

Gossner, Olivier; Kalai, Ehud; Weber, Robert

Working Paper

Information independence and common knowledge

Discussion Paper, No. 1453

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Gossner, Olivier; Kalai, Ehud; Weber, Robert (2007) : Information independence and common knowledge, Discussion Paper, No. 1453, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/31185>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

INFORMATION INDEPENDENCE AND COMMON KNOWLEDGE

OLIVIER GOSSNER, EHUD KALAI, AND ROBERT WEBER

ABSTRACT. Conditions of information independence are important in information economics and game theory. We present notions of partial independence in Bayesian environments, and study their relationships to notions of common knowledge.

KEYWORDS: Bayesian games, independent types, common knowledge

1. INTRODUCTION

This note deals with the private information available to different economic agents in a Bayesian environment, as described, for example, by the players' types in the terminology of Harsanyi [6]. In particular, we are interested in a condition of subjective independence and related notions. This condition leads to strong consequences in several important models of Bayesian games as illustrated by the following examples.

First, in perfect-monitoring Bayesian repeated games (with arbitrary discount parameter), at every Bayesian equilibrium the play converges to the play of a Nash equilibrium of the repeated game in which the realized types are common knowledge. Second, in Bayesian repeated games there is a full folk theorem even if monitoring is imperfect. And

Kalai's research is partially supported by the National Science Foundation, Grant No. SES-0527656.

third, the equilibria of one-shot Bayesian games with many players are ex-post Nash and structurally robust. The Appendix provides more details on all three models.

Motivated by the strong implication of subjective independence, this note offers deeper understanding of this condition and related notions. In particular, we discuss a weaker condition, called “independence under common knowledge,” but show that in the case of three or more players the two conditions are equivalent. Continuing with the case of three or more players, we elaborate on the environments in which these conditions hold, and how they relate to common knowledge.

2. ILLUSTRATIVE EXAMPLES

The following examples help illustrate the concepts.

Example 1. Independence under common knowledge of weather.

In a two-person Bayesian environment, the weather ω may be either *sunny* or *rainy* and, conditional on the realized weather, player i 's mood μ_i is either *happy* or *depressed*. Each player privately learns both the state of the weather and his own mood (for example, “It is sunny and I am depressed”) and hence each player may be any one of four possible types.

Let $w(\cdot)$ be the probability distribution over the two forms of weather (with positive probability on each), and let $m_i(\cdot|\omega)$ be player i 's (marginal) distribution of moods given the weather. Assume that the likelihood of any pair of types (ω, μ_1) and (ω, μ_2) is $w(\omega)m_1(\mu_1|\omega)m_2(\mu_2|\omega)$ and that all these distributions are commonly known to the players.

In the example above the players' types are not independent. For example, knowledge of one's own type precludes two possible opponent types. But the weather *factors the types* in the sense that conditional on every form of weather the types are independent. Moreover, the weather is *common knowledge*. In situations like the above, where a common-knowledge variable factors the types, we say that types are *independent under common knowledge*.

Example 2. Independence under ozone levels.

The possible moods of three players (μ_1, μ_2, μ_3) are correlated in the following way: either each $\mu_i = -1$ or 1 , or each $\mu_i = 2$ or 4 . All possible $16 (= 2^3 + 2^3)$ triples are known to be equally likely, but whatever triple is realized, every player i is informed only of his own mood, μ_i .

This example illustrates a situation of *subjective independence*: when a player knows his own type, he assesses his opponents' types to be *independent* of each other. For example, conditional on $\mu_1 = 2$, μ_2 and μ_3 are independent of each other and each takes the value 2 or 4 with equal probabilities.

One possible explanation for the situation above is that there is a variable that the players cannot observe directly, say, the ozone level, which affects all their moods. When the ozone level is *low*, the moods are independently drawn to be 2 or 4 each, but when it is *high*, the moods are independently drawn to be -1 or 1 each.

Even though the players in the second example may not be aware of the existence of ozone (let alone its level), the ozone level factors the moods (i.e. conditional on it, the types are independent). Moreover, there are many common-knowledge variables, equivalent to the ozone level, that factor the moods.

For example, consider a general mood variable G that takes on the value *up* when all three moods are in the set $\{2, 4\}$ and the value *down* when all three moods are in the set $\{-1, 1\}$. Notice that G is common knowledge (under the usual assumption that the prior probability distribution is common knowledge). And, in a similar way to the weather example, G factors the types.

The existence of such a G is not particular to this example, but is a consequence of Theorem 1 below. For three or more players, whenever types are subjectively independent, there must be some common-knowledge variable (such as G above) that factors them.

Is G unique? It is, but only up to information equivalence. For example, the variable P which takes on the value *even* when all three moods are in the set $\{2, 4\}$ and the value *odd* when all three variables

are in the set $\{-1, 1\}$ is also common knowledge and under P , as under G , the types are independent.

But while G and P are technically different, they are informationally equivalent: knowledge of the value of one is equivalent to knowledge of the value of the other.

In general, a coarsest variable that factors the types is not unique, even up to equivalence (see Example 3). But a consequence of Theorem 3 below is that under subjective independence there is only one (up to equivalence) variable which satisfies these two conditions: it is common knowledge and it factors the types. Moreover, this variable is equivalent to the coarsest common-knowledge variable introduced by Aumann [1].

3. DEFINITIONS AND MAIN RESULTS

All random variables (or just variables) considered in the sequel are defined over a fixed finite probability space (Ω, p) . We assume without loss of generality that $p(\omega) > 0$ for every $\omega \in \Omega$.

In addition, there is a fixed finite set of n (≥ 2) players, $I = \{1, 2, \dots, n\}$, with a vector of random variables $T = (T_1, \dots, T_n)$, where each variable T_i represents the type (information) of player i . It is assumed that the vector T is publicly known. So when a state ω is realized, each player i is told that his type $t_i = T_i(\omega)$. Then, through his knowledge of the entire vector T , he can make further inferences about his opponents' types, about any inferences that the opponents may make about their opponents, etc.

The variables of T are *independent*, if for every vector of values $t = (t_1, \dots, t_n)$, $p(T = t) = \prod_i p(T_i = t_i)$.

The types are *independent conditional on a variable Z* if for all possible values t and z , $p(T = t|z) = \prod_i p(T_i = t_i|Z = z)$. When this is the case, we say that Z *factors* T and, summing over all possible values of Z , we may write:

$$p(T = t) = \sum_z \prod_i p(T_i = t_i|Z = z)p(Z = z).$$

A random variable Y *reveals* a random variable Z if there exists a function f such that $Z = f(Y)$. More precisely, f is a function from the range of Y to the range of Z , with $Z(\omega) = f(Y(\omega))$ for every $\omega \in \Omega$. (Equivalently, $Z(\omega) \neq Z(\omega')$ implies that $Y(\omega) \neq Y(\omega')$.) It is easy to see that the “ Y reveals Z ” relationship is a partial order on the set of random variables.

Two random variables, Y and Z , are (informationally) *equivalent*, $Y \approx Z$, if they reveal each other ($Z(\omega) = Z(\omega')$ if and only if $Y(\omega) = Y(\omega')$). Y *strictly reveals* Z if Y reveals Z and they are not equivalent.

As will be clear from the context of the statements that follow, when we discuss a variable Z , we are often concerned with its equivalence class, $[Z]$, rather than with the variable itself.

It is easy to see that under the equivalence above and the “ Y reveals Z ” relationship, the variables on the space form a (complete) lattice (there are only finitely many variables under the equivalence relationship). The maximal *all-revealing* variable is represented by the function

R , with range in the set of states Ω , defined by $R(\omega) = \omega$. The minimal *least-revealing* variable may be represented by any constant variable.

For a subgroup of players $J \subseteq I$ we say that a variable Z is J -*common-knowledge*, if Z is revealed by every T_j with $j \in J$, i.e., for every j , $Z = f_j(T_j)$ for some function f_j . For the special case where Z is I -common-knowledge, we simply say that it is *common knowledge*.¹

When a common-knowledge variable factors the types, we say that it is a common-knowledge factorization.

The following are three key notions of information (type) independence.

Full independence: the variables of T are independent.

Independence under common knowledge: there is a common-knowledge factorization of T .

Subjective independence: for every player i , the other player's types, T_{-i} ($= (T_j)_{j \neq i}$), are independent conditional on i 's type T_i (or T_i factors T_{-i}).

Since every constant variable is common knowledge, it is clear that full independence implies independence under common knowledge. Also, as follows directly from the definitions above, full independence implies subjective independence.

¹As we clarify below, this is equivalent to the standard definition of common knowledge described in Aumann [1]

The connection between independence under common knowledge and subjective independence is less straightforward. If the number of players is two, then subjective independence is vacuously true (whereas independence under common knowledge may fail). For three or more players, it is not surprising that independence under common knowledge implies subjective independence. But the implication turns out to go in both directions.

Theorem 1. *If the number of players $n \geq 3$, then the players' types are subjectively independent if and only if they are independent under common knowledge.*

The proof of Theorem 1 is postponed to Section 3.2. We first present some additional results including Theorems 2 and 3. These Theorems offer deeper understanding of the relationship between independence and common knowledge, and are used for the proof of Theorem 1.

3.1. Common Knowledge and Independence. Given the variables $T = (T_1, \dots, T_n)$, there are always common-knowledge variables (for example, every constant variable). Moreover, if Z is common knowledge and Z reveals Y , then Y is also common knowledge. So a common-knowledge variable remains common knowledge as information is removed.

But if one goes in the opposite direction, there is an essentially unique common-knowledge variable that is most informative in a strong sense.

A variable Y *reveals all common knowledge* if it reveals every common-knowledge variable Z , i.e., for every common-knowledge variable Z there is a function f_Z such that $Z = f_Z(Y)$. To construct a *common-knowledge* variable with this strong all-revealing property, we follow Aumann [1].

Fixing a group $J \subseteq I$ of players, let \sim be the equivalence relation defined by $\omega \sim \omega'$ if there exists a chain $\omega_0 = \omega, \omega_1, \dots, \omega_n = \omega'$ such that for every $k = 1, \dots, n$, there exists a player $j \in J$ who is the same type in states ω_{k-1} and ω_k : $T_j(\omega_{k-1}) = T_j(\omega_k)$. We define $A[J]$ by the equivalence classes for this relation: $A[J](\omega) = \{\omega', \omega' \sim \omega\}$. For notational simplicity we let $A = A[I]$.

Formally, as defined above, A is a random variable with a range in the set of all subsets of Ω . But it also describes the common-knowledge partition defined by Aumann [1].

The following lemma shows that A is common knowledge, reveals all common knowledge, and is essentially unique.

Lemma 1. *The variable A is common knowledge and it reveals all common knowledge. Moreover, if A' is another variable with these two properties, then $A' \approx A$.*

Proof. First we show that A is common knowledge: Under the construction of A , note that $T_i(\omega) = T_i(\omega')$ implies $A(\omega) = A(\omega')$. Thus, A is revealed by each T_i .

Now we show that A is fully revealing: Let Z be a common-knowledge variable. For any ω, ω' such that $A(\omega) = A(\omega')$, consider a chain $\omega_0 = \omega, \omega_1, \dots, \omega_n = \omega'$ such that for every $k = 1, \dots, n$, there exists i with $T_i(\omega_{k-1}) = T_i(\omega_k)$. For every k , $T_i(\omega_{k-1}) = T_i(\omega_k)$ implies $Z(\omega_{k-1}) = Z(\omega_k)$, so that $Z(\omega) = Z(\omega')$. Hence, A reveals Z .

Finally, if both A and A' are fully revealing, A reveals A' and A' reveals A . ■

The above lemma may be viewed as stating that there is a “strongest” common-knowledge variable under the partial order imposed by the “ Y reveals Z ” relationship. In general, this is not the case for the factorization property, as we discuss next.

It is true that the trivial all-revealing variable R factors the types and (being all-revealing) reveals every other type-factorization Z . But for factorization it is useful to go in the opposite direction. If Y and Z both factor the types and Y reveals Z , the dependencies of the types can be explained by Z , and the extra information revealed by Y is superfluous.

So one would like to find the factorization D that is minimal, i.e., any variable Y that is strictly revealed by D is not a factorization. However, the example below shows that, in general, minimal factorizations are not unique (even up to equivalence).

Example 3. Let Ω consist of all the pairs of integers (i, j) , with each taking the values 1, 2, or 3. With the exception of the pair $(1, 1)$,

which has probability zero, all other pairs are equally likely so that $p(i, j) = 1/8$. Two players are told that their types are the values of i and j , respectively.

Clearly, the types above are not independent, but consider the following two variables, S and H : $S(i, j) = \textit{weak}$ if $i = 1$ and $S(i, j) = \textit{strong}$ if $i = 2$ or 3 ; H is similarly defined through the second coordinate, j .

It is easy to see that both S and H are non-equivalent type factorizations and that any variable strictly revealed by either one of them is no longer a type factorization. Thus, we have two different minimal factorizations.

We will return to the issue of a unique minimal type factorization after the presentation of the next theorem. It shows that factoring types is a stronger condition than revealing common knowledge. The corollary that follows shows that if Z factors the types and is common knowledge, then it must have the desired property of being the unique minimal factorization discussed above.

Theorem 2. *If Z factors the types, then Z reveals all common knowledge.*

Proof. It suffices to show that Z reveals A . Assume that $\omega, \omega' \in \Omega$ with $Z(\omega) = Z(\omega')$ in order to show that $A(\omega) = A(\omega')$. Let $z = Z(\omega)$, $t_k = T_k(\omega)$, and $t'_k = T_k(\omega')$ for $k = 1, \dots, n$. Since $p(t_1, t_2, \dots, t_n, z) > 0$ and $p(t'_1, t'_2, \dots, t'_n, z) > 0$, the independence of $(T_i)_i$ conditional on

Z implies that for every $k = 1, \dots, n - 1$:

$$p(t_1, \dots, t_k, t'_{k+1}, \dots, t'_n) > 0$$

so that there exists $\omega_1, \dots, \omega_{n-1}$ such that $(t_1, \dots, t_k, t'_{k+1}, \dots, t'_n) = (T_1(\omega_k), \dots, T_n(\omega_k))$ for every k .

Hence, $A(\omega) = A(\omega')$ whenever $Z(\omega) = Z(\omega')$, which shows that Z reveals A . ■

The theorem above shows that independence under common knowledge can only occur by conditioning on A :

Corollary 1. *If Z is a common-knowledge factorization, then $Z \approx A$.*

Proof. By Theorem 2, Z reveals A . But being common knowledge, it is also revealed by A . ■

Remark. Combined with Theorem 1, the above corollary shows that the subjective independence condition is powerful. For example, consider an n -person game G with a vector of types T . Let Z be a common-knowledge factorization of T with $|Z|$ denoting the number of values in its range.

For any value z of Z , consider the game G_z in which the players' types are restricted to those compatible with z . Because Z is common knowledge, each G_z is a well defined game. Actually, G may be represented as a game in which first the realization z of Z is announced to all players, then the subgame G_z is played. Moreover, with Z being a

type factorization, each of the games G_z is a game with independent types.

The cardinality of the range of Z , $|Z|$, tells us how many such separate games must be considered to achieve the decomposition above.

The corollary above tells us that if the types are subjectively independent, there is only one way to get such common-knowledge factorization, namely, by conditioning on the all-revealing common-knowledge variable A . This is useful, because the identification of A is a relatively easy task that involves only intersections of sets (as in the construction of A above). Thus, the need to check multiplications, as required in studying independence, is eliminated.

In the ozone level example in the introduction, for instance, the general mood variable G is the only common-knowledge factorization; every other factorization must be equivalent to it or not be common knowledge.

As follows immediately from the definition, any variable which is common knowledge for all the players must be common knowledge for any subset of players. But as the next theorem shows, under subjective independence (and hence also when we have independence under common knowledge), the converse is also true: what is common knowledge to any subset of two or more players is common knowledge to all. It follows that any fact which is common knowledge to any group of players

must be common knowledge to any other group (wether overlapping or disjoint).

Theorem 3. No secret common knowledge. *Assume that the players' types are subjectively independent. Any variable Z is common knowledge to a group of players J if and only if it is common knowledge to a group of players K , where J and K are any two groups (overlapping or disjoint) with two or more players each.*

Proof. Let J be any group of at least two players, and i be any player. Since the family $(T_j)_{j \in J}$ is independent conditional on T_i , Theorem 2 implies that T_i reveals $A[J]$. Thus, $A[J \cup \{i\}] = A[J]$. Since this is true for any J , an induction argument shows that for any group J of at least two players, $A[J] = A$. ■

3.2. Proofs of Theorem 1. We start with the following:

Lemma 2. *If X_1 and (X_2, X_3) are independent conditional on X_4 , then X_1 and X_2 are independent conditional on (X_3, X_4) .*

Proof. First note that X_1 and X_3 are independent conditional on X_4 , so that $p(x_1|x_4) = p(x_1|x_3, x_4)$ when $p(x_3, x_4) > 0$. Now, for x_3, x_4 such that $p(x_3, x_4) > 0$, we have

$$\begin{aligned} p(x_1, x_2|x_3, x_4) &= \frac{p(x_1, x_2, x_3|x_4)}{p(x_3|x_4)} = \frac{p(x_1|x_4)p(x_2, x_3|x_4)}{p(x_3|x_4)} \\ &= p(x_1|x_3, x_4)p(x_2|x_3, x_4) \end{aligned}$$

■

Proof of Theorem 1. We start with the “if” part. Let i, j, k be three different players. Since T_k and $(T_l)_{l \neq k, i}$ are independent conditional on T_i , $p(t_k | (t_l)_{l \neq k, i}, t_i) = p(t_k | (t'_l)_{l \neq k, i}, t_i)$ if $p((t'_l)_{l \neq k, i}, t_i) p((t'_l)_{l \neq k, i}, t'_i) > 0$. Similarly, $p(t_k | (t_l)_{l \neq k, j}, t_j) = p(t_k | (t'_l)_{l \neq k, j}, t_j)$ if $p((t_l)_{l \neq k, j}, t_j) p((t'_l)_{l \neq k, j}, t'_j) > 0$. Thus for any a in the range of $A[\{i, j\}]$, $p(t_k | (t_l)_{l \neq k}, a) = p(t_k | (t'_l)_{l \neq k}, a)$ whenever $p((t_l)_{l \neq k}, a) p((t'_l)_{l \neq k}, a) > 0$. Hence, T_k and $(T_l)_{l \neq k}$ are independent conditional on $A[\{i, j\}]$. From Theorem 3, T_k and $(T_l)_{l \neq k}$ are independent conditional on A . Since this is true for any k , the family T is independent conditional on A .

For the “only if” part, assume that T_i and $(T_k)_{k \neq i}$ are independent conditional on common knowledge Z . For any $j \neq i$, applying lemma 2 with $X_1 = T_i$, $X_2 = (T_k)_{k \neq i, j}$, $X_3 = T_j$, and $X_4 = Z$ shows that T_i and $(T_k)_{k \neq i, j}$ are independent conditional on (T_j, Z) . This proves that T_i and $(T_k)_{k \neq i, j}$ are independent conditional on T_j since (T_j, Z) and T_j generate the same partition. Since this is true for every i, j , we have established subjective independence. ■

REFERENCES

1. R. J. Aumann, *Agreeing to disagree*, The Annals of Statistics **4** (1976), 1236–1239.
2. R. J. Aumann and L. S. Shapley, *Long-term competition—A game theoretic analysis*, Essays on game theory (N. Megiddo, ed.), Springer-Verlag, New York, 1994, pp. 1–15.
3. D. Fudenberg, D. K. Levine, and E. Maskin, *The folk theorem with imperfect public information*, Econometrica **62** (1994), no. 5, 997–1039.
4. D. Fudenberg and E. Maskin, *The folk theorem in repeated games with discounting or with incomplete information*, Econometrica **54** (1986), no. 3, 533–554.
5. O. Gossner and J. Hörner, *When is the individually rational payoff in a repeated game equal to the minmax payoff?*, DP 1440, CMS-EMS, 2006.

6. J.C. Harsanyi, *Games with incomplete information played by 'Bayesian' players, parts I, II, and III*, Management Science **14** (1967/68), 159–182, 320–334, and 486–502.
7. J. Hörner and W. Olszewski, *The folk theorem with private almost-perfect monitoring*, Econometrica **74** (2006), no. 6, 1499–1544.
8. E. Kalai, *Large robust games*, Econometrica **72** (2004), no. 6, 1631–1666.
9. E. Kalai and E. Lehrer, *Rational learning leads to Nash equilibrium*, Econometrica **61** (1993), no. 5, 1019–1945.

APPENDIX A. IMPLICATIONS OF SUBJECTIVE INDEPENDENCE

We discuss more in details several implications of subjective independence.

A.1. Learning in Bayesian repeated games. Prior to the start of an n -person Bayesian repeated game U , a vector T of n types is drawn according to a commonly known prior probability distribution. Each player i is informed only of his own realized type, t_i . With the realized vector of types t fixed, the players proceed to play repeatedly a finite stage game $G = (A = \times A_i, u = (u_i))$ in periods $j = 1, 2, \dots$. For every player i , A_i denotes his feasible set of actions, and $u_i(t_i, a)$ denotes his stage payoff as a function of his type, t_i , and the profile of actions of all the players, a . A player's objective is to maximize the expected value of the sum of his stage payoffs, discounted by a fixed positive parameter $\lambda < 1$.

A (behavioral) strategy of player i is a rule that prescribes a probability distribution over A_i at every stage. This distribution may depend on his own realized type t_i and on the history of past profiles of actions

chosen by all the players at all previous stages (i.e., perfect monitoring). A vector of individual strategies $\sigma = (\sigma_1, \dots, \sigma_n)$ is a Bayesian-Nash equilibrium of U if for every player i , σ_i is optimal (under the uncertainty about the realized opponents' types) relative to σ_{-i} .

As the game progresses, the players observe the actions chosen by their opponents. Can they learn over time to play optimally, as if they each know all the realized types of their opponents? The answer is positive, in the sense below, provided that T satisfies subjective independence.

For any vector of realized types t , consider the complete information repeated game C_t , in which every player knows from the beginning the entire vector of types. (Formally, think of the game above, but with prior probability 1 assigned to the vector of signals t .) A vector of strategies $f = (f_1, \dots, f_n)$ is a Nash equilibrium of C_t , if each f_i is optimal relative to f_{-i} .

For any Bayesian equilibrium σ of the Bayesian game U above, for every vector of realized types t , and for every period j , let $\sigma_{t,j}$ be the (random) strategy induced by σ and t on the game, starting at time j . Assuming subjectively independent types, Kalai and Lehrer [9] show that with probability one, as j becomes large, $\sigma_{t,j}$ approaches a Nash equilibrium of the complete information game C_t (see their Theorem 2.1 and follow up discussion). In other words, the players learn to play

optimally as if they know each other's realized types (and hence each other's realized induced strategies).

A.2. Tight folk theorems for repeated games with imperfect monitoring. The folk theorem shows that, if players in a repeated game are sufficiently patient, then every individually rational payoff vector can be sustained as a (perfect) equilibrium payoff. By definition, a payoff vector is individually rational when it provides each player with at least his min max payoff (in mixed strategies) of the one-shot game.

This result was first established for games with perfect monitoring ([2], [4]), in which actions chosen by the players are commonly observed after each stage of the repeated game. More recently, the folk theorem has been extended to several classes of repeated games with imperfect monitoring (see, e.g., [3], [7]), where each player gets to observe at each stage a (possibly partially informative) signal on the action profile chosen.

When do such results characterize the full set of equilibrium payoff vectors of the repeated game? With perfect monitoring, a simple argument shows that every equilibrium payoff vector must be individually rational.² This contrasts with games with imperfect monitoring, for which several examples have shown that a player cannot necessarily

²Each player can play, at every stage, a best response to the mixed action profile of the other players, and thus defend in the repeated game his min max payoff in mixed strategies of the one-shot game.

guarantee his min max payoff, and that not all equilibrium payoffs are individually rational.

In a recent paper, Gossner and Hörner [5] show that, for a given player i , if other players' signals are independent conditional on player i 's signal, then player i can guarantee his min max payoff in the repeated game.³

Our notion of subjective independence naturally applies to games with imperfect monitoring, where each player's signal plays the role of his type: subjective independence holds in this context whenever, for every player, conditional on this player's signal, all other players' signals are independent.

Under subjective independence of players' signals, every equilibrium payoffs of the repeated game is individually rational, and the folk theorem gives a full characterization of equilibrium payoffs of the repeated game.

A.3. Ex-post Nash stability in Bayesian games with many players. T is a vector of n types drawn according to a commonly known prior probability distribution. With private knowledge of their own realized types, the n players proceed to play just once a finite strategic game $G = (A = \times A_i, u = (u_i))$. Each A_i denotes the actions available to player i , and each $u_i(t, a)$ describes the payoff of player i

³Furthermore, this condition is also necessary in a precise sense: Consider a distribution of signals such that player i can guarantee his min max for any payoff function. Then, the other players' signals are independent conditional on some garbling of i 's signal.

under the vector of types t , when the profile of actions a is selected. We assume that the payoff functions u_i are anonymous in opponents: $u_i(t, a) = u_i(t', a')$ whenever $(t'_i, a'_i) = (t_i, a_i)$ and $(t'_{-i}, a'_{-i})_j$ can be obtained from $(t_{-i}, a_{-i})_j$ by permuting the indexes $j \neq i$.

Would the outcome of the game be ex-post stable? Or would the players have incentives to revise the choices they made (based solely on their own types) after they observe (partial or full) hindsight information about the realized types and actions of their opponents?

Assuming subjectively independent types and continuous utility functions, Kalai [8] shows that as the number of players increases, the outcome of the game becomes ex-post stable and fully information-proof: for any information revealed ex-post, the probability that some player will have significant incentives to revise his choice is negligible.

PARIS SCHOOL OF ECONOMICS, UMR CNRS-EHESS-ENPC-ENS 8545 AND
KSM-MEDS, NORTHWESTERN UNIVERSITY

E-mail address: ogossner@ens.fr

KSM-MEDS, NORTHWESTERN UNIVERSITY

E-mail address: kalai@kellogg.northwestern.edu

KSM-MEDS, NORTHWESTERN UNIVERSITY

E-mail address: rjweber@kellogg.northwestern.edu