that even under unfavorable beliefs, player one does not abandon the efficient action $\hat{i}$. The least favorable condition for action $\hat{i}$ is one where player two is omniscient, minimizes player one's payoff from action $\hat{i}$ and maximizes player one's payoff from every other action. This concern with extreme beliefs is necessary because player two's responses are mediated by messages and there may be only one (unknown) message that induces the desired response $\hat{j}$. In that vein, define a risk measure $\rho((\hat{i},\hat{j}),G)$ for the equilibrium $(\hat{i},\hat{j})$ that trades off player one's payoff from that equilibrium against the possibility of facing the kind of omniscient player two described above. Let

$$\rho((\hat{i},\hat{j}),G) := \max_{i \neq \hat{i}} \frac{\max_j u_1(i,j) - \min_j u_1(\hat{i},j)}{u_1(\hat{i},\hat{j}) - \max_j u_1(i,j)}.$$

This is a measure of the risk of player one at the equilibrium $(\hat{i},\hat{j})$. The risk of the equilibrium $(\hat{i},\hat{j})$ to player one decreases if his payoff at that equilibrium increases

or if the worst outcome from playing his equilibrium strategy $\hat{i}$ improves. Relative to any alternative strategy i, the risk increases if the maximum payoff from that strategy increases. Finally, only the maximal risk relative to any alternative strategy matters. Note that this measure is invariant with respect to positive affine transformations of the payoff function.

We introduce other risk measures below that are appropriate for the case of only two messages or can be used to compare risks among multiple strict Nash equilibria. However, the measure ρ remains central because it permits statements about the most general class of games. In that regard, it may appear as a limitation that ρ applies only to games in which there is a unique maximizer for player one's payoff and this is a strict Nash equilibrium. Thus, it is worth pointing out that outside this class of games no general results on effective communication are available. This follows because in every curb set of the communication game the communicating player's maximum payoff will be attained by some strategy combination. If in the underlying game this maximum payoff is not attained at a strict equilibrium then iterating best replies can lead to large curb sets which may include multiple equilibria with different payoffs. For example, take any version of *Dodo* and add a dominated strategy F (far right) for the silent (column) player such that (D,F) maximizes the communicating (row) player's payoff. In any corresponding communication game with at least two messages, there is a unique curb set, which contains of all strategies of the communicating player and nearly all strategies of the silent player (only the strategy which uses the dominated reply after every message is ruled out). This explains our restriction on the underlying class of games, which is reflected in our risk measure.

Our next result states that, for a given size of the message space, if risk is sufficiently small, then communication will be effective. We will construct a curb retract Q that for the communicating player is spanned by all strategies of the form $(m,\hat{i})$, and for the silent player is spanned by all strategies that respond with $\hat{j}$ to *some* message. A condition relating the risk of $(\hat{i},\hat{j})$ to the size of the message space ensures that strategies of the form (m,i), $i \neq \hat{i}$, are never best replies against any beliefs over Q. The max and min operators in the risk measure come from constructing a worst case scenario for beliefs concentrated on Q, in which player one believes that the message m induces player one's favorite reply given i, and alternative messages l induce player one's worst reply given $\hat{i}$ with high probability. Even in such a worst case scenario one can always find a message l that induces the reply $\hat{j}$ with probability of at least $1/\#(\hat{F}(G,M))$. The cardinality of the set $\hat{F}(G,M)$ is also the number of strategies of the silent player in Q and is an

increasing function of the number of messages; no belief concentrated on Q can assign weight less than the reciprocal of that number to all strategies in $\hat{F}(G,M)$. Thus, for any beliefs concentrated on Q, the communicating player can ensure that the weight of $u_1(\hat{i},\hat{j})$ in his expected payoffs is at least $1/\#(\hat{F}(G,M))$ for one of his strategies in Q, say $(l,\hat{i})$. Since $(\hat{i},\hat{j})$ is the unique profile that maximizes her payoff, if the possible payoffs from other strategy profiles (i,j), $i \neq \hat{i}$, in the underlying game are not "too high," this suffices to make $(l,\hat{i})$ a strictly better reply than (m,i). The risk measure makes the meaning of "too high" precise. The following proposition summarizes this discussion.

Proposition 3 *Let $\Gamma_1(G,M)$ be a one-sided communication game with underlying game G, and message space M, and risk $\rho((\hat{i},\hat{j}),G)$ at player one's preferred equilibrium. Suppose the following relation holds between $\rho((\hat{i},\hat{j}),G)$ and the cardinality $\#(\hat{F}(G,M))$ of player two's set of strategies that respond with $\hat{j}$ after some message.*

$$\frac{1}{\#(\hat{F}(G,M))-1} > \rho((\hat{i},\hat{j}),G).$$

Then (a) the retract $Q = Q_1 \times Q_2 := \{\mathrm{co}\{(m,\hat{i})\}_{m\in M}\} \times \{\mathrm{co}\{\hat{F}(G,M)\}\}$ is the unique curb retract in $\Gamma_1(G,M)$, and (b) the payoffs in all curb equilibria of the communication game $\Gamma_1(G,M)$ are $u_k(\hat{i},\hat{j})$, $k = 1,2$.

Proof: We begin showing that Q is closed under inclusion of best replies. Consider player two. By assumption, $\hat{j}$ is player two's unique best reply to $\hat{i}$ in the underlying game. Thus, against beliefs concentrated on Q_1, as defined in the statement of the proposition, any strategy f' with $f'(m) \neq \hat{j}$ $\forall m \in M$ has a strictly lower expected payoff than $\hat{f}$ where $\hat{f}(m) = \hat{j}$ $\forall m \in M$.

Turn to player one and suppose, to derive a contradiction, that (m,i), $i \neq \hat{i}$, is a best reply for player one against beliefs λ concentrated on Q_2. Let $\lambda(f)$ be the probability assigned to strategy f by λ. Let $\hat{F}(m,j) := \{f \in \hat{F}(G,M) | f(m) = j\}$; this is the set of all strategies in $\hat{F}(G,M)$ that respond to message m with action j. Then the payoffs from strategies (m,i) and $l,\hat{j}$ satisfy

$$\sum_{j\in J} \sum_{f\in\hat{F}(m,j)} u_1(i,j)\lambda(f) \geq \sum_{j\in J} \sum_{f\in\hat{F}(l,j)} u_1(\hat{i},j)\lambda(f) \;\; \forall l \neq m.$$

Note that finiteness of $\hat{F}(G,M)$ implies that at least one strategy in this set has probability bounded away from zero, i.e.

$$\max_{f\in\hat{F}(G,M)} \lambda(f) \geq \frac{1}{\#(\hat{F}(G,M))} > 0.$$

Let $\tilde{f} \in \arg\max_{f\in\hat{F}(G,M)} \lambda(f)$ be a strategy with maximum probability in $\hat{F}(G,M)$, and without loss of generality let $\tilde{f} \in \hat{F}(l,\hat{j})$. Then the above payoff relation between strategies (m,i) and $(l,\hat{i})$ implies that

$$u_1(i,\tilde{f}(m))\lambda(\tilde{f}) + (1-\lambda(\tilde{f}))\max_j u_1(i,j) \geq$$

$$u_1(\hat{i},\hat{j})\lambda(\tilde{f}) + (1-\lambda(\tilde{f}))\min_j u_1(\hat{i},j).$$

For this condition to be satisfied it is necessary that $\lambda(\tilde{f}) < 1$; thus we can rewrite it as

$$\frac{\lambda(\tilde{f})}{1-\lambda(\tilde{f})}(u_1(\hat{i},\hat{j}) - \max_j u_1(i,j)) \leq \max_j u_1(i,j) - \min_j u_1(\hat{i},j),$$

which implies

$$\frac{1}{\#(\hat{F}(G,M))-1} \leq \rho((\hat{i},\hat{j}),G),$$

in violation of the condition in the proposition. Thus, for a strategy (m,i) outside of Q_1 to be a best reply to beliefs concentrated on Q_2, the condition in the proposition must be violated. If it holds Q is closed under inclusion of best replies.

To see that Q is minimal among retracts closed under inclusion of best replies consider the following: The strategy $f(m) = \hat{j} \;\; \forall m \in M$ is a best reply to any strategy in Q_1. All strategies in Q_1 are best replies to $f(m) = \hat{j} \;\; \forall m \in M$. Any strategy in Q_2 is a best reply to some strategy in Q_1 by construction.

To establish uniqueness it suffices to show that for any set $\tilde{Q} \neq Q$ that is closed under best replies, $Q \cap \tilde{Q} \neq \emptyset$: If $(\tilde{m},i) \in \tilde{Q}_1$, $i \neq \hat{i}$, and j is a best reply to i, then there exists an $f \in \tilde{Q}_2$ such that $f(\tilde{m}) = j$, and $f(m) = \hat{j} \;\; \forall m \neq \tilde{m}$. Thus for all $m' \neq \tilde{m}$, $\{(m',\hat{i}),f\} \in Q \cap \tilde{Q}$. This shows that Q is the unique curb retract in the communication game.

It remains to show that all equilibria in the game restricted to Q have payoffs $u_k(\hat{i},\hat{j})$. Since $(\hat{i},\hat{j})$ is a strict Nash equilibrium in the underlying game, $u_2(\hat{i},\hat{j})$ is player two's maximal payoff in the communication game restricted to Q. Player

two can guarantee himself that payoff in Q. Whenever player two gets $u_2(\hat{i},\hat{j})$, player one gets $u_1(\hat{i},\hat{j})$ in the game restricted to Q. $\square$

The preceding two propositions concern the impact of message space size on the dynamic stability of efficient communication outcomes. And yet, it takes only two messages to destabilize inefficient outcomes. Moreover, we will show that with only two messages we can employ a less conservative risk measure. Define this alternative risk measure as

$$\tilde{\rho}((\hat{i},\hat{j}),G) := \max_{i\neq\hat{i}} \frac{u_1(i,\hat{j}) - \min_{j\neq\hat{j}} u_1(\hat{i},j)}{u_1(\hat{i},\hat{j}) - \max_{j\neq\hat{j}} u_1(i,j)}.$$

This measure turns out to be closely related Harsanyi and Selten's definition of risk dominance in two-player games. The motivation for the measure $\tilde{\rho}$ is similar as for ρ; the value $\tilde{\rho}((\hat{i},\hat{j}),G)$ is low whenever $\hat{i}$ is an "acceptable" reply not only to $\hat{j}$ but also to other strategies j. It is easily verified that $\rho((\hat{i},\hat{j}),G) \geq \tilde{\rho}((\hat{i},\hat{j}),G)$. In this sense the risk measure $\tilde{\rho}$ is less conservative than ρ. Like ρ this measure is invariant under positive affine transformations. The differences between the two measures are threefold. $u_1(i,\hat{j})$ replaces $\max_j u_1(i,j)$ in the numerator, which means that less weight is given to the best possible outcome under an alternative strategy i. The other two differences concern the ranges of j over which the min in the numerator and the max in the denominator are taken. Under $\tilde{\rho}$ in both cases $\hat{j}$ is excluded, which means that a (weakly) higher payoff is considered under the status quo and a weakly lower payoff under the alternative. In the symmetric 2×2 game discussed in Section 2 there is only one alternative and therefore there is no need to maximize over alternatives. Moreover in this case the limitations on the range of the max and min operators imply that these operators can be dropped. By doing so one arrives at what one might call the Harsanyi-Selten risk measure

$$\rho_{\text{HS}}((\hat{i},\hat{j}),G) := \frac{u_1(i,\hat{j}) - u_1(\hat{i},j)}{u_1(\hat{i},\hat{j}) - u_1(i,j)}$$

for 2×2 games. Note that in such a game the equilibrium $(\hat{i},\hat{j})$ is risk dominant in the Harsanyi-Selten sense if

$$u_1(\hat{i},\hat{j}) - u_1(i,\hat{j}) > u_1(i,j) - u_1(\hat{i},j),$$

which is equivalent to

$$\rho_{\text{HS}}((\hat{i},\hat{j}),G) < 1.$$

The following result resembles Proposition 3, with three differences. They concern the cardinality of the message space (two messages instead of finitely many), the risk measure ($\tilde{\rho}$ instead of ρ) and the standard of comparison (1 instead of $\frac{1}{\#(\hat{F}(G,M))-1}$).

Proposition 4 *For any one-sided communication game $\Gamma_1(G,M)$ with underlying game G and size of the message space $\#(M) = 2$, suppose that player one's risk, $\tilde{\rho}((\hat{i},\hat{j}),G)$, at his preferred equilibrium satisfies*

$$1 > \tilde{\rho}((\hat{i},\hat{j}),G).$$

Then (a) the retract $Q = Q_1 \times Q_2 :=$

$$\{\mathrm{co}\{(m,\hat{i})\}_{m\in M}\} \times \{\mathrm{co}\{\hat{F}(G,M)\}\}$$

is the unique curb retract in $\Gamma_1(G,M)$, and (b) the payoffs in all curb equilibria of $\Gamma_1(G,M)$ are $u_k(\hat{i},\hat{j})$, $k = 1,2$.

Proof: The proof of Proposition 2 is identical to the proof of Proposition 1 with one exception that concerns showing that Q is closed under inclusion of best replies.

Let $M := \{m_1, m_2\}$. Suppose, in order to derive a contradiction, that (m_1, i), $i \neq \hat{i}$, is a best reply against the belief λ concentrated on Q_2. With only two messages it is convenient to represent a strategy of player two as a vector (j^1, j^2) with the first (second) element being the response to message m_1, (m_2). Let $\lambda(j^1, j^2)$ be the probability assigned to strategy (j^1, j^2) by λ. Then, since (m_1, i), $i \neq \hat{i}$, must be at least as good a reply as $(m_1, \hat{i})$ to the belief λ, we have

$$\sum_{j\neq\hat{j}} u_1(i,\hat{j})\lambda(\hat{j},j) + \sum_{j\neq\hat{j}} u_1(i,j)\lambda(j,\hat{j}) \geq$$

$$\sum_{j\neq\hat{j}} u_1(\hat{i},\hat{j})\lambda(\hat{j},j) + \sum_{j\neq\hat{j}} u_1(\hat{i},j)\lambda(j,\hat{j}),$$

where we have used the fact that $u_1(\hat{i},\hat{j}) > u_1(i,\hat{j})$.

Therefore

$$[u_1(\hat{i},\hat{j}) - u_1(i,\hat{j})] \sum_{j\neq\hat{j}} \lambda(\hat{j},j) \leq$$

$$[\max_{j\neq\hat{j}} u_1(i,j) - \min_{j\neq\hat{j}} u_1(\hat{i},j)] \sum_{j\neq\hat{j}} \lambda(j,\hat{j}).$$

Together with the condition $1 > \tilde{\rho}((\hat{i},\hat{j}),G)$, this implies

$$\sum_{j\neq\hat{j}} \lambda(j,\hat{j}) > \sum_{j\neq\hat{j}} \lambda(\hat{j},j),$$

since at least one of these sums must be positive for (m_1,i), $i \neq \hat{i}$, to be a best reply.

Since (m_1,i), $i \neq \hat{i}$, must be at least as good a reply as $(m_2,\hat{i})$ we also have the following condition

$$\sum_{j\neq\hat{j}} u_1(i,\hat{j})\lambda(\hat{j},j) + \sum_{j\neq\hat{j}} u_1(i,j)\lambda(j,\hat{j}) \geq$$
$$\sum_{j\neq\hat{j}} u_1(\hat{i},j)\lambda(\hat{j},j) + \sum_{j\neq\hat{j}} u_1(\hat{i},\hat{j})\lambda(j,\hat{j}).$$

Therefore

$$[u_1(\hat{i},\hat{j}) - \max_{i\neq\hat{j}} u_1(i,j)] \sum_{j\neq\hat{j}} \lambda(j,\hat{j}) \leq$$
$$[u_1(i,\hat{j}) - \min_{j\neq\hat{j}} u_1(\hat{i},j)] \sum_{j\neq\hat{j}} \lambda(\hat{j},j).$$

Together with the condition $1 > \tilde{\rho}((\hat{i},\hat{j}),G)$, this implies

$$\sum_{j\neq\hat{j}} \lambda(\hat{j},j) > \sum_{j\neq\hat{j}} \lambda(j,\hat{j}).$$

Therefore we have reached a contradiction. Hence, (m_1,i), $i \neq \hat{i}$, cannot be a best reply to any beliefs λ concentrated on Q_2. An analogous argument works for (m_2,i), $i \neq \hat{i}$. □

Proposition 4 shows that with only two messages a less stringent risk measure, $\tilde{\rho}$, can be used to formulate a condition guaranteeing effective communication. We argued before that this measure is a generalization of the Harsanyi-Selten risk measure ρ_{HS} while for ρ, which we used for general message space sizes, the link to risk dominance is more tenuous. For that reason let me examine why ρ and not $\tilde{\rho}$ is the appropriate risk measure in games with more than two messages. Consider the following version of *Dodo*

	L	R
U	x,x	0,0
D	0,0	1,1

where $2 > x > 1$. We will analyze the communication game where only the row player can talk and the cardinality of the message space equals three. The reduced normal form of this game is given by

	LLL	LLR	LRL	LRR	RLL	RLR	RRL	RRR
m_1U	x,x	x,x	x,x	x,x	0,0	0,0	0,0	0,0
m_1D	0,0	0,0	0,0	0,0	1,1	1,1	1,1	1,1
m_2U	x,x	x,x	0,0	0,0	x,x	x,x	0,0	0,0
m_2D	0,0	0,0	1,1	1,1	0,0	0,0	1,1	1,1
m_3U	x,x	0,0	x,x	0,0	x,x	0,0	x,x	0,0
m_3D	0,0	1,1	0,0	1,1	0,0	1,1	0,0	1,1

It is evident that in this game any curb retract that supports the efficient outcome in the underlying game must contain all strategies that support the efficient outcome in the communication game. This means that for example the strategies *RLR* and *RRL* for player two must be included in the curb retract. However, if player one believes that player two uses only these two strategies and uses them with equal probability, then the strategy m_1D is a best reply, from which it follows easily that in fact all strategies will be in any curb retract. Only if we made x larger than 2 would this argument not work. It is this distinction that is captured by the risk measure ρ. Note that $\tilde{\rho}(x) = 0 \; \forall x > 1$ and $\rho(x) = \frac{1}{x-1} \; \forall x > 1$; the risk associated with the equilibrium UL gets large as x approaches one.

This example does not show why in addition to changing the risk measure we need to change our standard of comparison as we increase the number of messages since $\rho(x) < 1$ is equivalent to the condition that $x > 2$. The reader may check however that if we add another message, x needs to increase for it to be the unique curb equilibrium payoff in the communication game.

The risk measures introduced up to this point are all structurally similar to ρ. This measure relies on $(\hat{i}, \hat{j})$ being the unique strategy profile that maximizes player one's payoff in the underlying game. Therefore it is worth pointing out that instead of $\tilde{\rho}$ we could have considered an alternative measure in Proposition 4. For an *arbitrary* finite game and *any* strict equilibrium $(\hat{i}, \hat{j})$, this equilibrium's risk can be measured by

$$\rho_{\text{GHS}} := \max_{i \neq \hat{i}} \frac{\max_{j \neq \hat{j}} u_1(i, j) - \min_{j \neq \hat{j}} u_1(\hat{i}, j)}{u_1(\hat{i}, \hat{j}) - u_1(i, \hat{j})}.$$

The attraction of this measure is that it can both replace $\tilde{\rho}$ in Proposition 4 and be used to compare risk across different strict equilibria. We will refer to this as the Generalized Harsanyi-Selten measure of risk. Section 6 contains a brief discussion on the relation of this measure to some examples from the literature on stochastic evolutionary game dynamics.

5 Multi-Sided Communication

In this section we introduce multi-sided communication and compare it to one-sided communication. The result of this comparison turns out to depend on whether or not we allow the population learning dynamic to introduce permanent asymmetries into the message space.

So far we have examined a class of population learning dynamics in which individuals make only limited use of the desymmetrizing effects of history. They learn to distinguish messages according to payoff differences in their current environment. They do not however develop an affection for messages that have served them well for a long time or impute meaning (beyond the current use in the population) to messages that have had a long association with an equilibrium in the underlying game. This is a useful benchmark in the analysis of cheap-talk games but not very realistic. One ought to at least acknowledge a secondary, tie-breaking role for such behaviors.

For the central result of this section we leave the population learning dynamics unchanged and work with a message space that is already differentiated. This allows us to capitalize on our investment in establishing a link between the stable outcomes of population learning dynamics and curb retracts. Later, we propose a modification of the population learning dynamics with an endogenous process of message space differentiation.

To capture message space differentiation through minor alterations of the communication game, we assume that associations of message profiles and equilibria induce minute differences in payoffs. This link between message profiles and equilibria introduces *and formalizes* a small measure of *a priori* meaning of messages, which we refer to as *limited information content (LIC).* To formalize *LIC,* we focus on the recipients of messages. The small payoff changes that we introduce make players in their roles as recipients of messages slightly more "gullible;" if all messages sent agree on the same equilibrium, then the recipients become more favorably inclined toward their corresponding equilibrium actions.[7] It is useful to emphasize that the meaning appears via a minimal differentiation of the message space, which itself can be endogenized without much difficulty.

Two observations emerge. If costless messages have no a priori meaning whatsoever, multi-sided communication is *less* effective than one-sided communication. Multi-sided communication is *more* effective if all players communicate, have the same preferred equilibrium, and there is some a priori information content of messages, however small, that identifies message profiles that signal agreement on a strict equilibrium in the underlying game.

Consider the first of these claims. With multi-sided communication, and a priori meaningless messages, there is no result of the form: "If the risk of an efficient equilibrium in the stage game is sufficiently low, it will be the only curb equilibrium in the communication game." To see this consider the following example of a communication game derived from a version of *Dodo.*

[7]Gullibility of receivers is different from a "disincentive to lie" for senders and from the nominal message costs employed by Hurkens [1993]. The latter are concerned with *sender* preferences and have stronger implications than gullibility in games where only a subset of the players send messages.

	$m_1'LL$	$m_1'LR$	$m_1'RL$	$m_1'RR$	$m_2'LL$	$m_2'LR$	$m_2'RL$	$m_2'RR$
m_1UU	**x**,**x**	**x**,**x**	**0**,0	**0**,0	x,x	x,x	0,0	0,0
m_1UD	**x**,**x**	**x**,**x**	**0**,0	**0**,0	0,0	0,0	1,1	1,1
m_1DU	0,**0**	0,**0**	1,1	1,1	x,x	x,x	0,0	0,0
m_1DD	0,**0**	0,**0**	1,1	1,1	0,0	0,0	1,1	1,1
m_2UU	x,x	0,0	x,x	0,0	x,x	0,**0**	x,x	0,**0**
m_2UD	x,x	0,0	x,x	0,0	**0**,0	**1**,**1**	**0**,0	**1**,**1**
m_2DU	0,0	1,1	0,0	1,1	x,x	0,**0**	x,x	0,**0**
m_2DD	0,0	1,1	0,0	1,1	**0**,0	**1**,**1**	**0**,0	**1**,**1**

In this game two players simultaneously exchange messages before playing *Dodo;* with $x > 1$. The cardinality of each player's message space equals two. It is easily checked that any curb retract that supports the efficient outcome must include all strategies that are consistent with efficiency. However, the game induced by the retract that is formed by all such strategies contains an equilibrium in which the row player uses m_1UD and m_2DU with equal probability and the column player uses $m_1'LR$ and $m_2'RL$ with equal probability. Thus, for any x there is a curb equilibrium with payoffs $\frac{x+1}{2}$ for each player. Hence, within the class of population learning dynamics considered above, we cannot rule out inefficient outcomes.

Now assume instead that the message space is differentiated, that message exchanges alter, if only marginally, the players' perception of the underlying game. In the example assume that once players have exchanged the message pair (m_1, m_1'), they are somewhat more inclined to use actions U and L respectively in the underlying game. Model this increased inclination as an enlargement of the range of beliefs for which the actions U and L are optimal in the underlying game. That is, for the row player, for any given distribution over L and R, the action U yields a slightly higher payoff than it would without the prior exchange of messages m_1 and m_1'. Let there be a similar link between the message pair (m_2, m_2')

and the action pair (D,R) so that there is no bias in favor of either equilibrium.[8] In the example such a payoff boost affects the bold entries in the above representation of the communication game. It is easily checked that for any sufficiently small positive payoff boost, there is a unique curb retract in the communication game and it consists entirely of equilibria that support the efficient outcome.

For a general definition of *limited information content (LIC)* we want to introduce some minimal message space differentiation that links messages with equilibria in the underlying game. To this end, we let messages have a marginal effect on the players' perception of the game following the message exchange. This introduces a measure of effective communication when a message is sent that is independent of the current use of the message. It appears essential for effective communication to occur that the sender's message affect the receiver's perception. In contrast, while the act of sending a message may also have a self-committing effect on receivers, such an effect does not seem to be equally central to the notion of effective communication. The effect on senders of messages having been exchanged appears to be indirect. For a message to alter the receiver's perception, it has to be "believed." On the other hand, for a message to alter the sender's perception, the sender must "believe" that the message is "believed" by the receiver, and thus solve an *inference problem.* Therefore we take the position that the primary effect of the message exchange is on receivers' beliefs.

Another issue that we have to address is that players may be both senders and receivers of messages. Therefore, it can occur that a player sends a message linked to one equilibrium and at the same time receives a message linked to another equilibrium. In that case it appears reasonable that perceptions remain unaltered.

Accordingly, the formal definition of *LIC* reflects two principles, (1) that the primary effect of message exchange is on perceptions of receivers, and (2) that any effect on perceptions requires unanimity among senders.

These principles have strong implications for the impact of one- vs two-sided communication in two player games. If only one player communicates, only the receiving player will be swayed by an LIC message. If both players communicate, they will *both* be swayed by LIC messages, as long as there is no ambiguity. Two-sided communication avoids the above mentioned inference problem of senders.

We express the altered perception of the underlying game, following a message profile linked to an equilibrium in that game, by making the associated equi-

[8]This second link is present only to demonstrate that we can allow for a whole range of meanings. It does not affect our conclusions.

librium actions more attractive. An action in the underlying game becomes more attractive if its payoffs increase. Therefore we will express the *limited information content* of a message profile as an ε-increase in payoffs from corresponding equilibrium actions in the underlying game.

Since the main result refers to an P player game, we adopt a somewhat more general notation in this section than in the previous one. Let each player's set of pure strategies in the underlying game be denoted by S_p, $p \in \mathcal{P}$, and let s_p be a typical element of player p's strategy set. Analogous to the previous section assume that there is a strategy combination $\hat{s} = \{\hat{s}_p\}_{p=1}^{P}$ in the underlying game such that

$$\hat{s} = \arg\max_{s} u_p(s) \;\; \forall p \in \mathcal{P};$$

i.e., for every player $p \in \mathcal{P}$, $\hat{s}$ is the unique strategy combination that gives p's maximum payoff in the game. Let there be a set $C \subseteq \mathcal{P}$ of communicating players.

Let $\mathcal{E}$ denote the set of strict Nash equilibria in the underlying game, and for each communicating player p, let ι_p be an injective function

$$\iota_p : \mathcal{E} \to M_p,$$

which will be referred to as player p's *LIC* mapping. Thus, for each strict equilibrium, player p has exactly one message linked to that equilibrium. This implicitly assumes that a communicating player's message space is at least as large as the number of strict equilibria in the underlying game. Let $\iota(\cdot)$ be the vector of all *LIC* mappings for all the communicating players.

Using the *LIC* mappings, we can define *LIC* preferences with an ε-payoff boost in the communication game as follows. Denote payoffs in the communication game by $U_p(m, f)$ and let

$$U_p(m,f) = \begin{cases} u_p(f(m)) + \varepsilon, & \text{if } \exists e \in \mathcal{E},\ \iota(e) = m,\ e_p = f_p(m) \text{ and } \mathcal{P} \backslash p \subseteq C; \\ u_p(f(m)), & \text{otherwise.} \end{cases}$$

The *LIC* mapping establishes the meaning of messages but only through *LIC* preferences do these meanings become operational. Without *LIC* preferences, the game is completely unchanged, and even with *LIC* preferences, the meanings play a limited role because none of the equilibrium outcomes change and all messages can acquire all meanings, in equilibrium. According to this formulation of LIC preferences, communication makes an equilibrium more attractive for a player

if all other players communicate, and all communicating players, "agree" on the equilibrium in question.

It is easy to check that this way of introducing small a priori information content has no effect with one-sided communication. Even with only one-sided communication, an LIC message does alter perceptions. However, it does so only for the receiver. As a consequence, the sender's payoffs in the communication game (including all ties) are exactly as before and the results of the previous section go through virtually unchanged.

For the remainder of this section assume that $\mathcal{P} = C$. Denote by $\Gamma_P(G,M)$ the communication game in which the play of the underlying game G is preceded by one round of simultaneous communication in which all P players announce a message from their respective message spaces M_p.

The following proposition generalizes our example. If all players communicate and have LIC preferences, every strategy in the unique curb retract is an efficient equilibrium.

Proposition 5 *Let $\Gamma_P(G,M)$ be a multi-sided communication game with player set $\mathcal{P}$, underlying game G, unique efficient profile $\hat{s}$ in G, and message spaces M_p, $p \in \mathcal{P}$, where all players have* LIC *preferences. Then there exists a bound $\bar{\varepsilon} > 0$ on the LIC payoff boost ε such that the retract $Q = \times_{p=1}^{P} Q_p$, where $Q_p :=$ $\mathrm{co}\{(m_p, f_p) | m_p \in M_p,\ \iota(f(m)) = m,\ f(m) = \hat{s}\}$, is the unique curb retract in $\Gamma_P(G,M)$, for all $\bar{\varepsilon} > \varepsilon > 0$.*

Proof: Against any belief concentrated on Q_{-p} player p can achieve a payoff of $u_p(\hat{s}) + \varepsilon$ by using one of the strategies in Q_p. Any other strategy will at most yield a payoff of $u_p(\hat{s})$. This shows that Q is closed under inclusion of best replies.

Every strategy combination in Q gives player p his maximal payoff in the communication game. Therefore no strict subset of Q is closed under inclusion of best replies, which shows that Q is minimal.

It remains to show uniqueness. We will show that if $\tilde{Q}$ is curb, then $Q \cap \tilde{Q} \neq \emptyset$. Let $(m, f) \in \tilde{Q} \backslash Q$. Then there exists a strategy $(m', f') \in \tilde{Q}$ with $f'_p(\tilde{m}) = \hat{s}_p\ \forall p, \forall \tilde{m}_{-p} \neq m_{-p}$. Against (m', f') player p can guarantee that the other players will play $\hat{s}_{-p}$ in the underlying game by not sending message m_p. Hence, there exists $(m'', f'') \in \tilde{Q}$ with $f''(m'') = \hat{s}$, which implies that there exists $(m'', f''') \in \tilde{Q}$ with $f'''_p(\tilde{m}) = \hat{s}_p\ \forall p,\ \forall \tilde{m}_{-p}$. Hence, $(\hat{m}, f''') \in \tilde{Q}$, for $\hat{m} = \iota(\hat{s})$. But $(\hat{m}, f''') \in Q$. $\square$

Proposition 5 relies on message space differentiation already having been established. One could model this differentiation as arising endogenously. Here is an example of one message profile becoming distinct from all others: Consider a variation on our population learning dynamic applied to P-player games, in which (conforming to the central result of this section) all players communicate and the underlying game has a unique efficient strategy combination $\hat{s}$. Let each population member behave exactly as in our basic population learning dynamic, with one exception: For any strategy profile (m^*, f^*) with $f^*(m^*) = e^* \in \mathcal{E}$, define the set $S(m^*, f^*) := \{(m_p, f_p(\dot{})) | \forall p \in \mathcal{P},\ m_p = m_p^*,\ f_p(m^*) = f_p^*(m^*)\}$ of all strategies in which player p sends message m_p^* and takes the equilibrium action e_p^* when the profile m^* is sent. Suppose that individual p's sample consists entirely of strategies in $S(m^*, f^*)$. Then, for any message $m_p' \neq m_p^*$, let the probability of individual $i \in p$ adopting the strategy $(m_i', f_i'(\cdot))$ equal zero, if $u_p(f(m_p', m_{-p}^*)) \leq u_p(f^*(m^*))$, for all f such that $(m, f) \in S(m^*, f^*)$. That is, if an individual believes he can induce strict equilibrium e with certainty by sending the message he believes everyone in his population to send, he will send the same message unless sending an alternative message strictly raises his payoff. Consider any set of states in which all individuals use strategies in $S(m^*, f^*)$ for some (m^*, f^*). It is straightforward (although tedious) to check that (1) outside such sets of states, the dynamic behaves just like the original population learning dynamic, (2) the dynamic will leave any such set of states where $f^*(m^*)$ is an inefficient equilibrium, and will never leave such a set where $f^*(m^*)$ is the unique efficient strategy profile. Adapting the proof of Proposition 5 to explicit dynamics, these observation can be shown to imply that the altered class of dynamics converge to a set of efficient equilibria almost surely. In the limit players use a single endogenously determined message even if they believe other messages to lead to the same payoffs.

Evidently, the dramatic effect of differentiating the message space derives from its breaking of some of the numerous payoff ties in the communication game. Not all plausible tie-breaking rules will guarantee efficiency in the limit. Say, players develop a slight preference (in the above ε-sense) for strategies that ignore messages. It is easy to check that in that case every strict equilibrium of the underlying game gives rise to its own curb retract in the communication game. Other plausible behaviors give rise to stronger results than Proposition 5. For example, one can model a "disincentive to lie" by linking a player's actions to a subset his messages. Then the player gets ε-penalized for not following a given message with the appropriate action. In that case both one- and two-sided communication

lead to efficiency. The same strong efficiency result is derived by Hurkens [1993] who differentiates messages by nominal message costs.

Whether message space differentiation occurs and in which form is an empirical matter that should be addressed experimentally. The point of this section was to show that (1) without message space differentiation communication remains ineffective with multi-sided communication and (2) that message space differentiation can lead to an interesting reversion in the effectiveness of one- and multi-sided communication.

6 Related Literature

This section discusses three strands of literature relevant to this paper. It first looks at some experimental evidence on pre-play communication. Then follows an overview of the evolutionary approach to pre-play communication. The section concludes with comments on the literature on stochastic evolutionary game dynamics and its relation to generalized risk measures.

Cooper, DeJong, Forsythe and Ross (CDFR) [1992] report on experiments they conducted on a version of the *Stag Hunt* game ($a = 1000,\ b = 0,\ c = 800, d = 800$ in game $\mathbf{G_0}$). They repeatedly let players play one-shot communication games, where messages had an exogenously given meaning. They find that without communication the risk dominant equilibrium will be played. With one-way communication the frequency of the Pareto-efficient equilibrium increases but there are also a significant number of coordination failures. Two-way communication resolves the coordination problem; almost exclusively the Pareto-efficient equilibrium is played.

Analogous to LIC preferences, one can interpret CDFR's results as indicating that the unanimous assertion to play according to the preferred equilibrium is enough to reduce strategic uncertainty; on the other hand, if only one player communicates he maybe less confident of the response of the listener. It makes a difference whether one is informed about someone's intent, rather than having to make an inference about this intent. In the words of CDFR: "This doubt about the action of a receiver is overcome by the two-way communication design since both players receive information about the likely play of their opponents." [1992, p.757]

While refinements in the spirit of strategic stability have little power in games with pre-play communication the evolutionary approach yields sensible predic-

tions in common interest games. Several authors have used versions of Maynard Smith and Price's [1973] notion of an *evolutionarily stable strategy (ESS)* in these games. Roughly, in a symmetric game a strategy is evolutionarily stable if it is a symmetric Nash equilibrium and, if it is played by all members of a large population, it cannot be invaded by a small population of mutants who use a different strategy. An *ESS* must be a best reply to itself, and it must be a better reply to the post-entry population than any potential entrant.

Wärneryd [1991] studies *Dodo* preceded by one round of pre-play communication in which each player sends a message from a common finite message space that contains at least two messages. Wärneryd shows that any *neutrally stable strategy (NSS)* (a variant of em ESS) in *Dodo* preceded by one round of simultaneous pre-play communication leads to the efficient equilibrium in the underlying game. Wärneryd's analysis does not extend to more general games because there the use of unused messages may get penalized deterring players from introducing them. Matsui [1991] arrives at similar conclusions for a different solution concept, *cyclically stable sets (CSS)* [Gilboa and Matsui, 1991]. Unlike *NSSs, CSSs* are set valued and exist in every game. Kim and Sobel [1993] obtain similar results as Wärneryd for more general games than *Dodo* by using a set-valued solution concept, Swinkels' [1992] *Equilibrium Evolutionarily Stable (EES) Set*.

These evolutionary solutions assume that only a small fraction of the population moves at any given point in time. This reduces the possibility for coordination failures significantly. In the case of common interest games, once the evolutionary process reaches the efficient point, any potential invader must also play an efficient strategy in order to maximize its payoff against the post-entry population. The problem of a mutant in this situation is reduced to a simple optimization problem with no strategic component. Sobel [1993] has a particularly transparent version of the evolutionary argument. In his model only one player at a time adjusts his strategy.

With simultaneous adjustments as in the present paper, either one has to acknowledge the role of risk and strategic uncertainty, or one must introduce some other source of inertia if one wants to ensure that only efficient outcomes are stable. In Section 5 we examined one such form of inertia. Hurkens [1993] gets very strong result for the same solution concept, curb retracts, and a different form of inertia, namely nominal message costs. He shows that in two-player games the same result holds for persistent equilibria, which were defined by Kalai and Samet [1984]. Since messages are costly in Hurkens' paper his work also provides a link with the work on "burning money" by Ben-Porath and Dekel [1992]

and van Damme [1989]. The players who are given the opportunity to burn money can achieve their preferred outcome, and no money is actually burned in equilibrium. Ben-Porath and Dekel use iterative deletion of weakly dominated strategies as their solution concept and require that message costs be nonnegligible.

Risk becomes an issue in this paper because of the relative ease with which one can travel from one strategy combination to another. This is a consequence of the low entry requirements for new strategies in our dynamic and the absence of other sources of friction such as message costs, as in Hurkens' work.[9] One reason it is so easy to travel from one profile to another is that communication turns strict equilibria into weak ones. Thus communication creates an escape route from strict equilibria. An alternative escape route is analyzed in the literature on stochastic evolutionary game dynamics, as for example in Kandori, Mailath and Rob [1993], Young [1993] and Ellison [1993]. There too, risk plays a role.[10] Kandori, Mailath and Rob for example prove that in two-player two-strategy coordination games the limit of the stationary distributions of their dynamics as the noise vanishes puts all probability weight on the risk dominant equilibrium. Young [1993] examines a similar kind of dynamics and comes to similar conclusions. He also points that while in two-strategy games a characterization of stochastically stable equilibria in terms of risk dominance is possible no such characterization may be available for games with three or more strategies. He gives the following example.

	L	C	R
U	6,6	0,5	0,0
M	5,0	7,7	5,5
D	0,0	5,5	8,8

Young points out that while the equilibrium (D,R) pairwise risk dominates the other two pure strategy equilibria in this game, there are plausible dynamics under which (M,C) is the unique stochastically stable equilibrium. Against

[9]Hurkens' results holds also for the weaker solution concept "closed under inclusion of better replies." Such a solution concept poses even less stringent entry requirements than curb retracts. In Hurkens work this permissiveness is balanced by the cost of messages. For a discussion and characterization of this solution concept see Ritzberger and Weibull [1993].

[10]Risk dominance also plays a role in 2×2 games with small payoff uncertainty and incomplete information. This has been demonstrated in a recent paper by Carlsson and van Damme [1993].

this background it is perhaps interesting to calculate the Generalized Harsanyi-Selten measures for this game. They are $\rho_{\text{GHS}}(U,L) = 7$, $\rho_{\text{GHS}}(M,C) = 3/2$, and $\rho_{\text{GHS}}(D,R) = 7/3$. Therefore the equilibrium picked by Young's dynamics is also the one with the lowest GHS risk measure. This phenomenon does not generalize. Moreover, Ellison [1993] points out that variations in matching rules (e.g. more frequent matching with close neighbors vs. uniform matching) alter the long-run predictions in games similar to the example. Despite the weakness of the link to stochastic evolutionary game dynamics, we believe that the GHS risk measure identifies some of the dynamic forces acting on multiple Pareto-ranked strict Nash equilibria.

7 Conclusion

The chances for pre-play communication to allow players to coordinate on efficient equilibria are tied to the risk associated with these equilibria. This observation can be formally expressed through examining stable outcomes of a class of best reply dynamics that permit simultaneous strategy adjustments among a large fraction of the population. This contrasts with and complements results from the evolutionary literature on pre-play communication in which risk plays no role. One benefit of this approach is that it permits one to make a distinction between one- and two-sided communication, which parallels some experimental results.

References

Aumann, R. [1990], "Nash Equilibria are not Self-enforcing," in *Economic Decision Making: Games, Econometrics and Optimization*, edited by J.J. Gabszewicz, J.-F. Richard and L.A. Wolsey, Amsterdam: North-Holland.

Balkenborg, D. [1992], "The Properties of Persistent Retracts and Related Concepts," Ph.D. Thesis (University of Bonn).

Basu, K. and J.W. Weibull [1991], "Strategy Subsets Closed Under Rational Behavior," *Economics Letters*, **36**, 141-146.

Bhaskar, V. [1992], "Noisy Communication and the Evolution of Cooperation," Dehli University.

Ben-Porath, E. and Dekel, E. [1992], "Signaling Future Actions and the Potential for Self-Sacrifice," *Journal of Economic Theory*, **57**, 36-51.

Blume, A., Kim, Y.-G., and J. Sobel [1993], "Evolutionary Stability in Games of Communication," *Games and Economic Behavior*, **5**, 547-575.

Carlsson, H. and E. van Damme [1993], "Global Games and Equilibrium Selection," *Econometrica*, **61**, 989-1018.

Cooper, R., DeJong, D.V., Forsythe, R. and T.W. Ross [1992], "Communication in Coordination Games," *The Quarterly Journal of Economics*, **107**, 739-771.

Ellison, G. [1993], "Learning, Local Interaction and Coordination," *Econometrica*, **61**, 1047-1071.

Farrell, J. [1988], "Communication, Coordination and Nash Equilibrium," *Economics Letters*, **27**, 209-214.

Farrell, J. and M. Rabin [1996], "Cheap Talk," forthcoming in *The Journal of Economic Perspectives*.

Fudenberg, D. and E. Maskin [1991], "Evolution and Communication in Games," preliminary notes.

Gilboa, I., and A. Matsui [1991], "Social Stability and Equilibrium," *Econometrica,* **59**, 859-867.

Harsanyi, J.C. and R. Selten [1988], *A General Theory of Equilibrium Selection in Games,* Cambridge MA: MIT press.

Hurkens, S. [1993], "Multi-sided Pre-play Communication by Burning Money," CentER for Economic Research Discussion Paper.

Hurkens, S. [1995], "Learning by Forgetful Players," *Games and Economic Behavior,* **11**, 304-329.

Kalai, E. and D. Samet [1984], "Persistent Equilibria in Strategic Games," *International Journal of Game Theory,* **13**, 129-144.

Kandori, M., Mailath, G.J. and R. Rob [1993], "Learning, Mutation, and Long Run Equilibrium in Games," *Econometrica,* **61**, 29-56.

Kim, Y.-G. and J. Sobel [1993], "An Evolutionary Approach to Pre-Play Communication," UCSD Working Paper.

Matsui, A. [1991], "Cheap Talk and Cooperation in Society," *Journal of Economic Theory,* **54**, 245-258.

Maynard Smith, J. and G. Price [1973], "The Logic of Animal Conflict," *Nature,* **246**, 15-18.

Ritzberger, K. and J. Weibull [1993], "Evolutionary Selection in Normal Form Games," The Industrial Institute for Economic and Social Research Working Paper.

Sobel, J. [1993], "Evolutionary Stability in Communication Games," UCSD Working Paper.

Swinkels, J. [1992], "Evolutionary Stability with Equilibrium Entrants," *Journal of Economic Theory,* **57**, 306-332.

van Damme, E. [1989], "Stable Equilibria and Forward Induction," *Journal of Economic Theory,* **48**, 476-496.

Wärneryd, K. [1991], "Evolutionary Stability in Unanimity Games with Cheap Talk," *Economics Letters,* **36**, 375-378.

Young, H.P. [1993], "The Evolution of Conventions," *Econometrica,* **61**, 57-84.