Convergence of Behavior in Social Networks*

Dinah Rosenberg†
Laboratoire d'Analyse Géométrie
et Applications
Institut Galilée, Université Paris Nord
avenue Jean-Baptiste Clément
93430 Villetaneuse, France

Eilon Solan
The School of
Mathematical Sciences
Tel Aviv University
Tel Aviv 69978, Israel

Nicolas Vieille
Département Finance et Economie
HEC
1, rue de la Libération
78 351 Jouy-en-Josas, France

1 Introduction

The dissemination of private information, or knowledge, in a population has attracted much interest, first among sociologists and geographers (see references in Chamley (2004)), and more recently among economists and computer scientists. A question that has attracted a lot of attention is whether as time passes, information spreads through the entire population, and beliefs become more precise, consensus of some sort eventually arises.

Within economics, this work has developed independently in different directions, and several strands of literature can be recast under that heading. This includes getting to common knowledge (e.g., Geanakoplos and Polemarchakis (1982) and Parikh and Krasucki (1990)), learning in social networks (e.g., Goyal (2005)) and strategic experimentation (e.g., Bolton and Harris (1999), Keller et al. (2005), Rosenberg et al. (2005) and Murto and Valimaki (2006)).

We study a general model of information dissemination, that includes the above models as special cases, and study the limit behavior of the players. Our main findings are the following. (1) The number of stages in which players experiment is uniformly bounded. In particular, asymptotically all motives for experimenting disappear, and players play myopically. (2) Consensus need not arise. Even if we face a connected network, so that each pair of players are connected with a path, and each player on the path observes the actions of the player next to him, asymptotically one player may play actions which are perceived suboptimal by another player. (3) Nevertheless, if the network is connected then each player believes that her neighbors play asymptotically optimal.

Thus, player A may think that her neighbor B is using myopically optimal actions, know that B thinks that her neighbor C is using myopically optimal actions, and know that C thinks that A is using myopically optimal actions, yet, were A suggested to play an action that C plays, A might find out that this's action is not myopically optimal according to her information. This remains true *even* when players observe their own payoffs and their neighbors' payoffs.

The intuition for this result runs as follows. Suppose the three players A, B and C play along the game the actions a, b and c respectively infinitely many times. Denote by P_a (resp. P_b, P_c) the set of all beliefs over the state of the world for which the action a (resp. b, c) is optimal. Since player A believes that both her and player B's actions are optimal, the limit belief of player A is in both P_a and P_b. Similarly, the limit belief of player B is in both P_b and P_c, and the limit belief of player C is in both P_c and P_a. This does not imply that the limit belief of player A is in P_c, hence player A may view c, player C's action, as sub-optimal.

2 The Model

We consider games with incomplete information, in which identical players repeatedly choose an action, and receive a payoff that depends on their own action, and on the underlying state of the world.

The set of players is a finite set N. A directed graph $G = (N, E)$ with vertex set N describes the structure of the network: if the directed edge $i \to j \in E$ then player j is a neighbor of player i. The set of states of the world is a measurable space $(\Omega, \mathcal{A})$, endowed with a common prior $\mathbf{P}$. Time is discrete, and the set of stages is the set $\mathbf{N}$ of positive integers. At each stage n, each player i first receives a private signal s_n^i from some signal set S^i, then chooses an action a_n^i from her action set A^i, obtains a utility $u^i(\omega, a_n^i)$, and

*This research was supported by a grant from the Ministry of Science and Technology, Israel, and the Ministry of Research, France. We thank Ehud Lehrer for the comments he provided. This is an extended abstract of Rosenberg et al. (2006).

†Laboratoire d'Econométrie de l'Ecole Polytechnique, 1, rue Descartes 75005 Paris, France

finally observes the actions that his neighbors have just chosen $(a_n^j)_{j:\, i\to j\in E}$. Players discount future payoffs at the common rate $\delta \in [0,1)$. Players are identical, only in that they share the same action set $A := A^i$, the same signal set $S := S^i$, and the same utility function $u : \Omega \times A \to \mathbf{R}$. However, different players may receive different signals.

We impose the following technical assumptions:

- The common action set A is a compact metric space, endowed with the Borel σ-field.

- The common utility function $u : \Omega \times A \to \mathbf{R}$ is (jointly) measurable, and continuous over A for every fixed $\omega \in \Omega$. In addition, it satisfies the following boundedness condition: the highest payoff $\bar{u} : \omega \mapsto \max_{a\in A} u(\omega, a)$ is L_2-integrable.

- The signal set S is a measurable set. The signalling function maps past histories into probability distributions over S^N, the space of signal profiles. The past history at stage n is the complete list of the state of the world, and of the actions and signals of all players in all previous stages, hence, lies in $H_n := \Omega \times (S^N \times A^N)^{n-1}$. Technically, the signalling function at stage n is *any* transition probability[1] from H_n to S^N.

We emphasize that each player's utility function only depends on the underlying state of the world, and on one's own action, but does not depend on other players' actions. In that sense, the strategic interaction between players is purely informational: actions of player i may provide some information on player i's signals, and hence on the state of the world. Therefore, actions of player i are relevant to player j.

Example 2.1 ***Agree to disagree.*** *In the literature on agreeing to disagree, each player has a partition of the states of the world, and learns at the outset of the game the atom of the partition that contains the actual state of the world. The players then report the conditional probability of some fixed event E given their information, and update their belief using the reports of other players.*

This model can be embedded in our model as follows. Each player receives one signal at the outset of the game - the atom of the partition that contains the state of the world. Players are myopic, the action set is $[0,1]$, and the payoff function is defined so that the action p is optimal if and only if the conditional probability of E is p:

$$u(\omega, p) = -(1_E(\omega) - p)^2.$$

In our model players observe at every stage the actions of all their neighbors. This assumption can be weakened; if i is a neighbor of j, it is sufficient that in infinitely many stages i will observe the action of j, and j will know that i observes him. This allows us to represent the model of Geanakoplos and Polemarchakis (1982) in our setup.

In some models, as in Sebenius and Geanokoplos (1983), Parikh and Krasucki (1990), and Ménager (2006), at every stage each player has a given set of states of the world which are still possible, and each player reports some function of this set. If the partitions of the players satisfy some conditions, by properly defining the payoff function these models can be embedded in our model as well.

Example 2.2 ***Strategic Experimentation.*** *In the literature on strategic experimentation, see, e.g., Bolton and Harris (1999), Keller et al. (2005), Rosenberg et al. (2005) and Murto and Valimaki (2006), each player operates a multi-arm bandit. All bandit machines have the same unknown distribution. Players observe each other's actions, that is, which arm each player chose, and may or may not observe each other's payoffs.*

This model can be embedded in our model as follows. The state of the world ω represents the vector of types of the various arms. The set of actions available to each player is the set of arms. The signal each player observes is the payoff of the arm that he chose, and the payoff $u(\omega, a)$ is the expected payoff of the arm that he chose.

This example illustrates why it is critical to assume that payoffs may not be observed. Indeed, in some applications, such as strategic experimentation models, observing the payoff amounts to observing the *expected* payoff associated with (ω, a) – an assumption which is overly restrictive.

This example also illustrates how to incorporate random payoffs. We assume in this paper that the payoff to a player is a *deterministic* function $u(\omega, a)$ of the state of the world ω and of one own's action a, and that a player may only receive a noisy signal about this payoff. In some applications, such as strategic experimentation models, the payoff to a player is random, with an expectation that depends on ω and a. In such applications, it is typically assumed that the payoff is observed. Our model accommodates such situations by setting $u(\omega, a)$ to be the expectation of the

[1] A transition probability from X to Y is a function f that assigns for every $x \in X$ a probability distribution $f(x)$ over Y, such that for every measurable subset B of Y, the probability $f(x)[B]$ assigned to B is measurable in x.

payoff, and setting the signal of the player to include her actual (random) payoff.

The assumption that the set of possible signals is independent of the stage, and is the same for all players, is without loss of generality. Indeed, one may otherwise define S as the union of all the signal sets of all players in all stages.

2.1 Information and strategies

A play is an infinite list that contains the state of the world, the sequence of signals that all the players received, and the sequence of actions that all the players chose. The space of plays is $H_\infty := \Omega \times (S^N \times A^N)^{\mathbf{N}}$. A private history of player i at stage n is an element of $H_n^i := (S \times A)^{n-1} \times S$. The information at stage n can be described by a σ-algebra $\mathcal{H}_n^i$, which is spanned by sets of the form $(s_1^i, a_1^i, \ldots, s_{n-1}^i, a_{n-1}^i, s_n^i) \times (A \times S)^{\mathbf{N}}$.

A *(behavior) strategy* of player i is a sequence (σ_n^i), where σ_n^i assigns to every private history in H_n^i a probability distribution[2] over A. We denote by $\mathcal{H}_\infty^i$ the information of player i at the end of the game. It is the σ-algebra spanned by $(\mathcal{H}_n^i)_{n\in\mathbf{N}}$.

Any strategy profile σ, together with the common prior $\mathbf{P}$ on Ω, induces a probability distribution, $\mathbf{P}_\sigma$, over the set of plays H_∞. Expectation w.r.t. $\mathbf{P}_\sigma$ is denoted by $\mathbf{E}_\sigma$.

Given a strategy profile σ, and a stage $n \in \mathbf{N}$, we denote by q_n^i the conditional distribution over Ω given player i's information $\mathcal{H}_n^i$ at stage n. For a fixed (measurable) subset $F \subseteq \Omega$, the quantity $q_n^i(F)$ is the conditional probability that player i attached to the event F at stage n given his information. The sequence $q_n^i(F)$ is a bounded martingale, which converges, by the martingale convergence theorem, to $\mathbf{E}_\sigma[1_F \mid \mathcal{H}_\infty^i]$, $\mathbf{P}_\sigma$-a.s. We set $q_\infty^i(F) := \lim_{n\to\infty} q_n^i(F)$. One can verify that it is a probability distribution, to be interpreted as the limit belief of player i.

2.2 Main results

We focus on the asymptotic equilibrium behavior. For non-myopic players ($\delta > 0$), we use the Nash equilibrium notion. If $\delta = 0$, the Nash equilibrium criterion puts no restriction on the player's behavior beyond the first stage. In order to get asymptotic results for myopic players, we require that if $\delta = 0$ each player plays at *every* stage a myopically optimal action. In both cases, we will simply speak of equilibria and best-replies.

Given a strategy profile σ, a stage $n \geq 1$, and an action $a \in A$, we let $u(q_n^i, a) := \mathbf{E}_\sigma[u(\cdot, a) \mid \mathcal{H}_n^i]$ denote the expected payoff at stage n, when playing the action a. We also set $u_*(q_n^i) = \max_{a\in A} \mathbf{E}_\sigma[u(\cdot, a) \mid \mathcal{H}_n^i]$; similarly, $u_*(q_\infty^i) = \max_{a\in A} \mathbf{E}[u(\cdot, a) \mid \mathcal{H}_\infty^i]$. This is the highest (expected) payoff the player may obtain at stage n.

When playing optimally, a player faces a trade-off between *optimizing*, that is, playing an action which is myopically optimal in the light of the information accumulated so far, and *experimenting*, with the purpose of obtaining further information on ω.

Our first observation, Theorem 2.3 below, states that the number of stages in which a player experiments is bounded.

Given $\varepsilon > 0$, we denote by $N_i(\varepsilon)$ the number of stages in which a player plays an action which is suboptimal by at least ε:

$$N_i(\varepsilon) := \#\{n \in \mathbf{N} : u(q_n^i, a_n^i) \leq u_*(q_n^i) - \varepsilon\}.$$

As we show, $N_i(\varepsilon)$ is of the order $\delta/\varepsilon(1-\delta)$.

Theorem 2.3 *Suppose that $\delta > 0$, and let σ be an equilibrium. For every $\varepsilon > 0$, one has* $\mathbf{E}_\sigma[N_i(\varepsilon)] \leq 2\mathbf{E}[\bar{u}] \times \dfrac{\delta}{\varepsilon(1-\delta)}$.

In particular, the expected *discounted* number of stages in which player i plays a myopically ε-optimal action is bounded from above, independently of δ. That is, even taking discounting into account, there is always a non-negligible fraction of stages in which the player does not actively experiment.

Here is a rough intuition for Theorem 2.3. Whenever an ε-suboptimal action is played, the immediate loss of ε must be compensated for in the future, e.g., by a gain of $(1-\delta)\varepsilon/\delta$ in every future stage. Since the expected payoff is bounded by $\mathbf{E}[\bar{u}]$, the number of such compensations is at most $\mathbf{E}[\bar{u}]\delta/\varepsilon(1-\delta)$.

We next argue that the bound in Theorem 2.3 is tight.

Proposition 2.4 *Let $\delta > 0$ and $\varepsilon < 2\delta/(1-\delta)$ be given. There is a game and an equilibrium σ in this game such that* $\mathbf{E}_\sigma[N_i(\varepsilon)] \geq \mathbf{E}[\bar{u}](1-\varepsilon) \times \dfrac{\delta}{\varepsilon(1-\delta)}$.

By Theorem 2.3 the expected number of times a player plays an action which is not myopically ε-optimal is bounded. This implies that the players eventually play myopically optimal actions. This is the content of Theorem 2.5 below. We provide some definitions before stating this result.

Given a belief q, that is, a probability distribution over

[2] Formally, it is a transition probability from $(H_\infty, \mathcal{H}_n^i)$ to A.

Ω, the set of *myopically optimal* actions w.r.t. q is:[3]

$$BR(q) := \text{argmax}_{a\in A} \int u(\omega, a)q(d\omega) = \text{argmax}_{a\in A}\mathbf{E}_q[u(\cdot, a)].$$

An action $a \in A$ is a *limit action* of player i if it is a limit point of the sequence $(a_n^i)_{n\in\mathbf{N}}$ of the actions played by the player along the game.[4] We denote by A_*^i the set of limit actions. Since A is compact metric, A_*^i is compact and non-empty. Since the actions of the player depend on her information, A_*^i is a random variable[5] measurable w.r.t. the information of the player at infinity, $\mathcal{H}_\infty^i$.

Theorem 2.5 *Let σ be an equilibrium. Then for every player i,*

$$\mathbf{P}_\sigma(A_*^i \subseteq BR(q_\infty^i)) = 1.$$

According to Theorem 2.5, any action which is played infinitely often must be myopically optimal in the light of the information that is eventually available. All motives for experimentation eventually disappear. Plainly, that is not to say that the player then knows the state of the world.

Theorem 2.5 does not assert anything about the *existence* of optimal strategies. Indeed, without further continuity assumptions on the signalling functions, it is immediate to exhibit examples with no equilibrium. However, optimal strategies always exist if $\delta = 0$.

To obtain further results, we now impose some structure, and make assumptions on the observation structure. The graph G is *connected* if for every pair of players i and j there is a directed path from i to j.

Theorem 2.6 *Let σ be an equilibrium, and assume that G is connected. Then*

P1. $\mathbf{E}_\sigma[u_*(q_\infty^i)] = \mathbf{E}_\sigma[u_*(q_\infty^j)]$ *for every two players i and j.*

P2. $\mathbf{P}_\sigma(A_*^j \subseteq BR(q_\infty^i)) = 1$, *provided player j is a neighbor of player i.*

P3. *If each player observes his own payoffs, $u_*(q_\infty^i) = u_*(q_\infty^j)$ for every two players i and j.*

[3]By dominated convergence, the map $a \mapsto \mathbf{E}_q[u(\cdot, a)]$ is continuous. Hence, $BR(q)$ is non-empty for each q.

[4]That is, there exists an increasing sequence $(n_k)_{k\in\mathbf{N}}$ of stages such that $a = \lim_{k\to\infty} a_{n_k}^i$.

[5]The set of compact subsets of A is endowed with the usual, Hausdorf, distance.

If G is not connected, **P1** and **P2** still hold provided i and j belong to the same connected component of G.

Since the graph is connected, according to **P1** all players eventually perform equally well. Even if information is not divided equally, the relevant information spreads along the graph and guarantees asymptotic equality. The intuition behind **P1** relies on the so-called *imitation principle*. If player j is a neighbor of player i, he can do at least as well as player j does since he has the option of mimicking player j's behavior. Since the graph is connected, the result follows.

According to **P2**, player i eventually thinks that his neighbor j is playing in an optimal way: any action that j plays infinitely many times is optimal in player i's eyes. The intuition is that by the imitation principle, every limit action of j yields, in player i's eyes, at most as much as her own limit action yields. If with positive probability the limit action of the neighbor is strictly suboptimal in player i's eyes, the expectation $\mathbf{E}[u_*(q_\infty^j)]$ would be strictly below $\mathbf{E}[u_*(q_\infty^i)]$, which would violate **P1**.

P3 states that if the each player observes her payoffs, and if the graph is connected, then asymptotically all players receive the same payoff. The intuition is that since players observe their own payoffs, the imitation principle holds pathwise, so that the asymptotic payoff of player i's neighbor cannot exceed that of player i. Since the graph is connected the result follows.

P1 and **P3** hold for every pair of players, but **P2** holds only for neighbors. Indeed, it may well be that player i would think, if told, that a limit action of a player who is not her neighbor is sub-optimal. We stress that this negative result is not an artefact of strategic behavior, as it may happen even if all players are myopic.

Whereas **P2** and **P3** hold path-wise, **P1** holds in expected terms: if players do not observe their payoffs, *before* the beginning of the game each player can compute what her average stage payoff is going to be eventually; the result is the same for all players. It may happen that $u_*(q_\infty^i) \neq u_*(q_\infty^j)$ with positive probability. If the players observe their payoffs, according to **P3** the players eventually perform equally well pathwise: the limit payoff of all the players is the same.

3 Example

There are three players and two equally likely states of the world, ω_1 and ω_2. At stage 1, both players 2 and 3 receive an informative signal in $\{s_1, s_2\}$. The signal to player 2 reveals the state with probability 2/3: $\mathbf{P}(s_k \mid \omega_k) = 2/3$, for $k = 1, 2$, so that $\mathbf{P}(\omega_k \mid s_k) = 2/3$ as well. The signal to player 3 reveals the state with probability 5/6. No further information about ω is provided.

There are three actions, $A = \{a, b, c\}$. Denoting p the belief assigned to ω_1, the utility function u is such that action a is myopically optimal for $p \in [\frac{2}{7}, \frac{5}{7}]$, action b is myopically optimal for $p \in [0, \frac{2}{7}]$, and action c is myopically optimal for $p \in [\frac{5}{7}, 1]$.

At each stage $n > 1$, each player i observes only the decision of player $i+1$ (modulo 3) in the previous stage. We assume that players are myopic, and describe below one equilibrium profile (see Figure 2).[6]

signal	stage 1			stage 2			stage 3			stage 4		
	P1	P2	P3	P1	P2	P3	P1	P2	P3	P1	P2	P3
s_1s_1	$\frac{1}{2}$	$\frac{2}{3}$	$\frac{5}{6}$	$\frac{1}{2}$	$\frac{10}{11}$	$\frac{5}{6}$	$\frac{10}{11}$	$\frac{10}{11}$	$\frac{5}{6}$	$\frac{10}{11}$	$\frac{10}{11}$	$\frac{10}{11}$
	a	a	c	a	c	c	c	c	c	c	c	c
s_1s_2	$\frac{1}{2}$	$\frac{2}{3}$	$\frac{1}{6}$	$\frac{1}{2}$	$\frac{2}{7}$	$\frac{1}{6}$	$\frac{1}{2}$	$\frac{2}{7}$	$\frac{1}{6}$	$\frac{1}{2}$	$\frac{2}{7}$	$\frac{2}{7}$
	a	a	b	a	a	b	a	a	b	a	a	b
s_2s_1	$\frac{1}{2}$	$\frac{1}{3}$	$\frac{5}{6}$	$\frac{1}{2}$	$\frac{5}{7}$	$\frac{5}{6}$	$\frac{1}{2}$	$\frac{5}{7}$	$\frac{5}{6}$	$\frac{1}{2}$	$\frac{5}{7}$	$\frac{5}{7}$
	a	a	c	a	a	c	a	a	c	a	a	c
s_2s_2	$\frac{1}{2}$	$\frac{1}{3}$	$\frac{1}{6}$	$\frac{1}{2}$	$\frac{1}{11}$	$\frac{1}{6}$	$\frac{1}{11}$	$\frac{1}{11}$	$\frac{1}{6}$	$\frac{1}{11}$	$\frac{1}{11}$	$\frac{1}{11}$
	a	a	b	a	b	b	b	b	b	b	b	b

Figure 1: The strategies, and the evolution of the beliefs.

Stage 1: Player 1's prior belief assigns probability 1/2 to ω_1, hence she plays a. Player 2's posterior probability is either 1/3 or 2/3, depending on whether her signal is s_1 or s_2, hence she plays a. Player 3's posterior belief is either 1/6 or 5/6, hence she plays either b or c, depending on whether her signal is s_1 or s_2.

Stage 2: Players 1 and 3 hold the same belief as at stage 1, and therefore repeat their action. Player 2 infers player 3's signal from her action at stage 1, and she revises her belief accordingly.

If both signals are equal to s_1 (resp. s_2), player 2's posterior belief becomes

$$\frac{2}{3}\frac{5}{6} / \left(\frac{2}{3}\frac{5}{6} + \frac{1}{3}\frac{1}{6}\right) = \frac{10}{11} > \frac{5}{7}$$

(resp. equal to $1/11 < 2/7$). Hence player 2 switches to c (resp. to b).

If the signals of players 2 and 3 mismatch, player 2's new posterior belief is 5/7 if she received s_1, and 2/7 if she received s_2. In the former case, she is indifferent between a and c, whereas in the latter she is indifferent between a and b. In our equilibrium she plays a.

Stage 3: Players 2 and 3 hold the same belief as at stage 2. If the signals of players 2 and 3 match, the action of player 2 at stage 2 reveals the common signal, and player 1 revises her belief accordingly. If the two signals mismatch, the belief of player 1 remains 1/2.

Stage 4: Now only the beliefs of player 3 may change, but actions remain as at stage 3. After stage 3 beliefs and actions do not change.

In Figure 1, the signals received by players 2 and 3 appear in the left-most column. Subsequent columns describe the belief of each player at each stage and the players' actions.

Observe that the limit action of player 3 is either b or c. If the signals of players 2 and 3 differ, the limit belief of player 1 is $[\frac{1}{2}(\omega_1), \frac{1}{2}(\omega_2)]$. Hence, player 3's limit action is *not* optimal in the eyes of player 1. Moreover, in this case player 1's limit conditional payoff, $u_*(q^1_\infty)$, is 3/14, while the limit conditional payoff of players 2 and 3 is 3/7 (if the signals are s_1s_2) or 0 (if the signals are s_2s_1). Thus, there is a positive probability that $u_*(q^1_\infty) \neq u_*(q^2_\infty)$: the players need *not* agree about their limit payoff.

In this example the action chosen by players 2 and 3 when their conditional probability that ω_1 is the correct state is $\frac{2}{7}$ or $\frac{5}{7}$ differ. By adding additional states to Ω, one can construct an example in which players who have the same belief choose the same action.

Here player 1 would benefit if he observed the action of player 3; when the signals of players 2 and 3 differ his expected payoff would have increased. Thus, in this example players may benefit from having more neighbors. This phenomenon is not general, and there are examples in which adding more edges to the graph G make *all* players play an action b, which is inferior to some other action a that they would have played had those edges not been added to G.

[6] It can be checked that this profile is actually an equilibrium for every discount factor.

References

[1] Bala V. and S. Goyal (1998) Learning from Neighbors. *Review of Economic Studies*, **65**, 595-621.

[2] Bolton P. and C. Harris (1999) Strategic Experimentation. *Econometrica*, **67**, 349-374.

[3] DeMarzo P.M., D. Vayanos and J. Zwiebel (2003) Persuasion Bias, Social Influence, and Uni-Dimensional Opinions. *Quarterly Journal of Economics*, **118**, 909-968.

[4] Ellison G. and D. Fudenberg (1995) Word-of-Mouth Communication and Social Learning. *Quarterly Journal of Economics*, **110**(1), 93-125.

[5] Geanakoplos J. and H. Polemarchakis (1982) We Cant Disagree Forever. *Journal of Economic Theory*, **28**, 192-200.

[6] Goyal S. (2005) Learning in networks. In *Group Formation in Economics*, G. Demange and M. Wooders (eds), CUP.

[7] Keller G., S. Rady and M.W. Cripps (2005) Strategic Experimentation with Exponential Bandits. *Econometrica*, **73**, 39-68.

[8] Ménager L. (2006) Consensus, Communication and Knowledge: an Extension with Bayesian Agents. *Mathematical Social Science*, **51**(3),274-279.

[9] Murto P. and J. Valimaki (2006) Learning in a model of exit. Preprint.

[10] Parikh R. and P. Krasucki (1990) Communication, Consensus and Knowldge. *Journal of Economic Theory*, **52**, 178-189.

[11] Rosenberg D., E. Solan and N. Vieille (2005). Social Learning in One-Arm Bandit. *Econometrica*, forthcoming.

[12] Rosenberg D., E. Solan and N. Vieille (2006). Informational Externalities and Convergence of Behavior. Preprint.

[13] Sebenius J. and Geanakoplos J. (1983) Don't Bet On It: A Note on Contingent Agreements with Asymmetric Information. *Journal of the American Statistical Association*, **78**, 224-226.