

SONDERFORSCHUNGSBEREICH 303

„Information und die Koordination wirtschaftlicher Aktivitäten“



Discussion Paper No. A-304

Switching Away From Probability One Beliefs

Georg Nöldeke
Eric van Damme

July 1990

We would like to thank Dieter Balkenborg and Klaus Schmidt for many helpful comments on earlier versions of this paper. Financial support from the Deutsche Forschungsgemeinschaft through SFB 303 is gratefully acknowledged.

Switching Away From Probability One Beliefs

Abstract:

This paper considers a class of repeated signalling games to gain some intuitive insights into the effects and the desirability of modelling players in a dynamic game of incomplete information as being obstinate in the sense that their beliefs satisfy a support restriction. We demonstrate that such a restriction is rather dubious on a-priori grounds and in general imposes "too much" pooling on sequential equilibrium outcomes. Equilibria violating a support restriction should therefore not be dismissed in dynamic models of incomplete information and may actually reflect the possibility of reputation effects present in such a setting.

Georg Nöldeke
Department of Economics
University of Bonn
Adenauerallee 24-26
5300 Bonn 1
West-Germany

Eric van Damme
Center for Economic Research
Tilburg University
5000 Le Tilburg
Netherlands

1. Introduction

By requiring the specification of a belief for every information set of a game the sequential equilibrium concept of Kreps and Wilson [18] greatly expands the scope for applying backwards induction reasoning to the analysis of games of incomplete or imperfect information.

Consider for example the famous analysis of the chain-store game with incomplete information in Kreps and Wilson [19]: The game analysed in that paper has no proper subgames at all, but every information set of the entrant (the uninformed player) is the root of what Kreps [17] calls an "almost proper subgame". Once a belief is specified at such an information set, one can think of replacing all previous nodes in the game with a single chance move of nature generating the pre-specified beliefs at this information set and thereby defining a game of incomplete information. Determining the solutions of such games for all possible beliefs then allows the analysis to proceed at least partly by backwards induction by replacing such subgames by their solutions (conditional on the belief at the beginning of the subgame).

Proceeding in this fashion, which many analyses of games with incomplete information in the economics literature do, makes it necessary to deal with an issue which may appear to be a technicality: the definition of an extensive form given for example by Kreps and Wilson [18] requires that all moves of nature have strictly positive probability. The backwards induction procedure sketched above however makes it necessary to solve "subgames" for every possible belief, including those which assign probability zero to a subset of nodes in the information set. One apparently natural way to resolve this issue is to delete all those zero probability nodes and their successors from the game tree under consideration and solve the resulting game. In the well known incomplete information bargaining model of Rubinstein [23] for instance this kind of reasoning is used and implies that the solution of a "subgame" in which the uninformed party assigns probability one to the informed player being "strong" is given by the solution of the corresponding perfect information bargaining game.

In this paper we will argue that great care should be taken in applying this approach to solve games of incomplete information. Eliminating zero

probability nodes to solve subgames which are not well-defined restricts the analysis to sequential equilibria which satisfy a support restriction on the beliefs a player may hold. On an informal level such a restriction amounts to the following: If, at a given information set h , a player assigns probability zero to a subset of nodes X_h in this information set, then in the continuation of the game he never assigns strictly positive probability to any node which can only be reached by a path through X_h - unless this requirement would force him to assign probability zero to all nodes consistent with the observed history.

Notice that this requirement essentially models players as being obstinate. Once they ruled out a possible explanation for what they observed in the game so far, they will not reconsider this possibility, as long as they can possibly avoid doing so.¹

Clearly, given a Nash equilibrium, this property holds for all nodes in information sets which are reached with strictly positive probability on the equilibrium path. Imposing the support restriction as a condition on the "updating rule" a player may use (as, for example Grossman and Perry [14] do in their definition of perfect sequential equilibrium) is probably motivated by this fact. Identifying restrictions which have to hold for beliefs on the equilibrium path and requiring those restrictions also off the the equilibrium path seems like a sensible approach to rule out "threatening through beliefs". This, after all, is the main concern of the heuristic equilibrium refinement literature, as developed for example in Cho and Kreps [6].

Consequently one might be tempted to argue that assuming the support restriction is not only convenient in applying backwards induction arguments to solve games of incomplete information, but actually constitutes a sensible refinement of the sequential equilibrium concept in itself. As demonstrated by Madrigal, Tan and Werlang [22] it can however happen that the unique Nash equilibrium of a given extensive form game can not be supported by beliefs satisfying this restriction. This observation clearly points out that imposing the support restriction is problematic, however it is not quite clear what to conclude from this. Does this failure of existence imply that, while basically a sensible requirement, the restriction is "too strong" in some cases? Or does it mean that there is something more fundamentally wrong with insisting that

players should be obstinate?

Addressing this kind of question will be one of the main concerns of this paper. In the following section we try to argue by example that there is no compelling reason for imposing the support restriction on a-priori grounds. The example which will be analysed is similar to, but simpler, than the one considered by Madrigal, Tan and Werlang [22]. The simplicity of our example makes it more transparent why assuming that players are obstinate may lead to non-existence of equilibrium and that it restricts their updating rules in a rather non-intuitive way. Furthermore the example serves to show that the non-existence of a Nash equilibrium satisfying the support restriction may hinge on strategically irrelevant details of the extensive form representation. Thus the support restriction is not compatible with normal form reasoning. The kind of extensive form reasoning embodied in the support restriction may seriously lead the analysis astray in cases where arguments based on the normal form (iterated elimination of weakly dominated strategies) give nice and intuitive conclusions.

Nevertheless, faced with a multiplicity of equilibria and the intractability of either a normal form analysis or a full analysis of the extensive form, it could still be tempting to concentrate on equilibria satisfying the support restriction, if this requirement does not result in non-existence. As we noted above, any analysis using the perfect sequential equilibrium concept implicitly uses this simplification.² In addition there are of course the papers in which the support restriction is used directly; see Madrigal, Tan and Werlang [22] for references.³

We will argue in the following that this kind of response to the problem is inappropriate: it is possible to pinpoint a general consequence of the support restriction, which also lies at the heart of the non-existence problem, and forces a particular structure on the equilibria surviving this "refinement" - eliminating equilibria which may most adequately reflect our intuition about the game under consideration⁴. Violations of the support restriction may represent a sensible reasoning process which supports interesting equilibrium behaviour.

To make these points it is helpful to have a somewhat general model which at

the same time allows us to focus on the updating of beliefs as the main object of interest and is simple enough to allow for intuitive interpretation of the results. In almost all of the refinement literature the model of a signalling game has proven to be ideally suited for such purposes. In section 3 we therefore introduce the model of a repeated signalling game, which is a straightforward extension of the standard signalling model.

Section 4 presents an analysis of repeated signalling games, introducing the idea that violations of the support restriction can be thought of as arising from a "loss of reputation". Arguing that such equilibria should not be dismissed as plausible solution candidates, leads to some reflections on extending heuristic refinement concepts, developed for signalling games, to richer environments, such as the class of games we consider. In this way our analysis ties in with recent contributions by Cho [5], Cho and Sobel [7] and Vincent [26], also pointing out some yet to be resolved problems in this literature. Furthermore we demonstrate in Section 4 that for a restricted class of games full immediate revelation of information is the only (or the only plausible) equilibrium outcome.

As will be shown in Section 5, imposing the support restriction enforces some pooling until almost the end of the game in repeated signalling games satisfying a relatively weak assumption on the structure of payoffs. This property not only rules out a whole class of equilibria based on the reputation effect we introduced before, it can also be used to prove a simple non-existence result. These results provide a further illustration of the main message of this paper: Imposing the support restriction is not only problematic because this requirement is lacking a good justification, there is also a good chance that such a restricted analysis will ignore the most plausible equilibrium outcomes.

2. Violations of the Support Restriction: A Simple Example

Consider the signalling game given by Figure 1. This game will serve as the main building block in the construction of our example for the non-existence of equilibrium under the support restriction.

In this game nature (N) first chooses the type (t_1 or t_2) of the sender (S), who then either chooses m_1 or m_2 . In this game sending m_1 is an "outside option" - the sender (and the receiver as well) gets a payoff of zero if she chooses this message. (The convention we use in this and in all following examples is that the payoff of the sender is given in the upper part of each cell, the payoff of the receiver (R) is given in the lower part.) Sending m_2 is strictly dominated for the sender if her type is t_1 , if her type is t_2 however, sending this message strictly dominates the choice of the outside option. If he receives m_2 the receiver has two options: choosing a_1 is optimal if the true type he faces is t_1 .

If we assume that both types are drawn by nature with strictly positive probability, then independently of the prior probability of type 1, $p^0(t_1)$, this game clearly has a unique Nash (and sequential) equilibrium: The sender chooses the outside option if her type is t_1 , she sends m_2 if her type is t_2 ; the receiver responds to m_2 with a_2 (this equilibrium is indicated in Figure 1 by doubling the edges of the corresponding arrows).

The belief of the receiver at his information set in this equilibrium obviously is that he faces type t_2 with probability one, since t_1 does not send m_2 . Intuitively it is also clear that if we would ask the receiver about his belief after m_1 was send, he would assign probability one to facing t_1 . This belief however is not part of the formal description of the equilibrium for the extensive form given in Figure 1; this would require the existence of an information set for the receiver upon hearing m_1 . Such an information set can easily be incooperated into the extensive form, see Figure 2. The only difference between Figure 1 and Figure 2 is the inclusion of the move "do nothing" for the receiver after the sender opted out. Including this move in the game tree, $p(t_i|\cdot)$ as the belief of the receiver that after a given history he faces type t_i is well defined. In the sequential equilibrium given above we then have $p(t_1|m_1)=1$, $p(t_2|m_2)=1$.

To us it seems next to impossible to argue that there should be any difference in the game theoretic analysis of the games represented in Figures 1 and 2. Indeed we will follow standard practice throughout the other sections of this paper and define a belief for the receiver after every possible message by the sender, representing such a signalling game in matrix form as in Figure 3 for

the above example (Note that in this representation different rows in the payoff matrices correspond to different types of the sender).

Now let us suppose that in the above game before she sends m_1 or m_2 the sender sends another signal s_i ($i=1,2$) which does not affect the receiver's payoff. The sender t_1 gets an additional payoff of 1 if she sends the signal s_i ; otherwise the payoffs are as given in Figure 3. An extensive form for this game is given in Figure 4.

In this game it is a strictly dominating strategy for t_1 to send the sequence (s_1, m_1) . For t_2 it is still strictly dominated to choose the outside option in the signalling game and therefore the only strategies she could use with strictly positive probability in an equilibrium are (s_2, m_2) and (s_1, m_2) . If she would use the second strategy in equilibrium, the belief of the receiver after receiving this sequence, as an implication of Bayes' rule, would be $p(t_1 | (s_1, m_2)) = 0$. Consequently, the receiver would choose a_2 at this information set, resulting in a payoff of 1 for t_2 . However she can guarantee herself a payoff of 2 by choosing (s_2, m_2) , and this implies that the equilibrium strategy of t_2 can not use s_1 with positive probability. Therefore in the only Nash equilibrium outcome of this game t_1 sends (s_1, m_1) and after receiving (s_2, m_2) the receiver responds with a_2 .

Turning to the issue of the receiver's beliefs in this equilibrium, application of Bayes' rule implies $p(t_1 | (s_2, m_2))=0$. To make it an unprofitable deviation for t_2 to send (s_1, m_2) it has to be true that the receiver chooses a_1 with probability less than 1/2 in this case. This is a best response for the receiver only if $p(t_1 | (s_1, m_2)) \leq 0.5$. Beliefs satisfying these conditions completely specify a sequential equilibrium for the game in Figure 4. A sequential equilibrium with these beliefs trivially satisfies the support restriction, since along any path through the game tree there is at most one information set of the receiver.

However, it may appear more natural to include the receiver's belief after he has seen the first signal in the description of the game. Allowing for this by giving the receiver an information set after s_i has been sent, has no influence on the Nash equilibrium outcome and does not change the conditions on the equilibrium beliefs just stated. The corresponding extensive form is

given in Figure 5, where we also included information sets for R after he received m_1 . As an implication of Bayes' rule we get additionally $p(t_1|s_1)=1$, and $p(t_1|s_2)=0$ as the beliefs at the new information sets where R has the move "sit and wait": after seeing s_1 the receiver believes with probability 1 that he faces type t_1 .

Note that these beliefs are derived from the equilibrium strategies of the players and that their intuitive meaning is captured by a statement like the following for the receiver: "Given the equilibrium strategies, after observing s_1 , I expect the sender to choose m_1 next and therefore I believe now that I face type t_1 with probability one." Observing s_1 , say, however does not give the receiver any information about whether he actually observes equilibrium behaviour. Conditional on the fact that equilibrium strategies are being played, he should assign probability one to t_1 . If he would know at that point that one of the types is actually not playing according to her equilibrium strategy, the receiver could clearly have a different belief. Observing m_2 next (and therefore knowing that a deviation occurred) the receiver can look back and conclude that if he had had this knowledge earlier, he would have had a different belief after receiving s_1 . Therefore, at this point he should feel free to ignore his previous belief: there is no conclusion he can draw from the assumption of equilibrium behaviour to what he should believe after an event which has probability zero in this equilibrium.

Actually, if there is any sensible conclusion to draw from observing (s_1, m_2) , it is that the type of the sender is t_2 : Sending (s_1, m_2) is strictly dominated for t_1 , but not for t_2 . Elimination of dominated strategies therefore results in $p(t_1|(s_1, m_2))=0$.

This is just the opposite of what one gets from insisting on the support restriction in Figure 5: Given the belief $p(t_1|s_1)=1$ the support restriction implies $p(t_1|(s_1, m_2))=1$ and therefore any set of beliefs which supports the unique Nash equilibrium outcome of this game leads to a violation of the support restriction. If we impose the restriction t_2 would like to deviate by first pretending to be t_1 and then exploiting the fact that even if he observes m_2 next, the receiver still believes that he faces t_1 .

It should be apparent by now that the non-existence of a Nash equilibrium satisfying the support restriction in the game given in Figure 5 results from a rather unintuitive requirement on the updating rule. This restriction forces the receiver to continue to believe that only t_1 could have send s_1 , even though the assumption on which this conclusion was based, namely equilibrium behaviour, has just been contradicted and in addition he observed the signal m_2 , which only t_2 would like to send.

The implausibility of such a requirement is further illustrated by the fact that it depends on strategically irrelevant details of the extensive form. With the sequence of information sets as in Figure 5 the support restriction leads to the argument that "only t_1 sends s_1 in equilibrium, therefore every strategy involving s_1 implies that I face t_1 ". In the game in Figure 4 the same verbal argument could of course be made, but the support restriction has no bite in this game. Changing the order of the signals in the extensive form also changes the conclusions one gets from applying the support restriction without changing the normal form of the game.

Summing up the discussion of this example it seems to us that the support restriction is based on a wrong interpretation of the concept of a belief in the extensive form. An equilibrium belief at a given information set is always derived from and conditional on equilibrium behaviour in the whole game. In a given extensive form violations of the support restriction may very well reflect the fact that once a deviation from equilibrium behaviour has been observed, a reassessment of all previous beliefs - which were based on the assumption that equilibrium strategies are followed - is called for. In this light such "switching of beliefs" is not an unfortunate problem, which cannot be avoided in some cases, but actually is a natural consequence of observing a deviation.

It is especially this last point we wish to develop in more detail in the following sections.

3. Repeated Signalling Games

The model of a repeated signalling game we introduce in this section is a straightforward extension of the standard signalling game model, as defined

for example in Cho and Sobel [7]. A time dimension is included in this set-up by assuming that after the sender has received his private information at the beginning of the game he may send a fixed finite number of signals to the receiver. Contrary to the example we discussed in the previous section it will be assumed that the receiver chooses a response after each signal and that the signalling game to be played at each stage is the same throughout. So the only parameter of the signalling game to be played which changes from stage to stage is the belief of the receiver about the information of the sender, and it is this ingredient that the analysis in the next section will focus on.⁵

While it would be desirable to analyse the effect of the support restriction directly in more realistic models, doing so in the framework of repeated signalling games has the benefit of reducing the strategic complexities to a point, where (we hope) the restrictions on equilibrium outcomes due to the support restriction can be clearly identified. At the same time the insights from our analysis about the role violations of the support restriction play in allowing separating equilibria should be transferable to more complex models. The recent contribution by Vincent [26], who obtains a fully separating equilibrium in a repeated bargaining game, illustrates the potential for such generalizations of our analysis: Vincent basically introduces a small perturbation into his model to guarantee a full support at every information set of the uninformed party. At the limit his analysis results in a fully separating equilibrium in which violations of the support restriction play an essential role.

As in our previous example the support restriction has a very natural interpretation in the framework of a repeated signalling game, which should make it easy to develop some intuition about the issue: Once the uninformed player's belief is concentrated on a subset of the possible types of the informed player, he behaves as if he knew (actually: as if it were common knowledge) that he faces some type from this subset.

Since most of our arguments proceed by stating assumptions on a one-shot signalling game in order to derive conclusions for the repeated version of this game, it is helpful to remind the reader of the formal definition of a signalling game and to state a few assumptions which will be maintained throughout the analysis.

A one-shot signalling game G is defined by a six-tuple (T, p^0, M, A, u_1, u_2) in the following way:

There are two players, called "The Sender" (player 1) and "The Receiver" (player 2). The first stage in the game is that nature draws the type of the sender, t , from the set T according to the probability distribution p^0 over T . This probability distribution is common knowledge among the players. Then the sender, and only the sender, is informed about the realization of t and next sends a signal m from the set M to the receiver. Upon receiving this signal player 2 chooses an action a from the set A . This ends the game. Payoffs are given by $u_i: T \times M \times A \rightarrow \mathbb{R}$ for player i .

Note that the above description implicitly assumes that neither the set of signals available to the sender depends on his type, nor the set of actions of the receiver depends on the signal sent. To avoid trivialities we assume $\#T > 1$ and more substantially we restrict attention to games G satisfying the following technical

Assumption 1:

(a) The sets T , M , A are finite.

(b) For every t and for every m there exist a unique $a^*(t, m)$ maximizing $u_2(t, m, a)$; furthermore for every t there is a unique $m^*(t)$ maximizing $u_1(t, m, a^*(t, m))$.

To understand the role part (b) of this assumption will play in the later analysis, note that if we were to consider a finitely repeated game in which it is common knowledge that the type of the sender is t , then for any t this assumption would imply that there is a unique subgame perfect equilibrium in this game.

The main interest in studying signalling games lies in the analysis of the distortions caused by the desires of the various types of the sender to reveal/hide their information. To state assumptions guaranteeing that a signalling game is interesting in this sense, let us define:

Type s has an incentive to mimic type t ($s(ITM)t$) if $u_1(s, m, a^*(t, m)) >$

$u_1(s, m^*(s), a^*(s, m^*(s)))$ for some m in M .

Given a signalling game G the N -fold repetition of G is defined by assuming that at the beginning of the game - and only at the beginning of the game - nature draws the type of player 1 and then in each of the periods $n=1, \dots, N$ first player 1 chooses a message from the set M and then player 2 responds by choosing a from the set A . These actions can be made contingent on the relevant history of the game, observable to a player when he has to move. After period N the game stops and payoffs are realized. For simplicity we assume that the payoffs are given by the (undiscounted) sum of the payoffs resulting from the one-shot games.

Formally, let $h_2^0 = \emptyset$ and define the history of moves (excluding the move of nature) observable to the sender (resp. the receiver) at time n by $h_1^n = (h_2^{n-1}, a^{n-1})$ ($h_2^n = (h_1^n, m^n)$), where a^n (m^n) is the action chosen by the receiver (the sender) in period n . Denote by H_i^n the set of all possible histories for player i at time n . Then a behavioural strategy for the sender is a sequence of maps $\sigma_1: H_1^n \times T \rightarrow \Delta(M)$, where $\Delta(\cdot)$ denotes the unit simplex. A behavioural strategy for the receiver is a sequence of maps $\sigma_2: H_2^n \rightarrow \Delta(A)$.

Another bit of notation we will need is the definition of a belief for player 2, which gives the probability he assigns to the various types of the sender after observing a history h_2^n . Formally, a belief is a sequence of maps $p: H_2^n \rightarrow \Delta(T)$, where $p(t|h_2^n)$ denotes player 2's probability assessment of facing type t after observing the history h_2^n .

A tuple (σ_1, σ_2, p) is a sequential equilibrium of the N -fold repetition of a signalling game G if it satisfies the conditions given in Kreps and Wilson [18], which we will not restate here. In fact most of our arguments will be only based on the following conditions, which are clearly necessary for sequential equilibrium (see also Fudenberg and Tirole [11]):

(a) beliefs are consistent with Bayes' rule whenever possible, i.e. not only on the equilibrium path but also starting from every information set h_2^n given the belief $p(\cdot|h_2^n)$.

(b) strategies specify best responses at every information set given the

beliefs and the strategy of the other player.

Finally, an equilibrium satisfies the support restriction if for all h_2^n : $p(t|h_2^n) = 0 \Rightarrow p(t|h_2^m) = 0$, for all h_2^m s.t. h_2^n is a sub-history of h_2^m in the obvious sense.

4. Reputation and Separation in Repeated Signalling Games

Consider the signalling game given in Figure 6. Assuming $6 > c_1 > 3 > c_2 > 0$ this game is a very much simplified version of a Spence-type signalling game. Independent of the message sent, both types of the sender would prefer that the receiver responds with a_2 . The receiver however is only going to use this response if he assigns probability of more than 1/2 to facing t_2 . Sending m_2 is costly for both types, but more so for t_1 . For values of c_1 satisfying the given constraints this game has two sequential equilibrium outcomes: (a) both types choose m_1 and the receiver responds with a_1 (if m_2 were sent, the receiver would use response a_1 with sufficiently high probability to deter t_2 from sending this message). (b) t_1 sends m_1 and the receiver responds with a_1 ; t_2 sends m_2 and the receiver responds with a_2 .

Outcome (a) can be eliminated by noting that sending m_2 is a strictly dominated strategy for t_1 . Consequently the receiver should believe that he faces type t_2 if he receives this message and should respond with a_2 , making it profitable for t_2 to deviate from the pooling equilibrium by sending m_2 . We therefore conclude that the strategies given in (b) characterize the unique plausible equilibrium for the one-shot signalling game given in Figure 6; the corresponding equilibrium beliefs are completely determined by Bayes' rule and are given by $p(t_1|m_1) = 1$ and $p(t_2|m_2) = 1$.

Turning now to the analysis of repeated signalling games we will use the example we just developed to motivate and illustrate most of our formal results.

A natural starting point for our analysis is to note that the equilibrium we selected for the signalling game in Figure 6 is fully separating, i.e. in equilibrium the receiver can perfectly infer the type of the sender from the

message he received. Intuitively one might expect that such full separation should be more difficult to achieve in a repeated signalling game. The argument here is simple and might sound familiar from the analysis of repeated adverse selection problems (see for example Laffont and Tirole [20]): To have separation in a one-shot signalling game it is basically necessary to fulfil a set of incentive constraints: no type should find it worthwhile to mimic the strategy of some other type. In a repeated game, so it seems, it is no longer possible to provide such incentives once all information is revealed: why should t_2 in the game above continue to invest c_2 in sending m_2 if the receiver already knows her true type? But if she invests in the signal only once, t_1 would clearly like to mimic her strategy, if this would induce the receiver to respond with a_2 in all following rounds. As we will prove in the next section this kind of reasoning turns out to be correct if one insists on the support restriction, in general however it is false.

More specifically suppose we have a signalling game G which has a fully separating equilibrium. Let $p^*(\cdot|m)$ be the belief of the receiver in this equilibrium. In the N -fold repetition of G it is always possible for the receiver to ignore any information revealed in any previous periods of the game and to condition his belief just on the last signal he received in exactly the same way as in the equilibrium of the one-shot game. I.e. after any history in the repeated game such that the last message send was m , specify the belief of the receiver as $p^*(\cdot|m)$. Given these beliefs, it is optimal at every stage for both the sender and the receiver to choose exactly the same actions as in the one-shot game. It remains to note that, given these strategies, the specified beliefs are also consistent with Bayesian updating wherever possible.

This argument almost (but not quite) proves⁶

Proposition 1:

Suppose a signalling game G has a fully separating equilibrium. Then for every N , the N -fold repetition of this equilibrium is a sequential equilibrium of the N -fold repetition of G .

Applied to our example, Proposition 1 demonstrates that it is part of an

equilibrium in the N-fold repetition of this game for type t_2 to send the signal m_2 in each and every period, even though there is no additional information to be transmitted after the first period. A simple intuitive way to understand this equilibrium behaviour is to interpret it in terms of a reputation effect: Once she has send m_2 , t_2 has established a reputation for being the type she actually is. Continuing to send m_2 just serves the purpose of keeping that reputation. Deviating just once undoes all previous efforts since the receiver now concludes that he was mistaken in assuming equilibrium behaviour - and consequently there is no compelling reason left to assume that he really faces t_2 .

Given the utter simplicity and the somewhat surprising nature of the fully separating equilibrium in Proposition 1, it seems worthwhile to inquire if there are conditions under which this is the unique possible sequential equilibrium outcome in a repeated signalling game. The answer to this question is positive, even though the conditions needed to gurantee this result are rather strong:

Proposition 2:

Assume that in a given signalling game G for every type $t_i \in T$ there exist a message m_i which is strictly dominating for t_i , with $m_i \neq m_j$ for $i \neq j$. Assume furthermore $u_1(t_i, m_i, a^*(t_i, m_i)) = \min_{a \in A} u_1(t_i, m_i, a)$. Then the unique sequential equilibrium outcome of the N-fold repetition of G is the one given in Proposition 1.

Proof:

The proof is by a simple backwards induction argument. If the sender's type is t_i then on any possible realization of the equilibrium path she will send m_i in period N and the receiver will respond with $a^*(t_i, m_i)$ to this. Suppose this result is true for the last L periods of the game. Then it also has to hold in period $N-L$: By the induction hypothesis type t_i can ensure her continuation payoff by simply sending m_i in all following periods and she will therefore choose a message which is optimal in the one-shot signalling game in period $N-L$. By assumption, independent of the beliefs of the receiver this is m_i . Therefore only t_i sends m_i in period $N-L$ and given the uniqueness of the receiver's contiuation payoff and his equilibrium belief $p(t_i | \cdot) = 1$ he will respond with $a^*(t_i, m_i)$ and the result follows. ■

Proposition 2 could probably be proved under slightly different conditions. It is not sufficient, however, to assume that the one-shot signalling game has a unique, fully separating equilibrium outcome independent of the prior - a property which is already difficult enough to satisfy. To see that even given the first part of our assumption the second part is needed, too, consider the example in Figure 7.

Obviously the unique sequential equilibrium outcome in this game is for type t_1 to send m_1 , independently of the prior. In the twice repeated game, assuming $p^0(t_1) < 0.5$, it is however an equilibrium outcome for both types to send m_2 in the first period and then to continue with the equilibrium behaviour of the one-shot game in the second period. This equilibrium outcome is supported by the belief of the receiver that he assumes to face t_2 , if he should observe m_1 in the first period, and that he sticks to this belief, no matter what he observes in the second period. Given these beliefs and the corresponding best responses of the receiver it is indeed optimal for t_1 to stick to the specified strategy of sending m_2 in the first round.

This example is not only instructive in pointing out once again that requiring the support restriction might be quite counterintuitive: Even taking for granted that $p(t_1|m_1) = 0$ could be a sensible belief, shouldn't the receiver switch away from this belief after observing m_1 for the second time? It is also obvious in this example that pooling in the first period is only possible as an equilibrium outcome if the receiver bases his belief after observing a deviation on the assumption that some type (in this case t_2) is playing a strictly dominated strategy.

This observation raises an interesting question: Suppose that in a given signalling game there is a unique, fully separating equilibrium outcome satisfying some refinement of sequential equilibrium (like elimination of weakly dominated strategies or some stronger criterion, like the one used in Cho and Sobel [7] to guarantee uniqueness of the equilibrium outcome). Is it then true that, if we use an equivalent refinement in the repeated signalling game, we will also get a unique fully separating equilibrium outcome?

Of course, put this way, this question is not really well-defined, especially

since most heuristic refinement concepts are only defined for one-shot signalling games. Even the recent attempt by Cho [5] (see also Vincent [26]) to extend such a definition are not quite applicable to our situation, since it is not clear how the backwards induction procedure defined there would deal with "almost proper subgames" having a restricted support.

Nevertheless, considering the two-fold repetition of the game given in Figure 6, we can gain some intuition, why the answer to this question will be negative in the framework we consider here - suggesting also conditions under which the answer will be positive. In addition this exercise will also help us to understand under which conditions an equilibrium based on the "reputation effect" we outlined above will be a sensible solution to a repeated signalling game.

So let us again turn to our simplified Spence signalling game. First note that in any fully separating equilibrium of the twice repeated version of this game t_2 has to send m_2 on the equilibrium path in the second period: If she would send m_1 in the second period in equilibrium, the belief of the receiver would be given by Bayes' rule and he would respond with a_2 . Given this response t_1 would always find it worthwhile to mimic t_2 in the first round: The worst possible case for t_1 would be that she has to invest $c_1 < 6$ once to guarantee the response a_2 in both rounds, compared to a payoff of 0, which is not only the minimum payoff she can guarantee herself (by sending m_1 in both periods), but also her equilibrium payoff in any fully separating equilibrium⁷.

It follows from this argument that the only possible equilibrium path for a fully separating equilibrium is characterized by t_1 sending m_1 in both rounds and the receiver responding with a_1 to receiving m_1 along the equilibrium path. Proposition 1 shows that there is indeed a sequential equilibrium resulting in this outcome.

We will not even attempt to show here whether or not this equilibrium outcome is the unique outcome consistent with some refinement, instead we will demonstrate simply (and somewhat informally) that it depends on the values of the parameters c_1 and c_2 whether this outcome satisfies the independence of never weak best responses (INWBR) criterion defined in Kohlberg and Mertens [16]. As we have shown full separation is the only outcome compatible with

elimination of weakly dominated strategies - a criterion which is implied by INWBR - in the one-shot game. Consequently our argument shows that even though, according to INWBR, full separation is the only plausible equilibrium outcome in the one-shot game, full separation is not a plausible outcome in the repeated game for some values of c_1 and c_2 .

More precisely, note that given the equilibrium outcome the only deviation which could potentially be profitable for either t_1 or t_2 would be to send the sequence (m_2, m_1) . To deter t_1 from sending this sequence, the probability r_2 that the receiver responds with a_2 in the second round after observing this history has to satisfy

$$r_2 \leq c_1/3 - 1. \quad (4.1)$$

If (4.1) is satisfied then it is not worthwhile for t_1 to incur the costs of sending m_2 in the first round, because the chance that she will gain from this in the second round is sufficiently small. Correspondingly, it also has to be true that t_2 does not find it worthwhile to save the trouble of signalling in the second round and we therefore have

$$r_2 \leq 1 - c_2/3. \quad (4.2)$$

To have a sequential equilibrium outcome r_2 of has to satisfy both of these conditions simultaneously; for considering the INWBR criterion the relevant question however is which of these constraints is the binding one. Ignoring the borderline case in which both constraints are binding we have to distinguish the following two cases:

(a) $c_1 + c_2 < 6$.

In this case (4.1) is the binding constraint and therefore sending (m_2, m_1) is a never weak best response for t_2 . Eliminating this type as a candidate for sending this sequence results in a_1 as the unique best response of the receiver to the induced belief $p(t_1 | (m_2, a_2, m_1)) = 1$.

Intuitively the formal refinement supports our reputation story in this case: If both c_1 and c_2 are relatively low, there is not much to be gained for t_2 by not sending m_2 in the second period; t_1 however can try at relatively low cost

to pretend to be actually t_2 , hoping for a gain in the second period. Receiving m_2 only once should therefore make the receiver rather suspicious about the identity of the sender.

(b) $c_1 + c_2 > 6$.

In this case (4.2) is binding and all our conclusions for case (a) just get reversed. For these parameter values INWBR leads to the same result as insisting on the support restriction. The reasoning process, however, is quite a different one: t_1 is not ruled out a-priori as a candidate for sending (m_2, m_1) - it is the fact that sending m_2 is quite costly for her (and also for t_2) which induces the belief that she did not send this message pair.

Some general insights can be derived by considering a special case of the analysis just conducted. Specifically, suppose $c_1 = 3$. Under this assumption we are always in case (a) and the fully separating equilibrium will satisfy INWBR. This result is caused by the simple fact that t_1 is just indifferent between following her equilibrium strategy and mimicking t_2 . In judging a deviation in period 2 it is therefore only necessary to ask the following question: Which type is "more likely" to benefit by a deviation from (m_2, m_2) to (m_2, m_1) ? The answer is always t_1 , because by assumption she has the higher cost of sending m_2 . This kind of reasoning extends to the N-fold repetition of the underlying game for any N, indicating that the fully separating equilibrium will satisfy INWBR in these games.

Of course, $c_1 = 3$ is a very peculiar parameter value in our game. However, in the framework of Cho and Sobel [7], who consider games with a continuum of messages, it is a property of the unique equilibrium outcome satisfying their refinement that "low" types are just indifferent between following their equilibrium strategy and mimicking the next higher type. It is this kind of indifference we can only represent in our current setup by choosing $c_1 = 3$. The analysis of this case suggests that for a signalling game with a continuum of messages, satisfying the assumptions of Cho and Sobel, full separation can not be eliminated as a plausible equilibrium outcome in the repeated signalling game, whenever the same is true in the one-shot game.⁸

It should be noted that in the case where $c_1 = 3$ full separation is not the unique equilibrium outcome satisfying INWBR in the repeated game - this is not

even the case for the one-shot game. This is due to the fact that the best response of the receiver is not sensitive to small changes in his belief starting from $p(t_2|\cdot) = 1$. In fact taken together with the non-uniqueness of the receiver's best response at $p(t_1|\cdot) = 0.5$ this property can be exploited to construct semi-separating equilibria in the repeated game. By carefully manipulating the randomizations of the sender and the receiver it is possible to construct such equilibria for all values of the parameters c_i , such that they can not be eliminated by INWBR. This is once again reminiscent of the results in Cho and Sobel [7], who demonstrate that their uniqueness result may fail in games where the receiver's strategy set is discrete, making it impossible to eliminate pooling equilibria for particular priors of the receiver.

This suggests the following reformulation of the question with which we started our discussion of refinements: Suppose that in a given one-shot signalling game there is a unique, fully separating equilibrium outcome, satisfying some refinement of sequential equilibrium, independent of the prior p^0 (assuming full support). Is it then true that in the repeated signalling game there is also a unique (fully separating) outcome, satisfying an equivalent refinement?

Within our framework of games with finite strategy sets a positive answer to this question has unfortunately not too much applicability, since for any conceivable refinement the class of signalling games satisfying the stated condition is rather small. (see footnote 8, however). On the other hand, expecting uniqueness in a class of games like the one we consider here is rather strong, too, and our main goal is to convince the reader that full separation achieved through violations of the support restrictions is in many circumstances at least a sensible candidate for an equilibrium outcome. In this light it is helpful not only to establish existence, as we did in Proposition 1, but also to illustrate the conditions under which such a claim can be unambiguously made. We therefore state:

Proposition 3:

For a given signalling game G define M_i as the set of messages which is not strictly dominated for type t_i . Assume $M_i \cap M_j = \emptyset$ for all $i \neq j$. Then there is a unique equilibrium outcome of the N -fold repetition of G compatible with

iterated elimination of weakly dominated strategies. The strategy for the sender to send $m^*(t_1)$ in every period (independent of the history) if her type is t_1 is part of such an outcome.

Proof:

The obvious way to eliminate strategies in this repeated game is to start in the last period: Begin by eliminating t_1 as a candidate for sending any message not belonging to the set M_1 in the last period of the game. Having eliminated all those strategies, it is weakly dominated for the receiver not to respond with $a^*(t_1, m)$ for any $m \in M_1$ observed in the last period. Eliminating all other responses for signals in the set M_1 , it is now a strictly dominant strategy for the sender of type t_1 to choose $m^*(t_1)$ in the last period of the game. This completely determines the strategies of the sender and the receiver in the last period of the game, allowing the same dominance arguments for period $N-1$. Iterating this argument then leads to the desired result, since Proposition 6 of Kohlberg and Mertens [16] implies that independent of the order of elimination of weakly dominated strategies the corresponding outcome can not be eliminated ■

Before we turn to the effects of imposing the support restriction in repeated signalling games, it should be noted that the additional time dimension in such a repeated game may allow to support separation as an equilibrium outcome in cases where this is not possible in the one shot game. (This should be contrasted with our arguments so far, which basically demonstrated that full separation can be achieved by "ignoring" the dynamic nature of the interaction.) The idea here is simple and can be illustrated by the game in Figure 8, which is derived from the game in Figure 6 by setting $c_1 = 2$, $c_2 = 1$ and adding a third message m_3 which has cost 4 for t_2 and cost 8 for t_1 . In the one-shot game neither type will send m_3 - this message is strictly dominated for both types. Sending m_2 on the other hand is not sufficiently costly for t_1 to allow t_2 to separate herself. Consequently there is no fully separating equilibrium in the one-shot game.

In the 2-fold repetition of the signalling game however, t_2 can use the message m_3 as an investment in the first period to credibly signal her true type, eliminating the necessity of further investment into a costly signal in the second period. Mimicking such a strategy is not attractive for t_1 . It

follows without further difficulties that there is indeed a sequential equilibrium outcome in this game in which t_1 sends m_1 in both periods and t_2 separates herself by following the strategy outlined above.

5. The Support Restriction in Repeated Signalling Games

In most of the signalling game literature the emphasis has been on fully separating equilibria: Finding conditions under which such equilibria exist, developing refinements which single out such outcomes, etc... . This is a natural focus in the analysis of signalling games, which were essentially introduced in the economics literature to study the question of how information could be revealed through (costly) signals. As we tried to demonstrate in the previous section extending the framework of a signalling game to include a dynamic perspective still allows us to ask the same kind of questions and to extend results about the feasibility of full separation to this framework.

The reader will have noted that most of the equilibria constructed in the last section did involve violations of the support restriction. Indeed, as we will show now, any analysis which focuses on the question of information revelation and assumes the support restriction will arrive at the conclusion that a repeated signalling game differs substantially from a one-shot signalling game.⁹

To motivate the assumption we impose in the statement of the following proposition, note that in most economic applications it is assumed that a signalling game G to be analysed satisfies some form of a response monotonicity condition. This roughly requires that there is a complete strict ordering of types s .t. for all possible messages m and all possible types s , if $t' > t$ (according to this ordering) then $u(s, m, a^*(t', m)) > u(s, m, a^*(t, m))$. If response monotonicity holds $t' > t$ implies that t has an incentive to mimic type t' in order to convince the receiver that she is actually t' . This is captured in our formal definition of an incentive to mimic and we have $t(ITM)t'$. In a game satisfying the response monotonicity condition it is therefore clearly true that for any pair of types s, t either $s(ITM)t$ or $t(ITM)s$ holds.

It is this later, quite weak, property we require in the statement of the following

Proposition 4:

Suppose that in a given signalling game G for each pair of types (s, t) either s has an incentive to mimic t or t has an incentive to mimic s . Then there exist a constant $k(G)$, such that in the N -fold repetition of G full separation can not occur earlier than period $N - k(G)$ in any sequential equilibrium satisfying the support restriction.

Proof:

Consider the N -fold repetition of a given signalling game G and fix a period N^* . Suppose there exist a sequential equilibrium of this game satisfying the support restriction s.t. for any sub-history h_2 which is realized with strictly positive probability in equilibrium and satisfies $h_2 \in H_2^n$ for $n \geq N^*$ the following is true: $\exists t$ s.t. $p(t|h_2) = 1$, i.e. by period N^* for any realization of the equilibrium strategies the observed signals fully reveal the type of the sender to the receiver.

W.l.o.g we can choose N^* such that it is the first period for which this statement is true, i.e. there exist some history in $H_2^{N^*-1}$ appearing with strictly positive probability on the equilibrium path satisfying that the support of the receivers belief after this history includes at least two types. Pick any two types from the support, say t_1 and t_2 . Again w.l.o.g. we may assume that $t_1(ITM)t_2$. By construction, in the equilibrium t_1 and t_2 choose different actions in the next period with probability one and exploiting the support restriction and assumption 1 in each of the following periods t_1 's payoff will be given by $u(t_1, m^*(t_1), a^*(t_1, m^*(t_1)))$. Now consider the following deviation by t_1 : In period N^* he sends one of the messages type t_2 sends with strictly positive probability in equilibrium, given the history h_2 . This may result in a loss in the payoff this type of sender receives from period N^* . By assumption 1(a) this loss is bounded above due to the finite number of messages and responses available; let us denote by $L(t_1)$ such an upper bound. The belief of the receiver after this deviation of course puts probability one on type t_2 and by the support restriction this belief, once held, will not change, no matter what actions the sender chooses in the continuation. Exploiting once again assumption 1(b) it can be seen that the receiver will

therefore choose the response $a^*(t_2, \cdot)$ in all following periods. In each of the periods N^*+1, \dots, N t_1 can consequently send a signal m^* which maximizes $u_1(t_1, m, a^*(t_2, m))$ and she will actually receive the corresponding payoff. Compared to her equilibrium payoffs in each of these periods this will result in a (strictly positive) one-period gain of, say, $V(t_1, t_2)$.

Since such a deviation cannot be worthwhile for t_1 by the assumption that we have an equilibrium for the N -fold repetition of G , it follows that $(N-N^*) \leq L(t_1)/V(t_1, t_2)$. Taking the maximum of $L(t_1)/V(t_1, t_2)$ over all pair of types such that $t_1(ITM)t_2$ (note, that once again, by assumption 1(a) this maximum exists) and denoting the maximum by k , it follows by the above argument that for any perfect bayesian equilibrium $N^* \geq N-k$ holds, proving the statement of the Proposition. ■

An immediate consequence of Proposition 4 is that to prove non-existence of equilibrium under the support restriction for large enough N , it is sufficient to demonstrate that for a given game every sequential equilibrium has to be fully separating after a number of periods which is proportional to N . We did not succeed in providing transparent general conditions under which such a result would be true. Therefore we just give the following non-existence result, which is an immediate consequence of Proposition 2 and generalizes the example of Madrigal, Tan and Werlang [22]:

Corrolary 1:

Assume in addition to the assumptions stated in Proposition 2 that the following holds for all t_1 : $u_1(t_1, m_1, a^*(t_1, m_1)) < u_1(t_1, m_1, a^*(t_1, m_2))$ for some t_1 . Then for N large enough ($N > k(G)$) there exist no sequential equilibrium satisfying the support restriction in the N -fold repetition of G .

Proof:

This follows trivially by noting that (a) Proposition 2 implies that the unique sequential equilibrium outcome for such a game is fully separating, whereas (b) the added assumption implies that Proposition 4 is applicable, too, ruling out such an equilibrium outcome for a sufficiently high number of repetitions. ■

The careful reader will have noted that the conclusion stated in Proposition

4 actually looks very weak (even though, as the above Corrolary demonstrates, strong enough to cause non-existence and also strong enough to rule out all the equilibrium outcomes we considered in the previous section for games satisfying the stated assumptions). If we assume that N is large relative to $k(G)$, it says that in any equilibrium satisfying the support restriction it can happen with (an arbitrarily small, but strictly positive) probability that until "almost the end" of the game the receiver can not infer the type of the sender exactly from the observed actions.

That we only get this kind of result is due to the fact that our conclusions are just derived by working backwards from the assumption that after a given history appearing in the equilibrium the support of the belief is concentrated on a single type. As long as there is more than one type in the support of the receiver's belief there are no immediate restrictions on the response of the receiver in a given period, making it impossible to extend the kind of argument we use in our proof to such situations (In periods before the last one it is not even true that the receiver has to use a best response against some belief in his current support, cf. footnote 7).

However, it should be noted that the logic of the proof of Proposition 4 directly implies a number of additional substantial constraints on the way information is revealed in equilibrium.

First note that in any period before $N-k$, if the receiver is still uncertain about the true type of the sender, there is at most one type who will reveal her information in this period by choosing with positive probability an action which no other type is choosing: If there were two such types, then at least one of some could gain by mimicking the fully revealing signal of the other type and exploiting the obstinate beliefs of the receiver afterwards.

Next suppose that for a given type there is actually a positive probability that she will reveal her type in a given period, say period L . Then there exists a constant c , s.t. on the equilibrium path it is not possible that another type reveals her information before period $L + c(N-L)$. This follows by exactly the same argument as the last claim. The only additional fact to note there is a simple linear relationship which determines whether it is worthwhile to mimic another type: If it takes m periods of following the

strategy of some other type (whom one would like to mimic) to convince the receiver that one is actually this type, then this is worthwhile if there are $(m+1)k$ periods left in the game. In this sense revelation of information has to be slow and gradual on the equilibrium path.

Another instructive approach to see that the support restriction may have substantial bite (and few intuitive merits), is to consider the problems which will result if we want to construct an equilibrium in which some type, say t_1 , reveals her information in a given period L . Once we reach that point on the equilibrium path a ratchet effect will appear: From thereon the sender can do whatever she wants to do, the receiver will always respond with $a^*(t_1, m)$. This may allow types other than t_1 to achieve very high continuation payoffs, by choosing messages t_1 would actually never choose (she will send $m^*(t_1)$ in every round from L on). This in turn may make the separation of t_1 impossible to sustain as an equilibrium outcome. But why should the receiver continue to believe that he faces t_1 , if he does not observe $m^*(t_1)$ at a stage where he supposedly deals with t_1 with probability one? Allowing him to switch away from this probability one belief not only makes intuitive sense in this situation, it also makes it less attractive for the other types to mimic t_1 up to period L by reducing their freedom to optimize against the fixed belief of the receiver.

In the class of games we considered in Proposition 3 of the previous section the kind of reasoning we just presented is just what we need to support the separating outcome we singled out there: Whenever the receiver believes that he faces type t_1 and he does observe a message not belonging to the set M_1 he should be free to switch away from his probability one belief. However, as the opening example of Section 4 demonstrates this is only half the story. Given that the receiver continues to believe that he faces type t_1 as long as he receives $m^*(t_1)$ it might still be impossible to support a separating equilibrium, because some other type might actually be quite happy to send the same message as long as the receiver responds as if he were facing t_1 . But in such a situation receiving $m^*(t_1)$ should not necessarily convince the receiver that he faces t_1 . Dropping this constraint on the receiver's belief allows the construction of the kind of "reputational" equilibria we also discussed in Section 4. Consequently, it might be possible to support separation of t_1 by specifying a sequence of messages for her such that (a)

it is not desirable for any other type to mimic this sequence of messages and (b) the threat of a "loss of reputation" is sufficiently severe to stop t_1 from reverting to $m^*(t_1)$ (or any other message) once she revealed her type in equilibrium.

It should be apparent from these remarks that dropping the support restriction may greatly expand the set of possible equilibrium outcomes. Insisting on the support restriction to simplify the analysis and reduce the multiplicity of equilibrium outcomes is in our view not a solution to this problem. It would be more fitting to describe such an approach as one which ignores the problem and also ignores the rich possibilities for equilibrium behaviour inherent in many dynamic models of incomplete information.

This of course does not rule out the possibility that some equilibrium outcomes may appear more intuitive than others. As Section 4 of this paper demonstrates applying arguments in this spirit may help substantially in focusing the analysis on particular equilibrium outcomes in some repeated signalling games. Obviously, there is more work to be done in this area, especially in generalizing the formal definitions of heuristic refinement criteria to games in which almost proper subgames with restricted support necessarily arise. Once this is achieved, it seems reasonable to conjecture on the basis of the work presented here, that separation based on violations of the support restriction will often play a prominent role in such models.

References

1. A. Admati and M. Perry, Strategic Delay in Bargaining, *Rev. Econ. Stud.* 54 (1987), 345 - 364.
2. K. Bagwell and R. Staiger, The Role of Export Subsidies when Product Quality is Unknown, *J. Int. Econ.* 27 (1989), 69 - 89.
3. J.-P. Benoit and V. Krishna, Finitely Repeated Games, *Econometrica* 53 (1985), 905 - 922.
4. S. Bikhchandani, Market Games with Few Traders, PH.D. thesis, Stanford University, 1986.
5. I.K. Cho, Strategic Stability in Repeated Signaling Games, Mimeo, University of Chicago, 1989.
6. I.K. Cho and D. Kreps, Signaling Games and Stable Equilibria, *Quart. J. Econ.* 102 (1987), 179 - 221.
7. I.K. Cho and J. Sobel, Strategic Stability and Uniqueness in Signaling Games, *J. Econ. Theory* 50 (1990), 381 - 413.
8. D. Diamond, Reputation Acquisition in Debt Markets, *J. Polit. Econ.* 97 (1989), 828 - 862.
9. M. Fishman, A Theory of Preemptive Takeover Bidding, *Rand J. Econ.* 19 (1988), 88 - 101.
10. M. Fishman, Preemptive Bidding and the Role of the Medium of Exchange in Acquisitions, *J. Finance* 44 (1989), 41 - 57.
11. D. Fudenberg and J. Tirole, Perfect Bayesian and Sequential Equilibria, Mimeo, 1989.
12. I. Gale and J. Stiglitz, The Informational Content of Initial Public Offerings, *J. Finance* 44 (1989), 469 - 477.
13. R. Gertner, R. Gibbons and D. Scharfstein, Simultaneous Signalling to the Capital and Product Markets, *Rand J. Econ.* 19 (1988), 173 - 190.
14. S. Grossman and M. Perry, Perfect Sequential Equilibrium, *J. Econ. Theory* 39 (1986), 97 - 119.
15. S. Grossman and M. Perry, Sequential Bargaining under Asymmetric Information, *J. Econ. Theory* 39 (1986), 120 - 154.
16. E. Kohlberg and J.-F. Mertens, On the Strategic Stability of Equilibria, *Econometrica* 54 (1986), 1003 - 1038.
17. D. Kreps, A Course in Microeconomic Theory, Harvester Wheatsheaf, New York, 1990.
18. D. Kreps and R. Wilson, Sequential Equilibria, *Econometrica* 50 (1982), 863

19. D. Kreps and R. Wilson, Reputation and Imperfect Information, *J. Econ. Theory* 27 (1982), 280 - 312.
20. J.-J. Laffont and J. Tirole, The Dynamics of Incentive Contracts, *Econometrica* 56 (1988), 1153 - 1175.
21. N. Lutz, Warranties as Signals under Consumer Moral Hazard, *Rand J. Econ.* 20 (1989), 239 - 255.
22. V. Madrigal, T. Tan and S. Werlang, Support Restrictions and Sequential Equilibria, *J. Econ. Theory* 43 (1987), 329 - 334.
23. A. Rubinstein, A Bargaining Model with Incomplete Information, *Econometrica* 53 (1985), 1151 - 1172.
24. N. Stoughton, The Informational Content of Corporate Merger and Acquisition Offers, *J. Finan. Quant. Anal.* 23 (1988), 175 - 197.
25. D. Vincent, Dynamic Auctions, *Rev. Econ. Stud.* 57 (1990), 49 - 61.
26. D. Vincent, Bilateral Monopoly, Non-durable Goods and Dynamic Trading Relationships, Mimeo, Northwestern University, 1990.

1. Just restricting the analysis to subgame perfect equilibria also assumes that players are obstinate to some degree. For example they do not consider the possibility that they were mistaken about the extensive form game once they observe a deviation. Such explanations are ruled out by the assumption that the given extensive form game is common knowledge among the players, which makes it impossible to argue consistently within the model why deviations occur. Nevertheless, within the framework of a given extensive form it is at least clear what the set of possible strategies is which explains how it could have happened that a player is called upon to move at a given information set. It is at this level that the support restriction operates by dismissing some of the possibilities in the extensive form how an information set could have been reached.

2. Due to the nature of the games considered, the issue of the support restriction does not arise in most published applications appealing to the perfect sequential equilibrium concept. Notable exceptions are Grossman and Perry [15] and Vincent [25]. The widespread willingness to accept perfect sequential equilibrium as a sensible solution concept (see for example Bagwell and Staiger [2], Diamond [8], Fishman [9,10], Gertner, Gibbons and Scharfstein [13], Lutz [21], and Stoughton [24] nevertheless makes it worthwhile to point out this problem with the formal definition.

3. Further examples where the support restriction is used, include Bikchandani [4] and Gale and Stiglitz [12].

4. A similar position is taken in Admati and Perry [1]. They note that in the specific bargaining game they analyse, imposing the support restriction conflicts with other heuristic refinement ideas they use and consequently drop the support restriction as a requirement.

5. This description of a repeated signalling game is quite different from the one recently suggested by Cho [5]. In his formal model in every period a

different player signals his private information. The issue of the support restriction therefore does not arise.

6. What is missing from the argument is that the equilibrium is indeed sequential and not only perfect bayesian according to the weakest definition in Fudenberg and Tirole [11]. For most readers who really care about such technical details it should be sufficient to point out that the specified beliefs can easily shown to be structurally consistent: Whenever the receiver observes a deviation he explains this by assuming that up to the previous period all types would have chosen the observed sequence of messages. For the current period he assumes a strategy which generates the required belief given the prior p^0 at the start of the period and for the continuation of the game he expects to observe equilibrium behaviour.

7. This argument is not as straightforward as it may appear. The point is that in general it is not true that the receiver has to play a short-run best response against his current belief in equilibrium, if, given this belief, there are multiple continuation equilibria. By conditioning the equilibrium continuation on the response of the receiver, it is possible to induce him even to choose actions which would be strictly dominated in the one shot game for any possible prior. Basically the situation is very similiar to the one which allows the proof of a Folk theorem for finitely repeated games of complete information as in Benoit and Krishna [3]. In the situation we consider here, however, the receiver has to choose a short-run best response given his prior, which assigns probability one to t_2 . The reason is that the receiver actually does not care whether t_2 is going to choose m_1 or m_2 in the next round - he can gurantee himself the maximal possible payoff of 1 by choosing a_2 in either case. The same argument applied to a prior concentrated on t_1 serves to show that, once seperation has occured, the receiver will always choose a_1 .

8. Some preliminary formal analysis of such games we conducted indicates that such a result indeed holds. In fact we strongly conjecture that it will be

possible to prove a uniqueness result for repeated signalling games under essentially the same conditions Cho and Sobel [7] use for their uniqueness result. Outlining such results in any degree of detail is beyond the scope of the current paper, especially since such equilibria will in general have a more complicated structure than the ones given in Proposition 1. We hope to deal with this issue in future work.

9. A case in point is the analysis of Gale/Stiglitz [12]. Their model is very similar to a repeated signalling game and their main results concern the possibilities of information transmission in equilibrium. They argue that separating equilibria may fail to exist and single out a pooling equilibrium as the solution to their model. All their results depend critically on the assumption that "once the investors are certain of the entrepreneur's type, they do not subsequently revise their beliefs" [12, p.471]. None of their main results continues to hold once this assumption, which is never discussed in their paper, is dropped.

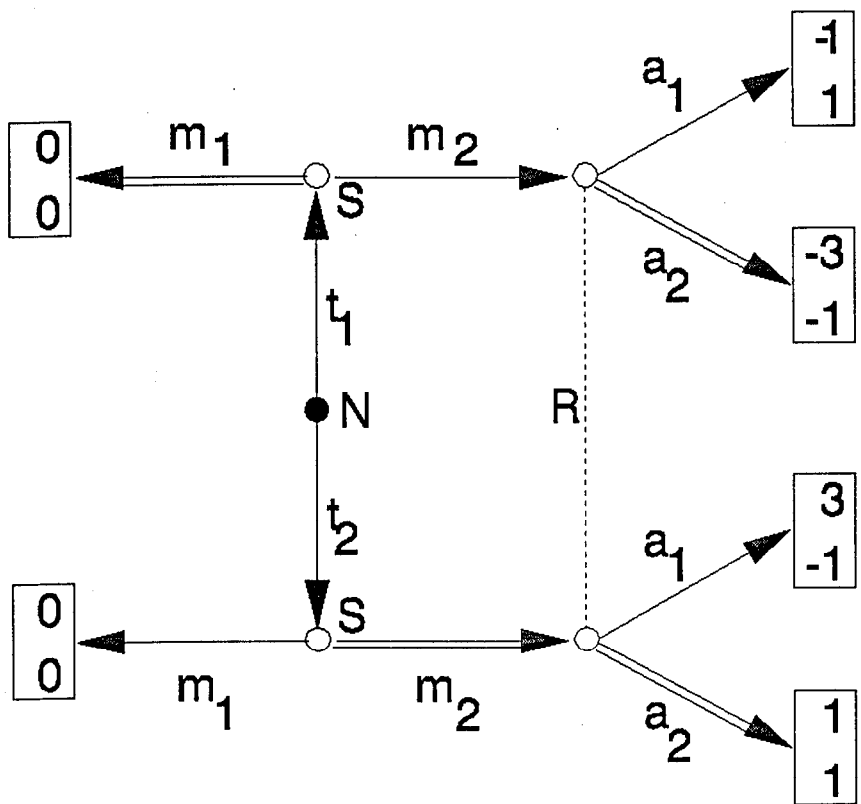


Figure 1

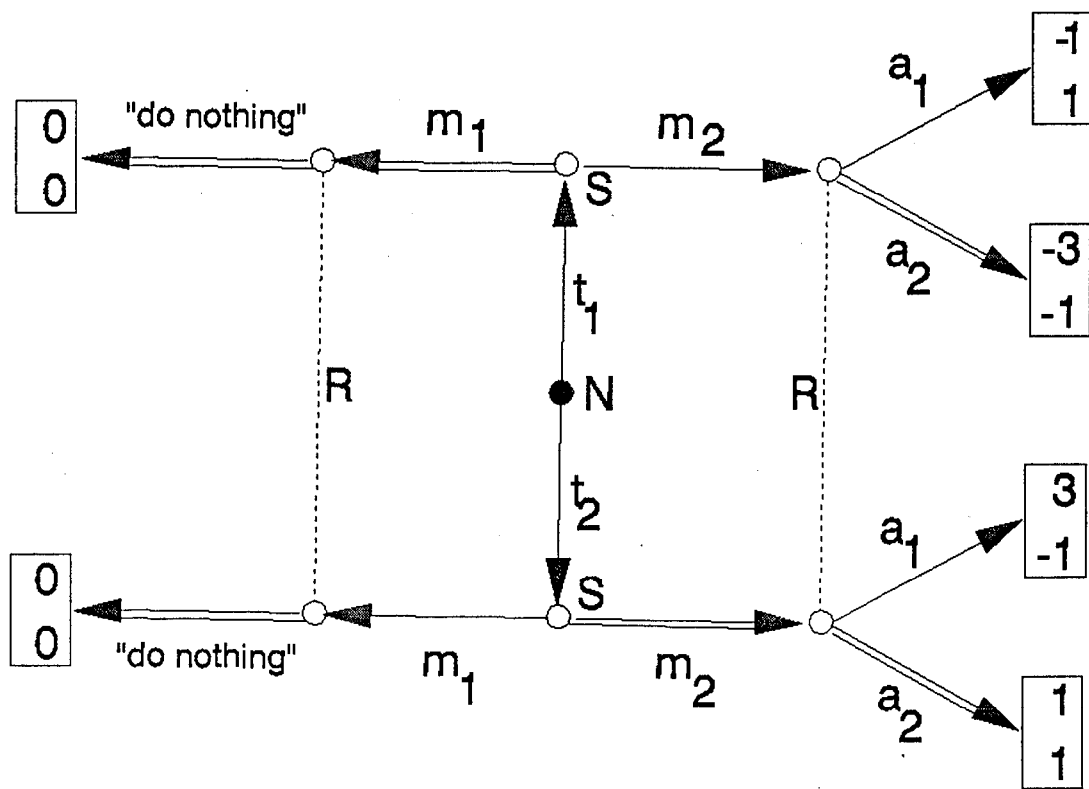


Figure 2

	a_1	a_2
t_1	0	0
t_2	0	0
	m_1	

	a_1	a_2
-1	1	-3
3	-1	1
	m_2	

Figure 3




Figure 4




Figure 5

	a_1	a_2
t_1	0	3
	1	-1
t_2	0	3
	-1	1
	m_1	

	a_1	a_2
	$-c_1$	$3-c_1$
	1	-1
	$-c_2$	$3-c_2$
	-1	1
	m_2	

$$p^0(t_1) > 0.5$$

Figure 6

	a_1	a_2
t_1	2	0
t_2	-1	-3

	m_1	m_2
t_1	1	-1
t_2	-1	1

	a_1	a_2
t_1	-1	-3
t_2	2	0

	m_1	m_2
t_1	-1	1
t_2	1	-1

Figure 7

	a_1	a_2
t_1	0	3
	1	-1
t_2	0	3
	-1	1

m_1

	a_1	a_2
t_1	-2	1
	1	-1
t_2	-1	2
	-1	1

m_2

	a_1	a_2
t_1	-8	-5
	1	-1
t_2	-4	-1
	-1	1

m_3

$$p^0(t_j) > 0.5$$

Figure 8