

ON VERSHIK'S STANDARDNESS CRITERION AND TSIRELSON'S NOTION OF COSINESS

M. Émery and W. Schachermayer

Abstract. — Building on work done by A. Vershik some thirty years ago, the insight into different types of filtrations has recently seen important progress, due in particular to B. Tsirelson, and L. Dubins, J. Feldman, M. Smorodinsky, B. Tsirelson. Key concepts are the notions of standard filtrations (due to A. Vershik) and cosy filtrations (due to B. Tsirelson). We investigate the relation between these two concepts and try to provide a comprehensive and self-contained presentation of the topic.

Part of this work is expository, and consists in translating into a probabilist's language Vershik's necessary and sufficient condition for standardness, and his theorem on lacunary isomorphism. There are also original results: Theorem 2 proves that standardness is in fact equivalent to a certain variant of the notion of cosiness, which we call *I-cosiness*; an example borrowed from Vershik and Smorodinsky then shows that I-cosiness is strictly stronger than another variant, D-cosiness, used in earlier works. Another new result is a (negative) answer to a question of H. von Weizsäcker: the last section gives an example of a filtration $(\mathcal{F}_n)_{n \in -\mathbb{N}}$ and a σ -field \mathcal{B} such that \mathcal{F}_0 and \mathcal{B} are “almost independent”, but nevertheless

$$\bigcap_{n \in -\mathbb{N}} (\mathcal{F}_n \vee \mathcal{B}) \not\cong \left(\bigcap_{n \in -\mathbb{N}} \mathcal{F}_n \right) \vee \mathcal{B}.$$

Many thanks to A. M. Vershik for enjoyable and fruitful conversations about his theory of filtrations. We are also grateful to the Schrödinger Institut in Vienna, where part of this work was done during the Mini-Symposium on the Classification of Filtrations held in December 1998.

Introduction

The objects of this study are filtrations. We shall not be interested in their set-theoretical properties, but in their probabilistic ones: we shall only consider filtrations on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$, and the notions relevant for our analysis, e.g. that of independence, are not invariant under changes of the measure \mathbb{P} . (In full rigor, we should speak of filtered probability spaces rather than filtrations.) We refer to the next section for a precise definition of an isomorphism between two filtrations in the present context.

In the late sixties and early seventies, A. Vershik [14] initiated a classification of filtrations. Consider filtrations $(\mathcal{F}_n)_{n \in -\mathbb{N}}$, where time is a negative integer, such that $\bigcap_n \mathcal{F}_n$ is degenerate and each \mathcal{F}_n is generated by \mathcal{F}_{n-1} and a random variable independent of \mathcal{F}_{n-1} , uniformly distributed on $[0, 1]$. A typical example of this situation is, of course, the filtration generated by an i.i.d. sequence $(X_n)_{n \in -\mathbb{N}}$ of random variables with uniform law on $[0, 1]$; a natural (and innocent-looking) question is whether this example already covers all cases of filtrations verifying the above properties. One of Vershik's results is the following, highly non-trivial, fact:

these filtrations are not all isomorphic to each other. More precisely, calling *standard* a filtration generated by an independent sequence of uniformly distributed random variables, he exhibited non standard filtrations satisfying the above conditions, and obtained his *standardness criterion*, a necessary and sufficient condition for a filtration to be standard.

Written in the language of ergodic theory, these ideas did not find their way into the probabilistic culture until 25 years later, when they were used by L. Dubins, J. Feldman, M. Smorodinsky and B. Tsirelson [5] to show that standard filtrations are not stable under equivalent changes of probability. They deduced therefrom that Brownian filtrations are not stable either under equivalent changes of probability.

A further step in the search of filtration invariants was made by B. Tsirelson, who showed in [13] that a Walsh process is not immersible into a Brownian filtration, that is, into a filtration generated by some (finite- or infinite-dimensional) Brownian motion. The strategy of his proof involves introducing a new property, *cosiness*, possessed by all Brownian filtrations (but more general than mere “Brownianness”); he then shows that a Walsh process is not immersible into a cosy filtration.

This strategy, establishing non-cosiness to deduce non-Brownianness, has been adapted to other situations: J. Warren proves in [15] that the filtration generated by sticky Brownian motion is not cosy; the non-Brownian change of probability constructed on Wiener space by Dubins, Feldman, Smorodinsky and Tsirelson in [5] is shown in [3] to be non-cosy; a non-cosy change of time on Wiener space is constructed in [7]. In the latter two articles, [3] and [7], Tsirelson’s definition of cosiness is slightly modified (weakened, and adapted to discrete time); what is used there is the variant of cosiness which we call *D-cosiness* below.

These two tools, the standardness criterion on the one hand, and cosiness and its variants on the other hand, are very efficient means of establishing that some given filtration is not standard (or not Brownian). The present article aims at bridging the gap between them, by establishing that standardness is equivalent to yet another variant of cosiness (we call it *I-cosiness*).

We shall first copy the proof of Vershik’s criterion in the language of stochastic processes, and graft thereupon the equivalence between standardness and I-cosiness (Theorem 2 and Corollary 5).

Then we shall test the efficiency of this new criterion on one of Vershik’s non standard, hence also non I-cosy, examples; somewhat unexpectedly, this particular non I-cosy example turns out to be D-cosy (Proposition 9).

Last, we shall show in Proposition 10 that the same example answers negatively a question raised by H. von Weizsäcker: if a filtration $(\mathcal{F}_n)_{n \leq 0}$ and a σ -field \mathcal{B} are almost independent (this will be explained), does the germ σ -field $\bigcap_n (\mathcal{F}_n \vee \mathcal{B})$ always equal $(\bigcap_n \mathcal{F}_n) \vee \mathcal{B}$?

Notation and definitions

All probability spaces $(\Omega, \mathcal{A}, \mathbb{P})$ will be \mathbb{P} -complete; by a sub- σ -field of \mathcal{A} , we always mean an $(\mathcal{A}, \mathbb{P})$ -complete sub- σ -field of \mathcal{A} . For instance, $\sigma(X)$ denotes the σ -field generated by the r.v. X and the null events; and when we consider a product $(\Omega' \times \Omega'', \mathcal{A}' \otimes \mathcal{A}'', \mathbb{P}' \times \mathbb{P}'')$ of probability spaces, $\mathcal{A}' \otimes \mathcal{A}''$ is completed for $\mathbb{P}' \times \mathbb{P}''$. Also,

all filtered probability spaces $(\Omega, \mathcal{A}, \mathbb{P}, \mathcal{F})$ will satisfy the usual hypotheses of the general theory of processes: each \mathcal{F}_t is $(\mathcal{A}, \mathbb{P})$ -complete and, if the time-parameter is continuous, the filtration is right-continuous.

Recall that the σ -field \mathcal{A} is essentially separable (resp. essentially finite) if it is generated by countably (resp. finitely) many events (and the null events; this is implicit). Equivalently, \mathcal{A} is generated by some random variable (resp. simple random variable). This is tantamount to saying that $L^1(\Omega, \mathcal{A}, \mathbb{P})$ is a separable Banach space (resp. is finite-dimensional), and entails that every sub- σ -field of \mathcal{A} is also essentially separable (resp. essentially finite).

DEFINITION. — An embedding of a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ into another one $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$ is a mapping Ψ from $L^0(\Omega, \mathcal{A}, \mathbb{P})$ to $L^0(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$ that commutes with Borel operations on countably many r.v.'s:

$$\Psi(f(X_1, \dots, X_n, \dots)) = f(\Psi(X_1), \dots, \Psi(X_n), \dots) \quad \text{for every Borel } f : \mathbb{R}^{\mathbb{N}} \rightarrow \mathbb{R}$$

and preserves the probability laws:

$$\bar{\mathbb{P}}[\Psi(X) \in E] = \mathbb{P}[X \in E] \quad \text{for every Borel } E \subset \mathbb{R}.$$

An embedding is always injective and transfers not only random variables, but also sub- σ -fields, filtrations, processes, etc. It is called an *isomorphism* if it is surjective; it then has an inverse. An embedding Ψ of $(\Omega, \mathcal{A}, \mathbb{P})$ into $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$ is always an isomorphism between $(\Omega, \mathcal{A}, \mathbb{P})$ and $(\bar{\Omega}, \Psi(\mathcal{A}), \bar{\mathbb{P}})$.

DEFINITION. — Given two filtered probability spaces $(\Omega, \mathcal{A}, \mathbb{P}, \mathcal{F})$ and $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}}, \bar{\mathcal{F}})$, the filtrations \mathcal{F} and $\bar{\mathcal{F}}$ are isomorphic if there exists an isomorphism Ψ from $(\Omega, \mathcal{F}_\infty, \mathbb{P})$ to $(\bar{\Omega}, \bar{\mathcal{F}}_\infty, \bar{\mathbb{P}})$ such that $\Psi(\mathcal{F}) = \bar{\mathcal{F}}$.

It would be more rigorous to say that the filtered probability spaces (and not only the filtrations) are isomorphic; this precision is necessary when there may be an ambiguity on the probabilities \mathbb{P} and $\bar{\mathbb{P}}$. In the sequel we shall never change probabilities, so we shall allow ourselves this abuse of language.

DEFINITIONS. — Let \mathcal{F} and \mathcal{G} be two filtrations on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$.

The filtration \mathcal{F} is included in \mathcal{G} if $\mathcal{F}_t \subset \mathcal{G}_t$ for each time t .

The filtration \mathcal{F} is immersed in \mathcal{G} if every \mathcal{F} -martingale is a \mathcal{G} -martingale (this is stronger than mere inclusion).

The filtrations \mathcal{F} and \mathcal{G} are jointly immersed¹ if each of them is immersed in $\mathcal{F} \vee \mathcal{G}$ (the smallest filtration where \mathcal{F} and \mathcal{G} are included; it can be defined by $(\mathcal{F} \vee \mathcal{G})_t = \bigcap_{\varepsilon > 0} (\mathcal{F}_{t+\varepsilon} \vee \mathcal{G}_{t+\varepsilon})$ for each time t).

As with the previous definition, the role played by \mathbb{P} should be stressed: The fact that \mathcal{F} is immersed in \mathcal{G} is in general not stable under a change of probability.

Note also that, by a density argument, \mathcal{F} is immersed in \mathcal{G} if and only if every bounded \mathcal{F} -martingale is a (bounded) \mathcal{G} -martingale; and, by stopping, if and only if every local martingale for \mathcal{F} is a local martingale for \mathcal{G} .

1. Or, more precisely, *jointly immersed in* $\mathcal{F} \vee \mathcal{G}$. But we shall see in Lemma 4 b) that this holds as soon as there exists a filtration \mathcal{H} such that both \mathcal{F} and \mathcal{G} are immersed in \mathcal{H} .

The notion of immersion is fundamental in many aspects of stochastic calculus; for instance, it is hidden inside the definition of a Brownian motion for a filtration, or of the Markov property with respect to a filtration. It has been used by many authors, sometimes implicitly, without giving it a name, sometimes explicitly, under various names; see [3] for more details and for some references. (The work [1] by D. Aldous and M. Barlow should be added to those references.)

LEMMA 1. — *Let \mathcal{F} and \mathcal{G} be two independent filtrations (that is, \mathcal{F}_t and \mathcal{G}_t are independent for each t). Then \mathcal{F} and \mathcal{G} are jointly immersed in $\mathcal{F}\vee\mathcal{G}$.*

PROOF. — An independent enlarging of a filtration preserves its martingales. ■

LEMMA 2. — *Let \mathcal{F} , \mathcal{G} , \mathcal{H} and \mathcal{K} be four filtrations on the same sample space, such that \mathcal{F} is immersed in \mathcal{H} and \mathcal{G} in \mathcal{K} . If \mathcal{H} and \mathcal{K} are independent, $\mathcal{F}\vee\mathcal{G}$ is immersed in $\mathcal{H}\vee\mathcal{K}$.*

PROOF. — It suffices to show that the product FG of a bounded \mathcal{F} -martingale F and a bounded \mathcal{G} -martingale G is an $\mathcal{H}\vee\mathcal{K}$ -martingale. This is obtained by taking any bounded, \mathcal{H}_t -measurable (respectively \mathcal{K}_t -measurable) random variable H_t (respectively K_t) and writing

$$\mathbb{E}[F_\infty G_\infty H_t K_t] = \mathbb{E}[F_\infty H_t] \mathbb{E}[G_\infty K_t] = \mathbb{E}[F_t H_t] \mathbb{E}[G_t K_t] = \mathbb{E}[F_t G_t H_t K_t]. \quad \blacksquare$$

REMARK. — If \mathcal{F} and \mathcal{G} are immersed in \mathcal{H} , it is not always true that $\mathcal{F}\vee\mathcal{G}$ is immersed in \mathcal{H} , even when \mathcal{F} and \mathcal{G} are independent. A very simple counter-example can be built from two independent random variables U and V with uniform law on $\{-1, 1\}$. Put $M_t = U \mathbb{1}_{\{t \geq 1\}}$ and $N_t = V \mathbb{1}_{\{t \geq 1\}}$. The filtrations \mathcal{F} and \mathcal{G} respectively generated by M and N are independent and immersed (by Lemma 1) in the filtration \mathcal{H} given by $\mathcal{H}_t = \sigma(UV)$ if $t < 1$ and $\mathcal{H}_t = \sigma(U, V)$ if $t \geq 1$; but the process $MN = UV \mathbb{1}_{\llbracket 1, \infty \rrbracket}$ is not an \mathcal{H} -martingale, though it is of course an $\mathcal{F}\vee\mathcal{G}$ -martingale.

In other words, if \mathcal{F} and \mathcal{G} are two independent filtrations, the product of an \mathcal{F} -martingale and a \mathcal{G} -martingale is always an $\mathcal{F}\vee\mathcal{G}$ -martingale; but if M and N are two independent martingales for a filtration \mathcal{H} , the product MN is not necessarily an \mathcal{H} -martingale. (Note, however, that the product of two independent *continuous* \mathcal{H} -martingales is always an \mathcal{H} -martingale, for in this case $\langle M, N \rangle = 0$.)

A sufficient condition for the product of two martingales to be a local martingale is that their covariation process is constant; this suggests the following statement: *Let two filtrations \mathcal{F} and \mathcal{G} be immersed in \mathcal{H} . Suppose there exists an \mathcal{H} -optional subset A of $\mathbb{R}_+ \times \Omega$ such that, for all \mathcal{F} -martingales M and \mathcal{G} -martingales N , the processes $\int \mathbb{1}_A d[M, M]$ and $\int \mathbb{1}_{A^c} d[N, N]$ are constant. Then the filtration $\mathcal{F}\vee\mathcal{G}$ is immersed in \mathcal{H} .* The simple proof of this statement is left to the reader; we shall only need the particular instance when A is the deterministic interval $\llbracket 0, t \rrbracket$:

LEMMA 3. — *Let two filtrations \mathcal{F} and \mathcal{G} be immersed in some filtration \mathcal{H} . Suppose that for some time t , \mathcal{F}_∞ is included in \mathcal{H}_t and \mathcal{G}_t is degenerate. The filtration $\mathcal{F}\vee\mathcal{G}$ is immersed in \mathcal{H} .*

PROOF. — It suffices to show that the product FG of a bounded \mathcal{F} -martingale F and a bounded \mathcal{G} -martingale G is an \mathcal{H} -martingale. The martingale equality may be checked separately on the intervals $[0, t]$ and $[t, \infty)$, since they have t in common. On $[t, \infty)$, this equality holds because $F = F_t$ is constant and \mathcal{H}_t -measurable; on $[0, t]$, it holds because $G = \mathbb{E}[G_\infty]$ is constant and deterministic. ■

LEMMA 4. — Let \mathcal{F} , \mathcal{G} and \mathcal{H} be three filtrations on some sample space $(\Omega, \mathcal{A}, \mathbb{P})$.

a) If \mathcal{F} is included in \mathcal{G} and \mathcal{G} in \mathcal{H} , and if \mathcal{F} is immersed in \mathcal{H} , then \mathcal{F} is immersed in \mathcal{G} .

b) if \mathcal{F} and \mathcal{G} are immersed in \mathcal{H} , they are jointly immersed (in $\mathcal{F} \vee \mathcal{G}$).

REMARK. — Lemma 4 b) can be rewritten as follows: two filtrations \mathcal{F} and \mathcal{G} on the same sample space are jointly immersed if and only if there exists on that space a filtration \mathcal{H} such that \mathcal{F} and \mathcal{G} are immersed in \mathcal{H} . This equivalence explains a posteriori the choice of the name for this property.

PROOF. — a) Any \mathcal{F} -martingale is an \mathcal{H} -martingale adapted to \mathcal{G} , whence a \mathcal{G} -martingale.

b) If both \mathcal{F} and \mathcal{G} are immersed in \mathcal{H} , applying a) to \mathcal{F} , $\mathcal{F} \vee \mathcal{G}$ and \mathcal{H} shows that \mathcal{F} is immersed in $\mathcal{F} \vee \mathcal{G}$; similarly for \mathcal{G} . ■

If a filtration \mathcal{F} is included in a filtration \mathcal{G} , each of the following three statements is a necessary and sufficient condition for \mathcal{F} to be immersed in \mathcal{G} :

for each t , the σ -fields \mathcal{F}_∞ and \mathcal{G}_t are conditionally independent given \mathcal{F}_t ;

for each t , the operators of conditional expectation verify $\mathbb{E}^{\mathcal{F}_\infty} \mathbb{E}^{\mathcal{G}_t} = \mathbb{E}^{\mathcal{F}_t}$;

for each t , the operators of conditional expectation verify $\mathbb{E}^{\mathcal{G}_t} \mathbb{E}^{\mathcal{F}_\infty} = \mathbb{E}^{\mathcal{F}_t}$.

These three characterizations of immersion can be found in Exercise V.4.16.1° of Revuz-Yor [10]. We shall not use them directly, but in a disguised form: Lemma 5 will rephrase them in terms of \mathcal{F} -saturation.

DEFINITION. — Let $(\Omega, \mathcal{A}, \mathbb{P}, \mathcal{F})$ be a filtered probability space. A sub- σ -field \mathcal{B} of \mathcal{A} is \mathcal{F} -saturated if $\mathcal{B} \subset \mathcal{F}_\infty$ and if $\mathbb{P}[B|\mathcal{F}_t]$ ($= \mathbb{E}[\mathbb{1}_B|\mathcal{F}_t]$) is \mathcal{B} -measurable for each $B \in \mathcal{B}$ and each time t .

LEMMA 5. — Let $(\Omega, \mathcal{A}, \mathbb{P}, \mathcal{F})$ be a filtered probability space. The map $\mathcal{E} \mapsto \mathcal{E}_\infty$ is a bijection between all filtrations immersed in \mathcal{F} and all \mathcal{F} -saturated sub- σ -fields of \mathcal{A} . Its inverse is the map $\mathcal{B} \mapsto \mathcal{E}$ defined by $\mathcal{E}_t = \mathcal{B} \cap \mathcal{F}_t$.

Consequently, if \mathcal{E} is a filtration immersed in \mathcal{F} , then $\mathcal{E}_t = \mathcal{E}_\infty \cap \mathcal{F}_t$ for every t , and \mathcal{E} is characterized by its end σ -field \mathcal{E}_∞ .

PROOF. — Let \mathcal{E} be immersed in \mathcal{F} and set $\mathcal{B} = \mathcal{E}_\infty$. Pick any $B \in \mathcal{B}$ and consider the \mathcal{E} -martingale $M_t = \mathbb{P}[B|\mathcal{E}_t]$; it is also an \mathcal{F} -martingale, and $\mathbb{P}[B|\mathcal{F}_t]$ is equal to $\mathbb{P}[B|\mathcal{E}_t]$, whence \mathcal{B} -measurable; so \mathcal{B} is \mathcal{F} -saturated. This equality also shows that if $B \in \mathcal{B} \cap \mathcal{F}_t$, $\mathbb{P}[B|\mathcal{E}_t] = \mathbb{P}[B|\mathcal{F}_t] = \mathbb{1}_B$, whence $B \in \mathcal{E}_t$, and $\mathcal{E}_t = \mathcal{B} \cap \mathcal{F}_t$.

Conversely, starting with any \mathcal{F} -saturated sub- σ -field \mathcal{B} of \mathcal{A} , define a filtration \mathcal{E} included in \mathcal{F} by $\mathcal{E}_t = \mathcal{B} \cap \mathcal{F}_t$. Clearly, $\mathcal{E}_\infty \subset \mathcal{B}$. For each $X \in L^\infty(\mathcal{B})$, the \mathcal{F} -martingale $M_t = \mathbb{E}[X|\mathcal{F}_t]$ is adapted to the smaller filtration \mathcal{E} , so it is an \mathcal{E} -martingale. Consequently, noticing that $\mathbb{E}[X|\mathcal{E}_\infty] = M_\infty = \mathbb{E}[X|\mathcal{F}_\infty] = X$ (because $\mathcal{B} \subset \mathcal{F}_\infty$), one sees that \mathcal{B} is included in \mathcal{E}_∞ , whence $\mathcal{B} = \mathcal{E}_\infty$. So M is the most general bounded \mathcal{E} -martingale, and \mathcal{E} is immersed in \mathcal{F} . ■

It is obvious that the intersection of two \mathcal{F} -saturated σ -fields is \mathcal{F} -saturated too; Lemma 5 translates this into a statement on immersed filtrations:

LEMMA 6. — *If two filtrations \mathcal{F} and \mathcal{G} are jointly immersed, their intersection $\mathcal{F} \cap \mathcal{G}$ (the filtration consisting of the intersections $\mathcal{F}_t \cap \mathcal{G}_t$) is immersed in each of them.*

PROOF. — Put $\mathcal{H} = \mathcal{F} \vee \mathcal{G}$; by hypothesis, both \mathcal{F} and \mathcal{G} are immersed in \mathcal{H} . It suffices to show that $\mathcal{F} \cap \mathcal{G}$ is immersed in \mathcal{H} , and the result will follow by applying Lemma 4 a) to $\mathcal{F} \cap \mathcal{G}$, \mathcal{F} and \mathcal{H} . By Lemma 5, the σ -fields \mathcal{F}_∞ and \mathcal{G}_∞ are \mathcal{H} -saturated; consequently, so is also $\mathcal{F}_\infty \cap \mathcal{G}_\infty$. Still by Lemma 5, the filtration \mathcal{J} defined by $\mathcal{J}_t = \mathcal{F}_\infty \cap \mathcal{G}_\infty \cap \mathcal{H}_t$ is immersed in \mathcal{H} . Applying Lemma 5 again gives $\mathcal{F}_\infty \cap \mathcal{H}_t = \mathcal{F}_t$ and $\mathcal{G}_\infty \cap \mathcal{H}_t = \mathcal{G}_t$, wherefrom $\mathcal{J}_t = \mathcal{F}_t \cap \mathcal{G}_t$. ■

LEMMA 7. — *Let $(\Omega, \mathcal{A}, \mathbb{P}, \mathcal{F})$ be a filtered probability space and \mathcal{B} an \mathcal{F} -saturated sub- σ -field of \mathcal{A} . For given t , \mathcal{B} is independent of \mathcal{F}_t if and only if $\mathcal{B} \cap \mathcal{F}_t$ is degenerate.*

If this holds, and if \mathcal{C} is any \mathcal{F} -saturated σ -field included in \mathcal{F}_t , then $\mathcal{B} \vee \mathcal{C}$ is \mathcal{F} -saturated.

PROOF. — Supposing $\mathcal{B} \cap \mathcal{F}_t$ is degenerate, take $B \in \mathcal{B}$. The random variable $\mathbb{P}[B|\mathcal{F}_t]$ is measurable with respect to \mathcal{B} (by saturation) and to \mathcal{F}_t , hence a.s. constant. Consequently, B is independent of \mathcal{F}_t ; so \mathcal{B} and \mathcal{F}_t are independent. Conversely, if \mathcal{B} and \mathcal{F}_t are independent, $\mathcal{B} \cap \mathcal{F}_t$ is independent of itself, that is, degenerate.

The second part of the lemma is a corollary of Lemmas 3 and 5. ■

From now on, the discussion will be restricted to filtrations indexed by the time-axis $-\mathbb{N} = \{\dots, -2, -1, 0\}$: the instants of time are negative integers. (In fact, only a neighbourhood of $-\infty$ is interesting; at the cost of a few minor changes, everything extends to the case when the time-axis is \mathbb{Z} .) All statements seen so far on immersion and saturation are still valid, with naturally \mathcal{F}_∞ being replaced by \mathcal{F}_0 . In this situation (and more generally whenever time is discrete), there is a very simple and useful instance of immersion:

LEMMA 8. — *Let $(\Omega, \mathcal{A}, \mathbb{P}, \mathcal{F})$ be a filtered probability space and, for each $n \leq 0$, let \mathcal{C}_n be a sub- σ -field of \mathcal{F}_n , independent of \mathcal{F}_{n-1} . The filtration \mathcal{E} defined by $\mathcal{E}_n = \sigma(\mathcal{C}_m, m \leq n)$ is immersed in \mathcal{F} , and the σ -field $\sigma(\mathcal{C}_n, n \leq 0)$ is \mathcal{F} -saturated.*

PROOF. — Every bounded, \mathcal{E}_0 measurable r.v. E has the form $\phi(\dots, C_{-1}, C_0)$ where each C_n is a \mathcal{C}_n -measurable r.v. For fixed $n \leq 0$, the r.v.'s \dots, C_{n-1}, C_n are \mathcal{F}_n -measurable and (C_{n+1}, \dots, C_0) is independent of \mathcal{F}_n ; hence $\mathbb{E}[E|\mathcal{F}_n]$ equals $\int \phi(\dots, C_{n-1}, C_n, c_{n+1}, \dots, c_0) d\gamma_{n+1}(c_{n+1}) \dots d\gamma_0(c_0)$, where γ_m is the law of C_m . As this is \mathcal{E}_n -measurable, $\mathbb{E}[E|\mathcal{F}_n]$ equals $\mathbb{E}[E|\mathcal{E}_n]$, showing immersion of \mathcal{E} in \mathcal{F} and \mathcal{F} -saturation of \mathcal{E}_0 . ■

DEFINITION. — *Two filtrations $\mathcal{F} = (\mathcal{F}_n)_{n \in -\mathbb{N}}$ and $\mathcal{G} = (\mathcal{G}_n)_{n \in -\mathbb{N}}$, defined on the same sample space $(\Omega, \mathcal{A}, \mathbb{P})$, are I-separate if there exists an $n \in -\mathbb{N}$ such that the σ -fields \mathcal{F}_n and \mathcal{G}_n are independent.*

The letter I in this name stands for Independence. Later on, we shall meet other separation conditions (D-separation, H-separation, ...).

DEFINITION. — A filtered probability space $(\Omega, \mathcal{A}, \mathbb{P}, \mathcal{F})$, where $\mathcal{F} = (\mathcal{F}_n)_{n \in -\mathbb{N}}$, is I-cosy if for every \mathcal{F}_0 -measurable r.v. R and every $\delta > 0$, there exists a probability space $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$ with two filtrations \mathcal{F}' and \mathcal{F}'' such that

- (i) each of \mathcal{F}' and \mathcal{F}'' is isomorphic to \mathcal{F} ;
- (ii) \mathcal{F}' and \mathcal{F}'' are jointly immersed;
- (iii) \mathcal{F}' and \mathcal{F}'' are I-separate;
- (iv) the copies $R' \in L^0(\bar{\Omega}, \mathcal{F}'_0, \bar{\mathbb{P}})$ and $R'' \in L^0(\bar{\Omega}, \mathcal{F}''_0, \bar{\mathbb{P}})$ of R by the isomorphisms in (i) are δ -close in probability: $\bar{\mathbb{P}}[|R' - R''| \geq \delta] < \delta$.

When there is no ambiguity on the underlying space $(\Omega, \mathcal{A}, \mathbb{P})$, we shall simply say that the filtration \mathcal{F} is I-cosy.

If the filtration \mathcal{F} were indexed by \mathbb{Z} instead of $-\mathbb{N}$, the σ -fields \mathcal{F}_0 , \mathcal{F}'_0 and \mathcal{F}''_0 in the above definition should be replaced with \mathcal{F}_∞ , \mathcal{F}'_∞ and \mathcal{F}''_∞ .

The definition of I-cosiness is inspired from two sources. The first one is Tsirelson's definition of cosiness in [13]; it is the same as I-cosiness, save the separation condition (iii) (Tsirelson works in continuous time and assumes that all martingales are continuous; in this framework, his separation condition is the existence of a constant $\rho < 1$ such that, for each \mathcal{F}' -martingale M' and each \mathcal{F}'' -martingale M'' , $[M', M'']_t \leq \rho [M', M']_t^{1/2} [M'', M'']_t^{1/2}$). The other source is a proof of non-standardness by Smorodinsky [12], who implicitly uses I-cosiness, without giving it an explicit name.

Observe that I-cosiness is invariant by isomorphisms: two isomorphic filtrations are either both I-cosy, or both non I-cosy.

PROPOSITION 1. — A filtration immersed in an I-cosy filtration is itself I-cosy.

PROOF. — If $(\Omega, \mathcal{A}, \mathbb{P})$ is endowed with two filtrations \mathcal{F} and \mathcal{G} , if \mathcal{F} is immersed in \mathcal{G} and if Ψ is an embedding of $(\Omega, \mathcal{G}_0, \mathbb{P})$ into some probability space, then the filtration $\Psi(\mathcal{F})$ is immersed in $\Psi(\mathcal{G})$. The proposition follows immediately from this remark, the definition of cosiness and the transitivity of immersions. ■

DEFINITION. — Let \mathcal{F} and \mathcal{G} be two filtrations, not necessarily on the same probability space. The filtration \mathcal{F} is immersible into \mathcal{G} if there exists a filtration immersed in \mathcal{G} and isomorphic to \mathcal{F} .

Using the invariance of I-cosiness under isomorphisms, Proposition 1 immediately bootstraps into a stronger result:

COROLLARY 1. — A filtration immersible into an I-cosy filtration is itself I-cosy.

PROPOSITION 2. — I-cosiness is stable by taking subsequences: Let $\mathcal{F} = (\mathcal{F}_n)_{n \in -\mathbb{N}}$ be an I-cosy filtration and $\sigma : -\mathbb{N} \rightarrow -\mathbb{N}$ a strictly increasing map. The filtration $\mathcal{G} = (\mathcal{G}_n)_{n \in -\mathbb{N}}$ defined by $\mathcal{G}_n = \mathcal{F}_{\sigma(n)}$ is I-cosy too.

PROOF. — Let R be a \mathcal{G}_0 -measurable r.v.; it is also \mathcal{F}_0 -measurable, so for $\delta > 0$ we have two isomorphic copies \mathcal{F}' and \mathcal{F}'' of \mathcal{F} , jointly immersed, I-separate, and verifying condition (iv). Put $\mathcal{G}'_n = \mathcal{F}'_{\sigma(n)}$, $\mathcal{G}''_n = \mathcal{F}''_{\sigma(n)}$ and $\mathcal{H} = \mathcal{G}' \vee \mathcal{G}''$. Plainly, the filtrations \mathcal{G}' and \mathcal{G}'' are isomorphic to \mathcal{G} , immersed in \mathcal{H} , and I-separate. ■

LEMMA 9. — Let (E, ρ) be a separable metric space, and \mathcal{F} an I-cosy filtration. The property defining I-cosiness still holds when the random variables R are taken E -valued (with $\rho(R', R'')$ replacing $|R' - R''|$).

PROOF. — Approximating in probability R by a simple r.v., we may suppose that R takes only finitely many values x_1, \dots, x_p in E . It then suffices to apply the definition of I-cosiness to the $\{1, \dots, p\}$ -valued r.v. S defined by $R = x_S$ and to $\delta' = \delta \wedge \frac{1}{2} \inf_{i \neq j} \rho(x_i, x_j)$. ■

PROPOSITION 3. — If two independent filtrations \mathcal{F} and \mathcal{G} are I-cosy, so is also $\mathcal{F} \vee \mathcal{G}$.

PROOF. — Every $\mathcal{F}_0 \vee \mathcal{G}_0$ -measurable r.v. can be approximated in probability by finite sums of the form $F^1 G^1 + \dots + F^k G^k$, where F^i are bounded and \mathcal{F}_0 -measurable and G^i bounded and \mathcal{G}_0 -measurable; so it suffices to check the definition of I-cosiness for such r.v.'s. By I-cosiness of \mathcal{F} and by Lemma 9 with $E = \mathbb{R}^k$, there exists $(\Omega^1, \mathcal{A}^1, \mathbb{P}^1)$ with two filtrations \mathcal{F}' and \mathcal{F}'' isomorphic to \mathcal{F} , jointly immersed in $\mathcal{F}' \vee \mathcal{F}''$, I-separate, such that $\mathbb{P}^1[|F^{1'} - F^{1''}| + \dots + |F^{k'} - F^{k''}| \geq \delta] < \delta$. Similarly, there exists $(\Omega^2, \mathcal{A}^2, \mathbb{P}^2)$ with two filtrations \mathcal{G}' and \mathcal{G}'' isomorphic to \mathcal{G} , jointly immersed in $\mathcal{G}' \vee \mathcal{G}''$, I-separate, such that $\mathbb{P}^2[|G^{1'} - G^{1''}| + \dots + |G^{k'} - G^{k''}| \geq \delta] < \delta$.

By replacing Ω^1 and Ω^2 with the product $\Omega = \Omega^1 \times \Omega^2$, we may suppose that $(\Omega^1, \mathcal{A}^1, \mathbb{P}^1) = (\Omega^2, \mathcal{A}^2, \mathbb{P}^2)$ and that $\mathcal{H}^1 = \mathcal{F}' \vee \mathcal{F}''$ and $\mathcal{H}^2 = \mathcal{G}' \vee \mathcal{G}''$ are independent. So they are jointly immersed in $\mathcal{H} = \mathcal{H}^1 \vee \mathcal{H}^2$ by Lemma 1, and \mathcal{F}' , \mathcal{F}'' , \mathcal{G}' and \mathcal{G}'' are all immersed in \mathcal{H} by transitivity of immersions.

(i) The filtrations $\mathcal{F}' \vee \mathcal{G}'$ and $\mathcal{F}'' \vee \mathcal{G}''$ are isomorphic to $\mathcal{F} \vee \mathcal{G}$.

(ii) By Lemma 2, these two filtrations are also immersed in \mathcal{H} ; by Lemma 4 b), they are jointly immersed.

(iii) We shall now see that they are I-separate. \mathcal{F}'_m and \mathcal{F}''_m are independent for some m , \mathcal{G}'_n and \mathcal{G}''_n are independent for some n ; replacing m and n by $m \wedge n$ allows us to suppose $m = n$. As $\mathcal{F}'_n \vee \mathcal{F}''_n$ and $\mathcal{G}'_n \vee \mathcal{G}''_n$ are independent (because \mathcal{H}^1 and \mathcal{H}^2 are), the four σ -fields \mathcal{F}'_n , \mathcal{F}''_n , \mathcal{G}'_n and \mathcal{G}''_n are independent, and $\mathcal{F}'_n \vee \mathcal{G}'_n$ is independent of $\mathcal{F}''_n \vee \mathcal{G}''_n$.

(iv) Outside an event with probability less than 2δ , $R' = F^{1'} G^{1'} + \dots + F^{k'} G^{k'}$ and $R'' = F^{1''} G^{1''} + \dots + F^{k''} G^{k''}$ verify $|R' - R''| \leq 2kM\delta$ where M is a common bound for all the F^i and G^i . The proof is over. ■

DEFINITION. — A filtration $\mathcal{F} = (\mathcal{F}_n)_{n \in -\mathbb{N}}$ is of product type if there exists an independent sequence $(\mathcal{C}_n)_{n \in -\mathbb{N}}$ of sub- σ -fields of \mathcal{A} such that $\mathcal{F}_n = \sigma(\mathcal{C}_m, m \leq n)$ for each $n \leq 0$.

This definition is borrowed from Feldman [8], who defines a *filtered probability space* $(\Omega, \mathcal{A}, \mathbb{P}, \mathcal{F})$ to be of product type when it satisfies the above property, thus stressing the role of the measure \mathbb{P} . But, as we already do for isomorphisms and immersions, we shall simply speak of *filtrations* of product type, keeping in mind that this notion is not invariant under changes of measure.

Notice that if \mathcal{F} is of product type, the σ -fields \mathcal{C}_n in the preceding definition are in general not uniquely determined. Consider for instance the natural filtration of a process $(\varepsilon_n)_{n \leq 0}$ made of i.i.d. r.v.'s uniform on $\{-1, 1\}$. This filtration is of product type, with $\mathcal{C}_n = \sigma(\varepsilon_n)$; replacing \mathcal{C}_0 by $\sigma(\varepsilon_{-1}\varepsilon_0)$ yields another family of σ -fields with the same property.

Of course, if $(Y_n)_{n \in -\mathbb{N}}$ is an independent sequence of random variables, the filtration generated by the process Y is of product type. Conversely, every filtration \mathcal{F} of product type and such that \mathcal{F}_0 is essentially separable, is the natural filtration of such an independent process.

PROPOSITION 4. — *Every filtration of product type is I-cosy.*

PROOF. — Let \mathcal{F} be of product type: there exists an independent sequence $(\mathcal{C}_n)_{n \in -\mathbb{N}}$ of sub- σ -fields of \mathcal{A} such that $\mathcal{F}_n = \sigma(\dots, \mathcal{C}_{n-1}, \mathcal{C}_n)$. Fix $R \in L^0(\mathcal{F}_0)$ and $\delta > 0$. Remark that the σ -fields $\mathcal{B}_n = \sigma(\mathcal{C}_{n+1}, \mathcal{C}_{n+2}, \dots, \mathcal{C}_0)$ form a monotone sequence $(\mathcal{B}_n)_{n \in -\mathbb{N}}$ of sub- σ -fields of \mathcal{A} , with limit $\bigvee_n \mathcal{B}_n = \mathcal{F}_0$ when $n \rightarrow -\infty$. By Doob's direct martingale convergence theorem, there exist an $n < 0$ and a r.v. $S \in L^0(\mathcal{B}_n)$ such that S is δ -close to R in probability. Fix these n and S .

On a suitable sample space $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$, for instance the product $(\Omega \times \Omega, \mathcal{A} \otimes \mathcal{A}, \mathbb{P} \times \mathbb{P})$, there exist two independent sub- σ -fields \mathcal{A}^1 and \mathcal{A}^2 of $\bar{\mathcal{A}}$ such that both $(\bar{\Omega}, \mathcal{A}^1, \bar{\mathbb{P}})$ and $(\bar{\Omega}, \mathcal{A}^2, \bar{\mathbb{P}})$ are isomorphic to $(\Omega, \mathcal{A}, \mathbb{P})$; call Ψ^1 and Ψ^2 the isomorphisms: $\mathcal{A}^i = \Psi^i(\mathcal{A})$ for $i \in \{1, 2\}$, and put $\mathcal{C}_m^i = \Psi^i(\mathcal{C}_m)$. Define three filtrations \mathcal{F}' , \mathcal{F}'' and \mathcal{G} on $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$ as follows:

$$\begin{aligned} \text{for every } m \leq 0, & \quad \mathcal{F}'_m = \sigma(\dots, \mathcal{C}_{m-1}^1, \mathcal{C}_m^1); \\ \text{for every } m \leq 0, & \quad \mathcal{G}_m = \sigma(\dots, \mathcal{C}_{m-1}^1, \mathcal{C}_{m-1}^2, \mathcal{C}_m^1, \mathcal{C}_m^2); \\ \text{for } m \leq n, & \quad \mathcal{F}''_m = \sigma(\dots, \mathcal{C}_{m-1}^2, \mathcal{C}_m^2) \\ \text{and for } n < m \leq 0, & \quad \mathcal{F}''_m = \sigma(\dots, \mathcal{C}_{n-1}^2, \mathcal{C}_n^2, \mathcal{C}_{n+1}^1, \dots, \mathcal{C}_{m-1}^1, \mathcal{C}_m^1). \end{aligned}$$

To show that \mathcal{F} is I-cosy, we shall check that \mathcal{F}' and \mathcal{F}'' verify the four conditions in the definition of I-cosiness.

(i) The restriction Ψ' of Ψ^1 to \mathcal{F}_0 is an isomorphism from \mathcal{F} to \mathcal{F}' . An isomorphism Ψ'' between \mathcal{F} and \mathcal{F}'' is given by the following algorithm: If C_m are \mathcal{C}_m -measurable r.v.'s and if ϕ is a Borel function, put

$$\begin{aligned} \Psi''(\phi(\dots, C_{n-1}, C_n, C_{n+1}, \dots, C_0)) \\ = \phi(\dots, \Psi^2(C_{n-1}), \Psi^2(C_n), \Psi^1(C_{n+1}), \dots, \Psi^1(C_0)). \end{aligned}$$

(ii) \mathcal{F}' and \mathcal{F}'' are immersed in \mathcal{G} by Lemma 1, because \mathcal{G} is an independent enlargement of each of them.

(iii) The σ -fields \mathcal{F}'_n and \mathcal{F}''_n are independent because they are respectively included in \mathcal{A}^1 and \mathcal{A}^2 .

(iv) Put $R' = \Psi'(R)$, $R'' = \Psi''(R)$, $S' = \Psi'(S)$ and $S'' = \Psi''(S)$. By isomorphic transfer, R' and S' (respectively R'' and S'') are δ -close in probability. Owing to the definitions of Ψ' and Ψ'' , these isomorphisms have the same restriction to \mathcal{B}_n ; so $S' = \Psi'(S) = \Psi''(S) = S''$. Consequently, R' and R'' are 2δ -close in probability. ■

COROLLARY 2. — *Any filtration immersible into a filtration of product type is I-cosy.*

PROOF. — Straightforward from Proposition 4 and Corollary 1. ■

As we shall now see, the converse is also true. We shall get it as a straightforward consequence of Vershik's criterion, more precisely of condition 3 in Theorem 3.2 of Vershik [14] (called condition (vii) in our Theorem 2 below). The next three sections are devoted to this topic; the reader already familiar with Vershik's theory can skip these sections and jump directly to the (short and easy) proof of (vi) \Rightarrow (vii) in Theorem 2.

Vershik's standardness criterion: Preliminary notions

Vershik's work on filtrations is written in Rohlin's language [11], where the idea of conditioning with respect to a sub- σ -field is expressed by quotienting the probability space. Sticking to a vocabulary more familiar to probabilists (at least, to us), the next proposition recalls what happens when the factor space is diffuse (that is, all equivalence classes in the quotient are isomorphic to the Lebesgue space $[0, 1]$).

PROPOSITION 5 AND DEFINITION. — *Given a sample space $(\Omega, \mathcal{A}, \mathbb{P})$, let \mathcal{B} be an essentially separable sub- σ -field of \mathcal{A} and \mathcal{C} a sub- σ -field of \mathcal{B} . The following four conditions are equivalent:*

- (i) *there exists a \mathcal{B} -measurable random variable Y such that, for every \mathcal{C} -measurable random variable Z , $\mathbb{P}[Y = Z] = 0$;*
- (ii) *there exists a \mathcal{B} -measurable random variable, independent of \mathcal{C} and having a diffuse law;*
- (iii) *there exists a \mathcal{B} -measurable random variable X , independent of \mathcal{C} , with uniform law on $[0, 1]$, and such that $\mathcal{C} \vee \sigma(X) = \mathcal{B}$;*
- (iv) *every random variable V that verifies $\mathcal{C} \vee \sigma(V) = \mathcal{B}$, has a diffuse law.*

When these hypotheses and conditions are met, we shall say that \mathcal{B} is conditionally non-atomic given \mathcal{C} , and every X satisfying condition (iii) will be called a complement to \mathcal{C} in \mathcal{B} .

PROOF. — Trivially, (iii) \Rightarrow (ii).

(ii) \Rightarrow (i). Let Y be \mathcal{B} -measurable, independent of \mathcal{C} , and with diffuse law η . For any \mathcal{C} -measurable Z , calling ζ the law of Z and using the fact that η is diffuse, one has $\mathbb{P}[Y = Z] = \iint \mathbb{1}_{\{y=z\}} \eta(dy) \zeta(dz) = \int \eta(\{z\}) \zeta(dz) = 0$.

(i) \Rightarrow (iii). Let Y be a \mathcal{B} -measurable random variable such that $\mathbb{P}[Y = Z] = 0$ for every \mathcal{C} -measurable Z . The essentially separable σ -field \mathcal{B} is generated by some random variable B ; its sub- σ -field \mathcal{C} is essentially separable too and is generated by some C . Call γ the law of C , and $(\beta_c)_{c \in \mathbb{R}}$ a regular version of the law of B given C : β_c is a probability well-defined for γ -almost every c , depending measurably on c , and the joint law of (C, B) is $\gamma(dc) \beta_c(db)$.

Remark that Y has the form $y \circ B$ for some measurable function y ; consequently, for any \mathcal{C} -measurable Z , $\mathbb{P}[B = Z] \leq \mathbb{P}[Y = y \circ Z] = 0$ by the choice of Y . This implies that almost all probabilities β_c are diffuse; indeed, for $\varepsilon > 0$, call $a(c)$ the smallest (i.e. leftmost) atom of β_c with mass at least ε , or $+\infty$ if there is no such atom; as $\mathbb{P}[B = a \circ C] = 0$, for almost every c , β_c has no atom weighing at least ε .

For each c , call F_c the distribution function of β_c , given by $F_c(x) = \beta_c((-\infty, x))$; as β_c has no atoms, for every $t \in [0, 1]$, β_c gives mass t to the interval $F_c^{-1}([0, t])$. Define a \mathcal{B} -measurable, $[0, 1]$ -valued random variable X by $X = F_C \circ B$.

For every bounded, measurable function ϕ and every $t \in [0, 1]$, one has

$$\begin{aligned} \mathbb{E}[\phi \circ C \mathbb{1}_{\{X \leq t\}}] &= \mathbb{E}[\phi \circ C \mathbb{1}_{\{F_C \circ B \leq t\}}] = \iint \phi(c) \mathbb{1}_{\{F_c(b) \leq t\}} \gamma(dc) \beta_c(db) \\ &= t \int \phi(c) \gamma(dc) = t \mathbb{E}[\phi \circ C]; \end{aligned}$$

thus X is independent of C (of \mathcal{C}) and uniformly distributed on $[0, 1]$. Calling G_c the right-continuous inverse of F_c , defined on $[0, 1]$, one almost surely has $B = G_C \circ X$, so B is $\mathcal{C} \vee \sigma(X)$ -measurable and $\mathcal{B} = \mathcal{C} \vee \sigma(X)$.

(i) \Rightarrow (iv). Suppose Y verifies (i) and C generates \mathcal{C} ; let V be such that $\mathcal{C} \vee \sigma(V) = \mathcal{B}$. For some Borel f , one has $Y = f(C, V)$, and for each $v \in \mathbb{R}$, one can write $\mathbb{P}[V = v] \leq \mathbb{P}[Y = f(C, v)] = 0$.

(iv) \Rightarrow (i). Assuming (iv), let B be a r.v. generating \mathcal{B} and such that $B \neq 0$ a.s. If Z is any \mathcal{C} -measurable r.v., $V = B \mathbb{1}_{\{B \neq Z\}}$ is \mathcal{B} -measurable and one has $\{V = 0\} = \{B = Z\}$, whence $B = V + Z \mathbb{1}_{\{V=0\}}$; so $\mathcal{C} \vee \sigma(V) = \mathcal{B}$. By (iv), V must be diffuse, wherefrom $\mathbb{P}[B = Z] = \mathbb{P}[V = 0] = 0$, and B fulfills (i). \blacksquare

COROLLARY 3. — *Let \mathcal{B} , \mathcal{C} and \mathcal{D} be three sub- σ -fields in an essentially separable sample space $(\Omega, \mathcal{A}, \mathbb{P})$; assume \mathcal{C} is included in \mathcal{B} .*

a) *If $\mathcal{B} \vee \mathcal{D}$ is conditionally non-atomic given $\mathcal{C} \vee \mathcal{D}$, then \mathcal{B} is conditionally non-atomic given \mathcal{C} .*

b) *If \mathcal{B} and \mathcal{D} are independent, and if \mathcal{B} is conditionally non-atomic given \mathcal{C} , then $\mathcal{B} \vee \mathcal{D}$ is conditionally non-atomic given $\mathcal{C} \vee \mathcal{D}$; moreover, every complement of \mathcal{C} in \mathcal{B} is also a complement of $\mathcal{C} \vee \mathcal{D}$ in $\mathcal{B} \vee \mathcal{D}$.*

PROOF. — a) If $\mathcal{B} \vee \mathcal{D}$ is conditionally non-atomic given $\mathcal{C} \vee \mathcal{D}$, for any V such that $\mathcal{C} \vee \sigma(V) = \mathcal{B}$, one also has $\mathcal{C} \vee \mathcal{D} \vee \sigma(V) = \mathcal{B} \vee \mathcal{D}$, so V is diffuse by condition (iv), and \mathcal{B} is conditionally non-atomic given \mathcal{C} by the same condition.

b) If \mathcal{B} and \mathcal{D} are independent and if \mathcal{B} is conditionally non-atomic given \mathcal{C} , every complement X to \mathcal{C} in \mathcal{B} is independent of $\mathcal{C} \vee \mathcal{D}$ (because X is independent of \mathcal{C} and \mathcal{D} is independent of $\mathcal{C} \vee \sigma(X)$); as $\mathcal{C} \vee \sigma(X) \vee \mathcal{D} = \mathcal{B} \vee \mathcal{D}$, X is also a complement to $\mathcal{C} \vee \mathcal{D}$ in $\mathcal{B} \vee \mathcal{D}$. \blacksquare

REMARK. — At this stage, it may be useful to warn the reader against two pitfalls:

a) *If \mathcal{B} is conditionally non-atomic given \mathcal{C} and if Y is \mathcal{B} -measurable, with diffuse law, and independent of \mathcal{C} , there may exist no complement X to \mathcal{C} in \mathcal{B} such that $\sigma(X) \supset \sigma(Y)$.* Consider for instance three independent random variables C , Y and Z with uniform law on $[0, 1]$. Call Γ the event $\{C < \frac{1}{2}\}$, define \mathcal{C} as the σ -field $\sigma(C)$ and \mathcal{B} as the σ -field $\sigma(C, Y, Z \mathbb{1}_\Gamma)$. Since Y satisfies condition (ii) of Proposition 5, \mathcal{B} is conditionally non-atomic given \mathcal{C} . Let X be any \mathcal{B} -measurable r.v. independent of \mathcal{C} and such that $\sigma(X) \supset \sigma(Y)$. We shall show $\sigma(X) = \sigma(Y)$, thus preventing $\mathcal{C} \vee \sigma(X)$ to equal \mathcal{B} (for clearly, $Z \mathbb{1}_\Gamma$ is not $\mathcal{C} \vee \sigma(Y)$ -measurable). There are two Borel functions f and g such that $Y = f(X)$ and $X = g(C, Y, Z \mathbb{1}_\Gamma)$ a.s. Putting $h(c, y) = g(c, y, 0)$, one has $X = h(C, f(X))$ a.s. on Γ^c . Consequently, for almost every $c \geq \frac{1}{2}$, one has $x = h(c, f(x))$ for almost all x . Fix such a c and put

$k(y) = h(c, y)$; equality $x = k \circ f(x)$ holds for almost all x , whence $X = k \circ f(X)$ a.s., giving $X = k \circ Y$ a.s., and $\sigma(X) = \sigma(Y)$.

b) If \mathcal{B} is conditionally non-atomic given \mathcal{C} and if \mathcal{D} is such that $\mathcal{C} \vee \mathcal{D} = \mathcal{B}$, there may exist no \mathcal{D} -measurable complement to \mathcal{C} in \mathcal{B} . Take for instance Ω equal to the union of the three rectangles $[0, \frac{1}{2}] \times [0, 1]$, $[\frac{1}{2}, 1] \times [0, \frac{1}{2}]$ and $[1, \frac{3}{2}] \times [\frac{1}{2}, 1]$, endowed with the restriction of the Lebesgue measure. Call C the first coordinate (it is $[0, \frac{3}{2}]$ -valued) and D the second one (it is $[0, 1]$ -valued); define \mathcal{C} as $\sigma(C)$, \mathcal{D} as $\sigma(D)$ and \mathcal{B} as $\sigma(C, D)$. Since $|D - \frac{1}{2}|$ is independent of C , \mathcal{B} is conditionally non-atomic given \mathcal{C} . Let X be any \mathcal{D} -measurable random variable such that $\mathcal{C} \vee \sigma(X) = \mathcal{B}$. We shall show $\sigma(X) = \mathcal{D}$, thus preventing X to be independent of \mathcal{C} . There are two Borel functions f and g such that $X = f(D)$ and $D = g(C, X)$ a.s.. Consequently, $D = g(C, f(D))$ a.s. and for almost every $c < \frac{1}{2}$, one has $d = g(c, f(d))$ for almost all d in $[0, 1]$. Fix such a c and put $h(x) = g(c, x)$; equality $d = h \circ f(d)$ holds for almost all $d \in [0, 1]$, whence $D = h \circ f(D)$ a.s., giving $D = h \circ X$ a.s., and $\sigma(X) = \mathcal{D}$.

In counter-example a) above, Y had a diffuse law. If, on the opposite, Y is discrete, a complement X always exists, as can easily be seen:

COROLLARY 4. — *Let \mathcal{B} and \mathcal{C} be two sub- σ -fields of \mathcal{A} , with $\mathcal{C} \subset \mathcal{B}$ and \mathcal{B} conditionally non-atomic given \mathcal{C} .*

If Y is a \mathcal{B} -measurable random variable taking finitely or countably many values, \mathcal{B} is conditionally non-atomic given $\mathcal{C} \vee \sigma(Y)$.

If furthermore Y is independent of \mathcal{C} , there exists a complement X to \mathcal{C} in \mathcal{B} such that $\sigma(X) \supset \sigma(Y)$.

PROOF. — For every random variable V such that $\mathcal{C} \vee \sigma(Y) \vee \sigma(V) = \mathcal{B}$, the pair (Y, V) has a diffuse law by Proposition 5 (iv). Since Y is discrete, this implies $\mathbb{P}[V=v] = \sum_y \mathbb{P}[V=v, Y=y] = 0$, so V is diffuse, and \mathcal{B} is conditionally non-atomic given $\mathcal{C} \vee \sigma(Y)$ by the same condition (iv).

If Y is independent of \mathcal{C} , let X' denote any complement to $\mathcal{C} \vee \sigma(Y)$ in \mathcal{B} , and X a random variable generating $\sigma(Y, X')$; notice that $\sigma(X) \supset \sigma(Y)$. Since \mathcal{C} , Y and X' are independent and $\mathcal{C} \vee \sigma(Y) \vee \sigma(X') = \mathcal{B}$, X is a complement to \mathcal{C} in \mathcal{B} . ■

Vershik's standardness criterion: First level

Vershik's theory of standard filtrations is a two-storied building. The first floor, which we now shall enter, gives a necessary and sufficient condition for a filtration to be of product type, in terms of \mathcal{F} -saturation. Theorem 1 below translates into our language the equivalence between 1 and 2 in Vershik's Theorem 3.2 of [14]. (His condition 3 uses a more sophisticated tool; this is postponed to the next section.) Besides changing the language, another difference is that, in this section, we deal with the non-atomic case only, whereas his statement covers the atomic case as well. Indeed, for the sake of simplification, in [14] some statements are given in the atomic case only; the general case has recently been written in full details (and in english) by J. Feldman [8].

DEFINITION. — A filtration $(\mathcal{F}_n)_{n \leq 0}$ is standard non-atomic if it is generated by a process $(X_n)_{n \leq 0}$, where X_n are independent random variables uniformly distributed on $[0, 1]$.

The important point is that the r.v.'s X_n are independent, with diffuse laws; that these laws can be chosen uniform on $[0, 1]$ is irrelevant, but shows that all standard non-atomic filtrations are isomorphic to each other. Clearly, a filtration isomorphic to a standard non-atomic filtration is standard non-atomic too.

THEOREM 1 (Vershik [14]). — Let $\mathcal{F} = (\mathcal{F}_n)_{n \leq 0}$ be a filtration on some probability space $(\Omega, \mathcal{A}, \mathbb{P})$. Suppose that \mathcal{F}_0 is essentially separable and that, for each $n \leq 0$, \mathcal{F}_n is conditionally non-atomic given \mathcal{F}_{n-1} . The following four statements are equivalent:

- (i) \mathcal{F} is standard non-atomic;
- (ii) \mathcal{F} is of product type;
- (iii) the tail σ -field $\mathcal{F}_{-\infty}$ is degenerate; and, for every \mathcal{F}_0 -measurable random variable R and every $\delta > 0$, there exist an essentially finite, \mathcal{F} -saturated σ -field \mathcal{B} and a \mathcal{B} -measurable random variable S , such that $\mathbb{P}[|R-S| \geq \delta] < \delta$;
- (iv) For every $\delta > 0$, $n \leq 0$ and every simple, \mathcal{F}_n -measurable r.v. R , there exist $m < n$, an \mathcal{F} -saturated \mathcal{B} independent of \mathcal{F}_m verifying $\mathcal{F}_m \vee \mathcal{B} = \mathcal{F}_0$, and a $\mathcal{B} \cap \mathcal{F}_n$ -measurable r.v. S such that $\mathbb{P}[S \neq R] < \delta$.

The gist of the equivalence (i) \Leftrightarrow (iv) in this theorem can be better understood with the help of the following lemma:

LEMMA. — Assume the hypotheses of Theorem 1 and fix $m \leq 0$. If an \mathcal{F} -saturated σ -field \mathcal{B} is independent of \mathcal{F}_m and if $\mathcal{F}_m \vee \mathcal{B} = \mathcal{F}_0$, there exist X_{m+1}, \dots, X_0 such that each X_ℓ is a complement to $\mathcal{F}_{\ell-1}$ in \mathcal{F}_ℓ and $\sigma(X_{m+1}, \dots, X_0) = \mathcal{B}$.

We mention this lemma only for the light it sheds on the theorem; we shall neither use it nor prove it. In the second paragraph of (iv) \Rightarrow (iv') in the proof of the theorem, we shall need and establish a slightly stronger property. (That paragraph contains a parameter n ; when $n = 0$ it reduces to a proof of the lemma.)

PROOF OF THEOREM 1. — (i) \Rightarrow (ii) is trivial: calling \mathcal{C}_n the σ -field generated by X_n , one has $\mathcal{F}_n = \sigma(\dots, \mathcal{C}_{n-1}, \mathcal{C}_n)$.

(ii) \Rightarrow (iii). Hypothesis (ii) says that $\mathcal{F}_n = \sigma(\dots, \mathcal{C}_{n-1}, \mathcal{C}_n)$, where the \mathcal{C}_n are independent; the degeneracy of $\mathcal{F}_{-\infty}$ follows by Kolmogorov's zero-one law. As \mathcal{F}_0 is essentially separable, so is also each \mathcal{C}_n ; hence there exists for each n an increasing sequence $(\mathcal{C}_n^k)_{k \in \mathbb{N}}$ of essentially finite sub- σ -fields such that $\mathcal{C}_n = \bigvee_{j \geq 0} \mathcal{C}_n^j$. For each $j \geq 0$, put

$$\mathcal{B}^j = \mathcal{C}_{-j}^j \vee \mathcal{C}_{-j+1}^j \vee \dots \vee \mathcal{C}_0^j;$$

this is an \mathcal{F} -saturated σ -field by Lemmas 8 and 5. The σ -fields \mathcal{B}^j form an increasing sequence whose limit $\sigma(\mathcal{B}^j, j \geq 0)$ contains every \mathcal{C}_n ; consequently $\sigma(\mathcal{B}^j, j \geq 0) = \mathcal{F}_0$, and $\bigcup_{j \geq 0} L^0(\mathcal{B}^j)$ is dense in $L^0(\mathcal{F}_0)$.

The proof of (iii) \Rightarrow (iv) will be made more clear by breaking it into two smaller steps. We shall introduce a new condition (iii'), and establish (iii) \Rightarrow (iii') \Rightarrow (iv).

Here is this intermediate statement (the letters SC stand for ‘saturated complement’):

(iii') Call $\text{SC}(m, \mathcal{F})$ the set of all \mathcal{F} -saturated sub- σ -fields \mathcal{B} of \mathcal{F}_0 that are independent of \mathcal{F}_m and verify $\mathcal{F}_m \vee \mathcal{B} = \mathcal{F}_0$. For every $n \leq 0$, the union $\bigcup_{m < n} \bigcup_{\mathcal{B} \in \text{SC}(m, \mathcal{F})} L^0(\Omega, \mathcal{B}, \mathbb{P})$ is dense in $L^0(\Omega, \mathcal{F}_0, \mathbb{P})$.

(iii) \Rightarrow (iii'). This is a straightforward consequence of the following fact: *Under the hypotheses of the theorem, and assuming $\mathcal{F}_{-\infty}$ is degenerate, for every $n \leq 0$ and every essentially finite, \mathcal{F} -saturated \mathcal{B} , there exist an $m < n$ and a $\mathcal{C} \in \text{CS}(m, \mathcal{F})$ such that $\mathcal{C} \supset \mathcal{B}$.* To establish this claim, remark first that if B is an event such that $0 < \mathbb{P}[B] < 1$, the degeneracy hypothesis implies $B \notin \mathcal{F}_m$ for m small enough. Since \mathcal{B} contains only finitely many events (modulo negligibility), there is an m (fixed in the sequel) smaller than n , such that $B \notin \mathcal{F}_m$ for every $B \in \mathcal{B}$ verifying $0 < \mathbb{P}[B] < 1$. So $\mathcal{B} \cap \mathcal{F}_m$ is degenerate, and, by Lemma 7, \mathcal{B} is independent of \mathcal{F}_m .

By Lemma 5, the filtration defined by $\mathcal{D}_\ell = \mathcal{B} \cap \mathcal{F}_\ell$ is immersed in \mathcal{F} ; \mathcal{D}_ℓ is degenerate for $\ell \leq m$. Noticing that each \mathcal{D}_ℓ is essentially finite, Corollary 4 asserts that \mathcal{F}_ℓ is conditionally non-atomic given $\mathcal{F}_{\ell-1} \vee \mathcal{D}_\ell$; call Z_ℓ a complement to $\mathcal{F}_{\ell-1} \vee \mathcal{D}_\ell$ in \mathcal{F}_ℓ , and define a filtration \mathcal{E} by taking \mathcal{E}_ℓ degenerate if $\ell \leq m$ and $\mathcal{E}_\ell = \mathcal{D}_\ell \vee \sigma(Z_{m+1}, \dots, Z_\ell)$ if $\ell > m$. The claim will be shown with $\mathcal{C} = \mathcal{E}_0$.

First, $\mathcal{F}_m \vee \mathcal{E}_\ell \supset \mathcal{F}_\ell$ for each $\ell \geq m$: This inclusion is trivial for $\ell = m$, and if it holds for ℓ , then $\mathcal{F}_m \vee \mathcal{E}_{\ell+1} \supset \mathcal{F}_m \vee \mathcal{E}_\ell \vee \mathcal{D}_{\ell+1} \vee \sigma(Z_{\ell+1}) \supset \mathcal{F}_\ell \vee \mathcal{D}_{\ell+1} \vee \sigma(Z_{\ell+1}) = \mathcal{F}_{\ell+1}$. So $\mathcal{F}_m \vee \mathcal{C} = \mathcal{F}_m \vee \mathcal{E}_0 = \mathcal{F}_0$.

Second, $\mathcal{C} = \mathcal{E}_0$ is \mathcal{F} -saturated. By Lemma 5, it suffices to show that \mathcal{E} is immersed in \mathcal{F} . For every $\ell \leq 0$ and every bounded, \mathcal{E}_ℓ -measurable U , we have to show that $\mathbb{E}[U | \mathcal{F}_{\ell-1}]$ is $\mathcal{E}_{\ell-1}$ -measurable. When $\ell \leq m$, U is deterministic; so we may suppose $\ell > m$. Without loss of generality, we may also suppose that U is a product $WV_\ell D_\ell$, where $W \in L^\infty(\sigma(Z_{m+1}, \dots, Z_{\ell-1}))$, $V_\ell \in L^\infty(\sigma(Z_\ell))$ and $D_\ell \in L^\infty(\mathcal{D}_\ell)$. Taking W out of the conditional expectation, it remains to show that $\mathbb{E}[V_\ell D_\ell | \mathcal{F}_{\ell-1}]$ is $\mathcal{E}_{\ell-1}$ -measurable. We may replace V_ℓ by $\mathbb{E}[V_\ell | \mathcal{F}_{\ell-1} \vee \mathcal{D}_\ell]$; but V_ℓ is independent of $\mathcal{F}_{\ell-1} \vee \mathcal{D}_\ell$ by definition of Z_ℓ ; so $\mathbb{E}[V_\ell | \mathcal{F}_{\ell-1} \vee \mathcal{D}_\ell]$ is a constant, and we are left with $\mathbb{E}[D_\ell | \mathcal{F}_{\ell-1}]$. This is $\mathcal{D}_{\ell-1}$ -measurable because \mathcal{D} is immersed in \mathcal{F} .

Last, by Lemma 5, $\mathcal{C} \cap \mathcal{F}_m = \mathcal{E}_0 \cap \mathcal{F}_m = \mathcal{E}_m$ is degenerate, and \mathcal{C} is independent of \mathcal{F}_m by Lemma 7.

So \mathcal{C} is in $\text{SC}(m, \mathcal{F})$; as $\mathcal{C} = \mathcal{E}_0 \supset \mathcal{D}_0 = \mathcal{B}$, the proof of (iii) \Rightarrow (iii') is complete.

(iii') \Rightarrow (iv). Assuming (iii'), fix $n \leq 0$, $\delta > 0$ and R measurable for \mathcal{F}_n and F -valued, where F is a finite subset of \mathbb{R} ; without loss of generality, we shall take $F = \{1, \dots, p\}$. Put $\delta' = \delta/p$. Hypothesis (iii') provides us with an $m < n$, a $\mathcal{B} \in \text{SC}(m, \mathcal{F})$ and a \mathcal{B} -measurable r.v. T such that $\mathbb{P}[|T - R| \geq \delta'] < \delta'$; by replacing if necessary T with $1 \vee T \wedge p$, we may further suppose $|T - R| \leq p-1$. This implies $\mathbb{E}[|T - R|] \leq \delta' + (p-1)\mathbb{P}[|T - R| \geq \delta'] \leq \delta' + (p-1)\delta' = \delta$. By L^1 -contractivity of conditional expectations, $T' = \mathbb{E}[T | \mathcal{F}_n]$ is also δ -close to R in L^1 . Since T is \mathcal{B} -measurable and \mathcal{B} is \mathcal{F} -saturated, T' is $\mathcal{B} \cap \mathcal{F}_n$ -measurable.

For $x \in \mathbb{R}$, call $\psi(x)$ the point in F closest to x (take the smallest such point if there are two of them). Among all F -valued r.v.'s, $S = \psi \circ T'$ is closest to T' in L^1 , whence $\mathbb{E}[|S - T'|] \leq \mathbb{E}[|T' - R|] \leq \delta$, and $\mathbb{E}[|R - S|] \leq 2\delta$. Since R and S are

F -valued, $\mathbb{P}[R \neq S] = \mathbb{P}[|R-S| \geq 1] \leq \mathbb{E}[|R-S|] \leq 2\delta$; as S is $\mathcal{B} \cap \mathcal{F}_n$ -measurable, (iv) is established.

The proof of (iv) \Rightarrow (i) will also be sliced into two smaller steps, by introducing a new statement (iv') and establishing (iv) \Rightarrow (iv') \Rightarrow (i). This intermediate step is:

(iv') *Suppose given $n < 0$ and X_{n+1}, \dots, X_0 such that each X_ℓ is a complement to $\mathcal{F}_{\ell-1}$ in \mathcal{F}_ℓ . For every $R \in L^0(\mathcal{F}_0)$ and $\delta > 0$, there exist some $m < n$, some X_{m+1}, \dots, X_n with the same property (each X_ℓ is a complement to $\mathcal{F}_{\ell-1}$ in \mathcal{F}_ℓ) and some r.v. $S \in L^0(\sigma(X_{m+1}, \dots, X_0))$ verifying $\mathbb{P}[|R-S| \geq \delta] < \delta$.*

(iv) \Rightarrow (iv'). Take R, δ, n and X_{n+1}, \dots, X_0 as in (iv'); by the assumption on the X_ℓ , \mathcal{F}_0 is equal to $\mathcal{F}_n \vee \sigma(X_{n+1}, \dots, X_0)$. Writing \mathcal{F}_n as the limit of an increasing sequence of essentially finite sub- σ -fields, one can δ -approximate R by a r.v. of the form $\phi(T, X_{n+1}, \dots, X_0)$, where ϕ is Borel and T is \mathcal{F}_n -measurable and simple. Applying (iv) to T , we obtain an $m < n$, a $\mathcal{B} \in \text{SC}(m, \mathcal{F})$ and a $\mathcal{B} \cap \mathcal{F}_n$ -measurable S verifying $\mathbb{P}[S \neq T] < \delta$. This gives $\mathbb{P}[\phi(S, X_{n+1}, \dots, X_0) \neq \phi(T, X_{n+1}, \dots, X_0)] < \delta$, and R is 2δ -close in probability to $\phi(S, X_{n+1}, \dots, X_0)$.

The filtration $\mathcal{E}_\ell = \mathcal{B} \cap \mathcal{F}_\ell$ associated to \mathcal{B} by Lemma 5 has the following properties: \mathcal{E} is immersed in \mathcal{F} , S is \mathcal{E}_n -measurable, and $\mathcal{F}_m \vee \mathcal{E}_0 = \mathcal{F}_0$; moreover, by definition of \mathcal{B} , \mathcal{E} is independent of the σ -field \mathcal{F}_m . According to Lemma 3, the filtration $\mathcal{F}' = \mathcal{F}^m \vee \mathcal{E}$, equal to \mathcal{F} up to time m and to $\mathcal{F}_m \vee \mathcal{E}$ from m on, is immersed in \mathcal{F} . Its end σ -field is $\mathcal{F}_m \vee \mathcal{E}_0 = \mathcal{F}_0$, so Lemma 5 gives $\mathcal{F}' = \mathcal{F}$ and one has $\mathcal{F}_m \vee \mathcal{E}_\ell = \mathcal{F}_\ell$ for all $\ell \in [m, 0]$. For $m < \ell \leq 0$, \mathcal{E}_ℓ is conditionally non-atomic given $\mathcal{E}_{\ell-1}$, for this property is inherited from \mathcal{F}_ℓ and $\mathcal{F}_{\ell-1}$ by Corollary 3 a). For $m < \ell \leq n$, choose a complement X_ℓ to $\mathcal{E}_{\ell-1}$ in \mathcal{E}_ℓ . By Corollary 3 b), X_ℓ is also a complement to $\mathcal{F}_{\ell-1}$ in \mathcal{F}_ℓ . As \mathcal{E}_m is degenerate, $\sigma(X_{m+1}, \dots, X_n) = \mathcal{E}_n$, and S is $\sigma(X_{m+1}, \dots, X_n)$ -measurable; consequently, R is 2δ -close to some $\psi(X_{m+1}, \dots, X_n, X_{n+1}, \dots, X_0)$.

(iv') \Rightarrow (i). Choose any r.v. R generating \mathcal{F}_0 , and a sequence $(\delta_j)_{j \in \mathbb{N}}$ tending to 0. Starting for instance with $n = -1$ and an arbitrary complement X_0 to \mathcal{F}_{-1} in \mathcal{F}_0 , and using repeatedly (iv') for each δ_j in turn, construct a sequence $(X_\ell)_{\ell \leq 0}$, a strictly decreasing sequence $(n_j)_{j \in \mathbb{N}}$ in $-\mathbb{N}$, and random variables S_j , respectively $\sigma(X_{n_j+1}, \dots, X_0)$ -measurable and δ_j -close to R in probability. Being the limit in probability of S_j , R is $\sigma(X_\ell, \ell \leq 0)$ -measurable, whence $\sigma(X_\ell, \ell \leq 0) = \mathcal{F}_0$. The filtration generated by the process X is immersed in \mathcal{F} by Lemma 8, and its value at time 0 is \mathcal{F}_0 , so it is equal to \mathcal{F} by Lemma 5. \blacksquare

Vershik's standardness criterion: Second level

Given a filtration $(\mathcal{F}_n)_{n \leq 0}$ (with \mathcal{F}_0 essentially separable) and an \mathcal{F}_0 -measurable r.v. R , all the information on R available at time -1 is contained in the conditional law $\pi_{-1}R = \mathcal{L}[R|\mathcal{F}_{-1}]$ of R given \mathcal{F}_{-1} . Similarly, all the information available at time -2 on the values of R is carried by the conditional law $\mathcal{L}[R|\mathcal{F}_{-2}]$. But this does not encapture everything that can be said about R at time -2 : it may miss some possible prediction at time -2 of how the values of R will progressively be revealed in the future. Specifically, the conditional law $\pi_{-2}R = \mathcal{L}[\pi_{-1}R|\mathcal{F}_{-2}]$ may contain more information than $\mathcal{L}[R|\mathcal{F}_{-2}]$. For an example, take a triple (N, X_{-1}, X_0)

of independent and non-degenerate r.v.'s, such that N takes values -1 and 0 , and X_{-1} and X_0 have the same law; take $\mathcal{F}_{-2} = \sigma(N)$, $\mathcal{F}_{-1} = \sigma(N, X_{-1})$ and $\mathcal{F}_0 = \sigma(N, X_{-1}, X_0)$, and choose $R = X_N$. The r.v. $\mathcal{L}[R|\mathcal{F}_{-2}]$ only tells you that R is independent of \mathcal{F}_{-2} , with the same law as X_{-1} and X_0 ; whereas the r.v. $\pi_{-2}R$ generates a bigger σ -field: it further tells you that if $N = -1$, R will actually be known at time -1 , and if $N = 0$, R will still be completely unknown at time -1 .

A key idea in Vershik's theory is to repeat this operation by putting $\pi_{-3}R = \mathcal{L}[\pi_{-2}R|\mathcal{F}_{-3}]$, and so on; he uses the full sequence $(\pi_n R)_{n \in -\mathbb{N}}$ of such iterated conditional laws. These $\pi_n R$ will be rigorously defined before Lemma 13, that characterizes the information they contain. They play the central role in the second part of Vershik's criterion, which says that \mathcal{F} is standard non-atomic if and only if for each R the iterated prediction $\pi_n R$ becomes closer and closer to being deterministic when n tends to $-\infty$. This should be compared to the well-known, much easier fact, that $\mathcal{F}_{-\infty}$ is degenerate if and only if for each R the conditional law of R given \mathcal{F}_n becomes closer and closer to being deterministic when n tends to $-\infty$.

As the successive $\pi_n R$ do not live in the same space, Vershik introduces for each n a distance ρ_n on the corresponding space, these distances being related to each other in a precise way; only then can the asymptotic condition be rigorously stated (condition (vii) in Theorem 2 below). It is also possible to give an equivalent statement that does not involve the distances ρ_n , namely I-cosiness; as we shall see, equivalence between I-cosiness and Vershik's second-level condition is easily established. We feel that I-cosiness may prove handier in some instances, because the ρ_n no longer appear; but it is essentially the same thing. (Vershik also gives another restatement, condition 4 in his theorem 3.2, in terms of his "tower of measures". This is a space where all the $\pi_n R$ can be made to live together; but the ρ_n are still implicitly there, in the very definition of the tower. We shall not elaborate further on this topic.)

Instead of working with real random variables, we shall take them K -valued, where (K, ρ) is a non-empty compact metric space; this will make some iterations easier, because the space K' of all probability measures on K is also compact and metrizable (for the weak convergence). The set of all a.s. defined, K -valued random variables will be called $L(K)$ (or $L(\mathcal{A}, K)$ to specify the σ -field), and endowed with the distance $\mathbb{E}[\rho(R, S)]$. Given an essentially separable sub- σ -field \mathcal{C} in a sample space $(\Omega, \mathcal{A}, \mathbb{P})$, the conditional law $\mathcal{L}[R|\mathcal{C}]$ of an $R \in L(K)$ given \mathcal{C} is a random variable belonging to $L(\mathcal{C}, K')$; it is almost surely well defined, for instance by disintegrating the joint law of (C, R) where C is any r.v. generating \mathcal{C} .

On the set K' of all probabilities on K , the Kantorovich-Rubinshtein distance ρ' is defined as

$$\rho'(\mu, \nu) = \inf_{\substack{\lambda \text{ has margins} \\ \mu \text{ and } \nu}} \int_{K \times K} \rho(r, s) \lambda(dr, ds),$$

where the infimum is taken over all probabilities λ on the product $K \times K$ with margins μ and ν . References on this definition are given by Vershik [14] and Dubins, Feldman, Smorodinsky and Tsirelson [5]; see also the survey by Belili [4], or the recent book [9] by Rachev and Rüschemdorf.

LEMMA 10. — *With these notations, ρ' is indeed a distance on K' . Moreover, the topology generated by ρ' is the topology of weak convergence; in particular, (K', ρ') is compact.*

Let $g : K \rightarrow \mathbb{R}$ be a c -Lipschitz function, that is, $|g(r) - g(s)| \leq c \rho(r, s)$. The function g' defined on K' by $g'(\mu) = \int_K g(r) \mu(dr)$ is c -Lipschitz too: $|g'(\mu) - g'(\nu)| \leq c \rho'(\mu, \nu)$.

PROOF. — First, to establish the triangle inequality, assume $\rho'(\mu_1, \mu_2) < \alpha$ and $\rho'(\mu_2, \mu_3) < \beta$. There exist a probability λ' with marginals μ_1 and μ_2 , such that $\int \rho(r, s) \lambda'(dr, ds) < \alpha$, and a λ'' with marginals μ_2 and μ_3 , such that $\int \rho(s, t) \lambda''(ds, dt) < \beta$. Disintegrating λ' and λ'' gives probabilities ν'_s and ν''_s , defined for μ_2 -almost all s , such that $\lambda'(dr, ds) = \mu_2(ds) \nu'_s(dr)$ and $\lambda''(ds, dt) = \mu_2(ds) \nu''_s(dt)$. Putting $\lambda(dr, dt) = \int_{s \in K} \nu'_s(dr) \nu''_s(dt) \mu_2(ds)$, it is a child's play to verify that λ is a probability with marginals μ_1 and μ_3 such that $\int \rho(r, t) \lambda(dr, dt) < \alpha + \beta$.

We now verify that a sequence $(\mu_j)_{j \in \mathbb{N}}$ in K' converges weakly to a limit $\nu \in K'$ if and only if $\rho(\mu_j, \nu)$ tends to 0; taking a constant sequence $\mu_j = \mu$ will by the same token give the separation condition and show that ρ is a distance.

If μ_j converges weakly to ν , there exist, on a suitable sample space, random variables R_j and S with these laws and such that $\rho(R_j, S)$ tends in probability to 0. As $\rho(R_j, S)$ is bounded by the diameter of K , convergence also holds in L^1 and $\mathbb{E}[\rho(R_j, S)] \rightarrow 0$. But, calling λ_j the joint law of R_j and S , one has $\rho'(\mu_j, \nu) \leq \int \rho(r, s) \lambda_j(dr, ds) = \mathbb{E}[\rho(R_j, S)]$, wherefrom $\rho'(\mu_j, \nu) \rightarrow 0$.

Conversely, supposing $\rho'(\mu_j, \nu) \rightarrow 0$, we have to verify that $\mu_j(f) \rightarrow \nu(f)$ for every continuous function f on K . By compactness, f is uniformly continuous, and given any $\varepsilon > 0$ there exists $\delta > 0$ such that

$$|f(r) - f(s)| \leq \begin{cases} \varepsilon & \text{if } \rho(r, s) < \delta \\ 2 \sup |f| & \text{if } \rho(r, s) \geq \delta \end{cases} \leq \varepsilon + \frac{2 \sup |f|}{\delta} \rho(r, s).$$

Now choose a probability λ_j with marginals μ_j and ν , such that

$$\int \rho(r, s) \lambda_j(dr, ds) \leq \rho'(\mu_j, \nu) + \frac{\delta \varepsilon}{2 \sup |f|}$$

and write

$$\begin{aligned} \left| \int f d\mu_j - \int f d\nu \right| &= \left| \int f(r) \lambda_j(dr, ds) - \int f(s) \lambda_j(dr, ds) \right| \\ &\leq \int |f(r) - f(s)| \lambda_j(dr, ds) \\ &\leq \varepsilon + \frac{2 \sup |f|}{\delta} \int \rho(r, s) \lambda_j(dr, ds) \\ &\leq \varepsilon + \frac{2 \sup |f|}{\delta} \rho'(\mu_j, \nu) + \varepsilon. \end{aligned}$$

This implies $\limsup_j \left| \int f d\mu_j - \int f d\nu \right| \leq 2\varepsilon$, and $\int f d\mu_j$ tends to $\int f d\nu$.

Last, if g is a real, c -Lipschitz function on K , given μ and ν in K' a similar computation yields for every λ with marginals μ and ν

$$|g'(\mu) - g'(\nu)| \leq \int |g(r) - g(s)| \lambda(dr, ds) \leq c \int \rho(r, s) \lambda(dr, ds);$$

taking the infimum over all such λ yields $|g'(\mu) - g'(\nu)| \leq c \rho'(\mu, \nu)$. ■

The next two lemmas show how the definition of ρ' is tailor-made to transfer distance estimates all the way down or up a filtration. Going down the ladder is immediate:

LEMMA 11. — *Let R and S be two K -valued random variables. If \mathcal{C} is an essentially separable sub- σ -field, one has $\mathbb{E}[\rho'(\mathcal{L}[R|\mathcal{C}], \mathcal{L}[S|\mathcal{C}])] \leq \mathbb{E}[\rho(R, S)]$.*

PROOF. — For almost all ω the conditional law $\mathcal{L}[(R, S)|\mathcal{C}](\omega)$ is a probability on K with marginals $\mathcal{L}[R|\mathcal{C}](\omega)$ and $\mathcal{L}[S|\mathcal{C}](\omega)$. Calling it $\lambda(\omega)$ and inserting it into the definition of ρ' yields the almost sure inequality

$$\rho'(\mathcal{L}[R|\mathcal{C}](\omega), \mathcal{L}[S|\mathcal{C}](\omega)) \leq \int \rho(r, s) \lambda(\omega)(dr, ds) = \mathbb{E}[\rho(R, S)|\mathcal{C}](\omega),$$

whence $\mathbb{E}[\rho'(\mathcal{L}[R|\mathcal{C}], \mathcal{L}[S|\mathcal{C}])] \leq \mathbb{E}[\rho(R, S)]$. ■

Climbing up the ladder is a little more arduous; this is done in the next lemma. To make things technically easier, we shall deal with simple random variables, that is, random variables that take only finitely many values.

LEMMA 12. — *On some $(\Omega, \mathcal{A}, \mathbb{P})$, let \mathcal{B} and \mathcal{C} be sub- σ -fields such that $\mathcal{C} \subset \mathcal{B}$ and \mathcal{B} is conditionally non-atomic given \mathcal{C} ; suppose $R \in \mathcal{L}(\mathcal{B}, K)$ and $L \in \mathcal{L}(\mathcal{C}, K')$ are simple and the values (in K') taken by the random variable L are probability measures on K with finite support. There exists a r.v. $S \in \mathcal{L}(\mathcal{B}, K)$ verifying $\mathcal{L}[S|\mathcal{C}] = L$ and $\mathbb{E}[\rho(R, S)] = \mathbb{E}[\rho'(\mathcal{L}[R|\mathcal{C}], L)]$.*

PROOF. — There exists a finite set $F \subset K$ such that the values of R are points of F and the (finitely many) values taken by L are probabilities with supports in F . By weak compactness of the set of all probabilities on $K \times K$, the infimum in the definition of ρ' is reached for some λ ; writing this for $\mu = \mathcal{L}[R|\mathcal{C}](\omega)$ and $\nu = L(\omega)$ shows for almost every ω the existence of a probability $\lambda(\omega)$ on $F \times F$, verifying

$$\sum_{s \in F} \lambda(\omega, \cdot, s) = \mathcal{L}[R|\mathcal{C}](\omega) \quad \sum_{r \in F} \lambda(\omega, r, \cdot) = L(\omega)$$

and
$$\sum_{(r,s) \in F \times F} \rho(r, s) \lambda(\omega, r, s) = \rho'(\mathcal{L}[R|\mathcal{C}](\omega), L(\omega));$$

moreover, $\lambda(\omega)$ can be taken \mathcal{C} -measurable, for instance by the measurable section theorem, or by choosing λ as a continuous function of the vectors $\mathcal{L}[R|\mathcal{C}]$ and L in \mathbb{R}^F .

Put $\Sigma(\omega) = \sum_{s \in F} \lambda(\omega, R(\omega), s) = \sum_{r \in F} \mathbb{1}_{\{R=r\}} \mathbb{P}[R=r|\mathcal{C}]$ and deduce from

$$\mathbb{P}[\Sigma = 0 \text{ and } R = r|\mathcal{C}] = \mathbb{P}[R = r|\mathcal{C}] \mathbb{1}_{\{\mathbb{P}[R=r|\mathcal{C}]=0\}} = 0$$

that $\mathbb{P}[\Sigma = 0] = 0$ and $\Sigma > 0$ a.s. Call s_1, \dots, s_p the points of F and for $0 \leq j \leq p$ set

$$Q_j(\omega) = \frac{\lambda(\omega, R(\omega), s_1) + \dots + \lambda(\omega, R(\omega), s_j)}{\Sigma(\omega)};$$

these r.v.'s are $\mathcal{C} \vee \sigma(R)$ -measurable and verify $0 = Q_0 \leq Q_1 \leq \dots \leq Q_p = 1$. Now R is simple, so, by Corollary 4, \mathcal{B} is conditionally non-atomic given $\mathcal{C} \vee \sigma(R)$, and by Proposition 5 there exists a complement X to $\mathcal{C} \vee \sigma(R)$ in \mathcal{B} . The F -valued r.v. S defined by

$$S(\omega) = s_j \iff Q_{j-1}(\omega) < X(\omega) \leq Q_j(\omega)$$

verifies

$$\begin{aligned}\mathbb{P}[S = s | \mathcal{C} \vee \sigma(R)] &= \frac{\lambda(\omega, R(\omega), s)}{\Sigma(\omega)} ; \\ \mathbb{P}[R = r \text{ and } S = s | \mathcal{C} \vee \sigma(R)] &= \mathbb{1}_{\{R=r\}} \frac{\lambda(\omega, r, s)}{\mathbb{P}[R=r | \mathcal{C}]} ; \\ \mathbb{P}[R = r \text{ and } S = s | \mathcal{C}] &= \mathbb{P}[R = r | \mathcal{C}] \frac{\lambda(\omega, r, s)}{\mathbb{P}[R=r | \mathcal{C}]} = \lambda(\omega, r, s)\end{aligned}$$

(everything vanishes if the denominator is null). This shows that the conditional law of (R, S) given \mathcal{C} is λ . It implies on the one hand $\mathcal{L}[S | \mathcal{C}] = L$ and on the other hand

$$\mathbb{E}[\rho(R, S) | \mathcal{C}] = \sum_{(r,s) \in F \times F} \rho(r, s) \lambda(r, s) = \rho'(\mathcal{L}[R | \mathcal{C}], L) ,$$

whence $\mathbb{E}[\rho(R, S)] = \mathbb{E}[\rho'(\mathcal{L}[R | \mathcal{C}], L)]$. ■

Fixing a filtered space $(\Omega, \mathcal{A}, \mathbb{P}, (\mathcal{F}_n)_{n \leq 0})$ such that \mathcal{F}_0 is essentially separable, the (a.s. well defined) conditional law $\mathcal{L}[R | \mathcal{F}_{-1}]$ of an $R \in L(K)$ given \mathcal{F}_{-1} is a random variable belonging to $L(\mathcal{F}_{-1}, K')$. As mentioned at the beginning of this section, this will be iterated by considering the conditional law of this r.v. given \mathcal{F}_{-2} , and so on. To do so, we shall use the following notation: $(K_0, \rho_0) = (K, \rho)$, $\pi_0 R = R$; and for each $n \leq 0$, $(K_{n-1}, \rho_{n-1}) = (K_n', \rho_n')$ and $\pi_{n-1} R$ is the conditional law of $\pi_n R$ given \mathcal{F}_{n-1} . Notice that each (K_n, ρ_n) is a compact metric space, and the random variable $\pi_n R$ belongs to $L(\mathcal{F}_n, K_n)$.

What information is carried by the r.v. $\pi_n R$? It contains the conditional law of R given \mathcal{F}_n , but also (see the example at the beginning of this section) predictions about how these conditional laws may evolve from the present time n to the future time 0 (call this *second-order information*). And it also contains predictions about how this second-order information itself may evolve in the future (this is third order), and so on, up to order $|n|$ (no more information is added to $\pi_n R$ by iterating beyond order $|n|$). This procedure is an informal description from inside; the next lemma gives a (simpler) characterization from outside.

LEMMA 13. — Assume \mathcal{F}_0 is essentially separable and fix $R \in L(\mathcal{F}_0, K)$. The filtration generated by the process $(\pi_n R)_{n \leq 0}$ is the smallest filtration \mathcal{E} immersed in \mathcal{F} such that R is \mathcal{E}_0 -measurable. The σ -field $\sigma(\pi_n R, n \leq 0)$ is the smallest \mathcal{F} -saturated σ -field making R measurable.

PARTIAL PROOF. — We shall only show the immersion and saturation properties. Minimality is not needed in the sequel; its proof is left to the reader.

To see that the filtration generated by the process $(\pi_n R)_{n \leq 0}$ is immersed in \mathcal{F} , it suffices to verify that $\mathbb{E}[\phi(\pi_m R, m \leq 0) | \mathcal{F}_n]$ is $\sigma(\pi_m R, m \leq n)$ -measurable for each $n \leq 0$. Writing $\mathbb{E}[\phi(\dots, \pi_{-1} R, \pi_0 R) | \mathcal{F}_{-1}]$ as $\int \phi(\dots, \pi_{-1} R, r) d(\pi_{-1} R)(r)$ shows it for $n = -1$; the general case follows by iteration.

Consequently, by Lemma 5, the σ -field $\sigma(\pi_n R, n \leq 0)$ is \mathcal{F} -saturated. ■

The filtration \mathcal{F} is still kept fixed; in Lemmas 14 to 18 we assume that \mathcal{F}_0 is essentially separable, and in Lemmas 15 to 18, that each \mathcal{F}_n is conditionally non-atomic given \mathcal{F}_{n-1} .

LEMMA 14. — *If R and S are two K -valued, \mathcal{F}_0 -valued random variables,*

$$\mathbb{E}[\rho_n(\pi_n R, \pi_n S)] \leq \mathbb{E}[\rho(R, S)].$$

PROOF. — Immediate by iterating Lemma 11. ■

LEMMA 15. — *Let R be an \mathcal{F}_0 -measurable r.v. with values in a finite set $F \subset K$. For every $\delta > 0$, there exists an \mathcal{F}_0 -measurable, F -valued r.v. S such that $\pi_{-1}S$ is simple and $\mathbb{E}[\rho(R, S)] < \delta$.*

PROOF. — Call F' the set of all probabilities on F ; it is compact, so there exist a finite $G \subset F'$ and a Borel $f : F' \rightarrow G$ such that $\rho_{-1}(\mu, f(\mu)) < \delta$ for every $\mu \in F'$. Lemma 12 with $L = f(\pi_{-1}R)$ gives a K -valued S such that $\pi_{-1}S = L$ (hence S is F -valued) and that $\mathbb{E}[\rho(R, S)] = \mathbb{E}[\rho_{-1}(\pi_{-1}R, L)] = \mathbb{E}[\rho_{-1}(\pi_{-1}R, f(\pi_{-1}R))] < \delta$. ■

DEFINITION. — *The filtration $(\mathcal{F}_m)_{m \leq 0}$ and $n \leq 0$ being fixed, a r.v. $R \in L(\mathcal{F}_0, K)$ is n -simple if the $|n|+1$ r.v.'s $\pi_n R, \pi_{n+1} R, \dots, \pi_0 R = R$ are simple.*

The next lemma says that n -simplicity of R is in fact a property of $\pi_n R$.

NOTATION. — Set $K_0^S = K$ and for $n < 0$ call K_n^S the set of all probabilities carried by finitely many points of K_{n+1}^S . Observe that K_n^S is included in K_n , but K_n^S is not compact (unless $n = 0$ or $|K| = 1$).

LEMMA 16. — *Given $n \leq 0$, a random variable $R \in L(\mathcal{F}_0, K)$ is n -simple if and only if $\pi_n R$ is simple and takes its values in K_n^S .*

PROOF. — For $n = 0$, this is just the definition of a simple r.v. Assuming the lemma holds for some n , we shall prove it for $n - 1$.

If R is $(n-1)$ -simple, it is also n -simple, and, by induction hypothesis, $\pi_n R$ is simple and takes its values in K_n^S ; these values belong to some finite $F \subset K_n^S$. The set F' of all probabilities on F is included in K_{n-1}^S ; $\pi_{n-1}R$ has values in F' and a fortiori in K_{n-1}^S . It is also simple, by definition of $(n-1)$ -simplicity.

Conversely, if $\pi_{n-1}R$ is simple and takes its values in K_{n-1}^S , these values are finitely many probabilities finitely supported in K_n^S , so $\pi_n R$ is simple and K_n^S -valued, and R is n -simple (induction hypothesis). As $\pi_{n-1}R$ is simple too, R is $(n-1)$ -simple. ■

LEMMA 17. — *For some $n \leq 0$, let $R \in L(\mathcal{F}_0, K)$ and L be a simple, \mathcal{F}_n -measurable r.v. with values in K_n^S . There exists an n -simple $S \in L(\mathcal{F}_0, K)$ such that $\pi_n S = L$ and $\mathbb{E}[\rho(R, S)] = \mathbb{E}[\rho_n(\pi_n R, L)]$.*

PROOF. — If $n = 0$, take $S = L$. If $n < 0$, writing L_n instead of L and working by induction, it suffices to show that our hypothesis implies the existence of a simple, K_{n+1}^S -valued $L_{n+1} \in L(\mathcal{F}_{n+1}, K_{n+1})$ verifying $\mathcal{L}[L_{n+1} | \mathcal{F}_n] = L_n$ and $\mathbb{E}[\rho_{n+1}(\pi_{n+1}R, L_{n+1})] = \mathbb{E}[\rho_n(\pi_n R, L_n)]$. But this is just Lemma 12 with K_{n+1} instead of K , $\mathcal{B} = \mathcal{F}_{n+1}$ and $\mathcal{C} = \mathcal{F}_n$. ■

LEMMA 18. — For fixed $n \leq 0$, the set of all n -simple r.v.'s is dense in $L(\mathcal{F}_0, K)$.

PROOF. — For $n = 0$, this just recalls that simple r.v.'s are dense. To prove the lemma, it suffices to show that the set of all $(n-1)$ -simple r.v.'s is dense in the set of all n -simple ones. So let R be n -simple; $\pi_n R$ is simple and K_n^S -valued (Lemma 16), its values belong to a finite subset F of K_n^S . Applying Lemma 15 to $\pi_n R$ yields an F -valued L_n such that $\mathcal{L}[L_n | \mathcal{F}_{n-1}]$ is simple and $\mathbb{E}[\rho_n(\pi_n R, L_n)]$ arbitrarily small; then Lemma 17 gives a n -simple S such that $\pi_n S = L_n$ and $\mathbb{E}[\rho(R, S)]$ is small; as $\pi_{n-1} S = \mathcal{L}[L_n | \mathcal{F}_{n-1}]$ is simple, S is $(n-1)$ -simple. \blacksquare

With Lemmas 17 and 18 at our disposal, the proof of Vershik's second-level criterion is within hand reach. We now take for (K, ρ) the interval $[0, 1]$ with the usual distance.

THEOREM 2 (Vershik [14]). — Let $\mathcal{F} = (\mathcal{F}_n)_{n \leq 0}$ be a filtration on some probability space $(\Omega, \mathcal{A}, \mathbb{P})$. As in Theorem 1, suppose \mathcal{F}_0 is essentially separable and \mathcal{F}_n is conditionally non-atomic given \mathcal{F}_{n-1} for each $n \leq 0$. The following three statements are equivalent to each other and to statements (i), (ii), (iii) and (iv) of Theorem 1:

(v) \mathcal{F} is immersible into a filtration of product type;

(vi) \mathcal{F} is I -cosy;

(vii) for every $\delta > 0$ and every \mathcal{F}_0 -measurable, $[0, 1]$ -valued r.v. R , there exist an $n \leq 0$ and a $\mu \in K_n$ such that $\mathbb{E}[\rho_n(\pi_n R, \mu)] < \delta$.

PROOF. — (ii) \Rightarrow (v) is trivial, and (v) \Rightarrow (vi) has already been seen in Corollary 2.

(vi) \Rightarrow (vii). Given an \mathcal{F}_0 -measurable, $[0, 1]$ -valued R and a $\delta > 0$, the cosiness hypothesis provides us with an n , an $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$ and two isomorphic copies \mathcal{F}' and \mathcal{F}'' of \mathcal{F} on $\bar{\Omega}$, jointly immersed (in $\bar{\mathcal{F}} = \mathcal{F}' \vee \mathcal{F}''$), independent at time n , and such that $\bar{\mathbb{P}}[|R' - R''| \geq \delta] < \delta$. As $|R' - R''| \leq 1$, $\bar{\mathbb{E}}[|R' - R''|] < 2\delta$, whence $\bar{\mathbb{E}}[\rho_n(\pi_n R', \pi_n R'')] < 2\delta$ by Lemma 14. In this formula, the K_n -valued r.v. $\pi_n R'$ is defined in the filtration $\bar{\mathcal{F}}$; by immersion, it remains the same when computed in \mathcal{F}' . By isomorphic transfer, $\pi_n R'$ and $\pi_n R''$ have the same law as $\pi_n R$; call q this law. As $\pi_n R'$ and $\pi_n R''$ are independent, the measurable function ϕ defined on K_n by $\phi(\lambda) = \int \rho_n(\lambda', \lambda) q(d\lambda') = \mathbb{E}[\rho_n(\pi_n R, \lambda)]$ verifies

$$\begin{aligned} \mathbb{E}[\phi(\pi_n R)] &= \int \phi(\lambda'') q(d\lambda'') = \iint \rho_n(\lambda', \lambda'') q(d\lambda') q(d\lambda'') \\ &= \mathbb{E}[\rho_n(\pi_n R', \pi_n R'')] < 2\delta; \end{aligned}$$

hence $\mathbb{P}[\phi(\pi_n R) \geq 3\delta] \leq \frac{2}{3}$, and there exists an $\omega_0 \in \Omega$ such that $\phi((\pi_n R)(\omega_0)) < 3\delta$. Taking $\mu = (\pi_n R)(\omega_0)$, one has $\mathbb{E}[\rho_n(\pi_n R, \mu)] = \phi(\mu) < 3\delta$.

(vii) \Rightarrow (iii). Define a sequence of functions $g_n : K_n \rightarrow [0, 1]$ by $g_0 = \text{Id}$ and $g_{n-1}(\mu) = \int_{K_n} g_n(\lambda) \mu(d\lambda)$. Given any \mathcal{F}_0 -measurable, $[0, 1]$ -valued r.v. R , equality $\mathbb{E}[R | \mathcal{F}_n] = g_n(\pi_n R)$ holds for $n = 0$, and if it holds for some n , the definition of $\pi_{n-1} R$ implies

$$\mathbb{E}[R | \mathcal{F}_{n-1}] = \mathbb{E}[g_n(\pi_n R) | \mathcal{F}_{n-1}] = \int_{K_n} g_n(\lambda) \pi_{n-1} R(d\lambda) = g_{n-1}(\pi_{n-1} R);$$

so $\mathbb{E}[R | \mathcal{F}_n] = g_n(\pi_n R)$ for all n . Now, by Lemma 10, each g_n is 1-Lipschitz. Consequently, $|\mathbb{E}[R | \mathcal{F}_n] - g_n(\mu)| \leq \rho_n(\pi_n R, \mu)$ for any $\mu \in K_n$. Combining this

estimate with hypothesis (vii), we obtain for every $\delta > 0$ an n and a constant c such that $\mathbb{E}[R|\mathcal{F}_n]$ is δ -close to c in L^1 . As this holds for some n , it also holds for all n small enough, so there are constants c_n such that $\mathbb{E}[R|\mathcal{F}_n] - c_n$ tends to 0 in L^1 when $n \rightarrow -\infty$. This implies that $\mathbb{E}[R|\mathcal{F}_n] \rightarrow \mathbb{E}[R]$, which in turn shows that $\mathcal{F}_{-\infty}$ is degenerate.

It remains to show that every $R \in L^0(\mathcal{F}_0)$ can be approximated to any given accuracy by a r.v. belonging to some essentially finite, \mathcal{F} -saturated σ -field. By truncation, we may suppose that R is bounded, and by an affine transformation, that $0 \leq R \leq 1$. Take $K = [0, 1]$ and fix $\delta > 0$. Hypothesis (vii) gives an n (fixed in the sequel) and a $\mu \in K_n$ such that $\mathbb{E}[\rho_n(\pi_n R, \mu)] < \delta$. According to Lemma 18, there exists an n -simple r.v. S such that $\mathbb{E}[\rho(R, S)] < \delta$. Lemma 14 gives $\mathbb{E}[\rho_n(\pi_n R, \pi_n S)] < \delta$, whence $\mathbb{E}[\rho_n(\pi_n S, \mu)] < 2\delta$, and $\mathbb{P}[\rho_n(\pi_n S, \mu) \geq 3\delta] < \frac{2}{3}$. This estimate, and the fact that $\pi_n S$ is a.s. K_n^S -valued (apply Lemma 16 to S), imply that for some ω_0 one has $\rho_n((\pi_n S)(\omega_0), \mu) < 3\delta$ and $(\pi_n S)(\omega_0) \in K_n^S$. Putting $\nu = (\pi_n S)(\omega_0)$, one gets $\rho_n(\mu, \nu) < 3\delta$ and $\mathbb{E}[\rho_n(\pi_n S, \nu)] < 5\delta$. Taking now L deterministic and equal to ν in Lemma 17, we get an n -simple T such that $\mathbb{E}[\rho(S, T)] < 5\delta$ and $\pi_n T = \nu$.

The σ -field $\mathcal{B} = \sigma(\pi_m T, m \leq 0)$ is \mathcal{F} -saturated by Lemma 13. For each $m \leq n$, $\pi_m T$ is deterministic (easy induction, starting from $\pi_n T = \nu$); and for $m \geq n$, $\pi_m T$ is simple because T is n -simple; so \mathcal{B} is essentially finite. Since $\mathbb{E}[\rho(R, T)] < 6\delta$ and T is \mathcal{B} -measurable, (iii) is established. ■

To lift the assumptions on \mathcal{F} in theorems 1 and 2, we need a definition of standardness suitable for the general case (atomic or not, homogeneous or not).

DEFINITION — A filtration \mathcal{F} is standard² if it is immersible into a standard non-atomic filtration.

This name is not misleading, because every standard non-atomic filtration is standard (it is immersed in itself), and because a standard filtration such that \mathcal{F}_n is conditionally non-atomic given \mathcal{F}_{n-1} for every n , is standard non-atomic (by condition (v) of Theorem 2).

COROLLARY 5. — Let $(\Omega, \mathcal{A}, \mathbb{P}, \mathcal{F})$ be a filtered probability space. The following four conditions are equivalent:

- (i) \mathcal{F} is standard;
- (ii) \mathcal{F} is immersible into a standard filtration;
- (iii) if $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}}, \bar{\mathcal{F}})$ is a standard non-atomic filtered sample space, the independent product of \mathcal{F} and $\bar{\mathcal{F}}$ is standard non-atomic too;
- (iv) \mathcal{F}_0 is essentially separable and \mathcal{F} is I -cosy.

The independent product in condition (iii) just means a filtration of the form $\mathcal{F}' \vee \bar{\mathcal{F}}'$, where \mathcal{F}' and $\bar{\mathcal{F}}'$ are defined on the same sample space, independent and respectively isomorphic to \mathcal{F} and $\bar{\mathcal{F}}$. The product of \mathcal{F} with a standard non-atomic filtration is well-defined, up to isomorphism; so if (iii) holds for some standard non-atomic $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}}, \bar{\mathcal{F}})$, it holds for every standard non-atomic $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}}, \bar{\mathcal{F}})$.

2. We follow Vershik's terminology; Feldman [8] says *prestandard*. The (equivalent) definition they use is condition (iii) of Corollary 5.

PROOF OF COROLLARY 5. — (i) \Rightarrow (iv). A standard non-atomic filtration satisfies (iv) by Theorem 2; and a filtration immersible into a filtration satisfying (iv) also satisfies (iv) by Corollary 1.

(iv) \Rightarrow (iii). Let \mathcal{F} satisfy (iv) and \mathcal{G} be a standard non-atomic filtration independent of \mathcal{F} (possibly defined on an enlargement of Ω); put $\mathcal{H} = \mathcal{F} \vee \mathcal{G}$. As both \mathcal{F}_0 and \mathcal{G}_0 are essentially separable, so is also \mathcal{H}_0 . For each n , \mathcal{G}_n is conditionally non-atomic given \mathcal{G}_{n-1} ; so $\mathcal{F}_{n-1} \vee \mathcal{G}_n$ is conditionally non-atomic given \mathcal{H}_{n-1} by Corollary 3, and a fortiori \mathcal{H}_n is conditionally non-atomic given \mathcal{H}_{n-1} . This allows us to apply Theorems 1 and 2 to \mathcal{H} . Now both \mathcal{F} and \mathcal{G} are I-cosy, \mathcal{H} is I-cosy too by Proposition 3, and standard non-atomic by (vi) \Rightarrow (i) in Theorems 1 and 2.

(iii) \Rightarrow (ii). By Lemma 1, a filtration is always immersible into its independent product with any other filtration.

(ii) \Rightarrow (i). If \mathcal{F} is immersible into a filtration \mathcal{G} which is in turn immersible into a standard non-atomic filtration \mathcal{H} , then \mathcal{F} is immersible into \mathcal{H} , whence standard. \blacksquare

Among the reasons that make standardness an interesting notion stands the following fact. *Let $(\mathcal{F}_n)_{n \leq 0}$ be a filtration, and $(E_n)_{n \leq 0}$ be a sequence of sets, each of which is either a finite set or the interval $[0, 1]$. Suppose there exists for each n an \mathcal{F}_n -measurable r.v. X_n , uniformly distributed on E_n , independent of \mathcal{F}_{n-1} , and such that $\mathcal{F}_{n-1} \vee \sigma(X_n) = \mathcal{F}_n$. The filtration \mathcal{F} is standard (if and) only if it is of product type.*

We shall not prove this statement. When each X_n has a diffuse law, it just repeats equivalence (i) \Leftrightarrow (v) in Theorems 1 and 2. In the general case, it is an immediate corollary of Theorem 3.2 of Vershik [14]. Vershik gives a complete proof only in the case when each E_n is finite, but the indices n such that E_n is infinite are easy to deal with, in the same way as above: approximate random variables by simple ones and use Corollary 4. Another proof is provided by Feldman [8]; a key step in his method consists in showing that, for some special K , ρ , μ and ν , it is possible to find a probability λ on $K \times K$, carried by a graph, and arbitrarily close to being optimal in the definition of the Kantorovich-Rubinshtein distance $\rho'(\mu, \nu)$.

The restriction that each X_n is uniformly distributed is essential. A very simple counter-example is attributed to Vinokurov by Vershik ([14] page 756; see also Feldman [8]): the natural filtration \mathcal{F} of “the” stationary Markov chain $(M_n)_{n \leq 0}$ with two states and transition matrix $\begin{pmatrix} p & 1-p \\ 1-p & p \end{pmatrix}$, where $0 < p < 1$ and $p \neq \frac{1}{2}$. An easy coupling argument³ shows that \mathcal{F} is I-cosy, hence standard (Corollary 5). On the other hand, there are (modulo negligibility) only two non-degenerate \mathcal{F}_n -events independent of \mathcal{F}_{n-1} , namely $\{M_n = M_{n-1}\}$ and $\{M_n \neq M_{n-1}\}$. So if \mathcal{F} were of product type, it would be generated by the process $Y_n = \mathbb{1}_{\{M_n = M_{n-1}\}}$; but this process determines M only up to one bit of information (interchange both states), a contradiction.

3. Consider two independent copies M^1 and M^2 of M , put $T_m = \inf \{n \geq m \mid M_n^1 = M_n^2\}$ and define another copy N of the process M by $N_n = M_n^1$ if $n \leq T_m$ and $N_n = M_n^2$ if $n \geq T_m$. The filtrations generated by M^1 and N are isomorphic to that of M and jointly immersed in that of (M^1, M^2) ; by taking m small enough, the processes M^1 and N have a large probability of being equal on a given interval $\{n, \dots, 0\}$.

Vershik's theorem on lacunary isomorphism

The same tools that were needed to establish Vershik's criterion will now be used to prove his theorem on lacunary isomorphism, a phenomenon that we find still much more mind-boggling than the existence of non standard filtrations. It says that a non standard filtration can always be made standard by a deterministic time-change, that is, by replacing $-\mathbb{N}$ with one of its subsequences. We keep following closely Vershik [14].

THEOREM 3 (lacunary isomorphism). — *Let \mathcal{F} be a filtration such that \mathcal{F}_0 is essentially separable and $\mathcal{F}_{-\infty}$ is degenerate. There exists a strictly increasing map $\sigma : -\mathbb{N} \rightarrow -\mathbb{N}$ such that the filtration \mathcal{G} defined by $\mathcal{G}_n = \mathcal{F}_{\sigma(n)}$ is standard.*

The argument will be split into several statements.

LEMMA 19. — *Let \mathcal{B} and \mathcal{C} be two sub- σ -fields of \mathcal{A} , with \mathcal{C} included in \mathcal{B} and \mathcal{B} conditionally non-atomic given \mathcal{C} . If R is a \mathcal{B} -measurable r.v. taking values in some finite set F , there exists a \mathcal{B} -measurable, F -valued r.v. S independent of \mathcal{C} and such that $\mathbb{P}[S \neq R] = \frac{1}{2} \sum_{r \in F} \mathbb{E}[|\mathbb{P}[R=r|\mathcal{C}] - \mathbb{P}[R=r]|]$.*

PROOF. — Remark first that if K is a finite set endowed with the distance $\rho(r, s) = \mathbb{1}_{\{r \neq s\}}$, then the Kantorovich-Rubinshtein distance ρ' on the compact K' is explicitly given by the formula $\rho'(\mu, \nu) = \frac{1}{2} \sum_{t \in K} |\mu(t) - \nu(t)|$. To see this, define

$$\begin{aligned} p(r) &= (\mu(r) - \nu(r))^+ = \mu(r) - \mu(r) \wedge \nu(r), \\ q(s) &= (\mu(s) - \nu(s))^- = \nu(s) - \mu(s) \wedge \nu(s), \\ C &= 1 - \sum_t \mu(t) \wedge \nu(t) = \sum_r p(r) = \sum_s q(s) = \frac{1}{2} \sum_t |\mu(t) - \nu(t)| \end{aligned}$$

and observe that $\lambda(r, s) = \mu(r) \wedge \nu(s) \mathbb{1}_{\{r=s\}} + p(r)q(s)/C$ is a probability on $K \times K$ verifying $\lambda(t, t) = \mu(t) \wedge \nu(t)$, thus achieving the infimum in the definition of $\rho'(\mu, \nu)$.

Now apply Lemma 12 to $K = F$ endowed with this ρ , to R , and to the constant r.v. $L = \ell$ where ℓ is the law of R . This gives an S with values in F such that $\mathcal{L}[S|\mathcal{C}] = \ell$ (so S is independent of \mathcal{C}) and $\mathbb{P}[S \neq R] = \frac{1}{2} \sum_{r \in F} \mathbb{E}[|\mathcal{L}[R|\mathcal{C}](r) - \ell(r)|]$. ■

LEMMA 20. — *Let \mathcal{F} be a filtration such that \mathcal{F}_0 is essentially separable, each \mathcal{F}_n is conditionally non-atomic given \mathcal{F}_{n-1} , and $\mathcal{F}_{-\infty}$ is degenerate. For every $\delta > 0$ and every \mathcal{F}_0 -measurable, simple r.v. R , there exist an $m < 0$ and an \mathcal{F}_0 -measurable r.v. S , independent of \mathcal{F}_m and verifying $\mathbb{P}[S \neq R] < \delta$.*

PROOF. — Call F a finite set where R takes its values. For each $r \in F$, the degeneracy of $\mathcal{F}_{-\infty}$ implies that $\mathbb{P}[R=r|\mathcal{F}_m] \rightarrow \mathbb{P}[R=r]$ in L^1 when $m \rightarrow -\infty$. So there is an $m < 0$ such that $\sum_{r \in F} \mathbb{E}[|\mathbb{P}[R=r|\mathcal{F}_m] - \mathbb{P}[R=r]|] < \delta$; and Lemma 19 gives the S sought for. ■

LEMMA 21. — *Let \mathcal{F} be a filtration verifying the same hypotheses as in the previous lemma. Suppose given an $n \leq 0$ and $|n|$ random variables X_{n+1}, \dots, X_0 such that each X_ℓ is a complement to $\mathcal{F}_{\ell-1}$ in \mathcal{F}_ℓ . For every $\delta > 0$ and every \mathcal{F}_0 -measurable r.v. R , there exist an $m < n$, a complement X' to \mathcal{F}_m in \mathcal{F}_n , and a r.v. S , $\sigma(X', X_{n+1}, \dots, X_0)$ -measurable and δ -close to R in probability.*

PROOF. — Writing $\mathcal{F}_0 = \mathcal{F}_n \vee \sigma(X_{n+1}, \dots, X_0)$ and approximating \mathcal{F}_n by essentially finite σ -fields, we may suppose R to be of the form $\phi(T, X_{n+1}, \dots, X_0)$ where ϕ is Borel and T is \mathcal{F}_n -measurable and simple. Lemma 20 applied to T and to the shifted filtration $(\dots, \mathcal{F}_{n-1}, \mathcal{F}_n)$ gives an $m < n$ and a simple, \mathcal{F}_n -measurable S , independent of \mathcal{F}_m and such that $\mathbb{P}[S \neq T] < \delta$; a fortiori, $\mathbb{P}[\phi(S, X_{n+1}, \dots, X_0) \neq R] < \delta$. According to Corollary 4, there exists a complement X' to \mathcal{F}_m in \mathcal{F}_n such that $\sigma(X') \supset \sigma(S)$; so $S = \psi(X')$ and $\phi(\psi(X'), X_{n+1}, \dots, X_0)$ is δ -close to R in probability. ■

PROPOSITION 6. — *Let \mathcal{F} be a filtration such that \mathcal{F}_0 is essentially separable, each \mathcal{F}_n is conditionally non-atomic given \mathcal{F}_{n-1} , and $\mathcal{F}_{-\infty}$ is degenerate. There exists a strictly increasing map $\sigma : -\mathbb{N} \rightarrow -\mathbb{N}$ such that the filtration \mathcal{G} defined by $\mathcal{G}_n = \mathcal{F}_{\sigma(n)}$ is standard non-atomic.*

PROOF. — Choose a r.v. R generating \mathcal{F}_0 and a sequence $(\delta_\ell)_{\ell \leq 0}$ such that $\delta_0 = 1$, $\delta_\ell > 0$ and $\delta_\ell \rightarrow 0$ when $\ell \rightarrow -\infty$. We shall first show the existence of a strictly increasing mapping $\sigma : -\mathbb{N} \rightarrow -\mathbb{N}$, a sequence $(X'_\ell)_{\ell \leq 0}$ such that each X'_ℓ is a complement to $\mathcal{F}_{\sigma(\ell-1)}$ in $\mathcal{F}_{\sigma(\ell)}$, and a sequence $(S_\ell)_{\ell \leq 0}$ of r.v.'s such that S_ℓ is $\sigma(X'_\ell, \dots, X'_0)$ -measurable and δ_ℓ -close to R in probability. Indeed, supposing $\sigma(\ell), \dots, \sigma(0)$, $X'_{\ell+1}, \dots, X'_0$ and $S_{\ell+1}, \dots, S_0$ have already been constructed, $\sigma(\ell-1)$, X'_ℓ and S_ℓ are obtained simply by applying Lemma 21 to the filtration $(\dots, \mathcal{F}_{\sigma(\ell+1)-2}, \mathcal{F}_{\sigma(\ell+1)-1}, \mathcal{F}_{\sigma(\ell+1)}, \mathcal{F}_{\sigma(\ell+2)}, \dots, \mathcal{F}_{\sigma(-1)}, \mathcal{F}_{\sigma(0)})$; this gives σ , X' and S by induction (the first step just starts with $\sigma(0) = 0$, no X' and no S ; applying Lemma 21 then yields $\sigma(-1)$, X'_0 and S_0 , and so on).

Now, Lemma 8 says that the (standard non-atomic) filtration \mathcal{G}' generated by the process X' is immersed in the filtration \mathcal{G} defined by $\mathcal{G}_n = \mathcal{F}_{\sigma(n)}$. But R is the limit in probability of S_ℓ , so it is \mathcal{G}'_0 -measurable, whence $\mathcal{G}'_0 = \mathcal{F}_0$ by the choice of R . As \mathcal{G}_0 is sandwiched between \mathcal{G}'_0 and \mathcal{F}_0 , one also has $\mathcal{G}'_0 = \mathcal{G}_0$; consequently, by Lemma 5, $\mathcal{G}' = \mathcal{G}$ (both these filtrations are immersed in \mathcal{G}), and \mathcal{G} is standard non-atomic. ■

PROOF OF THEOREM 3. — Let \mathcal{F} be any filtration such that \mathcal{F}_0 is essentially separable and $\mathcal{F}_{-\infty}$ is degenerate. By enlarging the sample space if necessary, we may suppose the existence of a standard non-atomic filtration \mathcal{H} independent of \mathcal{F} . The filtration $\mathcal{F} \vee \mathcal{H}$ satisfies the same hypotheses as \mathcal{F} and moreover its increments are conditionally non-atomic. So Proposition 6 can be applied to this filtration, giving a sub-sequence σ such that the filtration $\mathcal{K}_n = \mathcal{F}_{\sigma(n)} \vee \mathcal{H}_{\sigma(n)}$ is standard non-atomic. Set $\mathcal{G}_n = \mathcal{F}_{\sigma(n)}$. Being immersed in \mathcal{K} by Lemma 1, the filtration \mathcal{G} is standard. ■

Study of an example

To illustrate the notion of I-cosiness and how it can be used, we now turn to Vershik's Example 2 ([14], page 744). This will be done in two steps: we start with a modified, easier version of the example, due to Smorodinsky [12]; there the state spaces are finite. Then we come back to Vershik's version, which is slightly less simple (the state space is $[0, 1]$), but stationary.

DEFINITIONS. — *Given a filtered probability space $(\Omega, \mathcal{A}, \mathbb{P}, (\mathcal{F}_n)_{n \leq 0})$, a process $(\varepsilon_n)_{n \leq 0}$ is an \mathcal{F} -coin-tossing if, for each $n \leq 0$, ε_n is \mathcal{F}_n -measurable, independent of \mathcal{F}_{n-1} , and uniformly distributed on $\{-1, 1\}$.*

A filtration \mathcal{F} has the predictable representation property with respect to an \mathcal{F} -coin-tossing ε if $\mathcal{F}_{-\infty}$ is degenerate and, for each n , \mathcal{F}_n is generated by \mathcal{F}_{n-1} and ε_n .

The definition of an \mathcal{F} -coin-tossing is equivalent to demanding that the law of the process ε is that of a fair coin-tossing, and that the filtration it generates is immersed in \mathcal{F} .

The predictable representation property with respect to ε amounts to saying that every \mathcal{F} -martingale M has the form $M_n = \mathbb{E}[M_0] + \sum_{m \leq n} H_m \varepsilon_m$ for some \mathcal{F} -predictable process H .

THEOREM 4 (Vershik [14], Smorodinsky [12]). — *There exists a filtered probability space $(\Omega, \mathcal{A}, \mathbb{P})$ with a filtration $\mathcal{F} = (\mathcal{F}_n)_{n \leq 0}$ and an \mathcal{F} -coin-tossing $\varepsilon = (\varepsilon_n)_{n \leq 0}$ such that \mathcal{F} has the predictable representation property with respect to ε , but \mathcal{F} is not standard.*

In particular, this \mathcal{F} is not of product type (this would imply I-cosiness by Proposition 3), so it is not generated by any coin-tossing whatsoever; according to Corollary 2, is it not even immersible into a filtration of product type (for instance the filtration generated by some coin-tossing process).

We now describe Smorodinsky's construction of this paradoxical filtration.

Let A be a finite set, called the alphabet, with k elements, called letters ($k \geq 2$). An ordered sequence of m letters is called an m -word; the set of all m -words is A^m , its cardinality is k^m . An m -word w and an m' -word w' can be concatenated to form the $(m+m')$ -word $w \cdot w'$. For $n \leq 0$, put $B_n = A^{2^{|n|}}$.

Consider a Markov process $(X_n, \varepsilon_n)_{n \leq 0}$ with the following law: for each $n \leq 0$, (X_n, ε_n) is uniformly distributed on $B_n \times \{-1, 1\}$ (that is, the random $2^{|n|}$ -word X_n and the random sign ε_n are uniform and independent); the transition from $n-1$ to n is obtained by taking ε_n independent of $(X_{n-1}, \varepsilon_{n-1})$ and choosing X_n as the first half of the word X_{n-1} if $\varepsilon_n = 1$, and as the second half of X_{n-1} if $\varepsilon_n = -1$.

Clearly, this transition probability, when applied to a uniformly distributed $(X_{n-1}, \varepsilon_{n-1})$, yields a uniformly distributed (X_n, ε_n) ; this compatibility implies existence and uniqueness in law of the process (X, ε) . From now on, we suppose (X, ε) is realized on a sample space $(\Omega, \mathcal{A}, \mathbb{P})$, and we call \mathcal{F} the filtration it generates.

The Markov property and the independence of ε_n and $(X_{n-1}, \varepsilon_{n-1})$ imply that ε_n is independent of \mathcal{F}_{n-1} ; in other words, the process ε is an \mathcal{F} -coin-tossing.

THEOREM 4' (Smorodinsky [12]). — *With the above definitions, the filtration \mathcal{F} has the predictable representation property with respect to ε , but it is not standard.*

The proof is copied from Smorodinsky; the only difference is that we use the language of cosiness, but the mathematics are exactly the same.

PROOF OF THE PREDICTABLE REPRESENTATION PROPERTY. — As X_n is the first or second half of X_{n-1} according to the value of ε_n , X_n is a function of X_{n-1} and ε_n , and \mathcal{F}_n is generated by \mathcal{F}_{n-1} and ε_n . To establish the predictable representation property, it remains to verify that $\mathcal{F}_{-\infty}$ is degenerate, or, equivalently, that for each $Z \in L^1(\mathcal{F}_0)$, the conditional expectation $\mathbb{E}[Z|\mathcal{F}_n]$ tends to $\mathbb{E}[Z]$ when n tends to $-\infty$.

We know that X_n is one half of the word X_{n-1} ; call W_n the other half (which may happen to be equal to X_n). Knowing X_0 , the sequence (ε_n) and the sequence (W_n) , it is easy to recover all the X_n by backward induction; so \mathcal{F}_0 is generated by X_0 , the ε and the W . Consequently, we may suppose that Z is a function of X_0 , $(\varepsilon_{\ell+1}, \varepsilon_{\ell+2}, \dots, \varepsilon_0)$ and $(W_{\ell+1}, W_{\ell+2}, \dots, W_0)$ for some $\ell < 0$. The vector $(W_{\ell+1}, W_{\ell+2}, \dots, W_0, X_0)$ is itself a function of X_ℓ and $(\varepsilon_{\ell+1}, \varepsilon_{\ell+2}, \dots, \varepsilon_0)$; so $Z = f(X_\ell, \varepsilon_{\ell+1}, \dots, \varepsilon_0)$. Now, for $n \leq \ell$, to compute the conditional expectation $\mathbb{E}[Z|\mathcal{F}_n]$, we may replace Z by $\mathbb{E}[Z|\mathcal{F}_\ell]$; the latter is but $g(X_\ell)$, with $g(x) = \mathbb{E}[f(x, \varepsilon_{\ell+1}, \dots, \varepsilon_0)]$, for $(\varepsilon_{\ell+1}, \dots, \varepsilon_0)$ is independent of \mathcal{F}_ℓ . So it suffices to check that the conditional law of X_ℓ given \mathcal{F}_n tends to the uniform law on $A^{2^{|\ell|}}$. This conditional law can easily be described: the random $2^{|n|}$ -word X_n can be sliced into $2^{|n|-|\ell|}$ smaller words, each with length $2^{|\ell|}$; conditionally on \mathcal{F}_n , X_ℓ is uniformly chosen among those $2^{|n|-|\ell|}$ words. As they are independent and uniformly distributed on $A^{2^{|\ell|}}$, they constitute a $2^{|n|-|\ell|}$ -sample of the uniform law on $A^{2^{|\ell|}}$, and the conditional law of X_ℓ given \mathcal{F}_n is the empirical measure associated to this sample. When $n \rightarrow -\infty$, the size of the sample tends to infinity (it doubles at each step), and the empirical measure converges almost surely to the uniform law on $A^{2^{|\ell|}}$ by the law of large numbers. ■

Before we turn to non-standardness, some definitions and estimates will be useful. In the next lemma, the symmetric group \mathfrak{S}_m is identified with the group of all permutations of the set $I_m = \{1, 2, \dots, m\}$.

LEMMA 22. — *There exists a sequence $(G_n)_{n \leq 0}$ such that*

- (i) *for each $n \leq 0$, G_n is a sub-group of the symmetric group $\mathfrak{S}_{2^{|n|}}$, having $2^{2^{|n|}-1}$ elements;*
- (ii) *the permutation of $I_{2^{|n|+1}}$ defined as the translation by $2^{|n|}$ modulo $2^{|n|+1}$ is in G_{n-1} (it globally exchanges both halves of the interval $I_{2^{|n|+1}}$ without modifying the order inside them);*
- (iii) *if g_1 and g_2 are in G_n , then the permutation g of $I_{2^{|n|+1}}$ acting as g_1 on the first half $I_{2^{|n|}}$ of $I_{2^{|n|+1}}$, and as g_2 on its second half (identified with $I_{2^{|n|}}$ by a translation) is in G_{n-1} .*

The G_n are unique; this is easy to check but we shall not need it. They are the groups of automorphisms of the dyadic tree; Vershik calls them \mathcal{D}_n .

PROOF OF LEMMA 22. — By induction: define G_0 as $\{\text{Id}\}$ and G_{n-1} as the set of all permutations g of $I_{2^{|n|+1}}$ such that

either both halves of the interval $I_{2^{|n|+1}}$ are globally stable by g , and the restriction of g to each half is in G_n ;

or both halves of the interval are globally exchanged by g , and g is the product of the translation defined in condition (ii) by an element of the previous type.

It is easy to see that this set is indeed a group, whose order $|G_{n-1}|$ verifies $|G_{n-1}| = 2|G_n|^2$; this gives (i). \blacksquare

A word $w \in B_n$ can be written as $w_1 w_2 \dots w_{2^{|n|}}$. The group G_n acts on B_n in the obvious way, by permuting the $2^{|n|}$ letters: $g(w_1 w_2 \dots w_{2^{|n|}}) = w_{g(1)} w_{g(2)} \dots w_{g(2^{|n|})}$. Notice the following property: if g_1 and g_2 are in G_n , there exists $g \in G_{n-1}$ such that, for all words u and v in B_n , $g(u \cdot v) = g_1(u) \cdot g_2(v)$ and there exists $h \in G_{n-1}$ such that, for all words u and v in B_n , $h(v \cdot u) = g_1(u) \cdot g_2(v)$.

Endow the set B_n with the distance

$$\delta_n(w', w'') = \frac{1}{2^{|n|}} \sum_{i=1}^{2^{|n|}} \mathbb{1}_{w'_i \neq w''_i}.$$

Remark that the distance δ_n is invariant under the action of G_n (more generally, of the whole symmetric group) and define on $B_n \times B_n$ the symmetric function

$$d_n(w', w'') = \inf_{g \in G_n} \delta_n(w', g(w'')).$$

LEMMA 23. — *If u' , v' , u'' and v'' are four words in B_n ,*

$$\frac{1}{2} [(d_n(u', u'') + d_n(v', v'')) \wedge (d_n(u', v'') + d_n(v', u''))] \geq d_{n-1}(u' \cdot v', u'' \cdot v'').$$

(Actually, Smorodinsky establishes equality; only this inequality will be needed.)

PROOF OF LEMMA 23. — We first show $\frac{1}{2} (d_n(u', u'') + d_n(v', v'')) \geq d_{n-1}(w', w'')$, with $w' = u' \cdot v'$ and $w'' = u'' \cdot v''$. The left side is $\frac{1}{2} [\delta_n(u', g_1(u'')) + \delta_n(v', g_2(v''))]$, where g_1 and g_2 are in G_n ; this is also $\delta_{n-1}(u' \cdot v', g_1(u'') \cdot g_2(v''))$. But $g_1(u'') \cdot g_2(v'')$ is equal to $g(u'' \cdot v'')$, where g is the element of G_{n-1} acting as g_1 on the first half of the word and as g_2 on the second half. So we have on the left $\delta_{n-1}(w', g(w''))$, and this majorizes $d_{n-1}(w', w'')$ by definition of d_{n-1} .

The other minoration is similar, with an $h \in G_{n-1}$ such that $g_1(u'') \cdot g_2(v'')$ equals $h(v'' \cdot u'')$. \blacksquare

LEMMA 24. — *Let \mathcal{F}' and \mathcal{F}'' be two isomorphic copies of \mathcal{F} , both immersed in some filtered probability space $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}}, (\bar{\mathcal{F}}_n)_{n \leq 0})$; call X' and X'' the copies of X in \mathcal{F}' and \mathcal{F}'' . The process $d_n(X'_n, X''_n)$ is an $\bar{\mathcal{F}}$ -submartingale.*

PROOF OF LEMMA 24. — Fix $n \leq 0$ and write $X'_{n-1} = U' \cdot V'$ and $X''_{n-1} = U'' \cdot V''$, where U' is the first half of the random word X'_{n-1} and similarly for X'' . By immersion, the copies ε' and ε'' of ε in \mathcal{F}' and \mathcal{F}'' verify

$$\bar{\mathbb{P}}[\varepsilon'_n = 1 | \bar{\mathcal{F}}_{n-1}] = \bar{\mathbb{P}}[\varepsilon'_n = -1 | \bar{\mathcal{F}}_{n-1}] = \bar{\mathbb{P}}[\varepsilon''_n = 1 | \bar{\mathcal{F}}_{n-1}] = \bar{\mathbb{P}}[\varepsilon''_n = -1 | \bar{\mathcal{F}}_{n-1}] = \frac{1}{2}.$$

This implies

$$\begin{aligned} \bar{\mathbb{P}}[\varepsilon'_n = \varepsilon''_n = 1 | \bar{\mathcal{F}}_{n-1}] &= \bar{\mathbb{P}}[\varepsilon'_n = \varepsilon''_n = -1 | \bar{\mathcal{F}}_{n-1}] = \frac{1}{2} \bar{\mathbb{P}}[\varepsilon'_n = \varepsilon''_n | \bar{\mathcal{F}}_{n-1}]; \\ \bar{\mathbb{P}}[\varepsilon'_n = 1, \varepsilon''_n = -1 | \bar{\mathcal{F}}_{n-1}] &= \bar{\mathbb{P}}[\varepsilon'_n = -1, \varepsilon''_n = 1 | \bar{\mathcal{F}}_{n-1}] = \frac{1}{2} \bar{\mathbb{P}}[\varepsilon'_n \neq \varepsilon''_n | \bar{\mathcal{F}}_{n-1}]. \end{aligned}$$

By construction of X ,

$$X'_n = \begin{cases} U' & \text{if } \varepsilon'_n = 1 \\ V' & \text{if } \varepsilon'_n = -1 \end{cases} \quad X''_n = \begin{cases} U'' & \text{if } \varepsilon''_n = 1 \\ V'' & \text{if } \varepsilon''_n = -1; \end{cases}$$

so

$$\begin{aligned} \bar{\mathbb{E}}[d_n(X'_n, X''_n) | \bar{\mathcal{F}}_{n-1}] &= d_n(U', U'') \bar{\mathbb{P}}[\varepsilon'_n = \varepsilon''_n = 1 | \bar{\mathcal{F}}_{n-1}] \\ &\quad + d_n(U', V'') \bar{\mathbb{P}}[\varepsilon'_n = 1, \varepsilon''_n = -1 | \bar{\mathcal{F}}_{n-1}] \\ &\quad + d_n(V', U'') \bar{\mathbb{P}}[\varepsilon'_n = -1, \varepsilon''_n = 1 | \bar{\mathcal{F}}_{n-1}] \\ &\quad + d_n(V', V'') \bar{\mathbb{P}}[\varepsilon'_n = \varepsilon''_n = -1 | \bar{\mathcal{F}}_{n-1}] \\ &= \frac{1}{2} [d_n(U', U'') + d_n(V', V'')] \bar{\mathbb{P}}[\varepsilon'_n = \varepsilon''_n | \bar{\mathcal{F}}_{n-1}] \\ &\quad + \frac{1}{2} [d_n(U', V'') + d_n(V', U'')] \bar{\mathbb{P}}[\varepsilon'_n \neq \varepsilon''_n | \bar{\mathcal{F}}_{n-1}] \end{aligned}$$

and

$$\bar{\mathbb{E}}[d_n(X'_n, X''_n) | \bar{\mathcal{F}}_{n-1}] \geq \frac{1}{2} [(d_n(U', U'') + d_n(V', V'')) \wedge (d_n(U', V'') + d_n(V', U''))].$$

It now suffices to appeal to Lemma 23 to get

$$\bar{\mathbb{E}}[d_n(X'_n, X''_n) | \bar{\mathcal{F}}_{n-1}] \geq d_{n-1}(X'_{n-1}, X''_{n-1}). \quad \blacksquare$$

LEMMA 25. — Fix $n < 0$ and $m = 2^{|n|}$. If W' and W'' are two independent random m -words with uniform law on B_n ,

$$\mathbb{P}[d_n(W', W'') \leq \frac{1}{2}] < \frac{1}{2} \left(\frac{4}{\sqrt{k}} \right)^m.$$

In particular, if $k \geq 16$, $\mathbb{E}[d_n(W', W'')] \geq \frac{1}{4}$.

PROOF OF LEMMA 25. — We start with the distance δ_n instead of d_n :

$$\begin{aligned} \mathbb{P}[d_n(W', W'') \leq \frac{1}{2}] &= \mathbb{P}[W' \text{ and } W'' \text{ have at least } m/2 \text{ matching letters}] \\ &= \mathbb{P}[\exists J \subset \{1, \dots, m\} \quad |J| = \frac{1}{2}m \quad \forall j \in J \quad W'(j) = W''(j)] \\ &\leq \sum_{\substack{J \subset \{1, \dots, m\} \\ |J| = \frac{m}{2}}} \mathbb{P}[\forall j \in J \quad W'(j) = W''(j)] \\ &= \sum_{\substack{J \subset \{1, \dots, m\} \\ |J| = \frac{m}{2}}} \left(\frac{1}{k} \right)^{\frac{m}{2}} \leq \sum_{J \subset \{1, \dots, m\}} \left(\frac{1}{k} \right)^{\frac{m}{2}} = \left(\frac{2}{\sqrt{k}} \right)^m. \end{aligned}$$

Now, since for any g in G_n (or in \mathfrak{S}_m), gW'' is uniformly distributed and independent of W' ,

$$\begin{aligned} \mathbb{P}[d_n(W', W'') \leq \frac{1}{2}] &= \mathbb{P}[\exists g \in G_n \quad \delta_n(W', gW'') \leq \frac{1}{2}] \leq \sum_{g \in G_n} \mathbb{P}[\delta_n(W', gW'') \leq \frac{1}{2}] \\ &= 2^{m-1} \mathbb{P}[\delta_n(W', W'') \leq \frac{1}{2}] \leq \frac{1}{2} \left(\frac{4}{\sqrt{k}} \right)^m. \end{aligned}$$

When $k \geq 16$, the m -th power is at most 1, one has $\mathbb{P}[d_n(W', W'') \leq \frac{1}{2}] \leq \frac{1}{2}$, whence $\mathbb{E}[d_n(W', W'')] \geq \mathbb{E}[\frac{1}{2} \mathbb{1}_{\{d_n(W', W'') > \frac{1}{2}\}}] = \frac{1}{2} \mathbb{P}[d_n(W', W'') > \frac{1}{2}] > \frac{1}{4}$. \blacksquare

END OF THE PROOF OF THEOREM 4'. To establish that \mathcal{F} is not standard, we shall show it is not I-cosy.

First case : $k \geq 16$. — On some sample space $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$, let \mathcal{F}' and \mathcal{F}'' be any two filtrations isomorphic to \mathcal{F} , jointly immersed and I-separated. By triviality of G_0 and by Lemma 24, one has for every $n \leq 0$

$$\bar{\mathbb{P}}[X'_0 \neq X''_0] = \bar{\mathbb{E}}[\delta_0(X'_0, X''_0)] = \bar{\mathbb{E}}[d_0(X'_0, X''_0)] \geq \bar{\mathbb{E}}[d_n(X'_n, X''_n)] .$$

By I-separation, there is an $n < 0$ such that X'_n and X''_n are independent. For this n , Lemma 25 gives $\bar{\mathbb{E}}[d_n(X'_n, X''_n)] > \frac{1}{4}$, wherefrom $\bar{\mathbb{P}}[X'_0 \neq X''_0] > \frac{1}{4}$. This minoration shows that condition (iv) in the definition of I-cosiness cannot be satisfied for $R = X_0$, and \mathcal{F} is not I-cosy.

Second case : $k < 16$. — The new alphabet $\hat{A} = A^4$ has at least 16 letters; it consists of “new letters”, which are blocks of 4 old letters. Calling \hat{B}_n the space $\hat{A}^{2^{|n|}}$, there is a natural identification between B_{n-2} and \hat{B}_n , obtained by considering a word of $2^{|n|+2} = 4 \times 2^{|n|}$ old letters as a word of $2^{|n|}$ new letters. Putting $\hat{X}_n = X_{n-2}$ and $\hat{\varepsilon}_n = \varepsilon_{n-2}$ for $n \leq 0$, the natural filtration $\hat{\mathcal{F}}$ of the process $(\hat{X}, \hat{\varepsilon})$ is given by $\hat{\mathcal{F}}_n = \mathcal{F}_{n-2}$. The first case applies to the new, hatted, process and shows that $\hat{\mathcal{F}}$ is not I-cosy; according to Proposition 2, neither is \mathcal{F} . ■

Other forms of cosiness

As recalled in the introduction, two other definitions of cosiness can be found in the literature; all three definitions are identical but for the separation condition (iii). The genuine one, hereafter called II-cosiness, was introduced by Tsirelson [13] in a framework where all martingales are continuous; the separation condition it uses is a reinforcement of the Kunita-Watanabe inequality: Two jointly immersed filtrations \mathcal{F} and \mathcal{G} (such that all martingales are continuous) are II-separate if for some $r < 1$ and for all \mathcal{F} -martingales M and \mathcal{G} -martingales N started at 0, one has

$$[M, N]^2 \leq r [M, M] [N, N] .$$

When rewritten in discrete time, this is equivalent to a conditional correlation inequality: For \mathcal{F} and \mathcal{G} jointly immersed in \mathcal{H} and for all $F \in L^2(\mathcal{F}_n)$ and $G \in L^2(\mathcal{G}_n)$,

$$\text{Cov} [F, G \mid \mathcal{H}_{n-1}]^2 \leq r \text{Var} [F \mid \mathcal{H}_{n-1}] \text{Var} [G \mid \mathcal{H}_{n-1}] ;$$

with this separation condition, we do not know if the non I-cosy filtration of the above example is II-cosy or not.

The separation condition used in [2], [3] and [7] is slightly different: it does not involve time, but only the end σ -fields \mathcal{F}_∞ and \mathcal{G}_∞ of the filtrations (or \mathcal{F}_0 and \mathcal{G}_0 when time is $-\mathbb{N}$); this makes it in some sense a coarse tool when compared to II-separation or I-separation. We shall call it *D-separation* (D for diffuse). Its definition is quite simple: Two filtrations \mathcal{F} and \mathcal{G} are *D-separate* if for all random variables $F \in L^0(\mathcal{F}_\infty)$ and $G \in L^0(\mathcal{G}_\infty)$ with diffuse laws, one has $\mathbb{P}[F=G] = 0$. And *D-cosiness* is defined exactly as I-cosiness, but with D-separation instead of I-separation in condition (iii).

At the risk of adding a little more confusion to this admittedly already rather messy situation, we shall introduce yet another variant of cosiness. Not only by sheer pleasure of losing the reader in a maze of definitions, but also for a logical reason: we shall establish that the non I-cosy filtration of the previous section is cosy in this new sense, so it is worth stating this result with a definition of cosiness as strong as possible. (By contradistinction, D-cosiness was used in [2], [3] and [7] to show that some filtrations are *not* D-cosy, that is why its definition was made as weak as possible.)

This new definition formalizes an idea introduced and brilliantly used, without giving it a name, by Tsirelson in [13]; it consists in “rotating” a Gaussian processes and all associated filtrations. If $G = (G_\lambda)_{\lambda \in \Lambda}$ is a centered Gaussian process and G' an independent copy of G , define a new centered Gaussian process $(G_\lambda^\alpha)_{\alpha \in \mathbb{R}, \lambda \in \Lambda}$ by $G_\lambda^\alpha = G_\lambda \cos \alpha + G'_\lambda \sin \alpha$. Notice that for any fixed α , the process $G^\alpha = (G_\lambda^\alpha)_{\lambda \in \Lambda}$ has the same law as G ; so the σ -fields $\sigma(G)$ and $\sigma(G^\alpha)$ are isomorphic, with an isomorphism Ψ^α such that $\Psi^\alpha(G_\lambda) = G_\lambda^\alpha$.

DEFINITION. — *A filtration \mathcal{F} (not necessarily indexed by $-\mathbb{N}$) is G-cosy if there exist two independent centered Gaussian processes G and G' with the same law, defined on some sample space $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$, and an embedding Φ of $(\Omega, \mathcal{F}_\infty, \mathbb{P})$ into $(\bar{\Omega}, \sigma(G), \bar{\mathbb{P}})$, such that, with the above notation, for each α the filtrations $\Phi(\mathcal{F})$ and $\Psi^\alpha \circ \Phi(\mathcal{F})$ are jointly immersed.*

This simply means the existence (possibly on an extension of Ω) of a Gaussian process G verifying $\mathcal{F}_\infty \subset \sigma(G)$ and such that \mathcal{F} and any copy \mathcal{F}^α of \mathcal{F} obtained by rotating G are jointly immersed.

The simplest example of a G-cosy filtration is any filtration generated by a Gaussian process, for instance a Brownian filtration. This is the case Tsirelson considered when introducing cosiness; his proof in [13] that such filtrations are cosy is copied below, in Proposition 8 and in the remark following it. But we cannot restrict ourselves to this case: in the proof of Proposition 9, we shall need the definition of G-cosiness in its full extent.

LEMMA 26. — *A filtration immersible into a G-cosy filtration is itself G-cosy.*

PROOF. — Suppose that \mathcal{F} satisfies the above definition and \mathcal{E} is immersed in \mathcal{F} . One has $\Phi(\mathcal{E}_\infty) \subset \Phi(\mathcal{F}_\infty) \subset \sigma(G)$; and $\Phi(\mathcal{E})$ is immersed in $\Phi(\mathcal{F})$. By hypothesis, the filtrations $\Psi^\alpha \circ \Phi(\mathcal{F})$ and $\Phi(\mathcal{F})$ on $\bar{\Omega}$ are immersed in their supremum \mathcal{H} , so $\Phi(\mathcal{E})$ and $\Psi^\alpha \circ \Phi(\mathcal{E})$ are also immersed in \mathcal{H} , and jointly immersed by Lemma 4 b).

This shows that a filtration immersed in a G-cosy filtration is itself G-cosy, and the lemma follows by isomorphism. ■

PROPOSITION 7. — *Every standard filtration $(\mathcal{F}_n)_{n \leq 0}$ is G-cosy.*

PROOF. — A standard filtration is immersible into a standard, non-atomic filtration; so by the preceding lemma it suffices to verify that a (or “the”) standard non-atomic filtration is G-cosy. Now a standard non-atomic filtration is generated by an i.i.d. sequence $(G_n)_{n \leq 0}$ of standard Gaussian random variables; it suffices to enlarge Ω to accomodate both G and an independent copy G' of G , and the definition of G-cosiness is readily verified. ■

PROPOSITION 8. — *Every G-cosy filtration is D-cosy.*

PROOF. — Suppose a filtration \mathcal{F} is G-cosy. By replacing \mathcal{F} with the isomorphic filtration $\Phi(\mathcal{F})$, we may suppose that $\bar{\Omega} = \Omega$ and $\Phi = \text{Id}$ in the definition of G-cosiness.

Fix $\delta > 0$ and R in $L^0(\mathcal{F}_\infty)$; notice that $R \in L^0(\sigma(G))$. By Slutsky's lemma (see Théorème 1 of [2]), the map $\alpha \mapsto \Psi^\alpha(R)$ is continuous for the topology of convergence in probability; so, for α close enough to 0, the distance between $\Psi^\alpha(R)$ and $\Psi^0(R)$ is less than δ . For any α , the isomorphic copies $\mathcal{F}^\alpha = \Psi^\alpha(\mathcal{F})$ and \mathcal{F} of the filtration \mathcal{F} are jointly immersed by hypothesis. Last, for $\alpha \neq 0 \pmod{\pi}$, the filtrations \mathcal{F}^α and \mathcal{F} are D-separate; see for instance Proposition 2 of [3].

Consequently, for α close enough to 0 but not null, the filtrations \mathcal{F}^α and \mathcal{F}^0 fulfill all four conditions in the definition of D-cosiness. \blacksquare

REMARK. — Define two σ -fields \mathcal{B} and \mathcal{C} to be *H-separate* if, for some $p \in [1, 2)$,

$$\forall B \in L^p(\mathcal{B}) \quad \forall C \in L^p(\mathcal{C}) \quad \mathbb{E}[BC] \leq \|B\|_{L^p} \|C\|_{L^p}$$

(the letter H stands for Hypercontractivity, or for Hypo-independence). Define two filtrations \mathcal{F} and \mathcal{G} to be H-separate if their end σ -fields \mathcal{F}_∞ and \mathcal{G}_∞ are, and a filtration to be H-cosy if it satisfies the definition of D-cosiness (or I-cosiness, or Π -cosiness) with H-separation instead of D-separation (or I-separation, or Π -separation). It is not difficult to see that H-separation implies D-separation (see [2], or Proposition 1 of [3]); therefore H-cosiness implies D-cosiness. And the proof of Proposition 8 (or the proof of Proposition 2 of [3]) shows in fact a stronger result: G-cosiness \Rightarrow H-cosiness \Rightarrow D-cosiness.

PROPOSITION 9. — *The (non standard, hence also non I-cosy) filtration \mathcal{F} of Theorem 4' is G-cosy (hence also H-cosy and D-cosy).*

The proof of this proposition will use some properties of a very small, abelian subgroup of the group G_n described in Lemma 22. Tsirelson calls it the group of cube automorphisms (as opposed to the whole group G_n of tree automorphisms).

LEMMA 27. — *For $n \leq 0$ and $(\eta_1, \dots, \eta_{|n|}) \in \{-1, 1\}^{|n|}$, define an operation $\gamma_{\eta_1, \dots, \eta_{|n|}}^n$ on all $2^{|n|}$ -words as follows: γ_{\emptyset}^0 is the identity on A^1 , and for $n \leq -1$, writing every $2^{|n|}$ -word of $A^{2^{|n|}}$ as the concatenation $u \cdot v$ of two $2^{|n|-1}$ -words,*

$$\gamma_{\eta_1, \dots, \eta_{|n|}}^n(u \cdot v) = \begin{cases} \gamma_{\eta_2, \dots, \eta_{|n|}}^{n+1}(u) \cdot \gamma_{\eta_2, \dots, \eta_{|n|}}^{n+1}(v) & \text{if } \eta_1 = 1, \\ \gamma_{\eta_2, \dots, \eta_{|n|}}^{n+1}(v) \cdot \gamma_{\eta_2, \dots, \eta_{|n|}}^{n+1}(u) & \text{if } \eta_1 = -1. \end{cases}$$

These operations form a commutative group, called H_n , with $2^{|n|}$ elements. The map $(\eta_1, \dots, \eta_{|n|}) \mapsto \gamma_{\eta_1, \dots, \eta_{|n|}}^n$ is a group isomorphism between $\{-1, 1\}^{|n|}$ and H_n .

For $\gamma \in H_n \setminus \{\text{Id}\}$, there is a partition of $\{1, \dots, 2^{|n|}\}$ into $2^{|n|-1}$ pairs such that γ (considered as an element of \mathfrak{S}_n) acts by swapping both elements of each pair.

PROOF OF LEMMA 27. — Both properties

$$\begin{aligned} \gamma_{\eta_1, \dots, \eta_{|n|}}^n = \text{Id} & \Leftrightarrow (\eta_1, \dots, \eta_{|n|}) = (1, \dots, 1) \\ \gamma_{\eta'_1 \eta''_1, \dots, \eta'_{|n|} \eta''_{|n|}}^n & = \gamma_{\eta'_1, \dots, \eta'_{|n|}}^n \gamma_{\eta''_1, \dots, \eta''_{|n|}}^n \end{aligned}$$

are readily verified by induction on n .

For $1 \leq k \leq |n|$, divide the interval $\{1, \dots, 2^{|n|}\}$ into 2^k blocks (i.e. subintervals) of length $2^{|n|-k}$. If $\eta_1 = \dots = \eta_{k-1} = 1$ and $\eta_k = -1$, then $\gamma_{\eta_1, \dots, \eta_{|n|}}^n$ sends every element of the $(2j+1)$ st block into the $(2j+2)$ nd block and conversely (induction on n and k , with $|n| - k$ fixed, starting from $k = 1$ and $|n| \geq 1$). Consequently, if $(\eta_1, \dots, \eta_{|n|}) \neq (1, \dots, 1)$, the permutation $\gamma_{\eta_1, \dots, \eta_{|n|}}^n$ has no fixed point; and as its square is Id (isomorphism with $\{-1, 1\}^{|n|}$), its orbits form a partition of $\{1, \dots, 2^{|n|}\}$ into $2^{|n|-1}$ classes of 2 elements each. ■

PROOF OF PROPOSITION 9. — Consider a quadruply infinite supply $(E_n)_{n \leq 0}$, $(C_p)_{p \geq 1}$, $(E'_n)_{n \leq 0}$, $(C'_p)_{p \geq 1}$ of independent standard normal random variables; let f be a Borel function from \mathbb{R} to A transforming the standard normal law into the uniform law on the alphabet A . For $n \leq 0$, define random variables ε_n in $\{-1, 1\}$ and X_n in $B_n = A^{2^{|n|}}$ by

$$\varepsilon_n = \text{sgn } E_n ; \quad X_n = \gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n (f(C_1) \dots f(C_{2^{|n|}}))$$

where γ^n is the same as in Lemma 27.

The first step of the proof consists in justifying this notation by showing that the process (X, ε) has the same law as in the example studied in Theorem 4'. It is easy to see that for each n the law of (X_n, ε_n) is uniform on $B_n \times \{-1, 1\}$: just notice that the random word $W = f(C_1) \dots f(C_{2^{|n|}})$ is uniform on B_n and that X_n is obtained from W by a random permutation independent of W . It is also straightforward to get X_{n+1} as the first or second half of X_n according to the value 1 or -1 of ε_{n+1} : just replace $\gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n$ by its definition. To complete the first step, it only remains to verify that the r.v. ε_n is independent of $(X_k, \varepsilon_k)_{k < n}$. This will be done later; meanwhile, the first step is left uncompleted.

Call G (respectively G') the centered Gaussian process (E, C) (respectively (E', C')); notice that G' is an independent copy of G and that (X, ε) can be written as $\psi(G)$ for some Borel functional ψ . Define $G^\alpha = G \cos \alpha + G' \sin \alpha$, put $(X^\alpha, \varepsilon^\alpha) = \psi(G^\alpha)$, call \mathcal{F}^α the filtration generated by $(X^\alpha, \varepsilon^\alpha)$, and \mathcal{H} the filtration generated by all processes $(X^\beta, \varepsilon^\beta)$ when β takes all possible values. With the notation used in the definition of G-cosiness, \mathcal{F}^α is but $\Psi^\alpha(\mathcal{F})$. The second step is to establish that each \mathcal{F}^α is immersed in \mathcal{H} ; when both steps are done, the proposition will be proved. Trivially, \mathcal{F}^α is included in \mathcal{H} . As \mathcal{F}^α has the predictable representation property with respect to ε^α , immersion amounts to saying that ε_n^α is independent of \mathcal{H}_{n-1} . Showing it will a fortiori establish independence of ε_n^α and \mathcal{F}_{n-1}^α and, taking $\alpha = 0$, of ε_n and $\sigma((X_k, \varepsilon_k)_{k < n})$; so by the same token step 1 will also be completed.

Fixing n from now on, it only remains to prove independence of ε_n^α and \mathcal{H}_{n-1} ; this is equivalent to $\mathbb{E}[\varepsilon_n^\alpha | \mathcal{H}_{n-1}] = 0$, or to

$$\mathbb{E}[\varepsilon_n^\alpha h((X_k^{\beta_1})_{k < n}, (\varepsilon_k^{\beta_1})_{k < n}, \dots, (X_k^{\beta_q})_{k < n}, (\varepsilon_k^{\beta_q})_{k < n})] = 0$$

for all $q > 0$, $\beta_1, \dots, \beta_q \in \mathbb{R}$ and for all bounded, Borel h .

The random variable inside the expectation is measurable for $\sigma(E, E', C, C')$; therefore it is a functional

$$\Phi(E, E', C, C') = \Phi((E_k)_{k \leq 0}, (E'_k)_{k \leq 0}, (C_p)_{p \geq 1}, (C'_p)_{p \geq 1})$$

of infinitely many independent $N(0, 1)$ random variables. To prove the claim, we shall exhibit another Gaussian family $(\mathbf{E}, \mathbf{E}', \mathbf{C}, \mathbf{C}')$ of random variables, with the same law as (E, E', C, C') , but such that $\Phi(\mathbf{E}, \mathbf{E}', \mathbf{C}, \mathbf{C}') \equiv -\Phi(E, E', C, C')$; this will imply $\mathbb{E}[\Phi(E, E', C, C')] = \mathbb{E}[\Phi(\mathbf{E}, \mathbf{E}', \mathbf{C}, \mathbf{C}')] = -\mathbb{E}[\Phi(E, E', C, C')]$, whence $\mathbb{E}[\Phi(E, E', C, C')] = 0$.

Define \mathbf{E} and \mathbf{E}' by

$$\mathbf{E}_k = \begin{cases} E_k & \text{if } k < n \\ -E_k & \text{if } k \geq n \end{cases} \quad \mathbf{E}'_k = \begin{cases} E'_k & \text{if } k < n \\ -E'_k & \text{if } k \geq n \end{cases}$$

and put $\mathbf{C}_p = C_{s(p)}$ and $\mathbf{C}'_p = C'_{s(p)}$, where s is the *deterministic* permutation of $\{1, 2, \dots\}$ which globally preserves each interval $\{i 2^{|n|+1} + 1, \dots, (i+1) 2^{|n|+1}\}$, but completely reverses the order inside this interval. For instance, if $n = -2$, s is the permutation

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & 13 & 14 & 15 & 16 & 17 & 18 & \dots \\ 8 & 7 & 6 & 5 & 4 & 3 & 2 & 1 & 16 & 15 & 14 & 13 & 12 & 11 & 10 & 9 & 24 & 23 & \dots \end{pmatrix} .$$

Observe the effect of replacing (E, E', C, C') by $(\mathbf{E}, \mathbf{E}', \mathbf{C}, \mathbf{C}')$ in the argument of Φ . For $k < n$, $\varepsilon_k^\alpha = \text{sgn}(E_k \cos \alpha + E'_k \sin \alpha)$ does not change, whereas for $k \geq n$, $\varepsilon_k^\beta = \text{sgn}(E_k \cos \beta + E'_k \sin \beta)$ is replaced by $\text{sgn}(\mathbf{E}_k \cos \beta + \mathbf{E}'_k \sin \beta) = -\varepsilon_k^\beta$. And for $k < n$,

$$X_k^\beta = \gamma_{\varepsilon_{k+1}^\beta, \dots, \varepsilon_{n-1}^\beta, \varepsilon_n^\beta, \dots, \varepsilon_0^\beta}^k (f(C_1 \cos \beta + C'_1 \sin \beta) \dots f(C_{2^{|k|}} \cos \beta + C'_{2^{|k|}} \sin \beta))$$

is replaced by

$$\gamma_{\varepsilon_{k+1}^\beta, \dots, \varepsilon_{n-1}^\beta, -\varepsilon_n^\beta, \dots, -\varepsilon_0^\beta}^k (f(C_{s(1)} \cos \beta + C'_{s(1)} \sin \beta) \dots f(C_{s(2^{|k|})} \cos \beta + C'_{s(2^{|k|})} \sin \beta)) ;$$

the claim will be established and the proposition proved if we show that the latter is equal to X_k^β itself. This is a particular instance of the general identity

$$\gamma_{\eta_{k+1}, \dots, \eta_{n-1}, -\eta_n, \dots, -\eta_0}^k (w_{s(1)} \dots w_{s(2^{|k|})}) = \gamma_{\eta_{k+1}, \dots, \eta_{n-1}, \eta_n, \dots, \eta_0}^k (w_1 \dots w_{2^{|k|}}) ,$$

valid for all $k < n \leq 0$. To show this identity, notice first that, since

$$\gamma_{\eta_{k+1}, \dots, \eta_{n-1}, -\eta_n, \dots, -\eta_0}^k = \gamma_{\eta_{k+1}, \dots, \eta_{n-1}, \eta_n, \dots, \eta_0}^k \gamma_{1, \dots, 1, -1, \dots, -1}^k$$

as a consequence of the group homomorphism (Lemma 27), the identity reduces to

$$\gamma_{1, \dots, 1, -1, \dots, -1}^k (w_{s(1)} \dots w_{s(2^{|k|})}) = w_1 \dots w_{2^{|k|}} .$$

For $k = n - 1$, this is just the formula

$$\gamma_{-1, \dots, -1}^{n-1} (w_{2^{|n|+1}} \dots w_1) = w_1 \dots w_{2^{|n|+1}} ,$$

which can be verified by induction on n , using the definition of γ^{n-1} ; for fixed n and general $k < n$, it is proved by induction on $n-k$ using the definition of γ^k and the definition of s . The proof is now complete. \blacksquare

REMARK. — Proposition 9 and its proof remain true if the finite alphabet A is replaced with any separable probability space (A, \mathcal{A}, π) : it suffices to replace f by a Borel function transforming the normal law into the probability π . This remark is not interesting per se (the simpler the model, the better), but will be technically useful in the next section.

Vershik's Example 3

When showing that his example is not immersible into a standard filtration, Smorodinsky's aim was to explain the same property for Vershik's example; indeed, this property carries over immediately from the former to the latter. As we shall see, G-cosiness also transfers easily.

Vershik's example is a Markov process indexed by the signed integers \mathbb{Z} , with state space the set $A^{\mathbb{N}}$ made of all infinite words $v = v^0v^1v^2\dots$ and endowed with the product measure μ (the letters are independent and uniform). The transition probability is

$$v^0v^1v^2v^3v^4\dots \mapsto \begin{cases} v^0v^2v^4\dots & \text{with probability } \frac{1}{2} \\ v^1v^3v^5\dots & \text{with probability } \frac{1}{2}; \end{cases}$$

plainly, μ is invariant, so we have a stationary process $V = (V_n)_{n \in \mathbb{Z}}$ with this probability transition and with stationary law μ . Associated to it is the coin-tossing $(\varepsilon_n)_{n \in \mathbb{Z}}$ such that V_n is the "even half" of V_{n-1} if $\varepsilon_n = 1$ and the "odd half" if $\varepsilon_n = -1$. The process (V, ε) has the same natural filtration \mathcal{G} as V , since both halves of V_{n-1} are almost surely different.

For $n \leq 0$, consider the word $W_n = V_n^0V_n^1\dots V_n^{2^{|n|}}$ obtained from the infinite word V_n by keeping only the first $2^{|n|}$ letters. The process (W, ε) is very similar to Smorodinsky's (X, ε) -process: at each step, one half of the word is lost and one half is retained, but these halves are not the left half and right half; they are the halves consisting of letters in even or odd position. There exists for each n a (deterministic) permutation σ_n of $\{0, 1, \dots, 2^{|n|} - 1\}$ such that the words $X_n = W_n \circ \sigma_n$ form a process with the same law as in Smorodinsky's example. The existence of those σ_n is obvious by induction; an explicit description is also possible: $\sigma_n(i)$ and i , when written as binary numbers with $|n|$ digits, are mirror-image from each other. For instance, $X_{-3} = W_{-3}^0W_{-3}^4W_{-3}^2W_{-3}^6W_{-3}^1W_{-3}^5W_{-3}^3W_{-3}^7$. In this way, the process X is immersed in the process V . This really is an immersion, in the sense that the filtration \mathcal{F} generated by (X, ε) becomes a sub-filtration of \mathcal{G} , immersed in \mathcal{G} ; consequently, non-standardness of \mathcal{G} is a corollary of non-standardness of \mathcal{F} .

It is almost as easy to get G-cosiness of \mathcal{G} from Proposition 9. Given V , define for $n \leq 0$ and $0 \leq k < 2^{|n|}$ a random infinite word $Z_n^k \in A^{\mathbb{N}}$ by

$$Z_n^k = V_n^{\sigma_n(k)}V_n^{\sigma_n(k)+2^{|n|}}V_n^{\sigma_n(k)+2 \cdot 2^{|n|}}V_n^{\sigma_n(k)+3 \cdot 2^{|n|}} \dots$$

and put $Y_n = Z_n^0\dots Z_n^{2^{|n|}-1}$. This Y_n can be considered as a $2^{|n|}$ -word with letters Z_n^k in the new alphabet $A^{\mathbb{N}}$. This defines a process $(Y_n)_{n \leq 0}$ with two properties. First, as the finite word Y_n is just a rewriting of the infinite word V_n , the processes $(Y_n)_{n \leq 0}$ and $(V_n)_{n \leq 0}$ generate the same filtration. Second, by the choice of σ_n , Y_n is the first or second half of Y_{n-1} according to the value ± 1 of ε_n ; in other terms, Y is an instance of Smorodinsky's process, but with the infinite alphabet $A^{\mathbb{N}}$ instead of A . Consequently, by the remark following the proof of Proposition 9, the filtration $(\mathcal{G}_n)_{n \leq 0}$ is G-cosy. And so is also the whole filtration $(\mathcal{G}_n)_{n \in \mathbb{Z}}$, as can be seen by inserting another Gaussian sequence $(E_n)_{n \geq 1}$ in the proof of Proposition 9, and putting $\varepsilon_n = \text{sgn } E_n$ for $n \geq 1$ as well.

On a question by von Weizsäcker

The proof of Proposition 9 rests on a construction of the process X via certain random permutations of a sequence of random letters. The same construction will now be used to answer a question raised by H. von Weizsäcker during the Mini-Symposium on the Classification of Filtrations at the ESI, Vienna, December 1998.

It is well known that if $\mathcal{F} = (\mathcal{F}_n)_{n \leq 0}$ is a filtration and \mathcal{B} a σ -field, the equality $\bigcap_n (\mathcal{F}_n \vee \mathcal{B}) = \mathcal{F}_{-\infty} \vee \mathcal{B}$ does not always hold true, but it sometimes does, for instance when \mathcal{B} and \mathcal{F}_0 are independent; see von Weizsäcker [16] for a necessary and sufficient condition and for references. The question is, is equality always obtained when \mathcal{F}_0 and \mathcal{B} are “almost independent”? The hope for an affirmative answer relied on the fact that all previously known examples of situations where the above formula goes wrong, had a big intersection between the σ -fields \mathcal{F}_0 and \mathcal{B} , which therefore were far from independent.

By almost independence of \mathcal{F}_0 and \mathcal{B} , we mean either a hypercontractivity inequality:

$$\exists p \in (1, 2) \quad \forall B \in L^p(\mathcal{B}) \quad \forall F \in L^p(\mathcal{F}_0) \quad \mathbb{E}[BF] \leq \|B\|_{L^p} \|F\|_{L^p}$$

with p close to 1, or a correlation inequality:

$$\exists r \in (0, 1) \quad \forall B \in L^2(\mathcal{B}) \quad \forall F \in L^2(\mathcal{F}_0) \quad |\text{Cov}[B, F]| \leq r \sqrt{\text{Var}[B]} \sqrt{\text{Var}[F]}$$

with r close to 0. These inequalities are always satisfied with $p = 2$ and $r = 1$; they become stronger when p and r decrease, and for $p = 1$ and $r = 0$ each of them is equivalent to independence. The former was called *H-separation* in the remark after Proposition 8; the latter could be called C-separation, and C-cosiness could then be defined—but enough is enough!

Rotating a Gaussian process is an efficient way to generate σ -fields satisfying such inequalities. (We also find it efficient in providing some intuition about “almost independence”.) If $G = (G_\lambda)_{\lambda \in \Lambda}$ is a centered Gaussian process and if G' is an independent copy of G , recall that the rotated process is $G^\alpha = G \cos \alpha + G' \sin \alpha$. The σ -fields $\sigma(G)$ and $\sigma(G^\alpha)$ satisfy the above inequalities with exponent $p = 1 + |\cos \alpha|$ and correlation constant $r = |\cos \alpha|$; so G and G^α are almost independent when α is close to $\frac{\pi}{2}$. (On hypercontractivity, see for instance the references in and after the proof of Proposition 2 of [3]. On the correlation inequality, called Gebelein’s inequality in this case, see Exercise V.3.13 of Revuz-Yor [10], or Dym-McKean [6] page 66.)

The answer to von Weizsäcker’s question is negative:

PROPOSITION 10. — *For all $\delta \in (0, 1)$, there exist a filtration $(\mathcal{F}_n)_{n \leq 0}$ and a σ -field \mathcal{B} such that the σ -fields \mathcal{F}_0 and \mathcal{B} satisfy the hypercontractivity and correlation inequalities recalled above with $p = 1 + \delta$ and $r = \delta$, and such that the germ σ -field $\bigcap_n (\mathcal{F}_n \vee \mathcal{B})$ contains a non-constant r.v. independent of $\mathcal{F}_{-\infty} \vee \mathcal{B}$.*

PROOF. — It is inspired from that of Proposition 9. As in that proof, start with a quadruply infinite supply $(E_n)_{n \leq 0}$, $(C_q)_{q \geq 1}$, $(E'_n)_{n \leq 0}$, $(C'_q)_{q \geq 1}$ of independent standard normal random variables; our alphabet will have two letters: $A = \{-1, 1\}$. For $n \leq 0$, define random variables ε_n in $\{-1, 1\}$ and X_n in $B_n = A^{2^{|n|}}$ by

$$\varepsilon_n = \operatorname{sgn} E_n ; \quad X_n = \gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n (\operatorname{sgn}(C_1) \dots \operatorname{sgn}(C_{2^{|n|}})) ,$$

where $\gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n$ is the (random) cube automorphism associated to the (random) signs $\varepsilon_{n+1}, \dots, \varepsilon_0$ by Lemma 27. As was seen in the proof of Proposition 9, the process (X, ε) is a realization of the example studied in Theorem 4'; call $\mathcal{F} = (\mathcal{F}_n)_{n \leq 0}$ its natural filtration. Now rotate the Gaussian process (E, C) by putting

$$E_n^\alpha = E_n \cos \alpha + E'_n \sin \alpha \quad \text{and} \quad C_q^\alpha = C_q \cos \alpha + C'_q \sin \alpha$$

and fix $\alpha \in (0, \frac{\pi}{2})$ such that $\cos \alpha = \delta$. The properties of rotated Gaussian processes recalled earlier imply that, for α close to $\frac{\pi}{2}$, the σ -fields $\sigma(E, C)$ and $\sigma(E^\alpha, C^\alpha)$ are almost independent (hypercontractivity inequality with $p = 1 + \cos \alpha$ and correlation inequality with $r = \cos \alpha$). This property is immediately inherited by the sub- σ -field $\mathcal{B} = \sigma(\operatorname{sgn}(C_q^\alpha), q \geq 1)$ of $\sigma(E^\alpha, C^\alpha)$ and the sub- σ -field \mathcal{F}_0 of $\sigma(E, C)$.

From now on, the parameter α is fixed in $(0, \frac{\pi}{2})$. It should be considered as close to $\frac{\pi}{2}$ (to make p close to 1 and r close to 0), but the sequel does not depend on that.

We have seen in Theorem 4' that $\mathcal{F}_{-\infty}$ is degenerate; to prove the proposition, it remains to exhibit in the germ σ -field $\bigcap_n (\mathcal{F}_n \vee \mathcal{B})$ a non degenerate r.v. independent of \mathcal{B} . We shall show that the process $(\varepsilon_n)_{n \leq 0}$ is measurable with respect to the germ σ -field; this will prove the proposition, for this process is independent of (C, C') , and a fortiori of \mathcal{B} . (Notice that E' plays no rôle: all these σ -fields are included in $\sigma(E, C, C')$. The sequence E' was introduced only as a phantom, to lubricate the proof of hypercontractivity and correlation.)

So the rest of the proof will consist in showing that the knowledge of the sequence $\operatorname{sgn}(C_1^\alpha), \operatorname{sgn}(C_2^\alpha), \dots$ and of the remote past of (X, ε) is sufficient to recover the whole process ε .

Set $u_q = \operatorname{sgn}(C_q)$ and $v_q = \operatorname{sgn}(C_q^\alpha)$. The joint law of the random words $U_n = u_1 \dots u_{2^{|n|}}$ and $V_n = v_1 \dots v_{2^{|n|}}$ is easily described: u_q and v_q are positively correlated bits, with $\mathbb{P}[u_q \neq v_q] = \frac{\alpha}{\pi} < \frac{1}{2}$, and all pairs (u_q, v_q) are independent of each other. Observe that the word X_n is obtained from U_n by a random permutation $\gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n$ independent of (U_n, V_n) . The idea of the proof is that it is possible to recover this permutation $\gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n$ from the observation of the words X_n and V_n with a probability of error that tends to 0 when $n \rightarrow -\infty$. The recipe to get $\gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n$ is as follows: construct all the words γV_n , where γ ranges in H_n , match these words against X_n , and retain the γ giving the best fit, that is, the largest number of matching letters. This works well for large $|n|$ because, with high probability, U_n and V_n have more than half their letters in common, and the group H_n is not too large. The main quantitative estimate is the following lemma, that uses the distance δ_n defined above Lemma 23. (But all what follows remains valid if the factor 2^n in the definition of δ_n is dropped, that is, if $\delta_n(w', w'')$ is defined as the number of non-matching letters in the words w' and w'' .)

LEMMA 28. — Given $\alpha \in (0, \frac{\pi}{2})$, there exists a constant $c(\alpha) > 0$ such that, for all $n < 0$ and all $\gamma \in H_n \setminus \{\text{Id}\}$, one has

$$\mathbb{P}[\delta_n(U_n, \gamma V_n) \leq \delta_n(U_n, V_n)] \leq \exp[-2^{|n|}c(\alpha)].$$

PROOF OF LEMMA 28. — Put $m = 2^{|n|}$. The event to be evaluated is $\{S \geq 0\}$, where

$$S = \sum_{i=1}^m \mathbb{1}_{\{u_i \neq v_i\}} - \sum_{i=1}^m \mathbb{1}_{\{u_i \neq v_{\gamma^{-1}(i)}\}}.$$

Now, since γ is in $H_n \setminus \{\text{Id}\}$, Lemma 27 gives a partition of $\{1, \dots, m\}$ into $\frac{m}{2}$ pairs $\{i, j\}$ such that $\gamma(i) = j$ and $\gamma(j) = i$. For such a pair $\{i, j\}$, set

$$D_{ij} = \mathbb{1}_{\{u_i \neq v_i\}} + \mathbb{1}_{\{u_j \neq v_j\}} - \mathbb{1}_{\{u_i \neq v_j\}} - \mathbb{1}_{\{u_j \neq v_i\}}.$$

The sum S is also the sum of the $m/2$ r.v.'s D_{ij} when $\{i, j\}$ ranges over all these pairs. These r.v.'s are independent and identically distributed, their law depends on α only. This law could easily be computed; we shall not need it, we only retain that $\mathbb{E}[D_{ij}] = \frac{\alpha}{\pi} + \frac{\alpha}{\pi} - \frac{1}{2} - \frac{1}{2} < 0$. The rest of the proof is Chernoff's classical argument of large deviations: Since the derivative at $t = 0$ of the Laplace transform $\phi_\alpha(t) = \mathbb{E}[\exp(tD_{ij})]$ is $\mathbb{E}[D_{ij}] < 0$, there exists $t_\alpha > 0$ such that $0 < \phi_\alpha(t_\alpha) < 1$, and it suffices to write

$$\mathbb{P}[S \geq 0] = \mathbb{P}[\exp(t_\alpha S) \geq 1] \leq \mathbb{E}[\exp(t_\alpha S)] = (\phi_\alpha(t_\alpha))^{m/2}. \quad \blacksquare$$

END OF THE PROOF OF PROPOSITION 10. — If x and v are two $2^{|n|}$ -words, define $k_n(x, v)$ as the element $(\eta_1, \dots, \eta_{|n|}) \in \{-1, 1\}^{|n|}$ that minimizes the distance $\delta_n(x, \gamma_{\eta_1, \dots, \eta_{|n|}}^n v)$ if the minimum is uniquely achieved, and put for instance $k_n(x, v) = (1, \dots, 1)$ if the minimum is not unique.

Consider the event $\Omega_n = \{k_n(X_n, V_n) = (\varepsilon_{n+1}, \dots, \varepsilon_0)\}$ that the recipe described before Lemma 28 gives the correct answer. On Ω_n^c , there exists a (random) $\Gamma \in H_n$ such that $\Gamma \neq \gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n$ and $\delta_n(X_n, \Gamma V_n) \leq \delta_n(X_n, \gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n V_n)$. As the distance between two words is not modified when the same permutation is applied to both of them, we may let $\gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n$ act on all four words in this inequality; the word X_n becomes $\gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n X_n = U_n$, the word $\gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n V_n$ becomes V_n , and ΓV_n becomes $\Gamma' V_n$ with a (random) $\Gamma' = \gamma_{\varepsilon_{n+1}, \dots, \varepsilon_0}^n \Gamma \neq \text{Id}$. Consequently,

$$\Omega_n^c \subset \{\exists \gamma \in H_n \setminus \{\text{Id}\} \quad \delta_n(U_n, \gamma V_n) \leq \delta_n(U_n, V_n)\}.$$

So, using Lemma 28 and remembering that H_n has $2^{|n|}$ elements,

$$\mathbb{P}[\Omega_n^c] \leq \sum_{\gamma \in H_n \setminus \{\text{Id}\}} \mathbb{P}[\delta_n(U_n, \gamma V_n) \leq \delta_n(U_n, V_n)] \leq 2^{|n|} \exp[-2^{|n|}c(\alpha)];$$

hence the series $\sum_n \mathbb{P}[\Omega_n^c]$ is convergent, and equality $(\varepsilon_{n+1}, \dots, \varepsilon_0) = k_n(X_n, V_n)$ holds for all n less than some $N(\omega)$. Consequently, for each fixed $n < 0$, the whole process ε is a function of the sequence $(k_\ell(X_\ell, V_\ell), \ell \leq n)$; it is therefore measurable with respect to $\mathcal{F}_n \vee \mathcal{B}$. \blacksquare

REMARKS. — a) The filtration \mathcal{F} in this proof can be replaced, with exactly the same proof, by the smaller filtration \mathcal{X} generated by X . The σ -field $\mathcal{X}_{-\infty}$ is degenerate because it is included in $\mathcal{F}_{-\infty}$; and since $(\varepsilon_{n+1}, \dots, \varepsilon_0) = k_n(X_n, V_n)$ for all but finitely many n , the process ε is measurable for each of the σ -fields $\mathcal{X}_n \vee \mathcal{B}$.

b) Since $\mathcal{F}_n \vee \mathcal{B}$ contains X_n and the whole process ε , it also contains \mathcal{F}_0 , so $\mathcal{F}_n \vee \mathcal{B} = \mathcal{F}_0 \vee \mathcal{B}$ for each n and the germ σ -field $\bigcap_n (\mathcal{F}_n \vee \mathcal{B})$ too is equal to $\mathcal{F}_0 \vee \mathcal{B}$.

c) The filtration \mathcal{F} in the proof of Proposition 10 is not standard (Theorem 4'); in fact, it can be considered as an arch-example of a non standard filtration. One might ask whether the situation changes when one imposes the restriction that \mathcal{F} must be standard. By Vershik's lacunary isomorphism theorem (Theorem 3), the above construction easily carries over to this case, simply by taking a subsequence. And once this has been done, one can also get an example where \mathcal{F} is standard non-atomic, simply by replacing \mathcal{F} by its product with a standard non-atomic filtration.

d) The construction in the proof of Proposition 10 is a small modification of that of Proposition 9, which features a whole family of processes $(X^\alpha, \varepsilon^\alpha)$ isomorphic to (X, ε) , each of them immersed in the filtration \mathcal{H} generated by all of them. To be in a situation similar to Proposition 10, consider the setting of Proposition 9, but only in the case when the alphabet has two letters -1 and 1 , and $f(x) = \text{sgn } x$. Call \mathcal{F}^α the natural filtration of $(X^\alpha, \varepsilon^\alpha)$ and put $\mathcal{F}_n^{\alpha\beta} = \mathcal{F}_n^\alpha \vee \mathcal{F}_n^\beta$ and $\varepsilon_n^{\alpha\beta} = \varepsilon_n^\alpha \varepsilon_n^\beta$. If $\alpha - \beta$ is not a multiple of $\frac{\pi}{2}$, the same argument as above yields $(\varepsilon_{n+1}^{\alpha\beta}, \dots, \varepsilon_0^{\alpha\beta}) = k_n(X_n^\alpha, X_n^\beta)$ with overwhelming probability when n is small, so for each n , the whole process $\varepsilon^{\alpha\beta}$ is measurable with respect to the σ -field $\mathcal{F}_n^{\alpha\beta}$; in other words, *this process $\varepsilon^{\alpha\beta}$ is measurable with respect to the germ σ -field $\mathcal{F}_{-\infty}^{\alpha\beta}$ of the filtration $\mathcal{F}^{\alpha\beta}$* . Consequently, still for $\alpha - \beta$ not a multiple of $\frac{\pi}{2}$, the filtration $\mathcal{F}^{\alpha\beta}$ is dyadic:

$$\mathcal{F}_n^{\alpha\beta} = \mathcal{F}_{n-1}^{\alpha\beta} \vee \sigma(\varepsilon_n^\alpha) = \mathcal{F}_{n-1}^{\alpha\beta} \vee \sigma(\varepsilon_n^\beta),$$

and one has for all $m \leq n \leq 0$ the identity $\mathcal{F}_m^\alpha \vee \mathcal{F}_n^\beta = \mathcal{F}_m^\beta \vee \mathcal{F}_n^\alpha = \mathcal{F}_n^\alpha \vee \mathcal{F}_n^\beta$. But $\mathcal{F}^{\alpha\beta}$ does not have the predictable representation property with respect to ε^α , nor to ε^β , nor to any other coin-tossing, simply because $\mathcal{F}_{-\infty}^{\alpha\beta}$ is not degenerate: we just saw that it contains $\varepsilon^{\alpha\beta}$.

Thus, for $\alpha \neq \beta \pmod{\frac{\pi}{2}}$, we have an example of two filtrations \mathcal{F}^α and \mathcal{F}^β immersed in $\mathcal{F}^\alpha \vee \mathcal{F}^\beta$, arbitrarily close to independence as in Proposition 10, such that the three filtrations \mathcal{F}^α , \mathcal{F}^β and $\mathcal{F}^\alpha \vee \mathcal{F}^\beta$ are dyadic, $\mathcal{F}_{-\infty}^\alpha$ and $\mathcal{F}_{-\infty}^\beta$ are trivial, and $\bigcap_n (\mathcal{F}_n^\alpha \vee \mathcal{F}_n^\beta)$ is not trivial.

(When $|\alpha - \beta| = \frac{\pi}{2}$, the filtrations \mathcal{F}^α and \mathcal{F}^β are independent, $\mathcal{F}_{-\infty}^{\alpha\beta}$ is degenerate, and $\mathcal{F}_n^{\alpha\beta} = \mathcal{F}_{n-1}^{\alpha\beta} \vee \sigma(\varepsilon_n^\alpha, \varepsilon_n^\beta)$; this is just the product situation.)

Last, observe that the filtration \mathcal{H} generated by all the $(X^\alpha, \varepsilon^\alpha)$ is dyadic too, because $\mathcal{H}_n = \mathcal{H}_{n-1} \vee \sigma(\varepsilon_n^\alpha)$ for any α . Indeed, $\mathcal{H}_{-\infty}$ contains the products $\varepsilon_n^\alpha \varepsilon_n^\beta$ for all α, β and n (if $\alpha - \beta = k\frac{\pi}{2}$, write $\varepsilon^{\alpha\beta} = \varepsilon^{\alpha\gamma} \varepsilon^{\gamma\beta}$); the reason why \mathcal{H} can be dyadic and at the same time contain two independent coin-tossing processes ε^0 and $\varepsilon^{\pi/2}$, is simply that their product $\varepsilon^0 \varepsilon^{\pi/2}$ is $\mathcal{H}_{-\infty}$ -measurable.

Conversely, it looks likely that $\mathcal{H}_{-\infty} = \sigma(\varepsilon^{0\alpha}, \alpha \in \mathbb{R})$ and, for $\alpha - \beta \notin \mathbb{Z}\frac{\pi}{2}$, $\bigcap_n (\mathcal{F}_n^\alpha \vee \mathcal{F}_n^\beta) = \sigma(\varepsilon^{\alpha\beta})$; we have not investigated this question.

References

- [1] D. Aldous & M. Barlow. On countable dense random sets. *Séminaire de Probabilités XV*, Lecture Notes in Mathematics 850, Springer 1981.
- [2] M.T. Barlow, M. Émery, F.B. Knight, S. Song & M. Yor. Autour d'un théorème de Tsirelson sur des filtrations browniennes et non browniennes. *Séminaire de Probabilités XXXII*, Lecture Notes in Mathematics 1686, Springer 1998.
- [3] S. Beghdadi-Sakrani & M. Émery. On certain probabilities equivalent to coin-tossing, d'après Schachermayer. *Séminaire de Probabilités XXXIII*, Lecture Notes in Mathematics 1709, Springer 1999.
- [4] N. Belili. Dualité du problème des marges et ses applications. *Séminaire de Probabilités XXXIII*, Lecture Notes in Mathematics 1709, Springer 1999.
- [5] L. Dubins, J. Feldman, M. Smorodinsky & B. Tsirelson. Decreasing sequences of σ -fields and a measure change for Brownian motion. *Ann. Prob.* 24, 882–904, 1996.
- [6] H. Dym & H.P. McKean. Gaussian Processes, Function Theory, and the Inverse Spectral Problem. Academic Press, 1976.
- [7] M. Émery & W. Schachermayer. Brownian filtrations are not stable under equivalent time-changes. *Séminaire de Probabilités XXXIII*, Lecture Notes in Mathematics 1709, Springer 1999.
- [8] J. Feldman. Decreasing sequences of σ -fields: product type, standard, and prestandard. Preprint, Berkeley 1998.
- [9] S. Rachev & L. Rüschendorf. Mass Transportation Problems. Springer, 1998 (2 volumes).
- [10] D. Revuz & M. Yor. Continuous Martingales and Brownian Motion. Springer 1991, 1994, 1999.
- [11] V.A. Rohlin. On the fundamental ideas of measure theory. *Mat. Sb.* 25, 107–150, 1949. English translation: *A.M.S. Transl.* 71, 3–55, 1952.
- [12] M. Smorodinsky. Processes with no standard extension. *Israel J. Math.* 107, 327–331, 1998.
- [13] B. Tsirelson. Triple points: From non-Brownian filtrations to harmonic measures. *GAFa, Geom. funct. anal.* 7, 1096–1142, 1997.
- [14] A. M. Vershik. Approximation in measure theory. *Doctor Thesis*, Leningrad 1973. Expanded and updated english version: The theory of decreasing sequences of measurable partitions. *St. Petersburg Math. J.* 6, 705–761, 1995.
- [15] J. Warren. On the joining of sticky Brownian motion. *Séminaire de Probabilités XXXIII*, Lecture Notes in Mathematics 1709, Springer 1999.
- [16] H. von Weizsäcker. Exchanging the order of taking suprema and countable intersections of σ -algebras. *Ann. Institut H. Poincaré Section B*, 91–100, 1983.

Michel Émery
Université Louis Pasteur et C.N.R.S.
I.R.M.A.
7 rue René Descartes
F 67 084 Strasbourg Cedex
emery@math.u-strasbg.fr

Walter Schachermayer
Technische Universität Wien
Institut für Statistik, Wahrscheinlichkeitstheorie
und Versicherungsmathematik
Wiedner Hauptstraße 8–10 / E 107–5
A 1040 Wien
walter.schachermayer@fam.tuwien.ac.at