March 1986

Consistent Mechanism Design and the Noisy Revelation Principle*

by

Praveen Kumar
GSIA
Carnegie-Mellon University
Pittsburgh, PA-15213.
U.S.A

Abstract: We model the problem of *consistent* mechanism design in a finite horizon economy as an extensive form game of incomplete information. We identify sufficient conditions for the following version of the *Revelation Principle* to apply in our setting: for any sequential equilibrium in the repeated principal-agent game, there corresponds an incentive compatible direct revelation (sequential) equilibrium, where agents truthfully report their private information in the first period (or over time if it evolves). More generally, we show that it may not even be optimal for the principal to induce truthful revelation, and extend the Revelation Principle to the *Noisy Revelation Principle* : the optimal incentive mechanism in the class of *consistent* mechanisms is one where agents update the beliefs of other players in an incentive compatible fashion at every information set.

JEL Numbers 026, 020

* It is a pleasure to acknowledge the help and encouragement of Peter Hammond. All blemishes are my own.

1. Introduction

The theory of incentives, as developed in the past decade or so, has been concerned with the problem faced by a planner (or "principal") whose objective or payoff functions in some way depend on the private information and/or unobservable behavior of the agents (see Laffont and Maskin (1982) for a survey of major themes). The principal then pursues its objectives by precommiting to an *incentive scheme*, i.e., a rule that specifies in advance the principal's response on the basis of its perceptions of the agents' information and/or behavior.

The literature , however, has generally focussed on static or one-shot principal agent games, and therefore not addressed the issue of *enforceability* of these incentive schemes in the following sense: at every information set of the principal, the incentive scheme should be optimal. This issue becomes critical in repeated principal-agent games where the principal cannot credibly precommit to future policies.

The practical importance of such situations can hardly be exaggerated. Many, if not most, economically interesting situations beset with incentive problems involve repeated interactions between the principal and the agents, in the absence of costless precommitment on the part of the principal. Consider for example, the 'game' between a privately informed manager and its (shareholder) employers, or the analogous relationship between a planner and a plant manager in a command economy (Freixas, Guesnerie and Tirole (1986)); likewise, a repeated 'regulation' game between a regulator and a firm with private information on costs (Baron and Besanko (1984), Laffont and Tirole (1986)), or the repeated relation between a tax authority and its constituents who are privately informed of their endowments or preferences (Roberts(1985)).

The restriction of consistency on the principal's mechanism design has substantive consequences for the received theory on incentives. It is by now well known that there are problems with extending the Revelation Principle (Dasgupta, Hammond and Maskin (1979), Myerson (1979), Harris and Townsend (1981)), a basic and convenient result in static incentive theory, to sequential environments in the absence of long-term precommitment by the principal.[1] The Revelation Principle states that allocations attained by any abstract incentive mechanism can be replicated by a direct revelation mechanism where it is (weakly) optimal for agents to truthfully reveal their private information to the mechanism designer. This result has been convenient in expressing the incentive constraints imposed by private information, and has been central to advancing our understanding of the (economic) consequences of asymmetric information.

It then becomes desirable to study the following questions systematically: if the consistency (or the sequential rationality) constraint is imposed on the mechanism design by the planner, what are the conditions under which the Revelation Principle would apply for design of *sequential* incentive mechanisms? And if the Revelation Principle does not generally apply in such contexts, can one still characterize the set of enforceable allocations *usefully*? This latter aspect is important because there is no a priori intuitive restriction on the set of admissible message spaces that may be used in the design of incentive mechanisms. A characterization of the set of attainable allocations for all "admissible" message spaces, whatever the admissibility criteria, may then become quite intractable.

We develop a general finite horizon dynamic principal-agent game where the decision making framework is sequentially hierarchical: at every stage the principal moves first and chooses an incentive mechanism, and allocations are then made on the basis of agents' communications to the principal. A requirement of sequential rationality is imposed on the behavior of the agents *and* the mechanism choice of the principal. In other words, the principal can precommit to a one-period mechanism, but over time is required to choose these mechanisms in a consistent way. An adaptation of the sequential equilibrium concept (Kreps and Wilson (1982)), called the *sequential hierarchical equilibrium* (SHE) is used as the solution concept.

We first ask the following: given any SHE, does there exist another equivalent incentive compatible SHE? i.e., a mechanism where agents truthfully report their types in the first or early periods (or *sequentially* report these truthfully, if the types are evolving), and which yields to the players the same probability of consequences at every stage as in the original equilibrium? A sequential auction example is presented to demonstrate a SHE in which it is not even *desirable* for the seller (the principal in this case) to induce truthful revelation in the first period. Thus, in contrast to static incentive theory, the nonexistence of equivalent truthful mechanisms may be based on their undesirability from the point of view of the principal.

We identify, however, a sufficient condition for the Revelation Principle to apply for the context at hand: if the principal is a utilitarian planner who is (weakly) more patient than the agents in a risk-neutral exchange economy, and if intertemporal transfers are allowed, then the Revelation Principle is valid in the following sense: it loses no generality to restrict attention to mechanisms where it is the equilibrium strategy for agents to truthfully report their private information to the principal in the first period (or sequentially report their private information truthfully if the types are evolving), and receive the consistent (complete information) mechanisms from the second period onwards. These sufficient conditions are of some interest since much of the applications of the static Revelation Principle have been in risk neutral exchange environments. However, they also give some indication of the restrictions that need to be imposed in order for the Revelation Principle to go through without long term precommitment.

We therefore study the more general case where the Revelation Principle does not apply. We derive the *Noisy Revelation Principle (NRP)*, the appropriate extension of the Revelation Principle to the case at hand. The NRP states that in class of *consistent* incentive mechanisms, it loses no generality to consider only mechanisms where the message spaces are the space of (marginal) probability distributions on the agents' types, and where agents essentially announce, at every stage, the posterior beliefs of other players regarding their type in an incentive compatible fashion, i.e., where the randomly announced beliefs always equal the *Bayes consistent* beliefs of the other players.

The content of the NRP is then that to every SHE there corresponds a *noisy revelation* SHE, where the agents announce the marginals on their types (or equivalently the likelihood functions for their types), in which the probability of outcomes at every stage of the game is the same as in the original SHE. It is immediate that this mechanism gives every player the same expected utility at every information stage, as in the general equilibrium.

We argue that this result generalizes to the case where 'types' are evolving over time in a correlated way, and to the case where both moral hazard and adverse selection are present. The Noisy Revelation Principle generalizes in the latter case to the statement that among the class of consistent mechanisms, it loses no generality to consider mechanisms where, along with noisy message spaces, the agents are given incentive compatible probability distributions on their actions: the *Generalized Noisy Revelation Principle* is then both a *noisy revelation* and a *noisy obedience principle*.

The structure of the paper is as follows. Section 2 lays out the basic model, and sets out the solution concept. Section 3 presents the repeated auction example, and Section 4 proves a sufficient condition for the Revelation Principle to hold in the game at hand. Section 5 presents the Noisy Revelation Principle, Section 6 discusses the extensions, and Section 7 concludes.

Before moving on to the analysis, I state some basic notational conventions used throughout the paper: I will denote the profile $(a^1, \ldots, a^N)$ with a, and $\bar{A}$ will denote the Cartesian product of sets $\overset{N}{\underset{i=1}{X}} A^i$. a^{-i} will denote the profile a without the element a^i, and $\bar{A}^{-i}$, the product $\bar{A}$ without A^i. Finally, $\Delta(Z)$ will denote the space of probability measures on Z.

2.1 The Basic Model

I consider a *finite horizon* (T period) single good economy of n+1 players: n agents, and a principal denoted X. To further describe the economy concisely, I make the following definitions:

Table 1

Item	Description	Definition
Θ^i	Finite set of 'type' parameter of agent i=1,...n.	
N	Set of agents	$N=\{1,\ldots n\}$
θ^i	Generic element of Θ^i, $i\varepsilon N$	$\theta^i \varepsilon \Theta^i$, $i\varepsilon N$
C	Finite set of feasible allocations	$C \subseteq R^n$
K_t	Set of feasible allocations at time,	$K_t: C^{t-1} \twoheadrightarrow \Xi$, Ξ = the system of

	$t=1,\ldots T$	non-empty subsets of C
Λ	Prespecified class of message spaces	
M	Generic element of Λ	$M \varepsilon \Lambda$
u^j_t	Instantaneous payoff function of player j at t, $j \varepsilon \{N,X\}$, $t=1,\ldots T$	u^j_t: C x $\Theta \to R$, $j=1,..n,X$, $t=1,\ldots T$

Each agent has private information that is payoff relevant to other players. The private information of agent i is represented by the type parameter $\theta^i \varepsilon\ \Theta^i$, a summary of preferences, production possibilities and beliefs of the agent.[2] There is, however, complete information regarding the principal.[3]

The principal's decision problem is to choose *sequentially* an economic allocation k_t from the period feasible set K_t, $t=1,\ldots T$. This current feasible allocation set is a correspondence depending in the past allocations $(k_1,\ldots k_{t-1})$.[4] These allocations are chosen to maximize the lifetime expected utility of the principal (X). For simplicity, the objective functions for the players are held to be time-additive von-Neumann-Morgenstern expected utility functions, viz. $E \sum_{t=1}^{T} u^j_t\ (k_t,\ \theta),\ j\ \varepsilon\ \{N,X\}$.

The principal pursues his objectives by the following sequential procedure: at the beginning of every period, he asks agents for information regarding their types. He will obtain this information by providing the agents with a 'message' space, an entity best conceptualized as a language or framework of communication between the principal and the agents. On the basis of this information, he proceeds to select an allocation from K_t through a possibly randomized procedure.[5]

The message spaces will be drawn from a prespecified class or system of message spaces Λ. The only assumption made regarding Λ is that $\Delta(\Theta)$, the space of probability measures on Θ is included in Λ.

This procedure sets up an extensive form game of incomplete information with perfect recall. To describe the solution concept, I will use the following definitions.

Table 2

Item	Description	Definition
H_t^j	Space of observable histories of player j, j=1,...n,X	$H_t^j \subseteq (\Theta^j x\ \Lambda^N x\ C)^{t-1}$, t=1,...T
h_t^j	Generic element of H_t^j, j=1,...n,X	$h_t^j \in H_t^j$
$q_t(k_t \mid m_t, h_t^X)$	Randomized allocation rule at t	$q_t \in \Delta(K_t)$
I_t	Incentive mechanism at t	$I_t = \langle M_t, q_t \rangle$
$r_t^i(m_t \mid I_t, h_t^i)$	Agent i's mixed reporting strategy at t, j=1,...N	$r_t^i \in \Delta(M_t^i)$
$p_t^j(\theta \mid h_t^j)$	Beliefs or player j at information set h_t^j j=1,...n,X; t=1,...T.	$p_t^j \in \Delta(\Theta)$

Presented with the incentive mechanism I_t, agents choose their (possibly randomized) reporting or communication strategies r_t^i, independently and non-cooperatively. For simplicity, I will assume that *at the end of every period* agents' communications or messages, and the consequent allocations, become public knowledge ex-post. In other words, there is *confidential* and *simultaneous* communication by the agents (to the principal or a mediator) at every stage, but at the end of the period, these communications became publicly known. Over time, then, agents' knowledge of their types is their only private information.[6]

This assumption is irrelevant to the main results of the paper, as will become clear in the next section when we consider a repeated principal-agent example with only one agent. I will indicate at the appropriate juncture how the main results can be adapted to the more general case where agents observe only some idiosyncratic signals that are correlated with the messages actually sent to the principal.

2.2 The Solution Concept

We can usefully interpret the agents' decision problem in a dynamic programming framework. Let $I^t = (I_t, I_{t+1}, \ldots I_T)$ and $r^t = (r_t, \ldots r_T)$. Then from Bellman's principle of optimality there exist valuation functions $V_t^i(I^t, r^{t+1}, r_t^{-i}, h_t^i)$, $i \in N$, $t=1,\ldots T$; such that,

$$V_t^i(I^t, r^{t+1}, r_t^{-i}, h_t^i) = \underset{r_t^i \in \Delta(M_t^i)}{\mathrm{Max}} \left\{ \sum_\theta p_t^i(\theta | h_t^i) \int_{M_t} \left[\sum_{k_t} \left(\prod_{j=1}^n r_t^j(m_t^j | \theta^j) \right). \right.\right.$$

$$\left.\left. q_t(k_t | m_t, h_t^X) u_t^i(k_t, \theta) + V_{t+1}^i(I^{t+1}, r^{t+1}, h^i_{t+1}) \right] dm_t \right\} \qquad (1)$$

with

$$V_T^i(I_T, r_T^{-i}) = \underset{r_T^i \in \Delta(M_T^i)}{\mathrm{Max}} \left[\sum_\theta p_T^i \int_{M_T} \sum_{k_T} \left(\prod_{j=1}^n r_T^i(m_T^i | \theta^i) \right) q_T(k_T | m_T) u_T^i(k_T, \theta) dm_T \right] \qquad (2)$$

In (1) we have denoted the optimal reporting strategy as $r_t^i(m_t^i | \theta^i)$, rather than the more cumbersome $r_t^i(m_t^i | I_t, h_t^i)$. I will call $R_t^i(I^t, r^{t+1}, r_t^{-i}, h_t^i)$ the *rational* (optimal) *response correspondence* of i at t for the given state (I_t, h_t^i), and given the profile of future optimal incentive mechanisms and reporting strategies, (I^{t+1}, r^{t+1}).

We will also impose a statewise individual rationality constraint on the rational response correspendence: namely, that the valuation from the continuation game must be non-negative for each agent. In other words, the agents cannot precommit (or be forced) to stay in the 'economy'.

The consistent mechanism design problem can then be straightforwardly formulated in the dynamic programming framework developed above. Let $\Phi_\tau \subseteq \Lambda^N \times \Delta(K_\tau)$ be the class of *consistent* incentive mechanisms at time τ. $\{\Phi_t\}_{t=1}^T$ is generated as follows: fix any t, and let I^{t+1} be such that $I_\tau \in \Phi_\tau$, $\forall\ \tau \geq t+1$. Fix any $M_t \in \Lambda$. From the principle of optimality again,

$$V_t^X(I^{t+1}, r^{t+1}, M_t, h_t^X) =$$

$$\underset{q_t \in \Delta(K_t)}{\operatorname{argmax}} \sum_\theta p_t^X(\theta)\{\int_{M_t}[\sum_{k_t}(\prod_{i=1}^n r_t^i(m_t^i|\theta^i))q_t(k_t|m_t)u_t^X(k_t,\theta)+V_t^X(I^{t+1},r^{t+1},h_{t+1}^X)]dm_t\} \tag{3}$$

$$\text{s.t., (i) } r_t^j \in R_t^j(I^t, r^{t+1}, r_t^{-j}, h_t^j),\ \forall\ j \in N \tag{4}$$

$$\text{(ii) } V_t^j(I^t, r^{t+1}, r_t^{-j}, h_t^j) \geq 0,\ \forall\ j \in N \tag{5}$$

The optimal choice of message spaces may then be characterized as,

$$M_t^* \in \underset{M_t \in \Lambda^N}{\operatorname{argmax}}\ V_t^X(I^{t+1}, r^{t+1}, M_t, h_t^X) \tag{6}$$

Then all pairs (M_t^*, q_t^*) satisfying (3) to (6) define the *consistent* mechanism correspondence $\Phi_t(I^{t+1}, r^{t+1}, h_t^X)$. By backwards induction, this procedure then generates the sequence of consistent mechanism correspondences, $\{\Phi_t\}_{t=1}^T$.[7]

It now remains to set out the equilibrium concept for this sequential (hierarchical) game of incomplete information. The profile of beliefs and strategies $\sigma^* = \langle p_t^*, I_t^*, r_t^* \rangle_{t=1}^{T}$ will be called a *sequential hierarchical equilibrium* (SHE) if,[8]

(i) $I_\tau^* \in \Phi_\tau$, $\tau = 1,2,3...T.$

(ii) $r_\tau^{*i} \in R_\tau^i(I^{*\tau}, r^{*\tau+1}, r_\tau^{*-i}, h_\tau^i), \forall h_\tau^i \in H_\tau^i, i \in N, \tau = 1,2,3...T$

(iii) there exists a sequence $\langle p_\tau(\varepsilon), I_\tau(\varepsilon), r_\tau(\varepsilon)\rangle_{\tau=1}^{T}$, where $p(\varepsilon)$ is the system of beliefs generated by $r(\varepsilon)$ such that,

$$\lim_{\varepsilon \to 0} (p_\tau(\varepsilon), I_\tau(\varepsilon), r_\tau(\varepsilon)) = (p_\tau^*, I_\tau^*, r_\tau^*), \forall \tau.$$

I now present a two-period repeated auction example of a sequential game of incomplete information (with communication) of the kind modelled above. It is shown that for plausible parameter values, there exist SHE for which the *principal* does not desire an equivalent truth-revealing subgame perfect equilibrium.

3. Repeated Auction Example

In a two period repeated purchase model, a seller wishes to auction the same object in both periods. The seller is faced by a prospective buyer who will bid for the objects in both periods. The seller, however, is not certain of the subjective valuation of the buyer, v, regarding the object. It is assumed that the buyer's valuations are the same in both periods. Without loss of generality, the seller's subjective valuation of the object is taken to be zero. Finally, let β_b, β_s denote the discount factors of the buyer and seller respectively. These are presumed to be common knowledge.

To simplify matters, suppose there are only two possible valuations: a high valuation, v'', and a low valuation, v' i.e., $v'' > v'$. I will also assume that both the buyer and seller are risk neutral. The static version of this problem has been extensively studied (for example, Myerson (1980), Maskin and Riley (1982)). The sequential mechanism version of the static game is one where the seller moves first and sets up an auction mechanism A_t in each period t=1,2. The auction mechanism in each period consists of:

(i) the price each buyer type must pay contingent on its announced type for the period, and,

(ii) the probability of its acquisition of the good in the period, contingent again on its announced valuation for the period.

Let x_t' (x_t'') be the price that the buyer must pay if it announces v'(v'') in period t (t=1,2), and α_t' (α_t'') be the probability of acquiring the good for the corresponding announcement v'(v'').

Since the second period is the terminal period, the Revelation Principle applies. Letting $p_2(v')$ denote the seller's second period prior that the buyer is the v'-type, the optimal second period mechanism can be shown to be,

$$A_2^* = \begin{cases} \alpha_2'=1,\ x_2'=v',\ \alpha_2''=1,\ x_2''=v'', \text{ if } (1-p_2(v'))v'' < v' \\ \alpha_2'=0,\ x_2'=0,\ \alpha_2''=1,\ x_2''= v'' \quad \text{else} \end{cases} \tag{7}$$

Thus if $p_2(v')$ is high, the expected gain from price discrimination is low, and a flat price of v' is demanded. If $p_2(v')$ is low, the seller can separate by making the high valuation buyer indifferent between announcing v'' and v'.

Notice also that the low valuation seller receives zero surplus in either equilibrium in the second period. Hence, its reporting strategy in the first period is, using $r_1'(v'')$ to denote the probability that v' announces v'' in the first period,

$$r_1'^*(v'') = \begin{cases} 1 \text{ if } \alpha_1'' v'-x_1'' > \alpha_1'v'-x_1' \\ \epsilon[0,1] \text{ if } \alpha_1''v'- x_1' = \alpha_1'v'- x_1' \\ 0 \text{ else} \end{cases} \tag{8}$$

Since the high valuation buyer is concerned with the seller's second period's beliefs, it is also concerned with the v'-type buyer's first period strategy. Hence,

$$r_1''^*(v'') = \begin{cases} 1 \text{ if } r_1'^* (v'') = 1 \\ \underset{r_1''(v'')}{\text{Max}} \{(1-r_1''(v''))[\alpha_1'v'' - x_1' + Y_1(v')] + r_1''(v'')(\alpha_1''v''-x_1'' + Y_1(v''))\} \\ \qquad Y_1(z) = \begin{cases} 1 \text{ if } v_1' > [\frac{r_1''(z)p_1(v'')}{r_1''(z)p_1(v'')+r_1'(z)p_1(v')}]v'', \ z=v',v'' \\ 0 \text{ else} \end{cases} \end{cases} \tag{9}$$

The optimal first period mechanism then maximizes the seller's lifetime expected revenues taking as given A_2^*, $r_1'^*(v'')$, $r_1''^*(v'')$, and being subject to the participation constraint, viz.,

$$\underset{\{A_1\}}{\text{Max}} \ U^S(A_1;\ A_2^*,\ r_1'^*,\ r_1''^*)$$

s.t., (i) $r_1'^*$, $r_1''^*$ as in (8) and (9),

(ii) $r_1'^*(v')(\alpha_1'v' - x_1') + (1-r_1'^*(v'))(\alpha_1''v' - x_1'') \geq 0.$

(iii) $r_1''^*(v')(\alpha_1'v'' - x_1') + (1-r_1''^*(v'))(\alpha_1''v'' - x_1'') \geq 0.$

Suppose now that $\beta_s = 0.75$, $\beta_b = 0.8$, $v'' = 5.0$, $v' = 1.5$, $p_1(v') = 0.75$. For the particular parametization at hand, it can be shown that the first period optimal mechanism induces a pooling equilibrium, i.e., is of the form,

$$A_1^* = \langle \alpha_1^{*\prime} = 1,\ \alpha_1^{*\prime\prime} = 1,\ x_1^{*\prime} = 1.5,\ x_1^{*\prime\prime} = 5.0 \rangle .$$

Plugging A_1^* and A_2^* into the optimal response functions (8) and (9) yields, $r_1^{\prime *}(v') = 1$, $r_1^{\prime\prime *}(v') = 1$, $p_2(v') = p_1(v') = 0.75$. Next let, $U^i(A_1^*, A_2^*)$ denote the expected lifetime utility of player i, i = s,v',or v''. Then, from above, $U^s(A_1^*, A_2^*) = 2.625$, $U'(A_1^*, A_2^*) = 0$, $U''(A_1^*, A_2^*) = 6.3$.

The important question of whether there exists an *incentive compatible* SHE equivalent to the $\langle A_1^*, A_2^* \rangle$ equilibrium may now be addressed. Let $\tilde{A}_2$ be the optimal second period mechanism given that an incentive compatible mechanism obtains in the first period. Then, consistency requires that the seller extract the buyer's surplus in the next period, i.e., $\tilde{x}_2 = v''(v')$ if $v''(v')$ is announced in the first period.

The v''-type clearly has an incentive to be untruthful with respect to SHE (A_1^*, A_2^*).[9] In the spirit of the method of construction of equivalent incentive compatible mechanisms, an equivalent optimal first period incentive compatible mechanism is the solution to the following programme:

$$\max_{\{\tilde{A}_1\}} U_s(\tilde{A}_1, \tilde{A}_2). \tag{P}$$

$$\text{s.t., (i) } \tilde{\alpha}_1' v' - \tilde{x}_1' \geq U'(A_1^*, A_2^*) \quad \text{(ii) } \tilde{\alpha}_1'' v'' - \tilde{x}_1'' \geq U''(A_1^*, A_2^*),$$

where, $U^s(\tilde{A}_1, \tilde{A}_2) = p_1(v')(\tilde{x}_1' + \beta_s v') + (1 - p_1(v'))(\tilde{x}_1'' + \beta_s v'')$.

(i) and (ii) are a consequence of the constraint that expected payoffs from truthful revelation in the first period weakly dominate the corresponding payoffs in the original equilibrium.

The solution to the program P, is the following:

$$\tilde{A}_1 = \langle \tilde{\alpha}_1' = 1, \ \tilde{x}_1' = v', \ \tilde{\alpha}_1'' = 1, \ \tilde{x}_1'' = v'' - U''(A_1^*, A_2^*)\rangle.$$

For the particular parameterization at hand this entails,

$$U'(A_1^*, A_2^*) = U'(A_1^*, A_2^*); \ U''(A_1^*, A_2^*) = U''(\tilde{A}_1, \tilde{A}_2); \ U^S(\tilde{A}_1, \tilde{A}_2) = 2.58$$

Notice that $U^S(\tilde{A}_1,\tilde{A}_2) < U^S(A_1^*, A_2^*)$. In other words, *the seller is strictly worse off in the corresponding incentive compatible mechanism that gives equivalent utility to the agents.*[10] Thus a necessary condition for the construction of an equivalent sequentially rational incentive compatible mechanism does not hold: the principal is strictly worse off from the incentive compatible construction. The reason, as illustrated by this example, is that the less patient seller has to "compensate" the patient v"-type buyer for the loss of second period utility that follows from the truthful revelation in the first period. Players' subjective discount factors then play an important role in the *optimality* of equivalent (incentive compatible) mechanisms. It can be readily seen that $U^S(\tilde{A}_1,\tilde{A}_2)$ is strictly decreasing in β_b, since $\tilde{x}_1''$ is strictly decreasing in β_b. In fact, if $\beta_s \geq \beta_b$, then $U^S(\tilde{A}_1,\tilde{A}_2) \geq U^S(A_1^*,A_2^*)$, if $v' > (1-p_1')v''$.

Similar observations can be made regarding the subjective valuation difference, v" - v', for v" fixed. If this difference is low, the higher-valuation type has a lower net payoff in the mechanism A_2^*. Consequently, the expected loss from revealing itself in the first period is also lower. Thus, $U^S(\tilde{A}_1,\tilde{A}_2)$ is also decreasing in (v" - v'), for v" fixed. In fact,

$U^s(\tilde{A}_1, \tilde{A}_2)$ is increasing in any factor that decreases the expected payoffs of the v"-type in the (second period) incentive compatible mechanism, A_2^*.

Notice also that as T increases, the cost of inducing truthful revelation in the first period, (or in the early periods), increases. In general, for a given T, the problems with constructing incentive compatible mechanisms in early periods increase with the discount factors of the agents: if agents had incentives to be untruthful in all these periods, there is greater expected loss to revealing private information completely in the first period itself. The extreme case where agents live forever, do not discount future welfare, and have incentives to be untruthful in every period is considered by Roberts (1985). As would be suggested by the intuition of this example, *no* useful information is conveyed in any subgame perfect equilibrium in that model.

4. A Sufficient Condition for the Revelation Principle

The intuition that there is greater information revelation in the early periods if the principal is more patient then the agents can be formalised in certain contexts. In particular, in a risk neutral exchange economy (or when the allocation space is intertemporally independent) where intertemporal resource transfers are feasible, and the principal is a utilitarian planner, the following version of Revelation Principle goes through even *without* precommitment: to every SHE there corresponds a direct revelation SHE, where the agents reveal their private information truthfully to the principal in the first period.

The idea is that the principal can give 'upfront' payments for information revelation in the first period and give the complete information allocations from then on. This initial payment is designed to leave the agents' expected utility equal to the original SHE. To state this result concisely, let σ^* denote the SHE $\langle M_1^{i*} = \theta^i,\ i \in N,\ q_1^*,\ r_1^*,\ (c_2^*, \ldots c_N^*) \rangle$ where r_1^* *is the profile of incentive compatible* (truthful reporting) strategies, and $(c_2^*, \ldots c_T^*)$ are the complete information allocations for $t = 2, \ldots T$.

Theorem 1: *Suppose that the principal is a utilitarian planner who is at least as patient as the agents. If, (a) the allocation spaces are intertemporally independent, and (b) intertemporal resource transfers are feasible, then to every SHE* σ *there exists a corresponding incentive compatible SHE* σ^**, such that,* $V_t^j(\sigma^*) \geq V_t^j(\sigma)$, $\forall\ j=\{N,X\}$, $t=1,2\ldots T$.

For example, in a risk neutral exchange economy where the agents discount, and the utilitarian planner does not discount (say), there is no loss of generality in restricting attention to games with incentive compatible mechanisms in the first period, and first best allocations from then on. Of course, the first period incentive compatibility constraints will be based on *lifetime* expected utility, taking as given the consistent (complete information) allocations of the principal from the second period onwards.

It is easy to see that the conditions of Theorem 1 would also suffice for the Revelation Principle to apply even if the private information of the agents was evolving over time: it loses no generality to restrict attention to mechanisms where, in equilibrium, agents reveal their private information truthfully to the planner in every period.

These sufficient conditions, while of considerable economic interest, are obviously restrictive. I therefore move to the analysis of the general case where the Revelation Principle does not apply.

5. The Noisy Revelation Principle

The Noisy Revelation Principle (NRP) stems from the observation that at any $t=1,\dots T$, the feasible allocation set K_t, and the (profile of) beliefs p_t represent a sufficient statistic on the past history of mechanisms, reports and allocations. This is immediate from the fact that agents' are Bayesian decisionmakers, and from Bellman's principle of optimality (see section 2). Hence, for any given SHE, one can attempt to construct an equivalent SHE where at every stage agents update the priors of other players in an incentive compatible fashion, i.e., the updating is done in a manner such that that the probability of outcomes (allocations and beliefs) at every information set are the same as in the original equilibrium.

This construction, however, is not so straightforward as the 'equivalence' construction used to establish the Revelation Principle (see Myerson (1979)). I first present the Noisy Revelation Principle and then indicate the problematic issues.

An incentive mechanism $\tilde{I}_t = \langle \tilde{M}_t, \tilde{q}_t \rangle_{t=1}^{T}$ will be called a *noisy revelation mechanism* if,

$$\tilde{M}_t^i = \Delta\,(\Theta^i),\ i \epsilon N, \quad \text{and } \tilde{q}_t : K_t \times \Delta(\Theta) \to [0,1], \quad t=1,\dots T.$$

In other words, the noisy revelation mechanism has the space of marginal beliefs for each agent as its message space, and the allocation rule is conditional on the announced profile of marginal beliefs, i.e., on $(g^1(\theta^1),\dots,g^n(\theta^n))$, $g^i \;\epsilon\; \Delta(\Theta^i)$, $i \;\epsilon\; N$.

Next, let $p_t^{X,i}(.)$ be the principal's marginal beliefs regarding agent i's type, $i \;\epsilon\; N$. Then for any t and i, for any profile of prior beliefs p_t^i, the (mixed) reporting strategy $\tilde{r}_t^{i*} \;\epsilon\; \Delta(\Delta(\Theta^i))$, $i \;\epsilon\; N$, will be called an *incentive compatible noisy revelation (ICNR)* strategy for agent i if,

(i) $\tilde{r}_t^{i*} \varepsilon R_t^i(\tilde{I}^t, \tilde{r}^{*t+1}, \tilde{r}_t^{-i*}, p_t^i), \forall\ i \ \varepsilon\ N$, and,

(ii) for every $g^i(\theta^i)\ \varepsilon\ \Delta(\theta^i|\tilde{r}_t^*)$, the range of $\tilde{r}_t^*$,

$$p_{t+1}^{X,i}(\theta^i|\tilde{g}^i(\cdot),\tilde{r}_t^*) = \tilde{g}^i(\theta^{i\prime}), \forall\ \theta^{i\prime} \varepsilon\ \Theta^i, i\ \varepsilon\ N$$

Thus an ICNR strategy is the agent's dynamically consistent strategy for announcing the marginal probability distribution such that the announced beliefs are themselves the the *principal's Bayesian consistent marginal posterior beliefs regarding his type.* Finally, the pair (of profiles) $\langle \tilde{I}_t, (\tilde{r}_t^{1*},\ldots,\tilde{r}_t^{*n})\rangle$, will be called an incentive compatible noisy revelation (ICNR) mechanism. We are now ready to present the main result of the paper.

Theorem 2: *For each* $t=1,\ldots\ \varepsilon\ T$, *any* $h_t^X \varepsilon H_t^X$, *the optimal incentive mechanism in the class of all sequentially rational hierarchical mechanisms is an ICNR mechanism.*

An immediate corollary of Theorem 2 is that to every SHE with arbitrary message spaces, there exists a corresponding SHE with noisy mechanisms where the probability of outcomes - allocations and beliefs - at every information set is the same as in the original equilibrium. Theorem 2 then says that in designing *consistent* incentive mechanisms in sequential games of incomplete information with communication (with Bayesian decision makers) it loses no generality to restrict attention to ICNR mechanisms that use the space of probability distributions on types as the message spaces. Since agents are choosing independent (uncorrelated) announcement strategies, it is enough for them to announce the relevant marginal distributions on their types. The principal can solve for the likelihood functions (announcement strategies)

$\{r_t^{i*}\}_{i\varepsilon N}$, and construct the joint distribution on the agents' type associated with the original equilibrium. Furthermore, since the profile of messages is common knowledge, each agent can similarly reconstruct his beliefs associated with the original equilibrium (see Appendix for the relevant details).

Effectively, then, at each stage, agents now announce, in a consistent fashion, what they would prefer the principal and the other agents to believe as a basis for their future decisions. This notion of incentive compatible mechanisms includes, as a special case, the case where agents announce the degenerate beliefs,

$$g^i(\theta^i|\tilde{\theta}^i) = \begin{cases} 1 \text{ if } \theta^i = \tilde{\theta}^i. \\ 0 \text{ else.} \end{cases}$$

i.e., reveal themselves truthfully.

If agents observe only some idiosyncratic signals on the actual messages transmitted to the planner, then the notion of incentive compatibility implies that at each stage, agents announce in a consistent fashion what they would prefer *every other player* to believe as a basis for future actions.

There are two basic issues that have to be resolved in the construction of the equivalent noisy mechanisms: firstly, notice that in constructing the equivalent ICNRs we have to ensure that the *joint probability* on allocations and posterior beliefs at every stage is the same as in the original equilibrium, even if the original equilibrium had two different allocations from two messages that implied the same posterior beliefs. This equivalence construction will become clearer from the following example:

Consider the special case where: $n = 1$, $\theta^1 = (\theta_1, \theta_2)$. Let $T=2$, and suppose $I_1 = \langle M_1 = (m^1, m^2), q_1(k|m^i), i = 1,2\rangle$, with the equilibrium reporting strategies, $\langle r_1^1(m^1|\theta_1)=1/2,\ r_1^1(m^1|\theta_2)=2/3\rangle$. Finally, let $p_1^X=(3/4,1/4)$. Then it is easily seen that $p_2^X(\theta^1=\theta_1|m^1)=9/13$; and $p_2^X(\theta^1=\theta_1|m^2)=9/11$.

Let $g^1(\theta_1,\theta_2)=(9/13,\ 4/13)$, and $g^2(\theta_1,\theta_2)=(9/11,\ 2/11)$. If the agent is of type θ_1, the equivalent strategies are $\tilde{r}_1^1(g^1)=1/2$, $\tilde{r}_1^1(g^2)=1/2$; $\tilde{q}_1(k|g^1)=q_1(k|m^1)$, and $\tilde{q}_1(k|g^2)=q_1(k|m^2)$.

Suppose now that $p_1^X=(1/2,\ 1/2)$ and both types follow the same reporting strategies in the original equilibrium, viz. $r_1^1(m^1|\cdot)=1/2$, $r_1^1(m^2|\cdot)=1/2$. Then, $p_2^X=p_1^X$, and the equivalent strategies are, $\tilde{r}_1^1(g'|\cdot)=1/2$ for both types, where $g'=(1/2,\ 1/2)$, and $\tilde{q}_1(k|g')=\sum_{i=1}^{2}q_1(k|m^i)$.

A more subtle issue relating to the 'sparseness' of the message spaces also arises in the equivalent 'noisy' construction. Suppose, again, for simplicity that n=1, and $\Theta^1=\{\theta_1,\theta_2\}$. Further suppose that in the original SHE $M_1=\{m_1\}$. Then the principal could not have obtained any useful information in the original equilibrium, i.e., the posterior belief must be equal to the prior beliefs. In the noisy revelation mechanism, however, the agent is provided with a very rich message space, where the entire set of posterior beliefs is in the range of the beliefs induced by his reporting strategies. It is then possible that the original strategies are no longer an equilibrium with the richer message space.

This issue does not arise in the *static* Revelation Principle since, in constructing the equivalent mechanism, the principal can precommit to choose his allocations *as if* no information is revealed. Here the principal *cannot* precommit to arbitrary posterior beliefs in the future. This difficulty is resolved by the fact that in any *sequentially rational* mechanism, the principal must employ the most desirable message spaces (from (7)), and hence if it is worth inducing information revelation in the noisy mechanism, the same information must also have been induced in the original equilibrium.

6. Extensions

6.1 Agents with Evolving Characteristics

In models of incomplete information, of the kind considered up to this point, the agent's characteristic or type, θ^i is usually interpreted as a summary of all the relevant features that may affect the allocation, including preferences, endowments, abilities and beliefs. In a truly dynamic model, all these features are subject to change due, say, to the arrival of new information or the current effects of past agent decisions such as the investment in physical or human capital. The design of incentive mechanisms in intertemporal contexts should, in principle, then take into account that the agent characteristics may themselves evolve over time. In fact, in many important intertemporal incentive problems, the idea of designing mechanisms that over time would allow the principal to "know everything" about the agents is really an incongruous mechanical transplantation of a result that makes sense only in static models. I now informally argue that the framework developed in section 2 can be easily extended to handle this case so that the Noisy Revelation Principle would continue to apply.

Suppose that the profile of agent characteristics, θ, follows a Markov chain with a non-empty and finite support Θ, and with the stationary transition probability mass function, $g\colon \Theta \times \Theta \to [0,1]$, with a given initial period probability mass function $g_1\colon\Theta \to [0,1]$. Thus, current period realization of the characteristics or type yields some information on future type through the transition probability $g(\theta_t|\theta_{t-1})$. This characterization is without loss of generality if Θ is treated as an "agent" space of sufficient richness.[11]

The information structure of the model is the same as in section 2. At the beginning of every period, only each agent observes the realization of its own type for that period. The principal then presents an incentive mechanism I_t, sequentially, to obtain information on the *current realization* of the agents' types. As before, we shall assume that the players' messages are ex post common knowledge.

With the suitable expansion of the state space of agents to allow changing types over time, and using the appropriate updating procedure the structure of the game remains as given in section 2. For example, player i's prior on the second period type profile θ_2 is given by,

$$p_2^i(\theta_2|h_2^i) = \sum_{\theta \varepsilon \Theta} g(\theta_2|\theta_1)p_2^i(\theta_1|h_2^i)$$

where

$$p_2^i(\theta_1|h_2^i) = (\prod_{j=1}^{n} r_1^j(m_1^j|\theta_1^j))\ p_1^i(\theta_1)/\sum_{\theta_1 \varepsilon \Theta} (\prod_{j=1}^{n} r_1^j(m_1^j|\theta_1^j))p_1^i(\theta_1)$$

This procedure then inductively defines the updated beliefs for t =3,...T.

With these modifications, the definition of a SHE (section 2) serves the present case as well. From the discussion in section 3 we know that problem with the Revelation Principle will remain as long as current information transmission affects future payoffs. These linkages occur as long as the principal's *prior beliefs in the future are influenced by current information transmission*, i.e., as long as the $\{\theta(t)\}_{t=1}^{T}$ is not an independent process. One conjectures also that with changing types, the costs to agents of truthfully revealing their types in early periods increased with the *correlation* in the process generating the types. Finally, it is easily shown that with the appropriate updating procedure, the Noisy Revelation Principle extends to this general case as well, where each agent now uses ICNR strategies for other players' prior beliefs for his next period type.

6.2 The Generalized Principal-Agent Problem and the Noisy Revelation and Obedience Principles

The basic model of section 2 can also be extended to include the possibility that agents undertake actions that are only privately observable when the principal's welfare is dependent on these actions. The principal only observes signals that are stochastically related to the agents' actions.

In static models or in situations where the principal can credibly precommit (Myerson 1982), the Revelation Principle is extended to this "generalized principal - agent" model as follows: it loses no generality to restrict attention to mechanisms where agents truthfully reveal their private information and *obey* the directives of the principal regarding their actions. The Generalized Revelation Principle is then both a Revelation and an Obedience Principle, so to speak.

In a sequential framework without precommitment, however, the Obedience principle also poses problems when the principal can observe some signals (stochastically) related to the agents' actions. The reason is that the observed signals may yield information on the particular realization of the stochastic process that links the actions with the signals. In particular, if the agents were to obey the principal he could often precisely infer the realization of those stochastic shocks. If these are not independent over time, this inference could be used against the agent in the design of future incentive mechanisms, as was the case when only adverse selection problems were present.

In this case, it can be shown that the Generalized Revelation Principle extends to the Generalized Noisy Revelation Principle in the following way: the optimal sequentially rational mechanism is one where agents are proffered the space of probability distributions as their message space, *and* an incentive compatible probability distribution on their action space. In other words, there is both noisy revelation and noisy obedience. As earlier, to every SHE in the original game with arbitrary message spaces and action strategy of agent, there corresponds a SHE with noisy obedience and revelation mechanisms where the probability of outcomes at every information set is the same as in the original SHE, and the principal is at least as well off.

7. Some Concluding Comments

In analogous games with precommitment, the Revelation Principle is useful because it is usually much easier to characterize the set of direct incentive compatible mechanisms, than to characterize the set of all Bayesian equilibria. In like fashion, the Noisy Revelation Principle is useful here since it may be much easier to characterize the set of ICNR allocations than to characterize the set of allocations sustained as SHE. This is especially so since there is no intuitive restriction on the set of admissible message spaces. Put another way, the Noisy Revelation Principle provides a succinct characterization of the set of feasible or attainable allocations in sequential economies with incentive constraints but without precommitment.

APPENDIX

Proof of Theorem 1: Since the players are risk-neutral, let their period or current payoffs be as follows:

$$u_t^j(c,\theta) = \beta_j^t[A_j(\theta) + \sum_{i=1}^{n} \alpha_{ij}\theta_i c_i], \quad j=1,\ldots n, \ 0 < \beta_j < 1 \tag{a.1}$$

The principal's period felicity or utility is then,

$$u_t^X(c,\theta) = \beta_X^t \sum_{j=1}^{n} \gamma_j u^j(c,\theta) \qquad 0 < \beta_X < 1 \tag{a.2}$$

The players maximize their time-additive lifetime expected utility functions, $E\{\sum_{t=1}^{T} \beta_j^t u^j(c,\theta)\}$, $j=1,\ldots n,X$. Next, let, $(c_2^*,\ldots,c_T^*)$ denote the solution to the first best or complete information problem,

$$\underset{\langle c_2,\ldots c_T\rangle \in C^{T-1}}{\text{Max}} \sum_{t=2}^{T} \beta_X^t u^X(c_t,\theta) \tag{a.3}$$

Now fix any SHE $\sigma' = \langle p_t', I_t', r_t'\rangle_{t=1}^{T}$ and define, for any $m_1' \in M_1'$,

$$v^j(\sigma',\theta^j, m_1') = E_{\sigma',\theta}[\sum_{t=2}^{T}\beta_j^t u^j(c_t, \theta)|m_1'] - E_\theta[\sum_{t=2}^{T} \beta_j^t u^j(c_t^*, \theta)|m_1'] \tag{a.4}$$

$v^j(\sigma', \theta^j, m_1')$ is θ^j type's excess expected utility from the equilibrium σ', relative to the first best, from period two onwards, given that the message profile m_1' was sent in the first period. Now we construct, for any $\theta \in \Theta$, the profile of certainity equivalent consumptions, $\hat{c}^j(\sigma', \theta, m_1') \in TC$, $\forall j \in N$, from the system of linear equations,

$$\sum_j \alpha_{ji}\theta_j \hat{c}^j(\sigma', \theta, m_1') = v^i(\sigma', \theta^i, m_1'), \; i \in N \tag{a.5}$$

Let σ^* denote the mechanism profile $\langle I_1^*, c_2^*, \ldots c_T^*\rangle$, where I_1^* is a *direct revelation mechanism*, i.e., $I_1^* = \langle M_1^i = \theta^i, \; i \in N, \; q_1^*(c_1|\theta)\rangle$, and where q_1^* is constructed in the following way. For every $\tilde{c}^* \in TC$, $\theta \in \Theta$, let,

$$\phi(\tilde{c}^*, \theta) = \{y = (c_1', m_1') \in C \times M_1' | \; c_1' + \hat{c}(\sigma', \theta, m_1') = \tilde{c}^*\} \tag{a.6}$$

$\phi(\tilde{c}^*, \theta)$ is the set of first period consumption and message profiles in the original equilibrium that would imply a total certainity equivalent consumption of $\tilde{c}^*$. Now let,

$$q_1^*(\tilde{c}^*|\theta) = \int_{\phi(\tilde{c}^*, \theta)} (\prod_i r_1^{i\prime}(m_1^{i\prime}|\theta^i))q_1' \; (c_1'|m_1') \; dy \tag{a.7}$$

Then, by construction, and using the linearity of u^j, $j \in N$,

$$E_{\sigma', \theta} \sum_{t=1}^{T} \beta_j^t \; u^j(c_t', \theta) = E_{\sigma^*, \theta} \sum_{t=1}^{T} \beta_j^t \; u^j(c_t^*, \theta), \quad \forall \; j \in N, \; \theta \in \Theta \tag{a.8}$$

Hence, if σ' is SHE, then so also are the incentive compatible reporting strategies:

$$r_1^{i*}(\tilde{\theta}^i|I^*, \theta^i) = \begin{cases} 1 \text{ if } \tilde{\theta}^i = \theta^i \\ 0 \text{ else} \end{cases} \quad i \in N$$

Moreover, from (a.7) and (a.8), for the principal,

$$E_{\sigma^*, \theta} \sum_{t=1}^{T} \beta_X^t \; u^X(c_t^*, \theta) - E_{\sigma', \theta} \sum_{t=1}^{T} \beta_X^t \; u^X(c_t', \theta) \; =$$

$$\sum_{\theta} p_1(\theta) \sum_{t=2}^{T} \sum_{i=1}^{n} (\beta_X^t - \beta_i^t)\gamma_i u^i \; (c_t^*, \theta) \geq 0, \quad \text{if } \beta_X \geq \beta_i, \; \forall \; i \varepsilon N. \qquad (a.9)$$

Q.E.D.

Proof of Theorem 2: Consider any given SHE $\sigma^* = \langle p_t^*, I_t^*, r_t^* \rangle_{t=1}^{T}$, with arbitrary message spaces. For any t=1, ...T, let $p_{t+1}^{X,i}(\theta^i) \varepsilon \Delta(\Theta^i)$ denote the marginal with respect to θ^i, $i \varepsilon N$, associated with the joint posterior beliefs $p_{t+1}^X(\theta) \varepsilon \Delta(\Theta)$. Now consider any noisy revelation mechanism $\tilde{I}_t^*$, with $M_t^{i*} = \Delta(\Theta^i)$, $i \varepsilon N$. For any given noisy revelation message profile $(\tilde{p}_t^{X,i}(\theta^i))_{i \varepsilon N}$, we obtain the $n \times |\Theta|$ system of equations in $n \times |\Theta|$ unknowns $\{r_t^{i*}(m_t^{i*}|\theta^i), \theta^i \varepsilon \Theta^i, i \varepsilon N\}$, viz.,

$$\tilde{p}_t^{X,i}(\theta^i) = \sum_{\substack{\theta^j \\ j \neq i}} [(\prod_{z=1}^{n} r_t^{z*}(m_t^{z*}|\theta^z) p^X(\theta) / \sum_{\theta \varepsilon \Theta} (\prod_{z=1}^{n} r_t^{z*}) p^X(\theta)], \; \theta^i \varepsilon \Theta^i \qquad (a.10)$$

It can be straightforwardly checked through the application of the implicit function theorem that the solution to the sysytem (a.10) exists. Then by Bayes rule,

$$\tilde{\tilde{p}}_{t+1}^X \; (\theta \; |(\tilde{\tilde{p}}^{X,1},..,p^{\approx X,n})) = \frac{(\prod_i r_t^{i*}(m_t^{i*}|\theta^i)) \; p_t^X(\theta)}{\sum_{\theta \varepsilon \Theta} (\prod_i r_t^{i*}(m_t^{i*}|\theta^i)) \; p_t^X(\theta)} \qquad (a.11)$$

Next define $\phi_t^i(\tilde{\tilde{p}}^{X,i}|I_t^*, r_t^*)$ to be the set of messages sent by agent i at t which yield the marginal beliefs $\tilde{\tilde{p}}^{X,i}$ as the posteriers in the original equilibrium σ^* i.e.,

$$\phi_t^i(\tilde{\tilde{p}}^{X,i}|I_t^*, r_t^*) = \qquad (a.12)$$
$$\{m_t^{i*} \varepsilon M_t^{i*} | \; p_{t+1}^{X,i}(\theta^i|m_t^{i*}, m_t^{-i*}) = \tilde{\tilde{p}}^{X,i}(\theta^i), \; \forall \; m_t^{-i*} \varepsilon M_t^{-i*}, \; \theta^i \varepsilon \Theta^i\}.$$

Since the agents are choosing their reporting strategy independently, it follows that,

$$\text{Prob}(p^X_{t+1} = \tilde{\tilde{p}}^X(.)|I_t^*, r_t^*) = \sum_{m_t^* \varepsilon\ \phi_t(\tilde{\tilde{p}}_t^{X,i}|I_t^*, r_t^*)} \left(\prod_{j=1}^{n} r_t^{j*}(m_t^j)\right) \quad (a.13)$$

where,

$$\phi_t(\tilde{\tilde{p}}^X|I_t^*, r_t^*) = \prod_{i=1}^{n} \phi_t^i(\tilde{\tilde{p}}^{X,i}|I_t^*, r_t^*)$$

Now if we let each agent $i \varepsilon N$ announce the marginal posteriors associated with his type by the following rule,

$$\tilde{r}_t^{i*}(\tilde{\tilde{p}}^{X,i}(\theta^{i\prime}) \mid \theta^i) = \sum_{m_t^{i*} \varepsilon\ \phi_t^i(\tilde{\tilde{p}}^{X,i}|I_t^*, r_t^*)} [r_t^{i*}(m_t^{i*}|\theta^i)] \quad (a.14)$$

then the ex ante probability that the principal will hear the marginal profile, $(\tilde{\tilde{p}}^{X,1}(.),\ldots,\tilde{\tilde{p}}^{X,n}(.))$, and hence construct the joint distribution $\tilde{\tilde{p}}^X_{t+1}(.)$, is just,

$$\text{Prob.}(\tilde{\tilde{p}}^X_{t+1}(.)|\tilde{r}_t^*) = \prod_{j=1}^{n} \text{Prob}(\tilde{\tilde{p}}_t^{X,j}(.)|\tilde{r}_t^{j*}) = \prod_{j=1}^{n} \tilde{r}_t^{j*}(\tilde{\tilde{p}}_t^{X,j}) \quad (a.15)$$

Then from (a.13) and (a.15) we have that,

$$\text{Prob.}\ (p^X_{t+1}\big|I_t^*,\ _t^*) = \text{Prob.}\ (p^X_{t+1}\big|\tilde{I}_t, \tilde{r}_t^*),\ \forall\ p^X\ \varepsilon\ \Delta(\Theta|\sigma_t^*) \quad (a.16)$$

where $\Delta(\Theta|\sigma_t^*)$ is the support of the posteriors induced by the original equilibrium at t. Next, let,

$$\tilde{q}_t^*(k'|\tilde{\tilde{p}}^{X,1}(\theta^1),\ldots, \tilde{\tilde{p}}^{X,n}(\theta^n)) =$$

$$\sum_{m_t^* \in \phi_t(\tilde{\tilde{p}}^X | I_t^*, r_t^*)} (\prod_{j=1}^{n} r_t^{j*}(m_t^{j*} | \theta^j)) \; q_t^*(k' | m_t^*)\} / \; [\prod_{j=1}^{n} \tilde{r}_t^{j*}(\tilde{\tilde{p}}^{X,j}(\theta^j))]. \tag{a.17}$$

(a.16) and (a.17) together imply that,

$$\text{Prob.}(p_{t+1}^X, k_t | I_t^*, r_t^*) = \text{Prob.}(p_{t+1}^X, k_t | \tilde{I}_t^*, \tilde{r}_t^*) \; \forall \; (p_{t+1}^X, k_t) \in \Delta(\Theta | \sigma_t^*) \times K_t \tag{a.18}$$

Moreover, since the profile of messages announced is common knowledge, each agent can similarly derive the likelihood functions r_t^{j*} from the profile of announced marginal distributions, $\tilde{\tilde{p}}^{X,j}, j \in N$. And by applying Bayes' rule he can derive the joint updated beliefs, $p_{t+1}^i(\theta \mid (\tilde{\tilde{p}}^{X,1} \ldots \tilde{\tilde{p}}^{X,n}))$ for each i, as in the original equilibrium. Hence,

$$\text{Prob.}(h_{t+1}^i | I_t^*, r_t^*, h_t^i) = \text{Prob.}(h_{t+1}^i | \tilde{I}^*, \tilde{r}_t^*, h_t^i), \; \forall \; h_{t+1}^i \in H_{t+1}^i$$

It now remains to show that the noisy revelation reporting rule given in (a.14) is an ICNR strategy, given $(\tilde{I}^{*t+1}, \tilde{r}^{*t+1})$, for the agents. First notice that,

$$\text{Prob.}\ (\theta' | \tilde{\tilde{p}}^X(.),\ \tilde{r}_t^*) = \frac{(\prod_{j=1}^{n} \tilde{r}_t^{j*}(\tilde{\tilde{p}}^{X,j} | \theta^{j'})) \; p_t^X(\theta')}{\sum_{\theta' \in \Theta} (\prod_j^{n} r^{\sim j*}(\tilde{\tilde{p}}^{X,j} | \theta')) \; p_t^X(\theta)} \tag{a.19}$$

Substituting (a.14) for $\tilde{r}_t^{j*}$ in (a.19) yields,

$$\text{Prob.}(\theta' | \tilde{\tilde{p}}^X(.), \tilde{r}_t^*) = \frac{\tilde{\tilde{p}}^X(\theta') = [p_t^X(\theta') \sum_{m_t^* \in \phi_t^X(\tilde{\tilde{p}}^X | I_t^*, r_t^*)} [\prod_j r_t^{j*}]}{\sum_{\theta' \in \Theta} \tilde{\tilde{p}}^X(\theta')}$$

$$= \tilde{\tilde{p}}^X(\theta')$$

i.e., if the noisy revelation strategy is $\tilde{r}^*$, then the announced beliefs are themselves the Bayes consistent posteriors of X.

Now suppose that some player, say $j \in N$ defects from the reporting strategy $\tilde{r}_t^{j*}$, to some other strategy $\tilde{r}^{j'}_t \in \Delta(\Delta(\Theta^j))$ $j \in N$. For each $i = 1,\ldots n$, X, let,

$$P^{j,i}(\tilde{I}^*_t, \tilde{r}^{-j*}, \tilde{r}^{j'}_t) = \{p \in \Delta(\Theta) | p = p^i_{t+1}(\theta | p^i_t, \tilde{r}^{-j*}_t, \tilde{r}^{j'}_t)\} \tag{a.20}$$

$P^{j,i}$ is then the set of all posterior probabilities of player i that agent j can possibly induce by using the other strategy in $\Delta(\Delta(\Theta^j))$, when all other agents are using the prescribed strategies r^{i*}_t, $i \in N/\{j\}$. Further, let,

$$P^{j,i}(\tilde{I}^*_t, \tilde{r}^{-j*}_t) = \bigcup_{\tilde{r}^j \in \Delta(\Delta(\Theta^j))} [P^{j,i}(\tilde{I}^*_t, r^{-j*}_t, \tilde{r}^{j'}_t)]$$

Then, first, suppose that,

$$P^{j,i}(\tilde{I}^*_t, \tilde{r}^{-j*}_t) = P^{j,i}(M^*_t, r^{-j*}_t), \ \forall\, i \in \{N,X\}/\{j\} \tag{a.21}$$

Then no agent j can do better by defecting from $\tilde{r}^{j*}_t$ in the noisy mechanism $\tilde{I}^*_t$. If this were indeed possible, it would imply from (a.21) that the agent could have done better by defecting from r^{j*}_t when others are using r^{-j*}_t in the original equilibrium, and this would violate the assumption that r^*_t was an equilibrium reporting profile in the original mechanism.

Suppose, however, that

$$P^{j,X}(\tilde{I}^*_t, \tilde{r}^{-j*}_t) \subset P^{j,X}(M^*_t, r^{-j*}_t), \text{ for some } j \in N \tag{a.22}$$

Then notice that by presenting an noisy revelation mechanism, the principal can guarantee the same ex ante probability of period allocations as in the original equilibrium σ^*, while yet providing the agent with enhanced opportunities to transmit information, i.e.,

$$E_{P^{j,X}(\tilde{I}_t^*,.)}[u_t^X(k_t,\theta) + V_{t+1}^X(\tilde{I}^{t+1}, \tilde{r}^{t+1}, p_{t+1}^X, .)|q_t^*(k_t)] \tag{a.23}$$
$$\geq$$
$$E_{P^{j,X}(M^*(t), .)}[u_t^X(k_t,\theta) + V_{t+1}^X(I^{*t+1}, r^{*t+1}, p_t^X)|q_t^*(k_t)].$$

In (a.23) E_y implies that expectations are taken over the support y. Hence, in a Noisy Revelation mechanism, for the same probability of period allocations, q_t^*, the principal's future payoffs are defined over a larger support of possible information transfers, and hence he must be no worse off with such a mechanism than with the original mechanism. But given the assumption that $\Delta(\Theta) \in \Lambda$, the principal would lose nothing by choosing $\Delta(\Theta)$ as the message space in the original mechanism. Thus, (a.22) could not hold for any SHE σ^*, and we conclude that $\tilde{r}_t^*$ in (a.14) is an ICNR strategy for the mechanism $\tilde{I}^*(t)$ for all $\theta^i \varepsilon \Theta^i$ $i \varepsilon N$.

Finally, notice that we cannot have the case that, for any $j \varepsilon N$, $P^{j,X}(\tilde{I}_t^*, .) = P^{j,X}(M^*, .)$ but $P^{j,i}(\tilde{I}_t^*, .) \subset P^{j,i}(M^*, .)$, for some $i \varepsilon N$? The reason is that $P^{j,\cdot}$ is defined for any arbitrary priors, and $P_t^{j,X}$ and $P_t^{j,i}$ differ only because of the difference in parameters p_t^X and p_t^i.

Q.E.D.

NOTES

1. In various contexts Baron and Besanko, Laffont and Tirole and Roberts have shown that no truthful revelation equilibrium can be sustained, at least in the early periods, without precommitment. It has been shown by Myerson (1986) that when there is precommitment and agents' messages to the principal are private communications, the Revelation Principle generalizes to sequential contexts.

2. The finiteness assumption is made only to avoid some technical (measure theoretic) issues that have no bearing on the objective of the analysis. For an example of sequential incentive mechanism design with a continuum of types, see Baron and Besanko, and Laffont and Tirole.

3. An interesting extension may be to allow agents to have incomplete information regarding the principal, introducing "reputation" considerations for X.

4. More generally, K_t may also depend on the beliefs of the principal about the agents' private information. This extension can be straightforwardly accommodated in what follows.

5. Note that according to this procedure, it is entirely possible for X to receive different, and even conflicting, reports from the agents over time. But it may be consistent for X to ignore such discrepancies, and since ex ante threats based on possibly inconsistent strategies may not be credible, agents may in fact have a reporting strategy that over time includes discrepancies and conflicts.

6. Notice that the random allocation rule at t, q_t, is conditional on the principal's information set h_t^X due to this assumption. If the principal's information set is not commonly observable then q_t can only be conditional on m_t.

7. Note that the choice of message space has been modelled as a "pure" strategy: The principal first establishes the framework for communications (or the "language") through the choice of M, and then uses the random allocation rule.

8. It should be noted that in his analysis of a multistage incentive design problem with precommitment, Myerson (1986) is able to use the "coordinator" as a correlation device, and utilize a sequential correlated equilibrium solution concept. This cannot be done for our model since the principal's communications to the agents (for correlating their strategies), is vitiated by his own incentives in the game.

9. This is because if it reveals itself in the first period, it receives a net payoff of zero for sure in the second period, whereas he receives a positive expected payoffs if he has to be induced to reveal himself in the second period.

REFERENCES

[1] Baron, D. and Besanko, D. A theory of regulation and information in a continuing relationship, Research Paper No. 725(GSB), Stanford University, 1984.

[2] Dasgupta, P., Hammond, P. and Maskin, E. The implementation of social choice rules, Rev. of Econ. Studies, 46(1979), 185-216.

[3] Freixas, X., Guesnerie, R. and Tirole, J. Planning under incomplete information and the ratchet effect, IMSSS Working Paper No. 408, 1983, Stanford University.

[4] Grossman, S. and Perry, M. Sequential bargaining under asymmetric information, mimeo, 1984.

[5] Harsayani, J. Games with incomplete information played by bayesian players, Management Science, 14(1967-1968), 159-182, 320-334, 486-502.

[6] Kreps, D. and Wilson, R. Sequential equilibrium, Econometrica, 50(1982), 863-894.

[7] Kumar, P. Essays in the theory of intertemporal incentives and signalling, Unpublished Ph.D. Dissertation, 1985, Stanford University.

[8] ____________. Sequential incentive mechanism design and the incomplete revelation principle, mimeo, Carnegie-Mellon University, 1985.

[9] ____________. The sequential general principal-agent problem and the generalized incomplete revelation principle, mimeo, Carnegie-Mellon University, 1985.

[10] Laffont, J.J., and Maskin, E. The theory of incentives: an overview, in W. Hildenbrand(ed.), Advances in Economic Theory:Invited Papers for the IVth World Congress of the Econometric Society, 31-94, Cambridge University Press:Cambridge(1982).

[11] Laffont, J.J., and Tirole, J. Using cost observations to regulate firms, mimeo, MIT, 1985.

[12] Myerson, R. Incentive compatibility in a bargaining problem, Econometrica, 47(1979), 61-73.

[13] ____________. Optimal auction design, Math. of Operations Research, 6(1982), 58-73.

[14] ____________. Optimal coordination mechanisms in generalized principal agent problems, Journal Of Mathematical Economics, 10(1982), 67-81.

[15] ____________. Sequential correlated equilibria of multistage games, Discussion Paper No. 590, Northwestern University, 1984.

[16] Ramey, G. A note on structural consistency in sequential equilibrium, Mimeo, Stanford University, 1985.

[17] __________. Limit pricing equilibria, Mimeo, Stanford University.

[18] Roberts, K. The theoretical limits to redistribution, Rev. of Econ. Studies, 51(1984), 177-196.

[19] Selten,R. Reexamination of the perfectness concept for equilibrium points in extensive games, International Journal of Game Theory, 4(1975), 25-55.

[20] Townsend, R. Private information, volatility, and economic welfare, International Conference on Real Business Cycles, Lisbon, Portugal, 1986.