Obtaining a unique, efficient norm through negotiation by linking proposals

Nicolas Quérou

Obtaining a Unique, Efficient Norm Through Negotiation by Linking Proposals

N. Quérou[1]

Résumé

On considère le probème du choix d'une norme entre groupes de taille importante. Tous les membres ont une chance égale de participer à la décision par le biais d'un processus de négociation dans lequel un représentant est choisi dans chaque groupe à chaque étape. On suppose que les groupes ont des estimations subjectives de l'importance de leurs adversaires dans la décision finale. Dans cet article nous prouvons que lier les propositions peut permettre aux différentes parties prenantes de définir une norme unique et efficace. Certaines implications des résultats principaux sont discutées dans le contexte du choix de standardisation.

Mots clé: groupements commerciaux, négociation, définition d'un standard.

Abstract

We consider the problem of the choice of a norm between groups of significant size. All members have an equal chance to participate to the decision through a negotiation process where at each round a representative is chosen from each group. The groups are assumed to have subjective estimates of the importance of their opponents in the final decision. We show that linking the proposals of the different parties enables them to design a unique, efficient norm. The question of choice of standardization motivates this study: some implications of main results are discussed in this context.

Keywords: commercial groups, negotiation, standard setting.

JEL numbers: D71, D82, L15.

[1]INRA, UMR LAMETA, 2, Place Viala, 34060 Montpellier Cedex 1, France, and University of Maine; E-mail: querou@ensam.inra.fr

1 Introduction

The problem of designing a common standard is a very important economic question. In network economy for example (see Shapiro and Varian (1998) for an introduction), let us consider the situation of communication networks. The question of choices of standardization is very important for numerical services, because of interconnection choices and problems of compatibility. Another example is related to commercial associations. Lack of standards concerning quality or quantity is often the main reason for the inefficiency of such formal networks. In fact, the Internet works because private groups have set interface standards, and commercial associations will try to design a common standard of quality in order to improve benefits.

Most of the litterature on the subject consider extreme situations for the control of standard setting. A first part considers settings where government bodies have played the key role in adjucating between the desires of different parties about possible standards (Farrell and Shapiro (1992)). A second part attributes no role for a standard setting organization (Farrell and Saloner (1985)). The main point of these studies is the absence of design compromises. But the importance of intermediaries in certifying firms has been widely recognized (see Diamond (1984), Lizzeri (1999)): standard setting organizations seem very important because they act as intermediaries and promote actively design *compromises*. In the context of commercial associations, the existence of an authority of mediation may enable the design of a compromise too. The present paper gives some support to this idea.

Some works have tried to solve the problem of designing a common norm by negotiation between different groups. Most methods are based on a static approach; if agents agree to cooperate, then they use a predetermined rule to design the standard. Sometimes the process is represented as a war of attrition between different parties with their own proposed standards (Simcoe (2003)). It results most of the time in problems related to the efficiency of the resulting standard, because such collective decision-making is done without a complete knowledge of the characteristics (the types) of the different groups, and without a possibility of real compromises.

The objective of the present paper is to understand how groups may design a common and efficient norm (standard) through the use of negotiation (see Muthoo (1999) for an introduction to bargaining theory). More precisely, the key insight of this paper rests on a simple point: linking proposals of the different groups may create a correct incentive for designing a unique and efficient standard. This fact suggests that an intermediary may enable the parties to design a common and efficient norm if its influence is such that each party admits the importance of its opponents in the decision process. We will develop this idea and provide indications on the relevance of the intervention of a formal intermediary in the choice of standardization, such as an authority of mediation between commercial associations, or standard setting organizations (see Lerner and Tirole (2004)). Thus, unlike to most of the literature on standard setting, the present work does promote an active role for intermediaries such as standard setting organizations.

Concerning the main characteristics of the study, we model the negotiation as a dynamic bargaining process (see Muthoo (1999)) under private information. The two different groups are assumed to interact by bargaining over a continuous set of alternatives. At the beginning of the process each population is privately informed about its own type and has a prior belief about the type of the other. At each round a representative is drawn at random from each population: both representatives play by making simultaneous proposals. The different groups update their beliefs using available information. It is considered that each member of a given group has an equal chance to become the representative of it at each round to highlight the idea that each member has to be involved (in some way) in collective action if one wants the resulting norm to be efficient. This situation can be thought of as a bargain between different (large sized) interest groups (see Van Winden (2003) for a survey of this notion).

The different groups are assumed to be interested in quick agreements only (they behave myopically). This may be due to tensions within a given group (as there may be little differences between the views of the different members, each representative is concerned with trying to obtain immediate agreement). Refusing a proposal is costly (because each group may give a certain importance to its opponent). This cost is related to the difference between the proposals of the different groups. As accepting a current proposal is costless (the "no effort" property, it implies no effort to implement the proposal), we make use of an entropic-form cost (see Kullback (1959)). We will explain how this cost may be interpreted in the context of standard setting in section 4. The different groups are assumed to assign a weight to this cost: it may be understood as a subjective estimate of the abilities of the opponent in negotiation (a subjective "political weight"). The main result of this paper states that if each party gives a sufficient power to its opponent in the decision process, then the negotiation process leads to a common and efficient agreement. In the context of standardization, the parties succeed in setting a standard with good properties by adjusting their views through negotiation.

Even though the focus of this paper is on standard setting through negotiation, some results may be related to the litterature on the emergence of norms in game or bargaining models under incomplete information, and on costly bargaining. This has been the object of an important number of studies (for example Young (1993b), Bicchieri et al. (1997), Binmore (1998)). The main part deals with how *social* norms are formed and maintained, using the evolutionary approach (Agastya (1997), Kandori (1992), and Young (1993a)). This theory suggests that norms emerge due to a process of adaptation. In the evolutionary models it is assumed that types are unobservable, or that actions can not be linked to types (so the principle is to best respond to samples of past proposals). In the present paper it is assumed that the sets of types are commonly known, and the main question is on how a *unique and efficient* norm may be *set* (may one give a correct incentive to the different parties?).

The results of the present work are different from those of some litterature on costly bargaining (which states mainly that agreements do not exist, see Anderlini and Felli (2001)), but this kind of studies has assumed that agents are perfectly rational, and does not consider costs as an incentive for designing a common norm.

The paper is organized as follows: section 2 presents the main assumptions and the model in a general setting, and we provide arguments for studying the problem of standard setting within this framework. The convergence of beliefs and proposals in the negotiation process is discussed in section 3. Section 4 deals with the uniqueness and efficiency of the long run norm. A discussion about a more active role for an authority of mediation in standard setting is provided, and the case of a discrete set of alternatives is explained. Section 5 concludes.

2 The model

2.1 The foundations of the decision process

We consider two distinct groups of agents engaged in a negotiation process in order to select a norm (standard). These populations are large and of different types. Let us denote it by groups 1 and 2. All members of a given population have the same utility function. Each population i is characterized

by an attribute vector t_i in some set T_i, which specifies those parts of population i's utility function that may be unknown to group j. T_i is assumed to be a complete and separable metric space.
Each representative chooses a proposal from a common alternative space X, assumed to be a compact and convex subset of $\mathbf{R}^*_+$. The utility function of group i is the continuously differentiable, concave and strictly quasi concave function $u_i : X \times T_i \to \mathbf{R}$. It is assumed that agent i of type t_i has a reservation payoff u^{dflt}_{i,t_i}. It is the minimum payoff that i (of type t_i) must be proposed to be at least as well off as she would be at the disagreement payoff. We assume that there exist alternatives that yield a strictly better payoff than u^{dflt}_{i,t_i}. Let us define the application $r_i : T_i \to \mathbf{R}$, such that

$$r_i(t_i) = u^{dflt}_{i,t_i}.$$

In what follows we wish to think of higher types as more competitive groups. Therefore we make the following assumption that associates higher types t_i with higher levels of reservation amounts:

Assumption 1. $r_i : T_i \to \mathbf{R}$ is a strictly increasing function, $i = 1, 2$.

So, if type t_i of group i prefers the disagreement payoff to a proposal corresponding to alternative x, then any type $t'_i > t_i$ strictly prefers the disagreement payoff to the one resulting from x.
We will define in the subsequent sections the expected utility for i at each proposal round as a function of its own attribute vector $t_i \in T_i$, its own action and the actions of group j. It is assumed that the parameters t_i do not evolve during the process, and that each group is privately informed of its own type (at the beginning).
The process is dynamic: periods of proposals are followed by periods of acceptance/refusal. At each round of proposal k one agent is drawn at random from each population; this agent is "elected" as a representative of her own group. Representatives engage in negotiation by making simultaneous proposals. This way of making proposals enables one to avoid problems related to incentive compatibility. If proposals were made sequentially, agents might have an incentive to pretend to be of a different type than their true one.
In the following period $k+1$ acceptance is discussed. If at least one proposal is accepted, a standard has been found and the negotiation ends; otherwise, each representative is replaced by a member of the same population and a new round of proposal has to be held at period $k+2$... The assumption that each representative is replaced at each round reinforces the idea that each member has to be involved in the decision process if one wants to have a

chance to design an efficient standard; however it is restrictive, and a good possible extension to the present work would be to relax it in order to evaluate its impact on the results.
Before presenting the (general) model, we will indicate how the problem of standard setting may be presented in this framework.

2.2 Choosing a standard through negotiation

In the present paper it is considered that different groups engage in negotiation to design a norm. Section 4 will focus on the selection of a unique norm, a common standard. The present negotiation is not a war of attrition where each group either wins (its standard is accepted) or loses (the other standard is accepted), it is more about designing an efficient compromise. Some people may be surprised that the set of alternative be continuous, but in practice the different groups can choose, and the intermediary can try to alter, various characteristics of the proposed standards (see Lerner and Tirole (2004) for more on this issue in the context of standard setting organizations). So, this gradualism makes the present assumption relevant. However, in order to be as precise as possible we will discuss the case of a discrete set of alternatives in section 4.3. The role of standard setting organization is mainly restricted to no role in the design of standard in most studies; we will elaborate on a possible role in section 4.2. Now we will define the dynamics of negotiation by understanding the evolution of beliefs.

2.3 Definition and evolution of beliefs

From now on by "agent i" we mean "representative of group i in the current round". At the beginning of date k agent i will be characterized by her own type t_i and her beliefs about the type of the other agent $\mu_{i,k} \in \Delta(T_j)$ (where $\Delta(T_j)$ denotes the space of probability measures of support T_j). In fact we consider that, when there is no agreement at date $k+1$, each representative i is replaced by a member of the same population with the same belief over the agent j: combined with the assumption that both populations are large, this enables us to consider that each representative is myopic, there is no interest for the possibility of delaying agreement, they are only concerned by obtaining an agreement at the current period. Due to this simple framework, this belief uniquely specifies the belief over the date $k+2$ acceptance/refusal of agent j: it specifies the subjective probability that the proposal x made by agent i at period k would be accepted by agent j. We will denote this probability by $P_i(x, \mu_{i,k})$.
It is assumed that agents update their beliefs using Bayes' rule when possible.

Let us briefly explain the rules. At the beginning of the process, agent i is privately informed of her own type t_i. F_{ik} is the set (σ-algebra) representing the information of agent i at the beginning of date k (after both proposals of period $k-2$ have been refused at date $k-1$). This information would come from the observations of agent i from periods 0 through $k-1$: it would consist of a sequence of refused proposals by agent j, and a sequence of proposals made by agent j that agent i refused. Before making a proposal at date k, agent i's belief about the true type is represented by the value of the probability $\mu_{i,0}$ conditional on F_{ik}: we will have

$$\mu_{i,0}(.|F_{ik})(.) = \mu_{i,k}(.).$$

We shall assume that the collection of subjective prior beliefs of agents are such that $\mu_{1,0}$ and $\mu_{2,0}$ are mutually absolutely continuous, i.e.,

$$\forall D \subset \Omega, \mu_{1,0}(D) \;>\; 0 \; iff \; \mu_{2,0}(D) \;>\; 0.$$

This condition requires that agents agree ex ante about the events which have zero probability. We would like to emphasize that this condition is weaker than the common prior assumption, which is very unlikely to be found in a situation of limited rationality. The prior beliefs of each group might be quite different from the true distribution of types.

2.4 Expectations and costs

An agent i with type t_i and beliefs $\mu_{i,k}$ who makes a proposal x_i will expect a date k utility equal to

$$P_i(x_i, \mu_{i,k}) u_i(x_i, t_i),$$

where $P_i(x_i, \mu_{i,k})$ is the probability that x_i will be accepted according to agent i's beliefs over the type of agent j at round k. In the rest of the paper we will simply denote this probability by $P_{i,k}(x_i)$. The decision rule at each proposal round k is defined as follows. If no proposal has been accepted at period $k-1$, each agent undergoes a cost to make a new proposal at round k; the cost of agent i for making a new proposal $x(t_i)$ is a function of $x(t_i)$ and of the past proposal of agent j $x_{j,k-2}$. Let us denote it by $C_i(x(t_i), x_{j,k-2})$. So at each round k the proposal of agent i $x_{i,k}(t_i, \mu_{i,k})$ maximizes her expected utility U_i, i.e.

$$U_i(x_{i,k}(t_i), \mu_{i,k}) = \max_{x \in X} P_{i,k}(x) u_i(x, t_i) - C_i(x, x_{j,k-2}).$$

The cost C_i is assumed to be of the relative entropy form (see Kullback (1959)):

$$C_i(x, x_{j,k-2}) = \varepsilon_i x ln(\frac{x}{x_{j,k-2}}), \forall x \in X.$$

The parameter $\varepsilon_i > 0$ is used to represent the (subjective) estimate for agent i of the political weight of agent j in the negotiation process: the larger the parameter, the more important the cost for agent i when she chooses her proposal (so, the more she takes agent j into account). In section 4 we will provide a possible explanation of this cost in the context of choice of standardization, by considering the possibility of the intervention of an authority of mediation.
Refusing $x_{j,k-2}$ implies (for agent i) going to another round of negotiation, which is costly. We choose to model this cost by using an entropic form because entropy is commonly used in economics and possesses the important "no effort" property: $C_i(x, x_{j,k-2}) = 0$ if and only if $x = x_{j,k-2}$. Thus, no effort is needed from agent i to accept $x_{j,k-2}$, whereas making another proposal is costly.

Remark 2.1. *One may explicitly write the expected utility as*

$$U_i(x_i, x_{j,k-2}, t_i, \mu_{i,k});$$

however, since $x_{j,k-2}$ is not the current proposal of agent j, we will not use this notation before the section about the limiting behavior of agents (section 4).

Several points about the process and some technical assumptions have to be underlined.

Remark 2.2.
- *It is easily verified that the cost functions are strictly convex.*
- *As we will focus on the case where the subjective estimates are large, we may assume in the rest of the paper that the expected utility functions are strictly concave. Due to strict convexity of the cost functions, it is easily verified that, when ε_i is large enough, the second derivative of the expected utility function of agent i is strictly negative.*

Now a simple application of Berge's maximum theorem (see Sundaram (1996)) leads to the following result:

Lemma 2.1. *Under the assumptions of the problem, a unique solution to the decision problem of each agent exists. Moreover, the proposal of agent i at date k is a continuous function of her type and beliefs.*

Since the existence of a common prior is not assumed, the agents behave as being boundedly rational; they optimize, but given subjective beliefs, and using myopic rules. The assumption of a subjective estimate of "political weights" reinforces the idea of not perfectly rational agents. Now we will focus on the understanding of the dynamics of possible agreements.

3 The dynamics of agreements

3.1 A naive rule

The remark of the previous section about the dynamic process enables us to describe more precisely the negotiation at each date k. If an agreement has not been reached at date $k-1$, each representative is replaced by another member of her original population: the interaction relies on simple beliefs about acceptance of the opponent, and the process seems to be adaptive. It is no surprise that we obtain an immediate result that presents the very simple structure of the (Bayesian) equilibrium outcome of the process at some fixed date.

Proposition 3.1. *Under assumption 1 any Bayesian equilibrium outcome of the negotiation process at date k has the following structure: agent i makes a proposal $x(t_i, \mu_{i,k})$, where $x(t_i, \mu_{i,k})$ solves the decision problem of agent i at date k. Agent j accepts if her true type t_j is such that*

$$u_{j,t_j}^{dflt} < u_j(x(t_i, \mu_{i,k}), t_j) \tag{1}$$

and rejects if

$$u_{j,t_j}^{dflt} > u_j(x(t_i, \mu_{i,k}), t_j). \tag{2}$$

The proof of this statement is immediate due to the very simple structure of the bargaining process and is therefore omitted. The idea is very intuitive: agents act simply, accepting any proposal that would give them at least what they would get with their reservation proposal. An immediate consequence of this result is that rejections occur with positive probability: any agent would refuse a proposal that would give her strictly less than her reservation payoff.

Remark 3.1. *In fact, one can deduce from this result that, in order to have a chance to see her proposal accepted at round k, agent i has to propose $x(t_i, \mu_{i,k})$ such that the resulting payoff of agent j is greater than her expected reservation payoff with respect to $\mu_{i,k}$. That is, the belief of agent i at round k concerning the reservation payoff of agent j has to be greater than her true* *value. So in the dynamic process each agent is solely concerned with trying to learn the type of the other agent in order to be able to make the best possible proposal for herself (i.e., the proposal that would give her the highest payoff in the set of proposals that would get accepted with probability* 1*).*

An important question is to understand if negotiation will stabilize over time: if the beliefs and resulting proposals converge almost surely as time goes to infinity, we will be able to study the long-term equilibria of this process by simply studying the Nash equilibria of the limiting game. We will provide an answer to this question in the next subsection.

3.2 Stability of long-run behavior

We have seen previously that at each fixed date the structure of the bargaining game is very simple: agents make proposals with respect to their own type and their beliefs about the type of their opponents. The learning dynamics is related to types only: the beliefs about types vary over time because of the information received. However in the long run this information has smaller and smaller influence on the beliefs about the fixed distribution of types between the groups. We will show that each agent's beliefs about the type of her opponent converge over time; next we will see that in the long run agent's proposals are optimal with respect to their limiting beliefs, which is quite expectable. Things seems to happen in the right way: the different views of groups will finally adjust to long run standards.
We shall use the mutually absolute continuity of prior beliefs to obtain the convergence result. This result will hold almost surely, i.e., with probability one with respect to a measure μ which is mutually absolute continuous to each of the $\mu_{i,0}$. The true distribution μ^T generated by the ex ante distribution of types and the true behavior of the agents is an appropriate choice of such a measure. It is not claimed that beliefs converge to the "true" types of agents.
Let us denote by $\mu_{i,k}(w)$ the date k beliefs of group i over the type of group j in the sample path w, and $\mu_{i,\infty}(w)$ her limiting beliefs (i.e., beliefs under the limiting information field $F_{i\infty} = \cup_k F_{i,k}$) at w. We will consider the limit of probability measures in the weak topology. Let us define for each $i = 1, 2$ and each k,

$$C_{i,k} = \{w|\mu_{i,k}(w) \rightharpoonup \mu_{i,\infty}(w)\}, \qquad C = \cap_i \cap_k C_{i,k}.$$

The result is the following:

Theorem 1. *Under the assumptions of the study, we know that the sequence of beliefs of each group converges almost surely: we have* $\mu^T(C) = 1$.

Proof. See appendix. □

This result implies that the sequence $\{\mu_{i,k}\}_k$converges almost surely to the belief under the limiting information field. Next we will specify an important property concerning the set of proposals which are chosen infinitely often (long run proposals): this set is indeed optimal with respect to the limiting beliefs.

Proposition 3.2. *Let us consider w a sample path belonging to the set C.*
a) If x' is a proposal which is chosen infinitely often on w by group i then

$$x' \in \arg\max_{x \in X} U_i(x, t_i, \mu_{i,\infty}(w)).$$

b)The payoff of group i converges on each such sample path.

Proof. See appendix. □

To sum-up the results of the present section, we have established the very simple structure of the equilibrium outcomes of the game at each fixed date, and then we have shown that the dynamics stabilizes over time. Thus we deduce that the process converges to a limiting negotiation game. Now it is possible to study the long term equilibrium outcomes that are likely to be (Nash) equilibria of the limiting game. Indeed it is possible for the moment that the proposals made by each group lead to the emergence of long run norms that will be different according to each group. The study of this problem is the goal of next section.

4 Adjusting views to a unique, efficient norm

In the present section we will establish the main result of this study: if each group attributes sufficient importance to its opponent in the negotiation process, then in the long run the proposals and beliefs of the different groups result in a unique and efficient norm. Thus by linking proposals it is possible to help parties designing efficient standards. The main assumptions seem to be plausible ones: in real negotiation making a proposal has a cost, and the subjectivity of different interest groups plays an important role in the decision process.

4.1 Main result

Until the end of the section we consider the limiting problem, where the utility of agent i is function of the long run proposals $x_{i,\infty}$ and $x_{j,\infty}$. We will try to obtain the Nash equilibrium of the limiting problem, in order to be able to specify stable agreements (with respect to unilateral deviations).

For the proof of the main result we need condition (3):
there exists $x \in X$ such that, for any i, j in I, $i \neq j$, we have

$$u_i(x, t_i) \geq E_{\mu_{j,\infty}} u_i^{dflt} > u_{i,t_i}^{dflt}, \tag{3}$$

where $E_{\mu_{j,\infty}} u_i^{dflt}$ is the expected reservation value of agent i for agent j with respect to her limiting beliefs (the subjective estimate for agent j of the reservation value of agent i), and t_i her true type. This condition states that there exists an alternative which is better for each agent than the limiting estimation (for her opponent) of her reservation payoff. It is like a condition on the compatibility of limiting beliefs (it is not required that beliefs converge to the true distribution of types). Now we state the main result of the paper: the limiting equilibrium consists of a proposal shared by all agents and this alternative is efficient.

Theorem 2. *If ε_i is chosen large enough and if condition (3) is fulfilled, then there exists a unique (pure) steady-state (Nash) equilibrium, and it is $(x_{i,\infty}, x_{j,\infty})$. Moreover, $x_{i,\infty} = x_{j,\infty} = x_S$ and x_S is Pareto-optimal.*

Proof. See appendix. □

Basically we obtain that, if the different groups admit that their opponent has an important weight in bargaining then the negotiation process will lead to a common proposal that will be efficient. This is an important property because of two points: selection (of a unique proposal) and efficiency. If we think of this process as a negotiation between different associations in order to agree on a choice of standardization, now we may develop some ideas concerning a possible role for an intermediary in standard setting. This is the goal of section 4.2.

4.2 The need for a formal authority in standard setting

It was proven that an informational cost may create a correct incentive for groups to set a common and efficient norm through negotiation. The result seems to lead to the conclusion that an efficient standard may be set if each party admits the importance of its opponent in the negotiation process. The

existence of a real authority would enable to achieve this.
The presence of an authority of mediation accepted by all parties may enable this cost to be taken into account, and the importance of the other group to be more considered (bigger subjective political weights) during negotiation. By comparing the different proposals at each round, the authority may propose adjustments of their characteristics: taking account of this opinion, each group will undergo a cost in order to make a new proposal that will have a better chance to be accepted. This role is more significant than that of a "meeting room".
The results of the present study give some interesting elements on the possibility of a more active role of a formal intermediary (such as standard setting organizations) for designing a common and efficient standard through negotiation. There is little work on the role of intermediaries in standard setting. One could cite Lerner and Tirole (2004), who study another very interesting aspect of the question, how firms should choose between different competing standard setting organizations.

4.3 A finite set of alternatives

It was assumed that the set of alternatives was compact and convex. But in fact in the different results except theorem 2 it is easily verified that the only necessary assumption is compactness; thus, these results still hold in the context of a finite set of alternatives.
Concerning theorem 2, one has to be a little bit more restrictive on the statement of this result in the discrete case. But the same arguments can be used. In fact the only modification is the following: if one alternative Pareto dominates all other ones, then this alternative will be the final standard. Thus, even if the different groups have very different standards at the beginning, they will adjust to an efficient collective compromise.

5 Conclusion

In this paper new results are obtained regarding the choice of standardization when there is private information. The use of continued bargaining interaction over time enables agents to better design the long run norm, which results in efficiency. The fact that each member of the different groups is involved in the decision process seems to help reaching efficiency.
A very interesting information is that linking the different proposals may lead to the selection of efficient norm: this fact seems to yield interesting ideas about possible roles for intermediary in standard setting. The presence

of an authority recognized by all groups in negotiation may enable parties to admit the importance of their opponents in the decision process and to help designing a compromise.
The present study is appropriate more to the context of standard setting between large commercial associations. To describe a framework that will take account of a restricted number of members, one would have to consider the possibility that a particular member be the representative for the whole process (or that different members have different chances to become representatives). One will have to understand the impact of time preferences (and anticipations) on the present results. This point is left for future research.

References

[1] M. Agastya (1997), "Adaptive Play in Multiplayer Bargaining Situations", Review of Economic Studies 64, 411-426.

[2] Anderlini, L. and L. Felli (2001), "Costly Bargaining and Renegotiation", Econometrica 69, 377-411.

[3] Bicchieri, C., Skyms, B. and R. Jeffrey (ed.) (1997), *The Dynamics of Norms*, Cambridge University Press, Cambridge.

[4] K. Binmore (1998), *Game Theory and the Social Contract II - Just Playing*, The MIT Press, Cambridge Mass.

[5] D. Diamond (1984), "Financial Intermediation and Delegated Monitoring", Review of Economic Studies 51, 393-414.

[6] Farrell, J. and G. Saloner (1985), "Standardization, Compatibility, and Innovation", Rand Journal of Economics 16, 70-83.

[7] Farrell, J. and C. Shapiro (1992), "Standard Setting in High-Definition Television", Brookings Papers on Economic Activity: Microeconomics", 1-93.

[8] M. Kandori (1992), "Social Norms and Community Enforcement", Review of Economic Studies 59, 61-80.

[9] S. Kullback (1959), *Information Theory and Statistics*, Wiley, New York.

[10] Lerner, J. and J. Tirole (2004), "A Model of Forum Shopping, with Special Reference to Standard Setting Organizations", National Bureau of Economic Research Working Paper n^o 10664.

[11] A. Lizzeri (1999), "Information Revelation and Certification Intermediaries", Rand Journal of Economics 30, 214-231.

[12] A. Muthoo (1999), *Bargaining Theory*, Cambridge University Press, Cambridge.

[13] Shapiro, C. and H. Varian (1998), *Information Rules: A Strategic Guide to the Network Economy*, Harvard Business School Press.

[14] T. Simcoe (2003), "Committees and the Creation of Technical Standards", mimeo, University of California at Berkeley.

[15] R. K. Sundaram (1996), *A First Course in Optimization Theory*, Cambridge University Press, Cambridge, UK.

[16] F. Van Winden (2003), "Interest Group Behavior and influence", in C.K. Rowley and F. Schneider (eds.), Encyclopedia of Public Choices, Boston, Kluwer Academic Publishers, forthcoming.

[17] H. P. Young (1993a), "An Evolutionary Model of Bargaining", Journal of Economic Theory, 59, pp.145-68.

[18] H. P. Young (1993b), "The Evolution of Conventions," Econometrica 61 (1), 57-84.

6 Appendix

6.1 Proof of theorem 1

Proof. We proove the result by using a general reasoning on $T = T_1 \times T_2$: using an abuse of notations we will denote by $\mu_{i,k}$ the measure on T (and not on T_j) obtained by extending the belief $\mu_{i,k}$ to T in a natural (and unique) way (since agent i knows her true type t_i for sure at the beginning of the process).

Now using standard arguments it is enough easily verified that, as T is a complete and separable metric space, there exists a countable set, $F(T)$, of uniformly bounded continuous real-valued functions on T which are such that the weak convergence of a sequence of probability measures $\{p_k\}_k$ to a measure p holds whenever $\lim_{k\to\infty} \int f dp_k = \int f dp$ for each f in $F(T)$. Let us consider the set Ω which describes the randomness in the negotiation process, and $f : T \to \mathbf{R}$ a continuous function. Now $f(t_1, t_2)$ may be thought of as a bounded random variable on Ω; then a direct consequence of the Martingale Convergence Theorem leads to

$$\lim_{k\to\infty} \int f(t)\mu_{i,k}(dt) = \int f(t)\mu_{i,\infty}(dt).$$

This can be made to hold with probability one for any countable collection of such functions f and hence for all $f \in F(T)$. Then we have $\mu_{i,k} \rightharpoonup \mu_{i,\infty}$ with $\mu_{i,0}$ probability one, so $\mu_{i,0}(\cap_k C_{i,k}) = 1$ for each $i \in I$, and from the definition of μ^T, this implies that $\mu^T(\cap_k C_{i,k}) = 1$ for all i and hence $\mu^T(C) = 1$. □

6.2 Proof of proposition 3.2

Proof. Let $x \in X$. Since x' is chosen infinitely often on w, there exists a subsequence $\{k_n\}$ such that

$$U_i(x', t_i, \mu_{i,k_n}(w)) \succeq U_i(x, t_i, \mu_{i,k_n}(w)).$$

Taking limits and using the continuity of U_i, we get

$$U_i(x', t_i, \mu_{i,\infty}(w)) \succeq U_i(x, t_i, \mu_{i,\infty}(w)).$$

Since x is arbitrary, this proves statement (a). The second assertion follows from the maximum theorem and part (a). □

6.3 Proof of theorem 2

Proof. By proposition 3.2 we know that the limiting proposal $x_{i,\infty}$ of agent i is such that

$$x_{i,\infty} = argmax_{x \in X} P_{i,\infty}(x) u_i(x, t_i) - \varepsilon_i x ln(\frac{x}{x_{j,\infty}}), \tag{4}$$

where $P_{i,\infty}(x)$ is the limiting probability that proposal x would be accepted by agent j. From lemma 2.1 we deduce that $(x_{i,\infty}, x_{j,\infty})$ is a Nash equilibrium of the limiting problem.
Using condition (3) we deduce by concavity and continuity of the utility functions that the set

$$\{x \in X | u_j(x, t_j) \geq E_{\mu_{i,\infty}} u_j^{dflt} \forall i, j, i \neq j\}$$

is not empty, convex, compact. So, focusing on agreement obtained with probability 1, we deduce from proposition 3.1 (and remark 3.1) that $x_{i,\infty}$ is solution of

$$x_{i,\infty} = \operatorname{argmax}_{\{x \in X | u_j(x,t_j) \geq E_{\mu_{i,\infty}} u_j^{dflt} \forall i,j,i \neq j\}} u_i(x, t_i) - \varepsilon_i x ln(\frac{x}{x_{j,\infty}}). \tag{5}$$

Now one may consider the case of $x_{i,\infty} > x_{j,\infty}$ (otherwise we consider the problem for agent j). If ε_i is chosen large enough, one obtains necessarily that $x_{i,\infty} = x_{j,\infty}$. For example, one may consider ε_i such that

$$\varepsilon_i \geq \frac{u_i(x_i^*, t_i) - u_i(y_i^*, t_i)}{M ln(\frac{a}{b})},$$

where $M > 0$ is such that $x \leq M$ for any $x \in X$, (x_i^*, y_i^*) is the maximizer on X^2 of $(x, y) \to u_i(x, t_i) - u_i(y, t_i)$, which exists because u_i is continuous and X^2 is compact, and $a, b \in X$ are such that $ln(\frac{x}{y}) \leq ln(\frac{a}{b})$ for all $x, y \in X$. Then using characterization (5) and reasoning by contradiction lead quickly to the Pareto optimality of the resulting agreement. □