

Ergodicity, Decisions, and Partial Information

Ramon van Handel

Abstract In the simplest sequential decision problem for an ergodic stochastic process X , at each time n a decision u_n is made as a function of past observations X_0, \dots, X_{n-1} , and a loss $l(u_n, X_n)$ is incurred. In this setting, it is known that one may choose (under a mild integrability assumption) a decision strategy whose pathwise time-average loss is asymptotically smaller than that of any other strategy. The corresponding problem in the case of partial information proves to be much more delicate, however: if the process X is not observable, but decisions must be based on the observation of a different process Y , the existence of pathwise optimal strategies is not guaranteed. The aim of this paper is to exhibit connections between pathwise optimal strategies and notions from ergodic theory. The sequential decision problem is developed in the general setting of an ergodic dynamical system $(\Omega, \mathcal{B}, \mathbf{P}, T)$ with partial information $\mathcal{Y} \subseteq \mathcal{B}$. The existence of pathwise optimal strategies grounded in two basic properties: the conditional ergodic theory of the dynamical system, and the complexity of the loss function. When the loss function is not too complex, a general sufficient condition for the existence of pathwise optimal strategies is that the dynamical system is a conditional K -automorphism relative to the past observations $\bigvee_{n \geq 0} T^n \mathcal{Y}$. If the conditional ergodicity assumption is strengthened, the complexity assumption can be weakened. Several examples demonstrate the interplay between complexity and ergodicity, which does not arise in the case of full information. Our results also yield a decision-theoretic characterization of weak mixing in ergodic theory, and establish pathwise optimality of ergodic nonlinear filters.

1 Introduction

Let $X = (X_k)_{k \in \mathbb{Z}}$ be a stationary and ergodic stochastic process. A decision maker must select at the beginning of each day k a decision u_k depending on the past observations X_0, \dots, X_{k-1} . At the end of the day, a loss $l(u_k, X_k)$ is incurred. The decision maker would like to minimize her time-average loss

$$L_T(\mathbf{u}) = \frac{1}{T} \sum_{k=1}^T l(u_k, X_k).$$

How should she go about selecting a decision strategy $\mathbf{u} = (u_k)_{k \geq 1}$?

Ramon van Handel

Sherrerd Hall 227, Princeton University, Princeton, NJ 08544, USA. e-mail: rvan@princeton.edu

There is a rather trivial answer to this question. Taking the expectation of the time-average loss, we obtain for any strategy \mathbf{u} using the tower property

$$\begin{aligned} \mathbf{E}[L_T(\mathbf{u})] &= \mathbf{E}\left[\frac{1}{T}\sum_{k=1}^T \mathbf{E}[l(u_k, X_k)|X_0, \dots, X_{k-1}]\right] \\ &\geq \mathbf{E}\left[\frac{1}{T}\sum_{k=1}^T \min_u \mathbf{E}[l(u, X_k)|X_0, \dots, X_{k-1}]\right] = \mathbf{E}[L_T(\tilde{\mathbf{u}})], \end{aligned}$$

where $\tilde{\mathbf{u}} = (\tilde{u}_k)_{k \geq 1}$ is defined as $\tilde{u}_k = \arg \min_u \mathbf{E}[l(u, X_k)|X_0, \dots, X_{k-1}]$ (we disregard for the moment integrability and measurability issues, existence of minima, and the like; such issues will be properly addressed in our results). Therefore, the strategy $\tilde{\mathbf{u}}$ minimizes the *mean* time-average loss $\mathbf{E}[L_T(\mathbf{u})]$.

However, there are conceptual reasons to be dissatisfied with this obvious solution. In many decision problems, one only observes a single sample path of the process X . For example, if X_k is the return of a financial market in day k and $L_T(\mathbf{u})$ is the loss of an investment strategy \mathbf{u} , only one sample path of the model is ever realized: we do not have the luxury of averaging our investment loss over multiple “alternative histories”. The choice of a strategy for which the mean loss is small does not guarantee, *a priori*, that it will perform well on the one and only realization that happens to be chosen by nature. Similarly, if X_k models the state of the atmosphere and $L_T(\mathbf{u})$ is the error of a weather prediction strategy, we face a similar conundrum. In such situations, the use of stochastic models could be justified by some sort of ergodic theorem, which states that the mean behavior of the model with respect to different realizations captures its time-average behavior over a single sample path. Such an ergodic theorem for sequential decisions was obtained by Algoet [1, Theorem 2] under a mild integrability assumption.

Theorem 1.1 (Algoet [1]). *Suppose that $|l(u, x)| \leq \Lambda(x)$ with $\Lambda \in L \log L$. Then*

$$\liminf_{T \rightarrow \infty} \{L_T(\mathbf{u}) - L_T(\tilde{\mathbf{u}})\} \geq 0 \quad a.s.$$

for every strategy \mathbf{u} : that is, the mean-optimal strategy $\tilde{\mathbf{u}}$ is pathwise optimal.

The proof of this result follows from a simple martingale argument. What is remarkable is that the details of the model do not enter the picture at all: nothing is assumed on the properties of X or l beyond some integrability (ergodicity is not needed, and a similar result holds even in the absence of stationarity, cf. [1, Theorem 3]). This provides a universal justification for optimizing the mean loss: the much stronger pathwise optimality property is obtained “for free.”

In the proof of Theorem 1.1, it is essential that the decision maker has full information on the history X_0, \dots, X_{k-1} of the process X . However, the derivation of the mean-optimal strategy can be done in precisely the same manner in the more general setting where only partial or noisy information is available. To formalize this idea, let $Y = (Y_k)_{k \in \mathbb{Z}}$ be the stochastic process observable by the decision maker, and suppose that the pair (X, Y) is stationary and ergodic. The loss incurred at

time k is still $l(u_k, X_k)$, but now u_k may depend on the observed data Y_0, \dots, Y_{k-1} only. It is easily seen that in this setting, the mean-optimal strategy $\tilde{\mathbf{u}}$ is given by $\tilde{u}_k = \arg \min_u \mathbf{E}[l(u, X_k) | Y_0, \dots, Y_{k-1}]$, and it is tempting to assume that $\tilde{\mathbf{u}}$ is also pathwise optimal. Surprisingly, this is very far from being the case.

Example 1.2 (Weissman and Merhav [32]). Let $X_0 \sim \text{Bernoulli}(1/2)$ and let $X_k = 1 - X_{k-1}$ and $Y_k = 0$ for all k . Then (X, Y) is stationary and ergodic: $Y_k = 0$ indicates that we are in the setting of *no* information (that is, we must make blind decisions). Consider the loss $l(u, x) = (u - x)^2$. Then the mean-optimal strategy $\tilde{u}_k = 1/2$ satisfies $L_T(\tilde{\mathbf{u}}) = 1/4$ for all T . However, the strategy $u_k = k \bmod 2$ satisfies $L_T(\mathbf{u}) = 0$ for all T with probability $1/2$. Therefore, $\tilde{\mathbf{u}}$ is not pathwise optimal. In fact, it is easily seen that no pathwise optimal strategy exists.

Example 1.2 illustrates precisely the type of conundrum that was so fortuitously ruled out in the full information setting by Theorem 1.1. Indeed, it would be hard to argue that either \mathbf{u} or $\tilde{\mathbf{u}}$ in Example 1.2 is superior: a gambler placing blind bets u_k on a sequence of games with loss $l(u_k, X_k)$ may prefer either strategy depending on his demeanor. The example may seem somewhat artificial, however, as the hidden process X has infinitely long memory; the gambler can therefore beat the mean-optimal strategy by simply guessing the outcome of the first game. But precisely the same phenomenon can appear when (X, Y) is nearly memoryless.

Example 1.3. Let $(\xi_k)_{k \in \mathbb{Z}}$ be i.i.d. Bernoulli($1/2$), and let $X_k = (\xi_{k-1}, \xi_k)$ and $Y_k = |\xi_k - \xi_{k-1}|$ for all k . Then (X, Y) is a stationary 1-dependent sequence: $(X_k, Y_k)_{k \leq n}$ and $(X_k, Y_k)_{k \geq n+2}$ are independent for every k . We consider the loss $l(u, x) = (u - x_1)^2$. It is easily seen that X_k is independent of Y_1, \dots, Y_{k-1} , so that the mean-optimal strategy $\tilde{u}_k = 1/2$ satisfies $L_T(\tilde{\mathbf{u}}) = 1/4$ for all T . On the other hand, note that $\xi_{k-1} = (\xi_0 + Y_1 + \dots + Y_{k-1}) \bmod 2$. It follows that the strategy $u_k = (Y_1 + \dots + Y_{k-1}) \bmod 2$ satisfies $L_T(\mathbf{u}) = 0$ for all T with probability $1/2$.

Evidently, pathwise optimality cannot be taken for granted in the partial information setting even in the simplest of examples: in contrast to the full information setting, the existence of pathwise optimal strategies depends both on specific ergodicity properties of the model (X, Y) and (as will be seen later) on the complexity on the loss l . What mechanism is responsible for pathwise optimality under partial information is not very well understood. Weissman and Merhav [32], who initiated the study of this problem, give a strong sufficient condition in the binary setting. Little is known beyond their result, beside one particularly special case of quadratic loss and additive noise considered by Nobel [24].¹

¹ It should be noted that the papers [1, 32, 24], in addition to studying the pathwise optimality problem, also aim to obtain *universal* decision schemes that achieve the optimal asymptotic loss without any knowledge of the law of X (note that to compute the mean-optimal strategy $\tilde{\mathbf{u}}$ one must know the joint law of (X, Y)). Such strategies “learn” the law of X on the fly from the observed data. In the setting of partial information, such universal schemes cannot exist without very specific assumptions on the information structure: for example, in the blind setting (cf. Example 1.2), there is no information and thus universal strategies cannot exist. What conditions are required for the existence of universal strategies is an interesting question that is beyond the scope of this paper.

The aim of this paper is twofold. On the one hand, we will give general conditions for pathwise optimality under partial information, and explore some tradeoffs inherent in this setting. On the other hand, we aim to exhibit some connections between the pathwise optimality problem and certain notions and problems in ergodic theory, such as conditional mixing and individual ergodic theorems for subsequences. To make such connections in their most natural setting, we begin by rephrasing the decision problem in the general setting of ergodic dynamical systems.

1.1 The dynamical system setting

Let T be an invertible measure-preserving transformation of a probability space $(\Omega, \mathcal{B}, \mathbf{P})$. T defines the time evolution of the dynamical system $(\Omega, \mathcal{B}, \mathbf{P}, T)$: if the system is initially in state $\omega \in \Omega$, then at time k the system is in the state $T^k \omega$. The state of the system is not directly observable, however. To model the available information, we fix a σ -field $\mathcal{Y} \subseteq \mathcal{B}$ of events that can be observed at a single time. Therefore, if we have observed the system in the time interval $[m, n]$, the information contained in the observations is given by the σ -field $\mathcal{Y}_{m,n} = \bigvee_{k \in [m,n]} T^{-k} \mathcal{Y}$.

In this general setting, the decision problem is defined as follows. Let $\ell : U \times \Omega \rightarrow \mathbb{R}$ be a given loss function, where U is the set of possible decisions. At each time k , a decision u_k is made and a loss $\ell_k(u_k) := \ell(u_k, T^k \omega)$ is incurred. The decision can only depend on the observations: that is, a strategy $\mathbf{u} = (u_k)_{k \geq 1}$ is admissible if u_k is $\mathcal{Y}_{0,k}$ -measurable for every k . The time-average loss is given by

$$L_T(\mathbf{u}) := \frac{1}{T} \sum_{k=1}^T \ell_k(u_k).$$

The basic question we aim to answer is whether there exists a pathwise optimal strategy, that is, a strategy \mathbf{u}^* such that for every admissible strategy \mathbf{u}

$$\liminf_{T \rightarrow \infty} \{L_T(\mathbf{u}) - L_T(\mathbf{u}^*)\} \geq 0 \quad \text{a.s.}$$

The stochastic process setting discussed above can be recovered as a special case.

Example 1.4. Let (X, Y) be a stationary and ergodic stochastic process, where X_k takes values in the measurable space (E, \mathcal{E}) and Y_k takes values in the measurable space (F, \mathcal{F}) . We can realize (X, Y) as the coordinate process on the canonical path space $(\Omega, \mathcal{B}, \mathbf{P})$ where $\Omega = E^{\mathbb{Z}} \times F^{\mathbb{Z}}$, $\mathcal{B} = \mathcal{E}^{\mathbb{Z}} \otimes \mathcal{F}^{\mathbb{Z}}$, and \mathbf{P} is the law of (X, Y) . Let $T : \Omega \rightarrow \Omega$ be the canonical shift $(T(x, y))_n = (x_{n+1}, y_{n+1})$. Then $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is an ergodic dynamical system. If we choose the observation σ -field $\mathcal{Y} = \sigma\{Y_0\}$ and the loss $\ell(u, \omega) = l(u, X_1(\omega))$, we recover the decision problem with partial information for the stochastic process (X, Y) as it was introduced above. More generally, we could let the loss depend arbitrarily on future or past values of (X, Y) .

Let us briefly discuss the connection between pathwise optimal strategies and classical ergodic theorems. The key observation in the derivation of the mean-optimal strategy $\tilde{u}_k = \arg \min_u \mathbf{E}[\ell_k(u)|\mathcal{Y}_{0,k}]$ is that by the tower property

$$\mathbf{E} \left[\frac{1}{T} \sum_{k=1}^T \ell_k(u_k) \right] = \mathbf{E} \left[\frac{1}{T} \sum_{k=1}^T \mathbf{E}[\ell_k(u_k)|\mathcal{Y}_{0,k}] \right].$$

As the summands on the right-hand side depend only on the observed information, we can minimize inside the sum to obtain the mean-optimal strategy $\tilde{\mathbf{u}}$. Precisely the same considerations would show that $\tilde{\mathbf{u}}$ is pathwise optimal if we could prove the ergodic counterpart of the tower property of conditional expectations

$$\frac{1}{T} \sum_{k=1}^T \{\ell_k(u_k) - \mathbf{E}[\ell_k(u_k)|\mathcal{Y}_{0,k}]\} \xrightarrow{T \rightarrow \infty} 0 \quad \text{a.s.} \quad ?$$

The validity of such a statement is far from obvious, however.

In the special case of blind decisions (that is, \mathcal{Y} is the trivial σ -field) the “ergodic tower property” reduces to the question of whether, given $f_k(\omega) := \ell(u_k, \omega)$,

$$\frac{1}{T} \sum_{k=1}^T \{f_k - \mathbf{E}[f_k]\} \circ T^k \xrightarrow{T \rightarrow \infty} 0 \quad \text{a.s.} \quad ?$$

If the functions f_k do not depend on k , this is precisely the individual ergodic theorem. However, an individual ergodic theorem need not hold for arbitrary sequences f_k . Special cases of this problem have long been investigated in ergodic theory. For example, if $f_k = a_k f$ for some fixed function f and bounded sequence $(a_k) \subset \mathbb{R}$, the problem reduces to a weighted individual ergodic theorem, see [2] and the references therein. If $a_k \in \{0, 1\}$ for all k , the problem reduces further to an individual ergodic theorem along a subsequence (at least if the sequence has positive density), cf. [6, 2] and the references therein. A general characterization of such ergodic properties does not appear to exist, which suggests that it is probably very difficult to obtain necessary and sufficient conditions for pathwise optimality. The situation is better for mean (rather than individual) ergodic theorems, cf. [3] and the references therein, and we will also obtain more complete results in a weaker setting.

The more interesting case where the information \mathcal{Y} is nontrivial provides additional complications. In this situation, the “ergodic tower property” could be viewed as a type of *conditional* ergodic theorem, in between the individual ergodic theorem and Algoet’s result [1]. Our proofs are based on an elaboration of this idea.

1.2 Some representative results

The essence of our results is that, when the loss ℓ is not too complex, pathwise optimal strategies exist under suitable conditional mixing assumptions on the ergodic

dynamical system $(\Omega, \mathcal{B}, \mathbf{P}, T)$. To this end, we introduce conditional variants of two standard notions in ergodic theory: weak mixing and K -automorphisms.

Definition 1.5. An invertible dynamical system $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is said to be *conditionally weak mixing* relative to a σ -field \mathcal{Z} if for every $A, B \in \mathcal{B}$

$$\frac{1}{T} \sum_{k=1}^T |\mathbf{P}[A \cap T^k B | \mathcal{Z}] - \mathbf{P}[A | \mathcal{Z}] \mathbf{P}[T^k B | \mathcal{Z}]| \xrightarrow{T \rightarrow \infty} 0 \quad \text{in } L^1.$$

Definition 1.6. An invertible dynamical system $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is called a *conditional K -automorphism* relative to a σ -field $\mathcal{Z} \subset \mathcal{B}$ if there is a σ -field $\mathcal{X} \subset \mathcal{B}$ such that

1. $\mathcal{X} \subset T^{-1}\mathcal{X}$.
2. $\bigvee_{k=1}^{\infty} T^{-k}\mathcal{X} = \mathcal{B} \quad \text{mod } \mathbf{P}$.
3. $\bigcap_{k=1}^{\infty} (\mathcal{Z} \vee T^k\mathcal{X}) = \mathcal{Z} \quad \text{mod } \mathbf{P}$.

When the σ -field \mathcal{Z} is trivial, these definitions reduce² to the usual notions of weak mixing and K -automorphism, cf. [31]. Similar conditional mixing conditions also appear in the ergodic theory literature, see [26] and the references therein.

An easily stated consequence of our main results, for example, is the following.

Theorem 1.7. *Suppose that $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is a conditional K -automorphism relative to $\mathcal{Y}_{-\infty, 0}$. Then the mean-optimal strategy $\tilde{\mathbf{u}}$ is pathwise optimal for every loss function $\ell : U \times \Omega \rightarrow \mathbb{R}$ such that U is finite and $|\ell(u, \omega)| \leq \Lambda(\omega)$ with $\Lambda \in L^1$.*

This result gives a general sufficient condition for pathwise optimality when the decision space U is finite. In the stochastic process setting (Example 1.4), the conditional K -property would follow from the validity of the σ -field identity

$$\bigcap_{k=1}^{\infty} (\mathcal{Y}_{-\infty, 0} \vee \mathcal{X}_{-\infty, -k}) = \mathcal{Y}_{-\infty, 0} \quad \text{mod } \mathbf{P},$$

where $\mathcal{X}_{-\infty, k} = \sigma\{X_i : i \leq k\}$ (choose $\mathcal{X} := \mathcal{X}_{-\infty, 0} \vee \mathcal{Y}_{-\infty, 0}$ in Definition 1.6). In the Markovian setting, this is a familiar identity in filtering theory: it is precisely the necessary and sufficient condition for the optimal filter to be ergodic, see section 3.3 below. Our results therefore lead to a new pathwise optimality property of nonlinear filters. Conversely, results from filtering theory yield a broad class of (even non-Markovian) models for which the conditional K -property can be verified [14, 27]. It is interesting to note that despite the apparent similarity between the conditions for filter ergodicity and pathwise optimality, there appears to be no direct connection between these phenomena, and their proofs are entirely distinct. Let us also note that, in the full information setting ($Y_k = X_k$) the conditional K -property holds trivially, which explains the deceptive simplicity of Algoet's result.

² To be precise, our definitions are time-reversed with respect to the textbook definitions; however, T is a K -automorphism if and only if T^{-1} is a K -automorphism [31, p. 110], and the corresponding statement for weak mixing is trivial. Therefore, our definitions are equivalent to those in [31].

While the conditional ergodicity assumption of Theorem 1.7 is quite general, the requirement that the decision space U is finite is a severe restriction on the complexity of the loss function ℓ . We have stated Theorem 1.7 here in order to highlight the basic ingredients for the existence of a pathwise optimal strategy. The assumption that U is finite will be replaced by various complexity assumptions on the loss ℓ ; such extensions will be developed in the sequel. While some complexity assumption on the loss is needed in the partial information setting, there is a tradeoff between the complexity and ergodicity: if the notion of conditional ergodicity is strengthened, then the complexity assumption on the loss can be weakened.

All our pathwise optimality results are corollaries of a general master theorem, Theorem 2.6 below, that ensures the existence of a pathwise optimal strategy under a certain uniform version of the K -automorphism property. However, in the absence of further assumptions, this theorem does not ensure that the mean-optimal strategy $\tilde{\mathbf{u}}$ is in fact pathwise optimal: the pathwise optimal strategy constructed in the proof may be difficult to compute. We do not know, in general, whether it is possible that a pathwise optimal strategy exists, while the mean-optimal strategy fails to be pathwise optimal. In order to gain further insight into such questions, we introduce another notion of optimality that is intermediate between pathwise and mean optimality. A strategy \mathbf{u}^* is said to be weakly pathwise optimal if

$$\mathbf{P}[L_T(\mathbf{u}) - L_T(\mathbf{u}^*) \geq -\varepsilon] \xrightarrow{T \rightarrow \infty} 1 \quad \text{for every } \varepsilon > 0.$$

It is not difficult to show that if a weakly pathwise optimal strategy exists, then the mean-optimal strategy $\tilde{\mathbf{u}}$ must also be weakly pathwise optimal. However, the notion of weak pathwise optimality is distinctly weaker than pathwise optimality. For example, we will prove the following counterpart to Theorem 1.7.

Theorem 1.8. *Suppose that $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is conditionally weak mixing relative to $\mathcal{Y}_{-\infty, 0}$. Then the mean-optimal strategy $\tilde{\mathbf{u}}$ is weakly pathwise optimal for every loss function $\ell : U \times \Omega \rightarrow \mathbb{R}$ such that U is finite and $|\ell(u, \omega)| \leq \Lambda(\omega)$ with $\Lambda \in L^1$.*

There is a genuine gap between Theorems 1.8 and 1.7: in fact, a result of Conze [6] on individual ergodic theorems for subsequences shows that there is a loss function ℓ such that for a generic (in the weak topology) weak mixing system, a mean-optimal blind strategy $\tilde{\mathbf{u}}$ fails to be pathwise optimal.

While weak pathwise optimality may not be as conceptually appealing as pathwise optimality, the weak pathwise optimality property is easier to characterize. In particular, we will show that the conditional weak mixing assumption in Theorem 1.8 is not only sufficient, but also necessary, in the special case that \mathcal{Y} is an invariant σ -field (that is, $\mathcal{Y} = T^{-1}\mathcal{Y}$). Invariance of \mathcal{Y} is somewhat unnatural in decision problems, as it implies that no additional information is gained over time as more observations are accumulated. On the other hand, invariance of \mathcal{Z} in Definitions 1.5 and 1.6 is precisely the situation of interest in applications of conditional mixing in ergodic theory (e.g., [26]). The interest of this result is therefore that it provides a decision-theoretic characterization of the (conditional) weak mixing property.

1.3 Organization of this paper

The remainder of the paper is organized as follows. In section 2, we state and discuss the main results of this paper. We also give a number of examples that illustrate various aspects of our results. Our main results require two types of assumptions: conditional mixing assumptions on the dynamical system, and complexity assumptions on the loss. In section 3 we discuss various methods to verify these assumptions, as well as further examples and consequences (such as pathwise optimality of nonlinear filters). Finally, the proofs of our main results are given in section 4.

2 Main results

2.1 Basic setup and notation

Throughout this paper, we will consider the following setting:

- $(\Omega, \mathcal{B}, \mathbf{P})$ is a probability space.
- $\mathcal{Y} \subseteq \mathcal{B}$ is a sub- σ -field.
- $T : \Omega \rightarrow \Omega$ is an invertible measure-preserving ergodic transformation.
- (U, \mathcal{U}) is a measurable space.

As explained in the introduction, we aim to make sequential decisions in the ergodic dynamical system $(\Omega, \mathcal{B}, \mathbf{P}, T)$. The decisions take values in the decision space U , and the σ -field \mathcal{Y} represents the observable part of the system. We define

$$\mathcal{Y}_{m,n} = \bigvee_{k=m}^n T^{-k}\mathcal{Y} \quad \text{for } -\infty \leq m \leq n \leq \infty,$$

that is, $\mathcal{Y}_{m,n}$ is the σ -field generated by the observations in the time interval $[m, n]$. An admissible decision strategy must depend causally on the observed data.

Definition 2.1. A strategy $\mathbf{u} = (u_k)_{k \geq 1}$ is called *admissible* if it is $\mathcal{Y}_{0,k}$ -adapted, that is, $u_k : \Omega \rightarrow U$ is $\mathcal{Y}_{0,k}$ -measurable for every $k \geq 1$.

It will be convenient to introduce the following notation. For every $m \leq n$, define

$$\mathbb{U}_{m,n} = \{u : \Omega \rightarrow U : u \text{ is } \mathcal{Y}_{m,n}\text{-measurable}\}, \quad \mathbb{U}_n = \bigcup_{-\infty < m \leq n} \mathbb{U}_{m,n}.$$

Thus a strategy \mathbf{u} is admissible whenever $u_k \in \mathbb{U}_{0,k}$ for all k . Note that $\mathbb{U}_n \subsetneq \mathbb{U}_{-\infty, n}$: this distinction will be essential for the validity of our results.

To describe the loss of a decision strategy, we introduce a loss function ℓ .

- $\ell : U \times \Omega \rightarrow \mathbb{R}$ is a measurable function and $|\ell(u, \omega)| \leq \Lambda(\omega)$ with $\Lambda \in L^1$.

If $|\ell(u, \omega)| \leq \Lambda(\omega)$ with $\Lambda \in L^p$, the loss is said to be dominated in L^p . As indicated above, we will always assume³ that our loss functions are dominated in L^1 .

The loss function $\ell(u, \omega)$ represents the cost incurred by the decision u when the system is in state ω . In particular, the cost of the decision u_k at time k is given by $\ell(u_k, T^k \omega) = \ell_k(u_k)$, where we define for notational simplicity

$$\ell_n(u) : \Omega \rightarrow \mathbb{R}, \quad \ell_n(u)(\omega) = \ell(u, T^n \omega).$$

Our aim is to select an admissible strategy \mathbf{u} that minimizes the time-average loss

$$L_T(\mathbf{u}) = \frac{1}{T} \sum_{k=1}^T \ell_k(u_k)$$

in a suitable sense.

Definition 2.2. An admissible strategy \mathbf{u}^* is *pathwise optimal* if

$$\liminf_{T \rightarrow \infty} \{L_T(\mathbf{u}) - L_T(\mathbf{u}^*)\} \geq 0 \quad \text{a.s.}$$

for every admissible strategy \mathbf{u} .

Definition 2.3. An admissible strategy \mathbf{u}^* is *weakly pathwise optimal* if

$$\mathbf{P}[L_T(\mathbf{u}) - L_T(\mathbf{u}^*) \geq -\varepsilon] \xrightarrow{T \rightarrow \infty} 1 \quad \text{for every } \varepsilon > 0$$

for every admissible strategy \mathbf{u} .

Definition 2.4. An admissible strategy \mathbf{u}^* is *mean optimal* if

$$\liminf_{T \rightarrow \infty} \{\mathbf{E}[L_T(\mathbf{u})] - \mathbf{E}[L_T(\mathbf{u}^*)]\} \geq 0$$

for every admissible strategy \mathbf{u} .

These notions of optimality are progressively weaker: a pathwise optimal strategy is clearly weakly pathwise optimal, and a weakly pathwise optimal strategy is mean optimal (as the loss function is assumed to be dominated in L^1).

In the introduction, it was stated that $\tilde{u}_k = \arg \min_{u \in U} \mathbf{E}[\ell_k(u) | \mathcal{Y}_{0,k}]$ defines a mean-optimal strategy. This disregards some technical issues, as the arg min may not exist or be measurable. It suffices, however, to consider a slight reformulation.

Lemma 2.5. *There exists an admissible strategy $\tilde{\mathbf{u}}$ such that*

$$\mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{Y}_{0,k}] \leq \operatorname{ess\,inf}_{u \in U_{0,k}} \mathbf{E}[\ell_k(u) | \mathcal{Y}_{0,k}] + k^{-1} \quad \text{a.s.}$$

for every $k \geq 1$. In particular, $\tilde{\mathbf{u}}$ is mean-optimal.

³ Non-dominated loss functions may also be of significant interest, see [24] for example. We will restrict attention to dominated loss functions, however, which suffice in many cases of interest.

Proof. It follows from the construction of the essential supremum [25, p. 49] that there exists a countable family $(U^n)_{n \in \mathbb{N}} \subseteq \mathbb{U}_{0,k}$ such that

$$\operatorname{ess\,inf}_{u \in \mathbb{U}_{0,k}} \mathbf{E}[\ell_k(u) | \mathcal{Y}_{0,k}] = \inf_{n \in \mathbb{N}} \mathbf{E}[\ell_k(U^n) | \mathcal{Y}_{0,k}].$$

Define the random variable

$$\tau = \inf \left\{ n : \mathbf{E}[\ell_k(U^n) | \mathcal{Y}_{0,k}] \leq \operatorname{ess\,inf}_{u \in \mathbb{U}_{0,k}} \mathbf{E}[\ell_k(u) | \mathcal{Y}_{0,k}] + k^{-1} \right\}.$$

Note that $\tau < \infty$ a.s. as $\operatorname{ess\,inf}_{u \in \mathbb{U}_{0,k}} \mathbf{E}[\ell_k(u) | \mathcal{Y}_{0,k}] \geq -\mathbf{E}[\Lambda \circ T^k | \mathcal{Y}_{0,k}] > -\infty$ a.s. We therefore define $\tilde{u}_k = U^\tau$. To show that $\tilde{\mathbf{u}}$ is mean optimal, it suffices to note that

$$\mathbf{E}[L_T(\mathbf{u})] - \mathbf{E}[L_T(\tilde{\mathbf{u}})] = \frac{1}{T} \sum_{k=1}^T \mathbf{E} \left[\mathbf{E}[\ell_k(u_k) | \mathcal{Y}_{0,k}] - \mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{Y}_{0,k}] \right] \geq -\frac{1}{T} \sum_{k=1}^T k^{-1}$$

for any admissible strategy \mathbf{u} and $T \geq 1$. \square

In particular, we emphasize that a mean-optimal strategy $\tilde{\mathbf{u}}$ always exists. In the remainder of this paper, we will fix a mean-optimal strategy $\tilde{\mathbf{u}}$ as in Lemma 2.5.

2.2 Pathwise optimality

Our results on the existence of pathwise optimal strategies are all consequences of one general result, Theorem 2.6, that will be stated presently. The essential assumption of this general result is that the properties of the conditional K -automorphism (Definition 1.6) hold uniformly with respect to the loss function ℓ . Note that, in principle, the assumptions of this result do not imply that $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is a conditional K -automorphism, though this will frequently be the case.

Theorem 2.6 (Pathwise optimality). *Suppose that for some σ -field $\mathcal{X} \subset \mathcal{B}$*

1. $\mathcal{X} \subset T^{-1}\mathcal{X}$.
2. *The following martingales converge uniformly:*

$$\operatorname{ess\,sup}_{u \in \mathbb{U}_0} |\mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0} \vee T^{-n}\mathcal{X}] - \ell_0(u)| \xrightarrow{n \rightarrow \infty} 0 \quad \text{in } L^1,$$

$$\operatorname{ess\,sup}_{u \in \mathbb{U}_0} |\mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0} \vee T^n\mathcal{X}] - \mathbf{E}[\ell_0(u) | \bigcap_{k=1}^{\infty} (\mathcal{Y}_{-\infty,0} \vee T^k\mathcal{X})]| \xrightarrow{n \rightarrow \infty} 0 \quad \text{in } L^1.$$

3. *The remote past does not affect the asymptotic loss:*

$$L^* := \mathbf{E} \left[\operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] \right] = \mathbf{E} \left[\operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \bigcap_{k=1}^{\infty} (\mathcal{Y}_{-\infty,0} \vee T^k\mathcal{X})] \right].$$

Then there exists an admissible strategy \mathbf{u}^ such that for every admissible strategy \mathbf{u}*

$$\liminf_{T \rightarrow \infty} \{L_T(\mathbf{u}) - L_T(\mathbf{u}^*)\} \geq 0 \quad a.s., \quad \lim_{T \rightarrow \infty} L_T(\mathbf{u}^*) = L^* \quad a.s.,$$

that is, \mathbf{u}^* is pathwise optimal and L^* is the optimal long time-average loss.

The proof of this result will be given in section 4.1 below.

Before going further, let us discuss the conceptual nature of the assumptions of Theorem 2.6. The assumptions encode two separate requirements:

1. Assumption 3 of Theorem 2.6 should be viewed as a mixing assumption on the dynamical system $(\Omega, \mathcal{B}, \mathbf{P}, T)$ that is tailored to the decision problem. Indeed, $\mathcal{Y}_{-\infty,0}$ represents the information contained in the observations, while $\bigcap_{k=1}^{\infty} (\mathcal{Y}_{-\infty,0} \vee T^k \mathcal{X})$ includes in addition the remote past of the generating σ -field \mathcal{X} . The assumption states that knowledge of the remote past of the unobserved part of the model cannot be used to improve our present decisions.
2. Assumption 2 of Theorem 2.6 should be viewed as a complexity assumption on the loss function ℓ . Indeed, in the absence of the essential suprema, these statements hold automatically by the martingale convergence theorem. The assumption requires that the convergence is in fact uniform in $u \in \mathcal{U}_0$. This will be the case when the loss function is not too complex.

The assumptions of Theorem 2.6 can be verified in many cases of interest. In section 3 below, we will discuss various methods that can be used to verify both the conditional mixing and the complexity assumptions of Theorem 2.6.

In general, neither the conditional mixing nor the complexity assumption can be dispensed with in the presence of partial information.

Example 2.7 (Assumption 3 is essential). We have seen in Examples 1.2 and 1.3 in the introduction that no pathwise optimal strategy exists. In both these examples Assumption 2 is satisfied, that is, the loss function is not too complex (this will follow from general complexity results, cf. Example 3.6 in section 3 below). On the other hand, it is easily seen that the conditional mixing Assumption 3 is violated.

Example 2.8 (Assumption 2 is essential). Let $X = (X_k)_{k \in \mathbb{Z}}$ be the stationary Markov chain in $[0, 1]$ defined by $X_{k+1} = (X_k + \varepsilon_{k+1})/2$ for all k , where $(\varepsilon_k)_{k \in \mathbb{Z}}$ is an i.i.d. sequence of Bernoulli(1/2) random variables. We consider the setting of blind decisions with the loss function $\ell_k(u) = \lfloor 2^u X_k \rfloor \bmod 2$, $u \in U = \mathbb{N}$. Note that

$$X_k = \sum_{i=0}^{\infty} 2^{-i-1} \varepsilon_{k-i}, \quad \ell_k(u) = \varepsilon_{k-u+1}.$$

We claim that no pathwise optimal strategy can exist. Indeed, consider for fixed $r \geq 0$ the strategy \mathbf{u}^r such that $u_k^r = k + r$. Then $\ell_k(u_k^r) = \varepsilon_{1-r}$ for all k . Therefore,

$$\varepsilon_{1-r} - \limsup_{T \rightarrow \infty} L_T(\mathbf{u}^*) = \liminf_{T \rightarrow \infty} \{L_T(\mathbf{u}^r) - L_T(\mathbf{u}^*)\} \geq 0 \quad a.s. \quad \text{for all } r \geq 0$$

for every pathwise optimal strategy \mathbf{u}^* . In particular,

$$0 = \inf_{r \geq 0} \varepsilon_{1-r} \geq \limsup_{T \rightarrow \infty} L_T(\mathbf{u}^*) \geq \liminf_{T \rightarrow \infty} L_T(\mathbf{u}^*) \geq 0 \quad \text{a.s.}$$

As $|L_T(\mathbf{u}^*)| \leq 1$ for all T , it follows by dominated convergence that a pathwise optimal strategy \mathbf{u}^* must satisfy $\mathbf{E}[L_T(\mathbf{u}^*)] \rightarrow 0$ as $T \rightarrow \infty$. But clearly $\mathbf{E}[L_T(\mathbf{u})] = 1/2$ for every T and strategy \mathbf{u} , which entails a contradiction.

Nonetheless, in this example the dynamical system is a K -automorphism (even a Bernoulli shift), so that that Assumption 3 is easily satisfied. As no pathwise optimal strategy exists, this must be caused by the failure of Assumption 2. For example, for the natural choice $\mathcal{X} = \sigma\{X_k : k \leq 0\}$, Assumption 3 holds as $\bigcap_k T^k \mathcal{X}$ is trivial by the Kolmogorov zero-one law, but it is easily seen that the second equation of Assumption 2 fails. Note that the function $l(u, x) = \lfloor 2^u x \rfloor \bmod 2$ becomes increasingly oscillatory as $u \rightarrow \infty$; this is precisely the type of behavior that obstructs uniform convergence in Assumption 2 (akin to “overfitting” in statistics).

Example 2.9 (Assumption 2 is essential, continued). In the previous example, pathwise optimality fails due to failure of the second equation of Assumption 2. We now give a variant of this example where the first equation of Assumption 2 fails.

Let $X = (X_k)_{k \in \mathbb{Z}}$ be an i.i.d. sequence of Bernoulli(1/2) random variables. We consider the setting of blind decisions with the loss function $\ell_k(u) = X_{k+u}$, $u \in U = \mathbb{N}$. We claim that no pathwise optimal strategy can exist. Indeed, consider for $r = 0, 1$ the strategy \mathbf{u}^r defined by $u_k = 2^{r+n+1} - k$ for $2^n \leq k < 2^{n+1}$, $n \geq 0$. Then

$$L_{2^n-1}(\mathbf{u}^r) = \frac{1}{2^n-1} \sum_{m=0}^{n-1} \sum_{k=2^m}^{2^{m+1}-1} X_{k+u_k} = \frac{2^n}{2^n-1} \sum_{m=0}^{n-1} 2^{-(n-m)} X_{2^{r+m+1}}.$$

Suppose that \mathbf{u}^* is pathwise optimal. Then

$$\liminf_{T \rightarrow \infty} \mathbf{E}[L_T(\mathbf{u}^0) \wedge L_T(\mathbf{u}^1) - L_T(\mathbf{u}^*)] \geq \mathbf{E} \left[\liminf_{T \rightarrow \infty} \{L_T(\mathbf{u}^0) \wedge L_T(\mathbf{u}^1) - L_T(\mathbf{u}^*)\} \right] \geq 0.$$

But a simple computation shows that $\mathbf{E}[L_{2^n-1}(\mathbf{u}^0) \wedge L_{2^n-1}(\mathbf{u}^1)]$ converges as $n \rightarrow \infty$ to a quantity strictly less than $1/2 = \mathbf{E}[L_T(\mathbf{u}^*)]$, so that we have a contradiction.

Nonetheless, in this example Assumption 3 and the second line of Assumption 2 are easily satisfied, e.g., for the natural choice $\mathcal{X} = \sigma\{X_k : k \leq 0\}$. However, the first line of Assumption 2 fails, and indeed no pathwise optimal strategy exists.

It is evident from the previous examples that an assumption on both conditional mixing and on complexity of the loss function is needed, in general, to ensure existence of a pathwise optimal strategy. In this light, the complete absence of any such assumptions in the full information case is surprising. The explanation is simple, however: all assumptions of Theorem 2.6 are automatically satisfied in this case.

Example 2.10 (Full information). Let $X = (X_k)_{k \in \mathbb{Z}}$ be any stationary ergodic process, and consider the case of full information: that is, we choose the observation σ -field $\mathcal{Y} = \sigma\{X_0\}$ and the loss $\ell(u, \omega) = l(u, X_1(\omega))$. Then all assumptions of Theorem 2.6 are satisfied: indeed, if we choose $\mathcal{X} = \sigma\{X_k : k \leq 0\}$, then

$\mathcal{Y}_{-\infty,0} = \mathcal{Y}_{-\infty,0} \vee T^k \mathcal{X}$ for all $k \geq 0$, so that Assumption 3 and the second line of Assumption 2 hold trivially. Moreover, $\ell_0(u)$ is $T^{-k} \mathcal{X}$ -measurable for every $u \in \mathbb{U}_0$ and $k \geq 1$, and thus the first line of Assumption 2 holds trivially. It follows that in the full information setting, a pathwise optimal strategy always exists.

In a sense, Theorem 2.6 provides additional insight even in the full information setting: it provides an explanation as to why the case of full information is so much simpler than the partial information setting. Moreover, Theorem 2.6 provides an explicit expression for the optimal asymptotic loss L^* , which is not given in [1].⁴

However, it should be emphasized that Theorem 2.6 does not state that the mean-optimal strategy $\tilde{\mathbf{u}}$ is pathwise optimal; it only guarantees the existence of some pathwise optimal strategy \mathbf{u}^* . In contrast, in the full information setting, Theorem 1.1 ensures pathwise optimality of the mean-optimal strategy. This is of practical importance, as the mean-optimal strategy can in many cases be computed explicitly or by efficient numerical methods, while the pathwise optimal strategy constructed in the proof of Theorem 2.6 may be difficult to compute. We do not know whether it is possible in the general setting of Theorem 2.6 that a pathwise optimal strategy exists, but that the mean-optimal strategy $\tilde{\mathbf{u}}$ is not pathwise optimal. Pathwise optimality of the mean-optimal strategy $\tilde{\mathbf{u}}$ can be shown, however, under somewhat stronger assumptions. The following corollary is proved in section 4.2 below.

Corollary 2.11. *Suppose that for some σ -field $\mathcal{X} \subset \mathcal{B}$*

1. $\mathcal{X} \subset T^{-1} \mathcal{X}$.
2. *The following martingales converge uniformly:*

$$\operatorname{ess\,sup}_{u \in \mathbb{U}_0} |\mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0} \vee T^{-n} \mathcal{X}] - \ell_0(u)| \xrightarrow{n \rightarrow \infty} 0 \quad \text{in } L^1,$$

$$\operatorname{ess\,sup}_{u \in \mathbb{U}_0} |\mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0} \vee T^n \mathcal{X}] - \mathbf{E}[\ell_0(u) | \bigcap_{k=1}^{\infty} (\mathcal{Y}_{-\infty,0} \vee T^k \mathcal{X})]| \xrightarrow{n \rightarrow \infty} 0 \quad \text{in } L^1,$$

$$\operatorname{ess\,sup}_{u \in \mathbb{U}_{-n,0}} |\mathbf{E}[\ell_0(u) | \mathcal{Y}_{-n,0}] - \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}]| \xrightarrow{n \rightarrow \infty} 0 \quad \text{a.s.}$$

3. *The remote past does not affect the present:*

$$\mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] = \mathbf{E}[\ell_0(u) | \bigcap_{k=1}^{\infty} (\mathcal{Y}_{-\infty,0} \vee T^k \mathcal{X})] \quad \text{for all } u \in \mathbb{U}_0.$$

Then the mean-optimal strategy $\tilde{\mathbf{u}}$ (Lemma 2.5) satisfies $L_T(\tilde{\mathbf{u}}) \rightarrow L^$ a.s. as $T \rightarrow \infty$. In particular, it follows from Theorem 2.6 that $\tilde{\mathbf{u}}$ is pathwise optimal.*

⁴ In [1, Appendix II.B] it is shown that under a continuity assumption on the loss function l , the optimal asymptotic loss in the full information setting is given by $\mathbf{E}[\inf_u \mathbf{E}[l(u, X_1) | X_0, X_{-1}, \dots]]$. However, a counterexample is given of a discontinuous loss function for which this expression does not yield the optimal asymptotic loss. The key difference with the expression for L^* given in Theorem 2.6 is that in the latter the essential infimum runs over $u \in \mathbb{U}_0$, while it is implicit in [1] that the infimum in the above expression is an essential infimum over $u \in \mathbb{U}_{-\infty,0}$. As the counterexample in [1] shows, these quantities need not coincide in the absence of continuity assumptions.

The assumptions of Corollary 2.11 are stronger than those of Theorem 2.6 in two respects. First, Assumption 3 is slightly strengthened; however, this is a very mild requirement. More importantly, a third martingale is assumed to converge uniformly (pathwise!) in Assumption 2. The latter is not an innocuous requirement: while the assumption holds in many cases of interest, substantial regularity of the loss function is needed (see section 3.1 for further discussion). In particular, this requirement is not automatically satisfied in the case of full information, and Theorem 1.1 therefore does not follow in its entirety from our results. It remains an open question whether it is possible to establish pathwise optimality of the mean-optimal strategy $\tilde{\mathbf{u}}$ under a substantial weakening of the assumptions of Corollary 2.11.

A particularly simple regularity assumption on the loss is that the decision space U is finite. In this case uniform convergence is immediate, so that the assumptions of Corollary 2.11 reduce essentially to the $\mathcal{Y}_{-\infty,0}$ -conditional K -property. Therefore, evidently Corollary 2.11 implies Theorem 1.7. More general conditions that ensure the validity of the requisite assumptions will be discussed in section 3.

2.3 Weak pathwise optimality

In the previous section, we have seen that a pathwise optimal strategy \mathbf{u}^* exists under general assumptions. However, unlike in the full information case, it is not clear whether in general (without a nontrivial complexity assumption) the mean-optimal strategy $\tilde{\mathbf{u}}$ is pathwise optimal. In the present section, we will aim to obtain some additional insight into this issue by considering the notion of weak pathwise optimality (Definition 2.3) that is intermediate between pathwise optimality and mean optimality. This notion is more regularly behaved than pathwise optimality; in particular, it is straightforward to prove the following simple result.

Lemma 2.12. *Suppose that a weakly pathwise optimal strategy \mathbf{u}^* exists. Then the mean-optimal strategy $\tilde{\mathbf{u}}$ is also weakly pathwise optimal.*

Proof. Let $\Lambda_T = \frac{1}{T} \sum_{k=1}^T \Lambda \circ T^k$. As $|L_T(\mathbf{u})| \leq \Lambda_T$ for any strategy \mathbf{u} , we have

$$\mathbf{E}[(L_T(\tilde{\mathbf{u}}) - L_T(\mathbf{u}^*))_-] \leq \varepsilon \mathbf{P}[L_T(\tilde{\mathbf{u}}) - L_T(\mathbf{u}^*) \geq -\varepsilon] + \mathbf{E}[2\Lambda_T \mathbf{1}_{L_T(\tilde{\mathbf{u}}) - L_T(\mathbf{u}^*) < -\varepsilon}]$$

for any $\varepsilon > 0$. Note that the sequence $(\Lambda_T)_{T \geq 1}$ is uniformly integrable as $\Lambda_T \rightarrow \mathbf{E}[\Lambda]$ in L^1 by the ergodic theorem. Therefore, using weak pathwise optimality of \mathbf{u}^* , it follows that $\mathbf{E}[(L_T(\tilde{\mathbf{u}}) - L_T(\mathbf{u}^*))_-] \rightarrow 0$ as $T \rightarrow \infty$. We therefore have

$$\limsup_{T \rightarrow \infty} \mathbf{E}[|L_T(\tilde{\mathbf{u}}) - L_T(\mathbf{u}^*)|] = -\liminf_{T \rightarrow \infty} \{\mathbf{E}[L_T(\mathbf{u}^*)] - \mathbf{E}[L_T(\tilde{\mathbf{u}})]\} \leq 0$$

by mean-optimality of $\tilde{\mathbf{u}}$. It follows easily that $\tilde{\mathbf{u}}$ is also pathwise optimal. \square

While Theorem 2.6 does not ensure that the mean-optimal strategy $\tilde{\mathbf{u}}$ is pathwise optimal, the previous lemma guarantees that $\tilde{\mathbf{u}}$ is at least weakly pathwise optimal.

However, we will presently show that the latter conclusion may follow under considerably weaker assumptions than those of Theorem 2.6. Indeed, just as pathwise optimality was established for conditional K -automorphisms, we will establish weak optimality for conditionally weakly mixing automorphisms.

Let us begin by developing a general result on weak pathwise optimality, Theorem 2.13 below, that plays the role of Theorem 2.6 in the present setting. The essential assumption of this general result is that the conditional weak mixing property (Definition 1.5) holds uniformly with respect to the loss function ℓ . For simplicity of notation, let us define as in Theorem 2.6 the optimal asymptotic loss

$$L^* := \mathbf{E} \left[\operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] \right]$$

(let us emphasize, however, the Assumption 3 of Theorem 2.6 need not hold in the present setting!) In addition, let us define the modified loss functions

$$\bar{\ell}_0(u) := \ell_0(u) - \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}], \quad \bar{\ell}_0^M(u) := \ell_0(u) \mathbf{1}_{\Lambda \leq M} - \mathbf{E}[\ell_0(u) \mathbf{1}_{\Lambda \leq M} | \mathcal{Y}_{-\infty,0}].$$

The proof of the following theorem will be given in section 4.3.

Theorem 2.13. *Suppose that the uniform conditional mixing assumption*

$$\lim_{M \rightarrow \infty} \limsup_{T \rightarrow \infty} \left\| \frac{1}{T} \sum_{k=1}^T \operatorname{ess\,sup}_{u, u' \in \mathbb{U}_0} |\mathbf{E}[\{\bar{\ell}_0^M(u) \circ T^{-k}\} \bar{\ell}_0^M(u') | \mathcal{Y}_{-\infty,0}]| \right\|_1 = 0$$

holds. Then the mean-optimal strategy $\tilde{\mathbf{u}}$ is weakly pathwise optimal, and the optimal long time-average loss satisfies the ergodic theorem $L_T(\tilde{\mathbf{u}}) \rightarrow L^$ in L^1 .*

Remark 2.14. We have assumed throughout that the loss function ℓ is dominated in L^1 . If the loss is in fact dominated in L^2 , that is, $|\ell(u, \omega)| \leq \Lambda(\omega)$ with $\Lambda \in L^2$, then the assumption of Theorem 2.13 is evidently implied by the natural assumption

$$\frac{1}{T} \sum_{k=1}^T \operatorname{ess\,sup}_{u, u' \in \mathbb{U}_0} |\mathbf{E}[\{\bar{\ell}_0(u) \circ T^{-k}\} \bar{\ell}_0(u') | \mathcal{Y}_{-\infty,0}]| \xrightarrow{T \rightarrow \infty} 0 \quad \text{in } L^1,$$

and in this case $L_T(\tilde{\mathbf{u}}) \rightarrow L^*$ in L^2 (by dominated convergence). The additional truncation in Theorem 2.13 is included only to obtain a result that holds in L^1 .

Conceptually, as in Theorem 2.6, the assumption of Theorem 2.13 combines a conditional mixing assumption and a complexity assumption. Indeed, the conditional weak mixing property relative to $\mathcal{Y}_{-\infty,0}$ (Definition 1.5) implies that

$$\frac{1}{T} \sum_{k=1}^T |\mathbf{E}[\{f \circ T^{-k}\} g | \mathcal{Y}_{-\infty,0}] - \mathbf{E}[f \circ T^{-k} | \mathcal{Y}_{-\infty,0}] \mathbf{E}[g | \mathcal{Y}_{-\infty,0}]| \xrightarrow{T \rightarrow \infty} 0 \quad \text{in } L^1$$

for every $f, g \in L^2$ (indeed, for simple functions f, g this follows directly from the definition, and the claim for general f, g follows by approximation in L^2). Therefore,

in the absence of the essential supremum, the assumption of Theorem 2.13 reduces essentially to the assumption that the dynamical system $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is conditionally weak mixing relative to $\mathcal{Y}_{-\infty, 0}$. However, Theorem 2.13 requires in addition that the convergence in the definition of the conditional weak mixing property holds uniformly with respect to the possible decisions $u \in \mathbb{U}_0$. This will be the case when the loss function ℓ is not too complex (cf. section 3). For example, in the extreme case where the decision space U is finite, uniformity is automatic, and thus Theorem 1.8 in the introduction follows immediately from Theorem 2.13.

Recall that a pathwise optimal strategy is necessarily weakly pathwise optimal. This is reflected, for example, in Theorems 1.7 and 1.8: indeed, note that

$$\begin{aligned} & \|\mathbf{P}[A \cap T^k B | \mathcal{Z}] - \mathbf{P}[A | \mathcal{Z}] \mathbf{P}[T^k B | \mathcal{Z}]\|_1 \\ &= \|\mathbf{E}[\{\mathbf{1}_A - \mathbf{P}[A | \mathcal{Z}]\} \mathbf{1}_{T^k B} | \mathcal{Z}]\|_1 \\ &\leq \|\mathbf{E}[\{\mathbf{1}_A - \mathbf{P}[A | \mathcal{Z}]\} \mathbf{P}[T^k B | T^{k-n} \mathcal{X}] | \mathcal{Z}]\|_1 + \|\mathbf{1}_{T^k B} - \mathbf{P}[T^k B | T^{k-n} \mathcal{X}]\|_1 \\ &\leq \|\mathbf{P}[A | \mathcal{Z} \vee T^{k-n} \mathcal{X}] - \mathbf{P}[A | \mathcal{Z}]\|_1 + \|\mathbf{1}_B - \mathbf{P}[B | T^{-n} \mathcal{X}]\|_1 \end{aligned}$$

for any n, k , so that the conditional K -property implies the conditional weak mixing property (relative to any σ -field \mathcal{Z}) by letting $k \rightarrow \infty$, then $n \rightarrow \infty$. Along the same lines, one can show that a slight variation of the assumptions of Theorem 2.6 imply the assumption of Theorem 2.13 (modulo minor issues of truncation, which could have been absorbed in Theorem 2.6 also at the expense of heavier notation). It is not entirely obvious, at first sight, how far apart the conclusions of our main results really are. For example, in the setting of full information, cf. Example 2.10, the assumption of Theorem 2.13 holds automatically (as then $\bar{\ell}_0^M(u) \circ T^{-k}$ is $\mathcal{Y}_{-\infty, 0}$ -measurable for every $u \in \mathbb{U}_0$ and $k \geq 1$). Moreover, the reader can easily verify that in all the examples we have given where no pathwise optimal strategy exists (Examples 1.2, 1.3, 2.8, 2.9), even the existence of a weakly pathwise optimal strategy fails. It is therefore tempting to assume that in a typical situation where a weakly pathwise optimal strategy exists, there will likely also be a pathwise optimal strategy. The following example, which is a manifestation of a rather surprising result in ergodic theory due to Conze [6], provides some evidence to the contrary.

Example 2.15 (Generic transformations). In this example, we fix the probability space $(\Omega, \mathcal{B}, \mathbf{P})$, where $\Omega = [0, 1]$ with its Borel σ -field \mathcal{B} and the Lebesgue measure \mathbf{P} . We will consider the decision space $U = \{0, 1\}$ and loss function ℓ defined as

$$\ell(u, \omega) = -u(\mathbf{1}_{[0, 1/2]}(\omega) - 1/2) \quad \text{for } (u, \omega) \in U \times \Omega.$$

Moreover, we will consider the setting of blind decisions, that is, \mathcal{Y} is trivial.

We have not yet defined a transformation T . Our aim is to prove the following: *for a generic invertible measure-preserving transformation T , there is a mean-optimal strategy $\hat{\mathbf{u}}$ that is weakly pathwise optimal but not pathwise optimal.* This shows not only that there can be a substantial gap between Theorems 1.7 and 1.8, but that this is in fact the typical situation (at least in the sense of weak topology).

Let us recall some basic notions. Denote by \mathcal{T} the set of all invertible measure-preserving transformations of $(\Omega, \mathcal{B}, \mathbf{P})$. The weak topology on \mathcal{T} is the topology generated by the basic neighborhoods $B(T_0, B, \varepsilon) = \{T \in \mathcal{T} : \mathbf{P}[TB \Delta T_0 B] < \varepsilon\}$ for all $T_0 \in \mathcal{T}$, $B \in \mathcal{B}$, $\varepsilon > 0$. A property is said to hold for a *generic* transformation if it holds for every transformation T in a dense G_δ subset of \mathcal{T} . A well-known result of Halmos [13] states that a generic transformation is weak mixing. Therefore, for a generic transformation, any mean-optimal strategy $\tilde{\mathbf{u}}$ is weakly pathwise optimal by Theorem 1.8. This proves the first part of our statement.

Of course, in the present setting, $\mathbf{E}[\ell_k(u)|\mathcal{Y}_{0,k}] = \mathbf{E}[\ell_k(u)] = 0$ for every decision $u \in U$. Therefore, every admissible strategy \mathbf{u} is mean-optimal, and the optimal mean loss is given by $L^* = 0$, regardless of the choice of transformation $T \in \mathcal{T}$. It is natural to choose a stationary strategy $\tilde{\mathbf{u}}$ (for example, $\tilde{u}_k = 1$ for all k) so that $\lim_{T \rightarrow \infty} L_T(\tilde{\mathbf{u}}) = L^*$ a.s. We will show that for a generic transformation, the strategy $\tilde{\mathbf{u}}$ is not pathwise optimal. To this end, it evidently suffices to find another strategy \mathbf{u} such that $\liminf_{T \rightarrow \infty} L_T(\mathbf{u}) < L^*$ with positive probability.

To this end, we use the following result of Conze that can be read off from the proof of [6, Theorem 5]: there exists a sequence $n_k \uparrow \infty$ with $k/n_k \rightarrow 1/2$ such that for every $0 < \alpha < 1$ and $1/2 < \lambda < 1$, a generic transformation T satisfies

$$\mathbf{P} \left[\limsup_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N \mathbf{1}_{[0,1/2]} \circ T^{n_k} \geq \lambda \right] \geq 1 - \alpha.$$

Define the strategy \mathbf{u} such that $u_n = 1$ if $n = n_k$ for some k , and $u_n = 0$ otherwise. Then, for a generic transformation T , we have with probability at least $1 - \alpha$

$$\liminf_{T \rightarrow \infty} L_{n_T}(\mathbf{u}) = - \limsup_{T \rightarrow \infty} \frac{1}{n_T} \sum_{k=1}^T (\mathbf{1}_{[0,1/2]} \circ T^{n_k} - 1/2) \leq - \frac{2\lambda - 1}{4}.$$

In words, we have shown that for a generic transformation T , the time-average loss of the mean-optimal strategy $\tilde{\mathbf{u}}$ exceeds that of the strategy \mathbf{u} infinitely often by almost $1/4$ with almost unit probability. Thus the mean-optimal strategy $\tilde{\mathbf{u}}$ fails to be pathwise optimal in a very strong sense, and our claim is established.

Example 2.15 only shows that there is a mean-optimal strategy $\tilde{\mathbf{u}}$ that is weakly pathwise optimal but not pathwise optimal. It does not make any statement about whether or not a pathwise optimal strategy \mathbf{u}^* actually exists. However, we do not know of any mechanism that might lead to pathwise optimality in such a setting. We therefore conjecture that for a generic transformation a pathwise optimal strategy in fact fails to exist at all, so that (unlike in the full information setting) pathwise optimality and weak pathwise optimality are really distinct notions.

The result of Conze used in Example 2.15 originates from a deep problem in ergodic theory that aims to understand the validity of individual ergodic theorems for subsequences, cf. [6, 2] and the references therein. A general characterization of such ergodic properties does not appear to exist, which suggests that the pathwise optimality property may be difficult to characterize beyond general sufficient conditions such as Theorem 2.6. In contrast, the weak pathwise optimality property is

much more regularly behaved. The following theorem, which will be proved in section 4.4 below, provides a complete characterization of weak pathwise optimality in the special case that the observation field \mathcal{Y} is invariant.

Theorem 2.16. *Let $(\Omega, \mathcal{B}, \mathbf{P}, T)$ be an ergodic dynamical system, and suppose that $(\Omega, \mathcal{B}, \mathbf{P})$ is a standard probability space and that $\mathcal{Y} \subseteq \mathcal{B}$ is an invariant σ -field (that is, $\mathcal{Y} = T^{-1}\mathcal{Y}$). Then the following are equivalent:*

1. $(\Omega, \mathcal{B}, \mathbf{P}, T)$ conditionally weak mixing relative to \mathcal{Y} .
2. For every bounded loss function $\ell : U \times \Omega \rightarrow \mathbb{R}$ with finite decision space $\text{card } U < \infty$, there exists a weakly pathwise optimal strategy.

The invariance of \mathcal{Y} is automatic in the setting of blind decisions (as \mathcal{Y} is trivial), in which case Theorem 2.16 yields a decision-theoretic characterization of the weak mixing property. In more general observation models, invariance of \mathcal{Y} may be an unnatural requirement from the point of view of decisions under partial information, as it implies that there is no information gain over time. On the other hand, applications of the notion of conditional weak mixing relative to a σ -field \mathcal{Z} in ergodic theory almost always assume that \mathcal{Z} is invariant (e.g., [26]). Theorem 2.16 yields a decision-theoretic interpretation of this property by choosing $\mathcal{Y} = \mathcal{Z}$.

3 Complexity and conditional ergodicity

3.1 Universal complexity assumptions

The goal of this section is to develop complexity assumptions on the loss function ℓ that ensure that the uniform convergence assumptions in our main results hold regardless of any properties of the transformation T or observations \mathcal{Y} . While such universal complexity assumptions are not always necessary (for example, in the full information setting uniform convergence holds regardless of the loss function), they frequently hold in practice and provide easily verifiable conditions that ensure that our results hold in a broad class of decision problems with partial information.

The simplest assumption is Grothendieck's notion of equimeasurability [12].

Definition 3.1. The loss function $\ell : U \times \Omega \rightarrow \mathbb{R}$ on the probability space $(\Omega, \mathcal{B}, \mathbf{P})$ is said to be *equimeasurable* if for every $\varepsilon > 0$, there exists $\Omega_\varepsilon \in \mathcal{B}$ with $\mathbf{P}[\Omega_\varepsilon] \geq 1 - \varepsilon$ such that the class of functions $\{\ell_0(u)\mathbf{1}_{\Omega_\varepsilon} : u \in U\}$ is totally bounded in $L^\infty(\mathbf{P})$.

The beauty of this simple notion is that it ensures uniform convergence of almost anything. In particular, we obtain the following results.

Lemma 3.2. *Suppose that the loss function ℓ is equimeasurable. Then Assumption 2 of Corollary 2.11 holds, and thus Assumption 2 of Theorem 2.6 holds as well, provided that \mathcal{X} is a generating σ -field (that is, $\bigvee_n T^{-n}\mathcal{X} = \mathcal{B}$).*

Proof. Let us establish the first line of Assumption 2. Fix $\varepsilon > 0$ and Ω_ε as in Definition 3.1. Then there exist $N < \infty$ measurable functions $l_1, \dots, l_N : \Omega \rightarrow \mathbb{R}$ such that for every $u \in U$, there exists $k(u) \in \{1, \dots, N\}$ such that

$$\|\ell_0(u)\mathbf{1}_{\Omega_\varepsilon} - l_{k(u)}\mathbf{1}_{\Omega_\varepsilon}\|_\infty \leq \varepsilon$$

(and $u \mapsto k(u)$ can clearly be chosen to be measurable). It follows that

$$\begin{aligned} \operatorname{ess\,sup}_{u \in \mathbb{U}_0} |\mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0} \vee T^{-n}\mathcal{X}] - \ell_0(u)| &\leq \max_{1 \leq k \leq N} |\mathbf{E}[l_k \mathbf{1}_{\Omega_\varepsilon} | \mathcal{Y}_{-\infty,0} \vee T^{-n}\mathcal{X}] - l_k \mathbf{1}_{\Omega_\varepsilon}| \\ &\quad + 2\varepsilon + \mathbf{E}[\Lambda \mathbf{1}_{\Omega_\varepsilon^c} | \mathcal{Y}_{-\infty,0} \vee T^{-n}\mathcal{X}] + \Lambda \mathbf{1}_{\Omega_\varepsilon^c}. \end{aligned}$$

As \mathcal{X} is generating, the martingale convergence theorem gives

$$\limsup_{n \rightarrow \infty} \left\| \operatorname{ess\,sup}_{u \in \mathbb{U}_0} |\mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0} \vee T^{-n}\mathcal{X}] - \ell_0(u)| \right\|_1 \leq 2\varepsilon + \mathbf{E}[2\Lambda \mathbf{1}_{\Omega_\varepsilon^c}].$$

Letting $\varepsilon \downarrow 0$ yields the first line of Assumption 2. The remaining statements of Assumption 2 follow by an essentially identical argument. \square

Lemma 3.3. *Suppose that the following conditional mixing assumption holds:*

$$\lim_{M \rightarrow \infty} \limsup_{T \rightarrow \infty} \left\| \frac{1}{T} \sum_{k=1}^T |\mathbf{E}[\{\bar{\ell}_0^M(u) \circ T^{-k}\} \bar{\ell}_0^M(u') | \mathcal{Y}_{-\infty,0}]| \right\|_1 = 0 \quad \text{for every } u, u' \in U.$$

If the loss function ℓ is equimeasurable, then the assumption of Theorem 2.13 holds.

Proof. The proof is very similar to that of Lemma 3.2 and is therefore omitted. \square

As an immediate consequence of these lemmas, we have:

Corollary 3.4. *The conclusions of Theorems 1.7 and 1.8 remain in force if the assumption that U is finite is replaced by the assumption that ℓ is equimeasurable.*

We now give a simple condition for equimeasurability that suffices in many cases. It is closely related to a result of Mokobodzki (cf. [9, Theorem IX.19]).

Lemma 3.5. *Suppose that U is a compact metric space and that $u \mapsto \ell(u, \omega)$ is continuous for a.e. $\omega \in \Omega$. Then ℓ is equimeasurable.*

Proof. As U is a compact metric space (with metric d), it is certainly separable. Let $U_0 \subseteq U$ be a countable dense set, and define the functions

$$b_n = \sup_{u, u' \in U_0: d(u, u') \leq n^{-1}} |\ell_0(u) - \ell_0(u')|.$$

b_n is measurable, as it is the supremum of countably many random variables. Moreover, for almost every ω , the function $u \mapsto \ell(u, \omega)$ is uniformly continuous (being continuous on a compact metric space). Therefore, $b_n \downarrow 0$ a.s. as $n \rightarrow \infty$.

By Egorov's theorem, there exists for every $\varepsilon > 0$ a set Ω_ε with $\mathbf{P}[\Omega_\varepsilon] \geq 1 - \varepsilon$ such that $\|b_n \mathbf{1}_{\Omega_\varepsilon}\|_\infty \downarrow 0$. We claim that $\{\ell_0(u) \mathbf{1}_{\Omega_\varepsilon} : u \in U\}$ is compact in L^∞ . Indeed, for any sequence $(u_n)_{n \geq 1} \subseteq U$ we may choose a subsequence $(u_{n_k})_{k \geq 1}$ that converges to $u_\infty \in U$. Then for every r , we have $|\ell_0(u_{n_k}) - \ell_0(u_\infty)| \leq b_r$ for all k sufficiently large, and therefore $\|\ell_0(u_{n_k}) \mathbf{1}_{\Omega_\varepsilon} - \ell_0(u_\infty) \mathbf{1}_{\Omega_\varepsilon}\|_\infty \rightarrow 0$. \square

Let us give two standard examples of decision problems (cf. [1, 24]).

Example 3.6 (ℓ_p -prediction). Consider the stochastic process setting (X, Y) , and let f be a bounded function. The aim is, at each time k , to choose a predictor u_k of $f(X_{k+1})$ on the basis of the observation history Y_0, \dots, Y_k . We aim to minimize the pathwise time-average ℓ_p -prediction loss $\frac{1}{T} \sum_{k=1}^T |u_k - f(X_{k+1})|^p$ ($p \geq 1$). This is a particular decision problem with partial information, where the loss function is given by $\ell_0(u) = |u - f(X_1)|^p$ and the decision space is $U = [\inf_x f(x), \sup_x f(x)]$. It is immediate that ℓ is equimeasurable by Lemma 3.5.

Example 3.7 (Log-optimal portfolios). Consider a market with d securities (e.g., $d - 1$ stocks and one bond) whose returns in day k are given by the random variable X_k with values in \mathbb{R}_+^d . The decision space $U = \{p \in \mathbb{R}_+^d : \sum_{i=1}^d p_i = 1\}$ is the simplex: u_k^i represents the fraction of wealth invested in the i th security in day k . The total wealth at time T is therefore given by $\prod_{k=1}^T \langle u_k, X_k \rangle$. We only have access to partial information Y_k in day k , e.g., from news reports. We aim to choose an investment strategy on the basis of the available information that maximizes the wealth, or, equivalently, its growth $\frac{1}{T} \sum_{k=1}^T \log \langle u_k, X_k \rangle$. This corresponds to a decision problem with partial information for the loss function $\ell_0(u) = -\log \langle u, X_0 \rangle$.

In order for the loss to be dominated in L^1 , we impose the mild assumption $\mathbf{E}[\Lambda] < \infty$ with $\Lambda = \sum_{i=1}^d |\log X_0^i|$. We claim that the loss ℓ is then also equimeasurable. Indeed, as $\mathbf{E}[\Lambda] < \infty$, the returns must satisfy $X_0^i > 0$ a.s. for every i . Therefore, equimeasurability follows directly from Lemma 3.5.

As we have seen above, equimeasurability follows easily when the loss function possesses some mild pointwise continuity properties. However, there are situations when this may not be the case. In particular, suppose that $\ell(u, \omega)$ only takes the values 0 and 1, that is, our decisions are sets (as may be the case, for example, in predicting the shape of an oil spill or in sequential classification problems). In such a case, equimeasurability will rarely hold, and it is of interest to investigate alternative complexity assumptions. As we will presently explain, equimeasurability is almost necessary to obtain a universal complexity assumption for Corollary 2.11; however, in the setting of Theorem 2.6, the assumption can be weakened considerably.

The simplicity of the equimeasurability assumption hides the fact that there are two distinct uniformity assumptions in Corollary 2.11: we require uniform convergence of both martingales and reverse martingales, which are quite distinct phenomena (cf. [18, 17]). The uniform convergence of martingales can be restrictive.

Example 3.8 (Uniform martingale convergence). Let $(\mathcal{G}_n)_{n \geq 1}$ be a filtration such that each $\mathcal{G}_n = \sigma\{\pi_n\}$ is generated by a finite measurable partition π_n of the probability space $(\Omega, \mathcal{B}, \mathbf{P})$. Let $L : \mathbb{N} \times \Omega \rightarrow \mathbb{R}$ a bounded function such that $L(u, \cdot)$ is

\mathcal{G}_∞ -measurable for every $u \in \mathbb{N}$. Then $\mathbf{E}[L(u, \cdot) | \mathcal{G}_n] \rightarrow L(u, \cdot)$ a.s. for every u . We claim that if this martingale convergence is in fact uniform, that is,

$$\sup_{u \in \mathbb{N}} |\mathbf{E}[L(u, \cdot) | \mathcal{G}_n] - L(u, \cdot)| \xrightarrow{n \rightarrow \infty} 0 \quad \text{in } L^1,$$

then L must necessarily be equimeasurable. To see this, let us first extract a subsequence $n_k \uparrow \infty$ along which the uniform martingale convergence holds a.s. Fix $\varepsilon > 0$. By Egorov's theorem, there exists a set Ω_ε with $\mathbf{P}[\Omega_\varepsilon] \geq 1 - \varepsilon$ such that

$$\sup_{u \in \mathbb{N}} \|\mathbf{E}[L(u, \cdot) | \mathcal{G}_{n_k}] \mathbf{1}_{\Omega_\varepsilon} - L(u, \cdot) \mathbf{1}_{\Omega_\varepsilon}\|_\infty \xrightarrow{k \rightarrow \infty} 0.$$

Therefore, for every $\alpha > 0$, there exists k such that

$$\sup_{u \in \mathbb{N}} \|\alpha [\alpha^{-1} \mathbf{E}[L(u, \cdot) | \mathcal{G}_{n_k}] \mathbf{1}_{\Omega_\varepsilon}] - L(u, \cdot) \mathbf{1}_{\Omega_\varepsilon}\|_\infty \leq 2\alpha.$$

But as \mathcal{G}_n is finitely generated, we can write

$$\mathbf{E}[L(u, \cdot) | \mathcal{G}_n] \mathbf{1}_{\Omega_\varepsilon} = \sum_{P \in \pi_n} L_{n,u,P} \mathbf{1}_{P \cap \Omega_\varepsilon},$$

with $|L_{n,u,P}| \leq \|L\|_\infty$ for all n, u, P . In particular, $\{\alpha [\alpha^{-1} \mathbf{E}[L(u, \cdot) | \mathcal{G}_n] \mathbf{1}_{\Omega_\varepsilon}] : u \in \mathbb{N}\}$ is a finite family of random variables for every n . We have therefore established that the family $\{L(u, \cdot) \mathbf{1}_{\Omega_\varepsilon} : u \in \mathbb{N}\}$ is totally bounded in L^∞ .

In the context of Corollary 2.11, the previous example can be interpreted as follows. Suppose that the observations are finite-valued, that is, \mathcal{Y} is a finitely generated σ -field. Let us suppose, for simplicity, that the decision space U is countable (the same conclusion holds for general U modulo some measurability issues). Then, if the third line of Assumption 2 in Corollary 2.11 holds, then the conditioned loss $\mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}]$ is necessarily equimeasurable. While it is possible that the conditioned loss is equimeasurable even when the loss ℓ is not (e.g., in the case of blind decisions), this is somewhat unlikely to be the case given a nontrivial observation structure. Therefore, it appears that equimeasurability is almost necessary to obtain universal complexity assumptions in the setting of Corollary 2.11.

The situation is much better in the setting of Theorem 2.6, however. While the first line of Assumption 2 in Theorem 2.6 is still a uniform martingale convergence property, the σ -field \mathcal{X} cannot be finitely generated except in trivial cases. In fact, in many cases the loss ℓ will be $T^{-n}\mathcal{X}$ -measurable for some $n < \infty$, in which case the first line of Assumption 2 is automatically satisfied (in particular, in the stochastic process setting, this will be the case for *finitary* loss $\ell_0(u) = l(u, X_{n_1}, \dots, X_{n_k})$ if we choose $\mathcal{X} = \sigma\{X_k, Y_k : k \leq 0\}$). The remainder of Assumption 2 is a uniform reverse martingale convergence property, which holds under much weaker assumptions.

Definition 3.9. The loss $\ell : U \times \Omega \rightarrow \mathbb{R}$ on (Ω, \mathcal{B}) is said to be *universally bracketing* if for every probability measure \mathbf{P} and $\varepsilon, M > 0$, the family $\{\ell_0(u) \mathbf{1}_{A \leq M} : u \in U\}$ can be covered by finitely many brackets $\{f : g \leq f \leq h\}$ with $\|g - h\|_{L^1(\mathbf{P})} \leq \varepsilon$.

Lemma 3.10. *Let (Ω, \mathcal{B}) be a standard space, and let \mathcal{X}, \mathcal{Y} be countably generated. Suppose the loss ℓ is universally bracketing and finitary (that is, for some $n \in \mathbb{Z}$, $\ell_0(u)$ is $T^{-n}\mathcal{X}$ -measurable for all $u \in U$). Then Assumption 2 of Theorem 2.6 holds.*

Proof. The finitary assumption trivially implies the first line of Assumption 2. The second line follows along the lines of the proof of [17, Corollary 1.4(2 \Rightarrow 7)].⁵ \square

To show that universal bracketing can be much weaker than equimeasurability, we give a simple example in the context of set estimation.

Example 3.11 (Confidence intervals). Consider the stochastic process setting (X, Y) where X takes values in the set $[-1, 1]$, and fix $\varepsilon > 0$. We would like to pin down the value of X_k up to precision ε ; that is, we want to choose $u_k \in [-1, 1]$ as a function of the observations Y_0, \dots, Y_k such that $u_k \leq X_k < u_k + \varepsilon$ as often as possible. This is a partial information decision problem with loss function $\ell_0(u) = \mathbf{1}_{\mathbb{R} \setminus [u, u + \varepsilon]}(X_0)$.

The proof of the universal bracketing property of ℓ is standard. Given \mathbf{P} and $\varepsilon > 0$, we choose $-1 = a_0 < a_1 < \dots < a_n = 1$ (for some finite n) in such a way that $\mathbf{P}[a_i < X_0 < a_{i+1}] \leq \varepsilon$ for all i (note that every atom of X_0 with probability greater than ε is one of the values a_i). Put each function $\ell_0(u)$ such that $u = a_i$ or $u + \varepsilon = a_i$ for some i in its own bracket, and consider the additional brackets $\{f : \mathbf{1}_{\mathbb{R} \setminus [a_{i-1}, a_{j+1}]} \leq f \leq \mathbf{1}_{\mathbb{R} \setminus [a_i, a_j]}\}$ for all $1 \leq i \leq j < n$. Then evidently each of the brackets has diameter not exceeding 2ε , and for every $u \in U$ the function $\ell_0(u)$ is included in one of the brackets thus constructed.

On the other hand, whenever the law of X_0 is not purely atomic, the loss ℓ cannot be equimeasurable. Indeed, as $\|\ell_0(u)\mathbf{1}_{\Omega_\varepsilon} - \ell_0(u')\mathbf{1}_{\Omega_\varepsilon}\|_\infty = 1$ whenever $\ell_0(u)\mathbf{1}_{\Omega_\varepsilon} \neq \ell_0(u')\mathbf{1}_{\Omega_\varepsilon}$, it is impossible for $\{\ell_0(u)\mathbf{1}_{\Omega_\varepsilon} : u \in U\}$ to be totally bounded in L^∞ for any infinite set Ω_ε (and therefore for any set of sufficiently large measure).

In [17] a detailed characterization is given of the universal bracketing property. In particular, it is shown that a uniformly bounded, separable loss ℓ on a standard measurable space is universally bracketing if and only if $\{\ell_0(u) : u \in U\}$ is a universal Glivenko-Cantelli class, that is, a class of functions for which the law of large numbers always holds uniformly. Many useful methods have been developed in empirical process theory to verify this property, cf. [10, 29]. In particular, for a separable $\{0, 1\}$ -valued loss, a very useful sufficient condition is that $\{\ell_0(u) : u \in U\}$ is a Vapnik-Chervonenkis class. We refer to [17, 10, 29] for further details.

3.2 Conditional absolute regularity

In the previous section, we have developed universal complexity assumptions that are applicable regardless of other details of the model. In the present section, we

⁵ The pointwise separability assumption in [17, Corollary 1.4(2 \Rightarrow 7)] is not needed here, as the essential supremum can be reduced to a countable supremum as in the proof of Lemma 2.5.

will in some sense take the opposite approach: we will develop a sufficient condition for a stronger version of the conditional K -property (in the stochastic process setting) under which no complexity assumptions are needed. This shows that there is a tradeoff between mixing and complexity; if the mixing assumption is strengthened, then the complexity assumption can be weakened. An additional advantage of the sufficient condition to be presented is that it is in practice one of the most easily verifiable conditions that ensures the conditional K -property.

In the remainder of this section, we will work in the stochastic process setting. Let (X, Y) be a stationary ergodic process taking values in the Polish space $E \times F$. We define $\mathcal{Y}_{n,m} = \sigma\{Y_k : n \leq k \leq m\}$ and $\mathcal{X}_{n,m} = \sigma\{X_k : n \leq k \leq m\}$ for $n \leq m$, and we consider the observation and generating fields $\mathcal{Y} = \sigma\{Y_0\}$, $\mathcal{X} = \mathcal{X}_{-\infty,0} \vee \mathcal{Y}_{-\infty,0}$. In this setting, the conditional K -property relative to $\mathcal{Y}_{-\infty,0}$ reduces to

$$\bigcap_{k=1}^{\infty} (\mathcal{Y}_{-\infty,0} \vee \mathcal{X}_{-\infty,-k}) = \mathcal{Y}_{-\infty,0} \quad \text{mod } \mathbf{P}.$$

If \mathcal{Y} is trivial (that is, the observations Y are noninformative), this reduces to the statement that X has a trivial past tail σ -field, that is, X is regular (or purely nondeterministic) in the sense of Kolmogorov. This property is often fairly easy to check: for example, any Markov chain whose law converges weakly to a unique invariant measure is regular (cf. [28, Prop. 3]). When \mathcal{Y} is nontrivial, the conditional K -property is generally not so easy to check, however. We therefore give a condition, arising from filtering theory [27], that allows to deduce conditional mixing properties from their more easily verifiable unconditional counterparts.

We will require two assumptions. The first assumption states that the pair (X, Y) is absolutely regular in the sense of Volkonskiĭ and Rozanov [30] (this property is also known as β -mixing). Absolute regularity is a strengthening of the regularity property; assuming regularity of (X, Y) is not sufficient for what follows [16]. Many techniques have been developed to verify the absolute regularity property; for example, any Harris recurrent and aperiodic Markov chain is absolutely regular [22].

Definition 3.12. The process (X, Y) is said to be *absolutely regular* if

$$\left\| \mathbf{P}[(X_k, Y_k)_{k \geq n} \in \cdot \mid \mathcal{X}_{-\infty,0} \vee \mathcal{Y}_{-\infty,0}] - \mathbf{P}[(X_k, Y_k)_{k \geq n} \in \cdot] \right\|_{\text{TV}} \xrightarrow{n \rightarrow \infty} 0 \quad \text{in } L^1.$$

By itself, however, absolute regularity of (X, Y) is not sufficient for the conditional K -property, as can be seen in Example 1.3. In this example, the relation between the processes X and Y is very singular, so that things go wrong when we condition. The following nondegeneracy assumption rules out this possibility.

Definition 3.13. The process (X, Y) is said to be *nondegenerate* if

$$\mathbf{P}[Y_1, \dots, Y_m \in \cdot \mid \mathcal{Z}_{-\infty,0} \vee \mathcal{Z}_{m+1,\infty}] \sim \mathbf{P}[Y_1, \dots, Y_m \in \cdot \mid \mathcal{Y}_{-\infty,0} \vee \mathcal{Y}_{m+1,\infty}] \quad \text{a.s.}$$

for every $1 \leq m < \infty$, where $\mathcal{Z}_{n,m} := \mathcal{X}_{n,m} \vee \mathcal{Y}_{n,m}$.

The nondegeneracy assumption ensures that the null sets of the law of the observations Y do not depend too much on the unobserved process X . The assumption is often easily verified. For example, if $Y_k = f(X_k) + \eta_k$ where η_k is an i.i.d. sequence of random variables with strictly positive density, then the conditional distributions in Definition 3.13 have strictly positive densities and are therefore equivalent a.s.

Theorem 3.14 ([27]). *If (X, Y) is absolutely regular and nondegenerate, then*

$$\bigcap_{k=1}^{\infty} (\mathcal{Y}_{-\infty,0} \vee \mathcal{X}_{-\infty,-k}) = \mathcal{Y}_{-\infty,0} \quad \text{mod } \mathbf{P}.$$

Theorem 3.14 provides a practical method to check the conditional K -property. However, the proof of Theorem 3.14 actually yields a much stronger statement. It is shown in [27, Theorem 3.5] that if (X, Y) is absolutely regular and nondegenerate, then X is *conditionally* absolutely regular relative to $\mathcal{Y}_{-\infty,\infty}$ in the sense that

$$\left\| \mathbf{P}[(X_k)_{k \geq n} \in \cdot \mid \mathcal{X}_{-\infty,0} \vee \mathcal{Y}_{-\infty,\infty}] - \mathbf{P}[(X_k)_{k \geq n} \in \cdot \mid \mathcal{Y}_{-\infty,\infty}] \right\|_{\text{TV}} \xrightarrow{n \rightarrow \infty} 0 \quad \text{in } L^1.$$

Moreover, it is shown⁶ that under the same assumptions [27, Proposition 3.9]

$$\mathbf{P}[(X_k)_{k \leq 0} \in \cdot \mid \mathcal{Y}_{-\infty,0}] \sim \mathbf{P}[(X_k)_{k \leq 0} \in \cdot \mid \mathcal{Y}_{-\infty,\infty}] \quad \text{a.s.}$$

From these properties, we can deduce the following result.

Theorem 3.15. *In the setting of the present section, suppose that (X, Y) is absolutely regular and nondegenerate, and consider a loss function of the form $\ell_0(u) = l(u, X_0)$. Then the conclusions of Theorem 2.6 hold.*

The key point about Theorem 3.15 is that no complexity assumption is imposed: the loss function $l(u, x)$ may be an arbitrary measurable function (as long as it is dominated in L^1 in accordance with our standing assumption). The explanation for this is that the conditional absolute regularity property is so strong that the regular conditional probabilities $\mathbf{P}[X_0 \in \cdot \mid \mathcal{Y}_{-\infty,\infty} \vee \mathcal{X}_{-\infty,-n}]$ converge in total variation. Therefore, the corresponding reverse martingales converge uniformly over any dominated family of measurable functions. The strength of the conditional mixing property therefore eliminates the need for any additional complexity assumptions. In contrast, we may certainly have pathwise optimal strategies when absolute regularity fails, but then a complexity assumption is essential (cf. Example 2.8).

The proof of Theorem 3.15 will be given in section 4.5. The proof is a straightforward adaptation of Theorem 2.6; unfortunately, the fact that the conditional absolute regularity property is relative to $\mathcal{Y}_{-\infty,\infty}$ rather than $\mathcal{Y}_{-\infty,0}$ complicates a direct verification of the assumptions of Theorem 2.6 (while this should be possible along the lines of [27], we will follow the simpler route here). The results of [27] could also

⁶ Some of the statements in [27] are time-reversed as compared to their counterparts stated here. However, as both the absolute regularity and the nondegeneracy assumptions are invariant under time reversal (cf. [30] for the former; the latter is trivial), the present statements follow immediately.

be used to obtain the conclusion of Corollary 2.11 in the setting of Theorem 3.15 under somewhat stronger nondegeneracy assumptions.

3.3 Hidden Markov models and nonlinear filters

The goal of the present section is to explore some implications of our results to filtering theory. For simplicity of exposition, we will restrict attention to the classical setting of (general state space) hidden Markov models (see, e.g., [4]).

We adopt the stochastic process setting and notations of the previous section. In addition, we assume that (X, Y) is a hidden Markov model, that is, a Markov chain whose transition kernel can be factored as $\tilde{P}(x, y, dx', dy') = P(x, dx') \Phi(x', dy')$. This implies that the process X is a Markov chain in its own right, and that the observations Y are conditionally independent given X . In the following, we will assume that the observation kernel Φ has a density, that is, $\Phi(x, dy) = g(x, y) \varphi(dy)$ for some measurable function g and reference measure φ .

A fundamental object in this theory is the nonlinear filter Π_k , defined as

$$\Pi_k := \mathbf{P}[X_k \in \cdot | Y_0, \dots, Y_k].$$

The measure-valued process $\Pi = (\Pi_k)_{k \geq 0}$ is itself a (nonstationary) Markov chain [16] with transition kernel \mathcal{P} . To study the stationary behavior of the filter, which is of substantial interest in applications (see, for example, [15] and the references therein), one must understand the relationship between the ergodic properties of X and Π . The following result, proved in [16], is essentially due to Kunita [20].

Theorem 3.16. *Suppose that the transition kernel P possesses a unique invariant measure (that is, X is uniquely ergodic). Then the filter transition kernel \mathcal{P} possesses a unique invariant measure (that is, Π is uniquely ergodic) if and only if*

$$\bigcap_{k=1}^{\infty} (\mathcal{Y}_{-\infty, 0} \vee \mathcal{X}_{-\infty, -k}) = \mathcal{Y}_{-\infty, 0} \quad \text{mod } \mathbf{P}.$$

Evidently, ergodicity of the filter is closely related to the conditional K -property. We will exploit this fact to prove a new optimality property of nonlinear filters.

The usual interpretation of the filter Π_k is that one aims to track to current location X_k of the unobserved process on the basis of the observation history Y_0, \dots, Y_k . By the elementary property of conditional expectations, $\Pi_k(f)$ provides, for any bounded test function f , an optimal mean-square error estimate of $f(X_k)$:

$$\mathbf{E}[\{f(X_k) - \Pi_k(f)\}^2] \leq \mathbf{E}[\{f(X_k) - \hat{f}_k(Y_0, \dots, Y_k)\}^2] \quad \text{for any measurable } \hat{f}_k.$$

This interpretation may not be satisfying, however, if only one sample path of the observations is available (recall Examples 1.2 and 1.3): one would rather show that

$$\liminf_{T \rightarrow \infty} \left[\frac{1}{T} \sum_{k=1}^T \{f(X_k) - \hat{f}_k(Y_0, \dots, Y_k)\}^2 - \frac{1}{T} \sum_{k=1}^T \{f(X_k) - \Pi_k(f)\}^2 \right] \geq 0 \quad \text{a.s.}$$

for any alternative sequence of estimators $(\hat{f}_k)_{k \geq 0}$. If this property holds for any bounded test function f , the filter will be said to be *pathwise optimal*.

Corollary 3.17. *Suppose that the filtering process Π is uniquely ergodic. Then the filter is both mean-square optimal and pathwise optimal.*

Proof. Note that the filter $\Pi_k(f)$ is the mean-optimal policy for the partial information decision problem with loss $\ell_0(u) = \{f(X_0) - u\}^2$. As the latter is equimeasurable, the result follows directly from Theorem 3.16 and Corollary 2.11. \square

The interaction between our main results and the ergodic theory of nonlinear filters is therefore twofold. On the one hand, our main results imply that ergodic nonlinear filters are always pathwise optimal. Conversely, Theorem 3.16 shows that ergodicity of the filter is a sufficient condition for our main results to hold in the context of hidden Markov models with equimeasurable loss. This provides another route to establishing the conditional K -property: the filtering literature provides a variety of methods to verify ergodicity of the filter [14, 5, 7, 27]. It should be noted, however, that ergodicity of the filter is not necessary for the conditional K -property to hold, even in the setting of hidden Markov models.

Example 3.18. Consider the hidden Markov model (X, Y) where X is the stationary Markov chain such that $X_0 \sim \text{Uniform}([0, 1])$ and $X_{k+1} = 2X_k \bmod 1$, $Y_k = 0$ for all $k \in \mathbb{Z}$ (that is, we have noninformative observations). Clearly the tail σ -field $\bigcap_n \mathcal{X}_{-\infty, n}$ is nontrivial, and thus the filter fails to be ergodic by Theorem 3.16. Nonetheless, we claim that the conditional K -property holds, so that our main results apply for any equimeasurable loss; in particular, the filter is pathwise optimal.

The key point is that, even in the hidden Markov model setting, one need not choose the “canonical” generating σ -field $\mathcal{X} = \mathcal{X}_{-\infty, 0}$ in Definition 1.6. In the present example, we choose instead $\mathcal{X} = \sigma\{\mathbf{1}_{X_k > 1/2} : k \leq 0\}$. To verify the conditional K -property, note that $(\mathbf{1}_{X_k > 1/2})_{k \in \mathbb{Z}}$ are i.i.d. Bernoulli(1/2) random variables and

$$X_k = \sum_{\ell=0}^{\infty} 2^{-\ell-1} \mathbf{1}_{X_{k+\ell} > 1/2} \quad \text{a.s. for all } k \in \mathbb{Z}.$$

Thus $\mathcal{X} \subset T^{-1}\mathcal{X}$ by construction, $\bigvee_k T^{-k}\mathcal{X} = \sigma\{X_n : n \in \mathbb{Z}\}$ is a generating σ -field, and $\bigcap_k T^k\mathcal{X}$ is trivial by the Kolmogorov zero-one law.

Let us now consider the decision problem in the setting of a hidden Markov model with equimeasurable loss function $\ell_0(u) = l(u, X_0)$. If the filter is ergodic, then Corollary 3.4 ensures that the mean-optimal strategy $\tilde{\mathbf{u}}$ is pathwise optimal. In this setting, the mean-optimal strategy can be expressed in terms of the filter:

$$\tilde{u}_k = \arg \min_{u \in U} \mathbf{E}[l(u, X_k) | Y_0, \dots, Y_k] = \arg \min_{u \in U} \int l(u, x) \Pi_k(dx).$$

When X_k takes values in a finite set $E = \{1, \dots, d\}$, the filter can be recursively computed in a straightforward manner [4]. In this case, the mean-optimal strategy $\tilde{\mathbf{u}}$ can be implemented directly. On the other hand, when E is a continuous space, the conditional measure Π_k is an infinite-dimensional object which cannot be computed exactly except in special cases. However, Π_k can often be approximated very efficiently by recursive Monte Carlo approximations $\Pi_k^N = \frac{1}{N} \sum_{i=1}^N \delta_{Z_k^N(i)}$, known as particle filters [4], that converge to the true filter Π_k as the number of particles increases $N \rightarrow \infty$. This suggests to approximate the mean-optimal strategy $\tilde{\mathbf{u}}$ by

$$\tilde{u}_k \approx \tilde{u}_k^N := \arg \min_{u \in U} \int l(u, x) \Pi_k^N(dx) = \arg \min_{u \in U} \frac{1}{N} \sum_{i=1}^N l(u, Z_k^N(i)).$$

The strategy $\tilde{\mathbf{u}}^N$ is a type of sequential stochastic programming algorithm to approximate the mean-optimal strategy. In this setting, it is of interest to establish whether the strategy $\tilde{\mathbf{u}}^N$ is in fact approximately pathwise optimal, at least in the weak sense. To this end, we prove the following approximation lemma.

Lemma 3.19. *In the hidden Markov model setting with equimeasurable loss $\ell_0(u) = l(u, X_0)$, suppose that the filter is ergodic, and let Π_k^N be an approximation of Π_k . If*

$$\lim_{N \rightarrow \infty} \limsup_{T \rightarrow \infty} \mathbf{E} \left[\frac{1}{T} \sum_{k=1}^T \operatorname{ess\,sup}_{u \in \mathbb{U}_{0,k}} |\Pi_k^N(l(u, \cdot)) - \Pi_k(l(u, \cdot))| \right] = 0,$$

then the strategy $\tilde{\mathbf{u}}^N$ is approximately weakly pathwise optimal in the sense that

$$\lim_{N \rightarrow \infty} \liminf_{T \rightarrow \infty} \mathbf{P}[L_T(\mathbf{u}) - L_T(\tilde{\mathbf{u}}^N) \geq -\varepsilon] = 1 \quad \text{for every } \varepsilon > 0$$

holds for every admissible strategy \mathbf{u} .

Proof. We begin by noting that

$$\mathbf{P}[L_T(\mathbf{u}) - L_T(\tilde{\mathbf{u}}^N) < -\varepsilon] \leq \mathbf{P}[L_T(\mathbf{u}) - L_T(\tilde{\mathbf{u}}) < -\varepsilon/2] + \mathbf{P}[L_T(\tilde{\mathbf{u}}^N) - L_T(\tilde{\mathbf{u}}) > \varepsilon/2].$$

Under the present assumptions, the mean-optimal strategy $\tilde{\mathbf{u}}$ is (weakly) pathwise optimal. It follows⁷ as in the proof of Lemma 2.12 that $\mathbf{E}[(L_T(\tilde{\mathbf{u}}^N) - L_T(\tilde{\mathbf{u}}))_-] \rightarrow 0$ as $T \rightarrow \infty$, and we obtain for any admissible strategy \mathbf{u} and $\varepsilon > 0$

$$\limsup_{T \rightarrow \infty} \mathbf{P}[L_T(\mathbf{u}) - L_T(\tilde{\mathbf{u}}^N) < -\varepsilon] \leq \frac{2}{\varepsilon} \limsup_{T \rightarrow \infty} \mathbf{E}[L_T(\tilde{\mathbf{u}}^N) - L_T(\tilde{\mathbf{u}})].$$

To proceed, we estimate

⁷ As particle filters employ a random sampling mechanism, the strategy $\tilde{\mathbf{u}}^N$ is technically speaking not admissible in the sense of this paper: Π_k^N (and therefore \tilde{u}_k^N) depends also on auxiliary sampling variables ξ_0, \dots, ξ_k that are independent of Y_0, \dots, Y_k . However, it is easily seen that all our results still hold when such *randomized* strategies are considered. Indeed, it suffices to condition on $(\xi_k)_{k \geq 0}$, so that all our results apply immediately under the conditional distribution.

$$\begin{aligned}
\mathbf{E}[L_T(\tilde{\mathbf{u}}^N) - L_T(\tilde{\mathbf{u}})] &= \mathbf{E} \left[\frac{1}{T} \sum_{k=1}^T \int \{l(\tilde{u}_k^N, x) - l(\tilde{u}_k, x)\} \Pi_k(dx) \right] \\
&\leq \mathbf{E} \left[\frac{1}{T} \sum_{k=1}^T \int \{l(\tilde{u}_k^N, x) - l(\tilde{u}_k, x)\} \Pi_k^N(dx) \right] \\
&\quad + 2\mathbf{E} \left[\frac{1}{T} \sum_{k=1}^T \operatorname{ess\,sup}_{u \in \mathbb{U}_{0,k}} |\Pi_k^N(l(u, \cdot)) - \Pi_k(l(u, \cdot))| \right].
\end{aligned}$$

But note that by the definition of $\tilde{\mathbf{u}}^N$

$$\int \{l(\tilde{u}_k^N, x) - l(\tilde{u}_k, x)\} \Pi_k^N(dx) = \inf_{u \in \mathbb{U}} \int l(u, x) \Pi_k^N(dx) - \int l(\tilde{u}_k, x) \Pi_k^N(dx) \leq 0.$$

The proof is therefore easily completed by applying the assumption. \square

Evidently, the key difficulty in this problem is to control the time-average error of the filter approximation (in a norm determined by the loss function l) uniformly over the time horizon. This problem is intimately related with the ergodic theory of nonlinear filters. The requisite property follows from the results in [15] under reasonable ergodicity assumptions but under very stringent complexity assumptions on the loss (essentially that $\{l(u, \cdot) : u \in \mathbb{U}\}$ is uniformly Lipschitz). Alternatively, one can apply the results in [8], which require exceedingly strong ergodicity assumptions but weaker complexity assumptions. Let us note that one could similarly obtain a pathwise version of Lemma 3.19, but the requisite pathwise approximation property of particle filters has not been investigated in the literature.

3.4 The conditions of Algoet, Weissman, Merhav, and Nobel

The aim of this section is to briefly discuss the assumptions imposed in previous work on the pathwise optimality property due to Algoet [1], Weissman and Merhav [32], and Nobel [24]. Let us emphasize that, while our results cover a much broader range of decision problems, none of these previous results follow in their entirety from our general results. This highlights once more that our results are, unfortunately, nowhere close to a complete characterization of the pathwise optimality property.

3.4.1 Algoet

Algoet's results [1], which cover the full information setting only, were already discussed at length in the introduction and in section 2.2. The existence of a pathwise optimal strategy can be obtained in this setting under no additional assumptions from Theorem 2.6, which even goes beyond Algoet's result in that it gives an explicit

expression for the optimal asymptotic loss. However, Algoet establishes that in fact the mean-optimal strategy $\tilde{\mathbf{u}}$ is pathwise optimal in this setting, while our general Corollary 2.11 can only establish this under an additional complexity assumption. We do not know whether this complexity assumption can be weakened in general.

3.4.2 Weissman and Merhav

Weissman and Merhav [32] consider the stochastic process setting (X, Y) , where X_k takes values in $\{0, 1\}$ and Y_k takes values in \mathbb{R} for all $k \in \mathbb{Z}$, and where the loss function takes the form $\ell_0(u) = l(u, X_1)$ and is assumed to be uniformly bounded. As X is binary-valued, it is immediate that any loss function l is equimeasurable. Therefore, our results show that the mean-optimal strategy $\tilde{\mathbf{u}}$ is pathwise optimal whenever the model is a conditional K -automorphism relative to $\mathcal{Y}_{-\infty, 0}$.

The assumption imposed by Weissman and Merhav in [32] is as follows:

$$\sum_{k=1}^{\infty} \sup_{r \geq 1} \mathbf{E}[|\mathbf{P}[X_{r+k} = a | X_r = a, \mathcal{Y}_{0, r+k-1}] - \mathbf{P}[X_{r+k} = a | \mathcal{Y}_{0, r+k-1}]|] < \infty \text{ for } a = 0, 1.$$

Using stationarity, this condition is equivalent to

$$\sum_{k=0}^{\infty} \sup_{r \geq 1} \mathbf{E}[|\mathbf{P}[X_1 = a | X_{-k} = a, \mathcal{Y}_{-r-k, 0}] - \mathbf{P}[X_1 = a | \mathcal{Y}_{-r-k, 0}]|] < \infty \text{ for } a = 0, 1,$$

which readily implies

$$\sum_{k=0}^{\infty} \mathbf{E}[|\mathbf{P}[X_1 = a | \sigma\{X_{-k}\} \vee \mathcal{Y}_{-\infty, 0}] - \mathbf{P}[X_1 = a | \mathcal{Y}_{-\infty, 0}]|] < \infty.$$

If the σ -field $\sigma\{X_{-k}\} \vee \mathcal{Y}_{-\infty, 0}$ could be replaced by the larger σ -field $\mathcal{X}_{-\infty, -k} \vee \mathcal{Y}_{-\infty, 0}$ in this expression, then Assumption 3 of Corollary 2.11 would follow immediately. However, the smaller σ -field appears to yield a slightly better variant of the assumption imposed in [32]. This is possible because the result is restricted to the special choice of loss $\ell_0(u) = l(u, X_1)$ that depends on X_1 only. On the other hand, it is to be expected that in most cases the assumption of [32] is much more stringent than that of Corollary 2.11. Note that Assumption 3 of Corollary 2.11 is purely qualitative in nature: it states, roughly speaking, that two σ -fields coincide. This is a structural property of the model. On the other hand, the assumption of [32] is inherently *quantitative* in nature: it requires that a certain mixing property holds at a sufficiently fast rate (the mixing coefficients must be summable). A quantitative bound on the mixing rate is both much more restrictive and much harder to verify, in general, as compared to a purely structural property.

In a sense, the approach of Weissman and Merhav is much closer in spirit to the weak pathwise optimality results in this paper than it is to the pathwise optimality results. Indeed, if we replace the weak pathwise optimality property

$$\mathbf{P}[L_T(\mathbf{u}) - L_T(\mathbf{u}^*) < -\varepsilon] \xrightarrow{T \rightarrow \infty} 0 \quad \text{for every } \varepsilon > 0$$

by its quantitative counterpart

$$\sum_{T=1}^{\infty} \mathbf{P}[L_T(\mathbf{u}) - L_T(\mathbf{u}^*) < -\varepsilon] < \infty \quad \text{for every } \varepsilon > 0,$$

then pathwise optimality will automatically follow from the Borel-Cantelli lemma. In the same spirit, if in Theorem 2.13 we replace the uniform conditional mixing assumption by the corresponding quantitative counterpart

$$\sum_{k=1}^T \mathbf{E} \left[\text{ess sup}_{u, u' \in \mathbb{U}_0} |\mathbf{E}[\{\bar{\ell}_0^M(u) \circ T^{-k}\} \bar{\ell}_0^M(u') | \mathcal{Y}_{-\infty, 0}]| \right] = O(T^\alpha)$$

for some $\alpha < 1$ (that may depend on M), then we easily obtain a pathwise version of Lemma 4.7 below (using Etemadi's well-known device [11]), and consequently the conclusion of Theorem 2.13 is replaced by that of Theorem 2.6. It is unclear whether such quantitative mixing conditions provide a distinct mechanism for pathwise optimality as compared to qualitative structural conditions as in our main results.

3.4.3 Nobel

Nobel [24] considers the stochastic process setting (X, Y) with observations of the additive form $Y_k = X_k + N_k$, where $N = (N_k)_{k \in \mathbb{Z}}$ is an L^2 -martingale difference sequence independent of X . The loss function considered is the mean-square loss $\ell_0(u) = (u - X_1)^2$. This very special scenario is essential for the result given in [24]; on the other hand, it is not assumed that (X, Y) is even stationary or that the decision space U is a compact set (when $U = \mathbb{R}$, the quadratic loss is not dominated). In order to compare with our general results, we will additionally assume that (X, Y) is stationary and ergodic and that X_k are uniformly bounded random variables (so that we may choose $U = [-\|X_1\|_\infty, \|X_1\|_\infty]$ without loss of generality).

While this is certainly a decision problem with partial information, the key observation is that this special problem is in fact a decision problem with full information in disguise. Indeed, note that we can write for any strategy \mathbf{u}

$$L_T(\mathbf{u}) = \frac{1}{T} \sum_{k=1}^T (u_k - Y_{k+1})^2 + \frac{1}{T} \sum_{k=1}^T \{X_{k+1}^2 - Y_{k+1}^2\} + \frac{1}{T} \sum_{k=1}^T 2u_k N_{k+1}.$$

The last term of this expression converges to zero a.s. as $T \rightarrow \infty$ for any admissible strategy \mathbf{u} by the martingale law of large numbers, as $(u_k N_{k+1})_{k \in \mathbb{Z}}$ is an L^2 -martingale difference sequence. On the other hand, the second to last term of this expression does not depend on the strategy \mathbf{u} at all. Therefore,

$$\liminf_{T \rightarrow \infty} \{L_T(\mathbf{u}) - L_T(\tilde{\mathbf{u}})\} = \liminf_{T \rightarrow \infty} \left\{ \frac{1}{T} \sum_{k=1}^T (u_k - Y_{k+1})^2 - \frac{1}{T} \sum_{k=1}^T (\tilde{u}_k - Y_{k+1})^2 \right\} \quad \text{a.s.},$$

which corresponds to the decision problem with the full information loss $\ell_0(u) = (u - Y_1)^2$. Thus pathwise optimality of the mean-optimal strategy $\tilde{\mathbf{u}}$ follows from Algoet's result. (The main difficulty in [24] is to introduce suitable truncations to deal with the lack of boundedness, which we avoided here.)

Of course, we could deduce the result from our general theory in the same manner: reduce first to a full information decision problem as above, and then invoke Corollary 2.11 in the full information setting. However, a more relevant test of our general theory might be to ask whether one can deduce the result directly from Corollary 2.11, without first reducing to the full information setting. Unfortunately, it is not clear whether it is possible, in general, to find a generating σ -field \mathcal{X} such that Assumption 3 of Corollary 2.11 holds.

One might interpret the additive noise model as a type of "informative" observations: while X cannot be reconstructed from the observations Y , the law of X can certainly be reconstructed from the law of Y even if the former were not known *a priori* (this idea is exploited in [32, 24] to devise universal prediction strategies that do not require prior knowledge of the law of X). In the hidden Markov model setting, there is in fact a connection between "informative" observations and the conditional K -property. In particular, if (X, Y) is a hidden Markov model where X_k takes a finite number of values, and $Y_k = X_k + \xi_k$ where ξ_k are i.i.d. and independent of X , then the conditional K -property holds, and we therefore have pathwise optimal strategies for *any* dominated loss. This follows from observability conditions in the Markov setting, cf. [5, section 6.2] and the references therein. However, the ideas that lead to this result do not appear to extend to more general situations.

4 Proofs

4.1 Proof of Theorem 2.6

Throughout the proof, we fix a generating σ -field \mathcal{X} that satisfies the conditions of Theorem 2.6. In the following, we define the σ -fields

$$\mathcal{G}_k^n = \mathcal{Y}_{-\infty, k} \vee T^{n-k} \mathcal{X}, \quad \mathcal{G}_k^\infty = \bigcap_n \mathcal{G}_k^n.$$

Note that \mathcal{G}_k^n is decreasing in n and increasing in k .

We begin by establishing the following lemma.

Lemma 4.1. *For any admissible strategy \mathbf{u} and any $m, n \in \mathbb{Z}$*

$$\frac{1}{T} \sum_{k=1}^T \{ \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^m] - \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^n] \} \xrightarrow{T \rightarrow \infty} 0 \quad \text{a.s.}$$

Proof. Assume $m < n$ without loss of generality. Fix $r < \infty$, and define

$$\Delta_k^j = \mathbf{E}[\ell_k(u_k) \mathbf{1}_{\Lambda \circ T^k \leq r} | \mathcal{G}_k^j] - \mathbf{E}[\ell_k(u_k) \mathbf{1}_{\Lambda \circ T^k \leq r} | \mathcal{G}_k^{j+1}]$$

for $m \leq j < n$. Then it is easily seen that we have the inequality

$$\left| \frac{1}{T} \sum_{k=1}^T \{ \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^m] - \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^n] \} \right| \leq \sum_{j=m}^{n-1} \left| \frac{1}{T} \sum_{k=1}^T \Delta_k^j \right| + \frac{1}{T} \sum_{k=1}^T \{ \mathbf{E}[\Lambda \mathbf{1}_{\Lambda > r} | \mathcal{G}_0^m] + \mathbf{E}[\Lambda \mathbf{1}_{\Lambda > r} | \mathcal{G}_0^n] \} \circ T^k.$$

By the ergodic theorem, the second term on the right converges to $\kappa(r) := \mathbf{E}[2\Lambda \mathbf{1}_{\Lambda > r}]$ a.s. as $T \rightarrow \infty$. It remains to consider the first term.

To this end, note the inclusions $\mathcal{G}_k^{j+1} \subseteq \mathcal{G}_k^j \subseteq \mathcal{G}_{k+1}^{j+1}$. It follows that

$$\Delta_k^j \text{ is } \mathcal{G}_{k+1}^{j+1}\text{-measurable, } \mathbf{E}[\Delta_k^j | \mathcal{G}_k^{j+1}] = 0, \text{ and } |\Delta_k^j| \leq 2r$$

for $0 \leq j < n$. Thus $(\Delta_k^j)_{k \geq 1}$ is a uniformly bounded martingale difference sequence with respect to the filtration $(\mathcal{G}_{k+1}^{j+1})_{k \geq 1}$, and we consequently have

$$\frac{1}{T} \sum_{k=1}^T \Delta_k^j \xrightarrow{T \rightarrow \infty} 0 \quad \text{a.s.}$$

by the simplest form of the martingale law of large numbers (indeed, it is easily seen that $M_n = \sum_{k=1}^n \Delta_k^j / k$ is an L^2 -bounded martingale, so that the result follows from the martingale convergence theorem and Kronecker's lemma).

Putting together these results, we obtain

$$\limsup_{T \rightarrow \infty} \left| \frac{1}{T} \sum_{k=1}^T \{ \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^m] - \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^n] \} \right| \leq \kappa(r) \quad \text{a.s.}$$

for arbitrary $r < \infty$. Letting $r \rightarrow \infty$ completes the proof. \square

We can now establish a lower bound on the loss of any strategy.

Corollary 4.2. *Under the assumptions of Theorem 2.6, we have*

$$\frac{1}{T} \sum_{k=1}^T \{ \ell_k(u_k) - \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^\infty] \} \xrightarrow{T \rightarrow \infty} 0 \quad \text{a.s.}$$

for any admissible strategy \mathbf{u} . In particular,

$$\liminf_{T \rightarrow \infty} L_T(\mathbf{u}) \geq \mathbf{E} \left[\operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{G}_0^\infty] \right] = L^* \quad \text{a.s.}$$

Proof. We begin by noting that

$$\left| \frac{1}{T} \sum_{k=1}^T \{ \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^n] - \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^\infty] \} \right| \leq \frac{1}{T} \sum_{k=1}^T \operatorname{ess\,sup}_{u \in \mathbb{U}_k} | \mathbf{E}[\ell_k(u) | \mathcal{G}_k^n] - \mathbf{E}[\ell_k(u) | \mathcal{G}_k^\infty] |$$

$$\xrightarrow{T \rightarrow \infty} \mathbf{E} \left[\operatorname{ess\,sup}_{u \in \mathbb{U}_0} | \mathbf{E}[\ell_0(u) | \mathcal{G}_0^n] - \mathbf{E}[\ell_0(u) | \mathcal{G}_0^\infty] | \right] \quad \text{a.s.}$$

by the ergodic theorem. Similarly,

$$\left| \frac{1}{T} \sum_{k=1}^T \{ \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^m] - \ell_k(u_k) \} \right| \leq \frac{1}{T} \sum_{k=1}^T \operatorname{ess\,sup}_{u \in \mathbb{U}_k} | \mathbf{E}[\ell_k(u) | \mathcal{G}_k^m] - \ell_k(u) |$$

$$\xrightarrow{T \rightarrow \infty} \mathbf{E} \left[\operatorname{ess\,sup}_{u \in \mathbb{U}_0} | \mathbf{E}[\ell_0(u) | \mathcal{G}_0^m] - \ell_0(u) | \right] \quad \text{a.s.}$$

Therefore, using Lemma 4.1 and Assumption 2 of Theorem 2.6, the first statement of the Corollary follows by letting $n \rightarrow \infty$ and $m \rightarrow -\infty$.

For the second statement, it suffices to note that

$$\frac{1}{T} \sum_{k=1}^T \mathbf{E}[\ell_k(u_k) | \mathcal{G}_k^\infty] \geq \frac{1}{T} \sum_{k=1}^T \operatorname{ess\,inf}_{u \in \mathbb{U}_k} \mathbf{E}[\ell_k(u) | \mathcal{G}_k^\infty] \xrightarrow{T \rightarrow \infty} L^* \quad \text{a.s.}$$

by the ergodic theorem and Assumption 3 of Theorem 2.6. \square

As was explained in the introduction, a pathwise optimal strategy could easily be obtained if one can prove “ergodic tower property” of the form

$$\frac{1}{T} \sum_{k=1}^T \{ \ell_k(u_k) - \mathbf{E}[\ell_k(u_k) | \mathcal{Y}_{0,k}] \} \xrightarrow{T \rightarrow \infty} 0 \quad \text{a.s.} \quad ?$$

Corollary 4.2 establishes just such a property, but where the σ -field $\mathcal{Y}_{0,k}$ is replaced by the larger σ -field \mathcal{G}_k^∞ . This yields a lower bound on the asymptotic loss, but it is far from clear that one can choose a $\mathcal{Y}_{0,k}$ -adapted strategy that attains this bound.

Therefore, what remains is to show that there exists an admissible strategy \mathbf{u}^* that attains the lower bound in Corollary 4.2. A promising candidate is the mean-optimal strategy $\tilde{\mathbf{u}}$. Unfortunately, we are not able to prove pathwise optimality of the mean-optimal strategy in the general setting of Theorem 2.6. However, we will obtain a pathwise optimal strategy \mathbf{u}^* by a judicious modification of the mean-optimal strategy $\tilde{\mathbf{u}}$. The key idea is the following “uniform” version of the martingale convergence theorem, which we prove following Neveu [23, Lemma V-2-9].

Lemma 4.3. *The following holds:*

$$\operatorname{ess\,inf}_{u \in \mathbb{U}_{-k,0}} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-k,0}] \xrightarrow{k \rightarrow \infty} \operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] \quad \text{a.s. and in } L^1.$$

Proof. Using the construction of the essential supremum as in the proof of Lemma 2.5, we can choose for each $0 \leq k < \infty$ a countable family $\mathbb{U}_{-k,0}^c \subset \mathbb{U}_{-k,0}$ such that

$$\operatorname{ess\,inf}_{u \in \mathbb{U}_{-k,0}} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-k,0}] = \inf_{u \in \mathbb{U}_{-k,0}^c} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-k,0}] \quad \text{a.s.},$$

and a countable family $\mathbb{U}_0^c \subset \mathbb{U}_0$ such that

$$\operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] = \inf_{u \in \mathbb{U}_0^c} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] \quad \text{a.s.}$$

For every $0 \leq k < \infty$, choose an arbitrary ordering $(U_l^n)_{n \in \mathbb{N}}$ of the elements of the countable set $\mathbb{U}_{-k,0}^c \cup (\mathbb{U}_0^c \cap \mathbb{U}_{-k,0})$. Then we clearly have

$$M_k := \operatorname{ess\,inf}_{u \in \mathbb{U}_{-k,0}} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-k,0}] = \min_{0 \leq l < k} \inf_{n \in \mathbb{N}} \mathbf{E}[\ell_0(U_l^n) | \mathcal{Y}_{-k,0}] \quad \text{a.s.}$$

and

$$M := \operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] = \inf_{0 \leq l < \infty} \inf_{n \in \mathbb{N}} \mathbf{E}[\ell_0(U_l^n) | \mathcal{Y}_{-\infty,0}] \quad \text{a.s.}$$

Our aim is to prove that $M_k \rightarrow M$ a.s. and in L^1 as $k \rightarrow \infty$.

We begin by noting that $|M_k| \leq \mathbf{E}[\Lambda | \mathcal{Y}_{-k,0}]$. Therefore, the sequence $(M_k)_{k \geq 0}$ is uniformly integrable. Moreover, $(M_k)_{k \geq 0}$ is a supermartingale with respect to the filtration $(\mathcal{Y}_{-k,0})_{k \geq 0}$: indeed, we can easily compute

$$\mathbf{E}[M_{k+1} | \mathcal{Y}_{-k,0}] \leq \mathbf{E} \left[\min_{0 \leq l \leq k} \inf_{n \in \mathbb{N}} \mathbf{E}[\ell_0(U_l^n) | \mathcal{Y}_{-k-1,0}] \middle| \mathcal{Y}_{-k,0} \right] \leq M_k.$$

Thus $M_k \rightarrow M_\infty$ a.s. and in L^1 by the martingale convergence theorem for some random variable M_∞ . We must now show that $M_\infty = M$ a.s. Note that

$$M_\infty = \lim_{k \rightarrow \infty} M_k \leq \lim_{k \rightarrow \infty} \mathbf{E}[\ell_0(U_l^n) | \mathcal{Y}_{-k,0}] = \mathbf{E}[\ell_0(U_l^n) | \mathcal{Y}_{-\infty,0}] \quad \text{a.s.}$$

for every $n \in \mathbb{N}$ and $0 \leq l < \infty$, so $M_\infty \leq M$ a.s. To complete the proof, it therefore suffices to show that $\mathbf{E}[M_\infty] = \mathbf{E}[M]$.

To this end, define for $N \in \mathbb{N}$ and $0 \leq k \leq \infty$

$$M_k^N = \min_{l \leq N \wedge k} \min_{n \leq N} \mathbf{E}[\ell_0(U_l^n) | \mathcal{Y}_{-k,0}].$$

As $(M_k^N)_{k \geq 0}$ is again a supermartingale, clearly $\mathbf{E}[M_k^N]$ is doubly nonincreasing in k and N . The exchange of limits is therefore permitted, so that

$$\mathbf{E}[M_\infty] = \lim_{k \rightarrow \infty} \lim_{N \rightarrow \infty} \mathbf{E}[M_k^N] = \lim_{N \rightarrow \infty} \lim_{k \rightarrow \infty} \mathbf{E}[M_k^N] = \mathbf{E}[M].$$

This completes the proof. □

Corollary 4.4. *Suppose that Assumption 3 of Theorem 2.6 holds. Then*

$$\mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{G}_k^\infty] \circ T^{-k} \xrightarrow{k \rightarrow \infty} \operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{G}_0^\infty] \quad \text{in } L^1.$$

Proof. Define $\hat{u}_k = \tilde{u}_k \circ T^{-k} \in \mathbb{U}_{-k,0}$, so that

$$\mathbf{E}[\ell_k(\tilde{u}_k)|\mathcal{G}_k^\infty] \circ T^{-k} = \mathbf{E}[\ell_0(\hat{u}_k)|\mathcal{G}_0^\infty].$$

By stationarity and the definition of $\tilde{\mathbf{u}}$, we have

$$\begin{aligned} \mathbf{E}[\mathbf{E}[\ell_0(\hat{u}_k)|\mathcal{G}_0^\infty]] &= \mathbf{E}[\mathbf{E}[\ell_0(\hat{u}_k)|\mathcal{Y}_{-k,0}]] \\ &\leq \mathbf{E}\left[\text{ess inf}_{u \in \mathbb{U}_{-k,0}} \mathbf{E}[\ell_0(u)|\mathcal{Y}_{-k,0}]\right] + k^{-1}. \end{aligned}$$

Therefore, by Lemma 4.3, we have

$$\limsup_{k \rightarrow \infty} \mathbf{E}[\mathbf{E}[\ell_0(\hat{u}_k)|\mathcal{G}_0^\infty]] \leq \mathbf{E}\left[\text{ess inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u)|\mathcal{Y}_{-\infty,0}]\right] = L^*.$$

On the other hand, note that

$$\mathbf{E}[\ell_0(\hat{u}_k)|\mathcal{G}_0^\infty] \geq \text{ess inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u)|\mathcal{G}_0^\infty] \quad \text{a.s.}$$

Using Assumption 3, we therefore have

$$\limsup_{k \rightarrow \infty} \left\| \mathbf{E}[\ell_0(\hat{u}_k)|\mathcal{G}_0^\infty] - \text{ess inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u)|\mathcal{G}_0^\infty] \right\|_1 \leq 0.$$

This completes the proof. \square

We are now in the position to construct the pathwise optimal strategy \mathbf{u}^* . By Corollary 4.4, we can choose a (nonrandom) sequence $k_n \uparrow \infty$ such that

$$\mathbf{E}[\ell_{k_n}(\tilde{u}_{k_n})|\mathcal{G}_{k_n}^\infty] \circ T^{-k_n} \xrightarrow{n \rightarrow \infty} \text{ess inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u)|\mathcal{G}_0^\infty] \quad \text{a.s.}$$

Let us define

$$u_k^* = \tilde{u}_{k_n} \circ T^{k-k_n} \quad \text{for } k_n \leq k < k_{n+1}, \quad n \in \mathbb{N}.$$

Then clearly $\mathbf{u}^* = (u_k^*)_{k \geq 1}$ is an admissible strategy.

Lemma 4.5. *Suppose that the assumptions of Theorem 2.6 hold. Then*

$$\lim_{T \rightarrow \infty} L_T(\mathbf{u}^*) = L^* \quad \text{a.s.}$$

Proof. By construction,

$$\mathbf{E}[\ell_k(u_k^*)|\mathcal{G}_k^\infty] \circ T^{-k} \xrightarrow{k \rightarrow \infty} \text{ess inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u)|\mathcal{G}_0^\infty] \quad \text{a.s.}$$

Moreover,

$$\sup_{k \geq 1} |\mathbf{E}[\ell_k(u_k^*)|\mathcal{G}_k^\infty] \circ T^{-k}| \leq \mathbf{E}[\Lambda|\mathcal{G}_0^\infty] \in L^1.$$

Therefore, by Maker's generalized ergodic theorem [19, Corollary 10.8]

$$\frac{1}{T} \sum_{k=1}^T \mathbf{E}[\ell_k(\mathbf{u}_k^*) | \mathcal{G}_k^\infty] \xrightarrow{T \rightarrow \infty} \mathbf{E} \left[\operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{G}_0^\infty] \right] = L^* \quad \text{a.s.}$$

Thus $L_T(\mathbf{u}^*) \rightarrow L^*$ a.s. as $T \rightarrow \infty$ by Corollary 4.2. \square

The proof of Theorem 2.6 is now complete. Indeed, if \mathbf{u} is admissible, then

$$\liminf_{T \rightarrow \infty} \{L_T(\mathbf{u}) - L_T(\mathbf{u}^*)\} = \liminf_{T \rightarrow \infty} L_T(\mathbf{u}) - L^* \geq 0 \quad \text{a.s.}$$

by Lemma 4.5 and Corollary 4.2, so \mathbf{u}^* is pathwise optimal.

4.2 Proof of Corollary 2.11

To prove pathwise optimality, it suffices to show $L_T(\tilde{\mathbf{u}}) \rightarrow L^*$ a.s.

Lemma 4.6. *Under the assumptions of Corollary 2.11, the mean-optimal strategy $\tilde{\mathbf{u}}$ (Lemma 2.5) satisfies $L_T(\tilde{\mathbf{u}}) \rightarrow L^*$ a.s. as $T \rightarrow \infty$.*

Proof. By the definition of $\tilde{\mathbf{u}}$ and Lemma 4.3,

$$\mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{Y}_{0,k}] \circ T^{-k} \xrightarrow{k \rightarrow \infty} \operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] \quad \text{a.s.}$$

Therefore, the third part of Assumption 2 of Corollary 2.11 implies that

$$\mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{Y}_{-\infty,k}] \circ T^{-k} \xrightarrow{k \rightarrow \infty} \operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] \quad \text{a.s.}$$

But by Assumption 3 of Corollary 2.11 and stationarity, we obtain

$$\mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{G}_k^\infty] \circ T^{-k} \xrightarrow{k \rightarrow \infty} \operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] \quad \text{a.s.}$$

Moreover, we have

$$\sup_{k \geq 1} |\mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{G}_k^\infty] \circ T^{-k}| \leq \mathbf{E}[\Lambda | \mathcal{G}_0^\infty] \in L^1.$$

Maker's generalized ergodic theorem [19, Corollary 10.8] therefore yields

$$\frac{1}{T} \sum_{k=1}^T \mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{G}_k^\infty] \xrightarrow{T \rightarrow \infty} \mathbf{E} \left[\operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty,0}] \right] = L^* \quad \text{a.s.}$$

As the assumptions of Corollary 2.11 imply those of Theorem 2.6, the result as well as pathwise optimality of $\tilde{\mathbf{u}}$ now follow from Corollary 4.2. \square

4.3 Proof of Theorem 2.13

The proof of the Theorem is once again based on a variant of the “ergodic tower property” described in the introduction. In the present setting, the result follows rather easily from the conditional weak mixing assumption.

Lemma 4.7. *Suppose that the assumption of Theorem 2.13 holds. Then*

$$\frac{1}{T} \sum_{k=1}^T \{\ell_k(u_k) - \mathbf{E}[\ell_k(u_k) | \mathcal{Y}_{-\infty, k}]\} \xrightarrow{T \rightarrow \infty} 0 \quad \text{in } L^1$$

for every admissible strategy \mathbf{u} .

Proof. Define $\bar{\ell}_k^M(u) = \bar{\ell}_0^M(u) \circ T^k$ for $u \in U$. We begin by noting that

$$\mathbf{E} \left[\left(\frac{1}{T} \sum_{k=1}^T \bar{\ell}_k^M(u_k) \right)^2 \right] = \frac{1}{T^2} \sum_{n, m=1}^T \mathbf{E}[\bar{\ell}_n^M(u_n) \bar{\ell}_m^M(u_m)].$$

Suppose that $m \leq n$. Then by stationarity and as \mathbf{u} is admissible

$$\begin{aligned} \mathbf{E}[\bar{\ell}_n^M(u_n) \bar{\ell}_m^M(u_m)] &= \mathbf{E}[\bar{\ell}_0^M(u_n \circ T^{-n}) \{\bar{\ell}_0^M(u_m \circ T^{-m}) \circ T^{-(n-m)}\}] \\ &\leq \mathbf{E} \left[\text{ess sup}_{u, u' \in \mathbb{U}_0} |\mathbf{E}[\bar{\ell}_0^M(u') \{\bar{\ell}_0^M(u) \circ T^{-(n-m)}\} | \mathcal{Y}_{-\infty, 0}]| \right]. \end{aligned}$$

We can therefore estimate

$$\begin{aligned} \mathbf{E} \left[\left(\frac{1}{T} \sum_{k=1}^T \bar{\ell}_k^M(u_k) \right)^2 \right] &\leq \frac{2}{T^2} \sum_{n=1}^T \sum_{k=0}^{n-1} \mathbf{E} \left[\text{ess sup}_{u, u' \in \mathbb{U}_0} |\mathbf{E}[\bar{\ell}_0^M(u') \{\bar{\ell}_0^M(u) \circ T^{-k}\} | \mathcal{Y}_{-\infty, 0}]| \right] \\ &= \frac{2}{T^2} \sum_{k=0}^{T-1} (T-k) \mathbf{E} \left[\text{ess sup}_{u, u' \in \mathbb{U}_0} |\mathbf{E}[\bar{\ell}_0^M(u') \{\bar{\ell}_0^M(u) \circ T^{-k}\} | \mathcal{Y}_{-\infty, 0}]| \right] \\ &\leq \frac{2}{T} \sum_{k=0}^{T-1} \mathbf{E} \left[\text{ess sup}_{u, u' \in \mathbb{U}_0} |\mathbf{E}[\bar{\ell}_0^M(u') \{\bar{\ell}_0^M(u) \circ T^{-k}\} | \mathcal{Y}_{-\infty, 0}]| \right]. \end{aligned}$$

By the uniform conditional mixing assumption, it follows that

$$\lim_{M \rightarrow \infty} \limsup_{T \rightarrow \infty} \left\| \frac{1}{T} \sum_{k=1}^T \bar{\ell}_k^M(u_k) \right\|_2 = 0.$$

On the other hand, note that

$$\sup_{T \geq 1} \left\| \frac{1}{T} \sum_{k=1}^T \{\ell_k(u_k) - \mathbf{E}[\ell_k(u_k) | \mathcal{Y}_{-\infty, k}]\} - \frac{1}{T} \sum_{k=1}^T \bar{\ell}_k^M(u_k) \right\|_1 \leq \mathbf{E}[2\Lambda \mathbf{1}_{\Lambda > M}] \xrightarrow{M \rightarrow \infty} 0.$$

The result now follows by applying the triangle inequality. \square

Corollary 4.8. *Under the assumption of Theorem 2.13, we have*

$$\mathbf{P}\left[L_T(\mathbf{u}) - L^* \leq -\varepsilon\right] \xrightarrow{T \rightarrow \infty} 0 \quad \text{for every } \varepsilon > 0$$

for every admissible strategy \mathbf{u} .

Proof. Let \mathbf{u} be any admissible strategy. Then by Lemma 4.7

$$L_T(\mathbf{u}) - \frac{1}{T} \sum_{k=1}^T \mathbf{E}[\ell_k(u_k) | \mathcal{Y}_{-\infty, k}] \xrightarrow{T \rightarrow \infty} 0 \quad \text{in } L^1.$$

On the other hand, note that

$$\frac{1}{T} \sum_{k=1}^T \mathbf{E}[\ell_k(u_k) | \mathcal{Y}_{-\infty, k}] \geq \frac{1}{T} \sum_{k=1}^T \operatorname{ess\,inf}_{u \in \mathbb{U}_k} \mathbf{E}[\ell_k(u) | \mathcal{Y}_{-\infty, k}] \xrightarrow{T \rightarrow \infty} L^* \quad \text{in } L^1$$

by the ergodic theorem. The result follows directly. \square

In view of Corollary 4.8, in order to establish weak pathwise optimality of $\tilde{\mathbf{u}}$ it evidently suffices to prove that $\tilde{\mathbf{u}}$ satisfies the ergodic theorem $L_T(\tilde{\mathbf{u}}) \rightarrow L^*$ in L^1 . However, most of the work was already done in the proof of Theorem 2.6.

Lemma 4.9. *Under the assumption of Theorem 2.13, $L_T(\tilde{\mathbf{u}}) \rightarrow L^*$ in L^1 .*

Proof. By the definition of $\tilde{\mathbf{u}}$, we have

$$\mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{Y}_{0, k}] \circ T^{-k} \leq \operatorname{ess\,inf}_{u \in \mathbb{U}_{-k, 0}} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-k, 0}] + k^{-1} \quad \text{a.s.}$$

Therefore, by Lemma 4.3, we obtain

$$\limsup_{k \rightarrow \infty} \mathbf{E}[\mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{Y}_{-\infty, k}] \circ T^{-k}] \leq \mathbf{E}\left[\operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty, 0}]\right] = L^*.$$

On the other hand,

$$\mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{Y}_{-\infty, k}] \circ T^{-k} \geq \operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty, 0}] \quad \text{a.s.}$$

for all $k \in \mathbb{N}$. It follows that

$$\begin{aligned} \limsup_{k \rightarrow \infty} \left\| \mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{Y}_{-\infty, k}] \circ T^{-k} - \operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty, 0}] \right\|_1 &= \\ \limsup_{k \rightarrow \infty} \mathbf{E}[\mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{Y}_{-\infty, k}] \circ T^{-k}] - \mathbf{E}\left[\operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty, 0}]\right] &\leq 0. \end{aligned}$$

Therefore, by Maker's generalized ergodic theorem [19, Corollary 10.8]

$$\frac{1}{T} \sum_{k=1}^T \mathbf{E}[\ell_k(\tilde{u}_k) | \mathcal{Y}_{-\infty, k}] \xrightarrow{T \rightarrow \infty} L^* \quad \text{in } L^1.$$

The result now follows using Lemma 4.7. \square

Combining Corollary 4.8 and Lemma 4.9 completes the proof of Theorem 2.13.

4.4 Proof of Theorem 2.16

The implication $1 \Rightarrow 2$ of Theorem 2.16 follows immediately from Theorem 1.8. In the following, we will prove the converse implication $2 \Rightarrow 1$: that is, we will show that if $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is *not* conditionally weak mixing relative to \mathcal{Y} , then we can construct a bounded loss function ℓ with some finite decision space U for which there exists no weakly pathwise optimal strategy.

We begin by providing a “diagonal” characterization of conditional weak mixing.

Lemma 4.10. *$(\Omega, \mathcal{B}, \mathbf{P}, T)$ is conditionally weak mixing relative to \mathcal{Z} if and only if*

$$\frac{1}{T} \sum_{k=1}^T |\mathbf{E}[\{h \circ T^{-k}\} h | \mathcal{Z}] - \mathbf{E}[h \circ T^{-k} | \mathcal{Z}] \mathbf{E}[h | \mathcal{Z}]| \xrightarrow{T \rightarrow \infty} 0 \quad \text{in } L^1$$

for every $h \in L^2$, provided that $\mathcal{Z} \subseteq T^{-1}\mathcal{Z}$.

Proof. It suffices to show that if the equation display in the lemma holds, then $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is conditionally weak mixing relative to \mathcal{Z} . To this end, let us fix $h \in L^2$ and denote by \mathcal{A} the class of all functions $g \in L^2$ such that

$$\frac{1}{T} \sum_{k=1}^T |\mathbf{E}[\{g \circ T^{-k}\} h | \mathcal{Z}] - \mathbf{E}[g \circ T^{-k} | \mathcal{Z}] \mathbf{E}[h | \mathcal{Z}]| \xrightarrow{T \rightarrow \infty} 0 \quad \text{in } L^1.$$

Clearly \mathcal{A} is closed linear subspace of L^2 . Note that \mathcal{A} certainly contains every random variable of the form $h \mathbf{1}_B \circ T^m$ or $\mathbf{1}_B \circ T^m$ for $m \in \mathbb{Z}$ and $B \in \mathcal{Z}$. Therefore, the closed linear span K of all such random variables is included in \mathcal{A} . On the other hand, suppose that $g \in K^\perp$. Then for every $k \in \mathbb{Z}$, we have

$$\mathbf{E}[\mathbf{E}[\{g \circ T^{-k}\} h | \mathcal{Z}] \mathbf{1}_B] = \mathbf{E}[g \{h \mathbf{1}_B \circ T^k\}] = 0$$

for all $B \in \mathcal{Z}$. It follows that $\mathbf{E}[\{g \circ T^{-k}\} h | \mathcal{Z}] = 0$ a.s. for all $k \in \mathbb{N}$. Similarly, we find that $\mathbf{E}[g \circ T^{-k} | \mathcal{Z}] = 0$ a.s. for all $k \in \mathbb{N}$. Thus evidently $K^\perp \subseteq \mathcal{A}$ also. Therefore, \mathcal{A} contains $K \oplus K^\perp = L^2$, and the proof is complete. \square

In the remainder of this section, we suppose that $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is not conditionally weakly mixing relative to \mathcal{Y} . By Lemma 4.10, there is a function $h \in L^2$ such that

$$\limsup_{T \rightarrow \infty} \mathbf{E} \left[\frac{1}{T} \sum_{k=1}^T |\mathbf{E}[\{H \circ T^{-k}\} H | \mathcal{Y}]| \right] \geq \varepsilon > 0$$

where $H := h - \mathbf{E}[h|\mathcal{Y}]$. By approximation in L^2 , we may clearly assume without loss of generality that h takes values in $[0, 1]$, so that H takes values in $[-1, 1]$. We will fix such a function in the sequel, and consider the loss function

$$\ell(u, \omega) = uH(\omega)$$

where we initially choose decisions $u \in [-1, 1]$ (the decision space will be discretized at the end of the proof as required by Theorem 2.16). We claim that for the loss function ℓ there exists no weakly pathwise optimal strategy. This will be proved by a randomization procedure that will be explained presently.

In the following $([0, 1], \mathcal{J})$ denotes the unit interval with its Borel σ -field.

Lemma 4.11. *Suppose that $(\Omega, \mathcal{B}, \mathbf{P})$ is a standard probability space. Then there exists a $(\mathcal{Y} \otimes \mathcal{J})$ -measurable map $\iota : \Omega \times [0, 1] \rightarrow \Omega$ such that*

$$\mathbf{E}[X|\mathcal{Y}](\omega) = \int_0^1 X(\iota(\omega, \lambda)) d\lambda \quad \mathbf{P}\text{-a.e. } \omega \in \Omega.$$

for any bounded $(\mathcal{B}\text{-})$ measurable function $X : \Omega \rightarrow \mathbb{R}$.

Proof. As $(\Omega, \mathcal{B}, \mathbf{P})$ is a standard probability space, this is [19, Lemma 3.22] together with the existence of regular conditional probabilities [19, Theorem 6.3]. \square

Consider the quantity

$$A_T^\lambda(\omega) = \frac{1}{T} \sum_{k=1}^T H(T^k \iota(\omega, \lambda)) H(T^k \omega).$$

Then we can compute

$$\begin{aligned} \int_0^1 (A_T^\lambda)^2 d\lambda &= \frac{1}{T^2} \sum_{m,n=1}^T H(T^m \omega) H(T^n \omega) \int_0^1 H(T^m \iota(\omega, \lambda)) H(T^n \iota(\omega, \lambda)) d\lambda \\ &= \frac{1}{T^2} \sum_{m,n=1}^T H(T^m \omega) H(T^n \omega) \mathbf{E}[\{H \circ T^m\} \{H \circ T^n\} | \mathcal{Y}](\omega). \end{aligned}$$

In particular, using the invariance of \mathcal{Y} , we have

$$\begin{aligned} \left[\int_0^1 \mathbf{E}[(A_T^\lambda)^2] d\lambda \right]^{1/2} &= \mathbf{E} \left[\frac{1}{T^2} \sum_{m,n=1}^T \mathbf{E}[\{H \circ T^m\} \{H \circ T^n\} | \mathcal{Y}]^2 \right]^{1/2} \\ &\geq \mathbf{E} \left[\frac{1}{T^2} \sum_{m,n=1}^T |\mathbf{E}[\{H \circ T^m\} \{H \circ T^n\} | \mathcal{Y}]| \right] \\ &\geq \mathbf{E} \left[\frac{1}{T^2} \sum_{n=1}^T \sum_{m=1}^n |\mathbf{E}[\{H \circ T^{m-n}\} H | \mathcal{Y}]| \right] \end{aligned}$$

$$\begin{aligned}
&= \mathbf{E} \left[\frac{1}{T^2} \sum_{k=0}^{T-1} (T-k) |\mathbf{E}[\{H \circ T^{-k}\} H | \mathcal{Y}]| \right] \\
&\geq \mathbf{E} \left[\frac{1}{2T} \sum_{k=0}^{\lfloor T/2 \rfloor} |\mathbf{E}[\{H \circ T^{-k}\} H | \mathcal{Y}]| \right].
\end{aligned}$$

By our choice of H , it follows that

$$\limsup_{T \rightarrow \infty} \mathbf{E}[(A_T^\lambda)^2] \geq \frac{\varepsilon^2}{16}$$

for some $\lambda = \lambda_0 \in [0, 1]$. Define

$$u_k(\omega) = H(T^k \mathbf{t}(\omega, \lambda_0)).$$

Then u_k is \mathcal{Y} -measurable for all k (and is therefore admissible if we choose, for the time being, the continuous decision space $U = [-1, 1]$), and $L_T(\mathbf{u}) = A_T^{\lambda_0}$. Moreover,

$$\frac{\varepsilon^2}{16} \leq \limsup_{T \rightarrow \infty} \mathbf{E}[(L_T(\mathbf{u}))^2] \leq \frac{\varepsilon^2}{64} + \limsup_{T \rightarrow \infty} \mathbf{P} \left[L_T(\mathbf{u}) > \frac{\varepsilon}{8} \right] + \limsup_{T \rightarrow \infty} \mathbf{P} \left[L_T(\mathbf{u}) < -\frac{\varepsilon}{8} \right]$$

implies that we may assume without loss of generality that

$$\limsup_{T \rightarrow \infty} \mathbf{P} \left[L_T(\mathbf{u}) < -\frac{\varepsilon}{8} \right] > 0$$

(if this is not the case, simply substitute $-\mathbf{u}$ for \mathbf{u} in the following). But note that the strategy $\tilde{\mathbf{u}}$ defined by $\tilde{u}_k = 0$ for all k is mean-optimal (indeed, $\mathbf{E}[\ell_k(u) | \mathcal{Y}] = u \mathbf{E}[H | \mathcal{Y}] \circ T^k = 0$ for all u by construction). Thus evidently

$$\limsup_{T \rightarrow \infty} \mathbf{P} \left[L_T(\mathbf{u}) - L_T(\tilde{\mathbf{u}}) < -\frac{\varepsilon}{8} \right] > 0,$$

so $\tilde{\mathbf{u}}$ is not weakly pathwise optimal. It follows from Lemma 2.12 that no weakly pathwise optimal strategy can exist if we choose the decision space $U = [-1, 1]$.

To complete the proof of Theorem 2.16, it remains to show that this conclusion remains valid if we replace $U = [-1, 1]$ by some finite set. This is easily attained by discretization, however. Indeed, let $U = \{k\varepsilon/16 : k = -\lfloor 16/\varepsilon \rfloor, \dots, \lfloor 16/\varepsilon \rfloor\}$, and construct a new strategy \mathbf{u}' such that u'_k equals the value of u_k (which takes values in $[-1, 1]$) rounded to the nearest element of U . Clearly $\tilde{\mathbf{u}}$ and \mathbf{u}' both take values in the finite set U , and we have $|L_T(\mathbf{u}) - L_T(\mathbf{u}')| \leq \varepsilon/16$. Therefore,

$$\limsup_{T \rightarrow \infty} \mathbf{P} \left[L_T(\mathbf{u}') - L_T(\tilde{\mathbf{u}}) < -\frac{\varepsilon}{16} \right] > 0,$$

and it follows again by Lemma 2.12 that no weakly pathwise optimal strategy exists.

4.5 Proof of Theorem 3.15

By stationarity, we can rewrite the conditional absolute regularity property as

$$\left\| \mathbf{P}[(X_k)_{k \geq 0} \in \cdot | \mathcal{X}_{-\infty, -n} \vee \mathcal{Y}_{-\infty, \infty}] - \mathbf{P}[(X_k)_{k \geq 0} \in \cdot | \mathcal{Y}_{-\infty, \infty}] \right\|_{\text{TV}} \xrightarrow{n \rightarrow \infty} 0 \quad \text{in } L^1.$$

Using a simple truncation argument (as the loss is dominated in L^1), this implies

$$\text{ess sup}_{u \in \mathbb{U}_0} \left| \mathbf{E}[l(u, X_0) | \mathcal{X}_{-\infty, -n} \vee \mathcal{Y}_{-\infty, \infty}] - \mathbf{E}[l(u, X_0) | \mathcal{Y}_{-\infty, \infty}] \right| \xrightarrow{n \rightarrow \infty} 0 \quad \text{in } L^1.$$

If only we could replace $\mathcal{Y}_{-\infty, \infty}$ by $\mathcal{Y}_{-\infty, 0}$ in this expression, all the assumptions of Theorem 2.6 would follow immediately. Unfortunately, it is not immediately obvious whether this replacement is possible without additional assumptions.

Remark 4.12. In general, it is not clear whether a conditional K -automorphism relative to $\mathcal{Y}_{-\infty, \infty}$ is necessarily a conditional K -automorphism relative to $\mathcal{Y}_{-\infty, 0}$. In this context, it is interesting to note that the corresponding property does hold for conditional weak mixing. We briefly sketch the proof. Suppose that $(\Omega, \mathcal{B}, \mathbf{P}, T)$ is conditionally weakly mixing relative to $\mathcal{Y}_{-\infty, \infty}$. We claim that then also

$$\frac{1}{T} \sum_{k=1}^T \left| \mathbf{E}[\{f \circ T^{-k}\} g | \mathcal{Y}_{-\infty, 0}] - \mathbf{E}[f \circ T^{-k} | \mathcal{Y}_{-\infty, 0}] \mathbf{E}[g | \mathcal{Y}_{-\infty, 0}] \right| \xrightarrow{T \rightarrow \infty} 0 \quad \text{in } L^1$$

for every $f, g \in L^2$. Indeed, the conclusion is clearly true whenever f is $\mathcal{Y}_{-\infty, n}$ -measurable for some $n \in \mathbb{Z}$. By approximation in L^2 , the conclusion holds whenever f is $\mathcal{Y}_{-\infty, \infty}$ -measurable, and it therefore suffices to consider $f \in L^2(\mathcal{Y}_{-\infty, \infty})^\perp$. But in this case we have $\mathbf{E}[f \circ T^{-k} | \mathcal{Y}_{-\infty, \infty}] = \mathbf{E}[f \circ T^{-k} | \mathcal{Y}_{-\infty, 0}] = 0$ for all k , and

$$\left\| \frac{1}{T} \sum_{k=1}^T \left| \mathbf{E}[\{f \circ T^{-k}\} g | \mathcal{Y}_{-\infty, 0}] \right| \right\|_1 \leq \left\| \frac{1}{T} \sum_{k=1}^T \left| \mathbf{E}[\{f \circ T^{-k}\} g | \mathcal{Y}_{-\infty, \infty}] \right| \right\|_1 \xrightarrow{T \rightarrow \infty} 0 \quad \text{in } L^1$$

by Jensen's inequality and the conditional weak mixing property relative to $\mathcal{Y}_{-\infty, \infty}$.

As we cannot directly replace $\mathcal{Y}_{-\infty, \infty}$ by $\mathcal{Y}_{-\infty, 0}$, we take an alternative approach. We begin by noting that, using the conditional absolute regularity property as described above, we obtain the following trivial adaptation of Corollary 4.2.

Lemma 4.13. *Under the assumptions of Theorem 3.15, we have*

$$\frac{1}{T} \sum_{k=1}^T \{l(u_k, X_k) - \mathbf{E}[l(u_k, X_k) | \mathcal{Y}_{-\infty, \infty}]\} \xrightarrow{T \rightarrow \infty} 0 \quad \text{a.s.}$$

for any admissible strategy \mathbf{u} .

We will now proceed to replace $\mathcal{Y}_{-\infty, \infty}$ by $\mathcal{Y}_{-\infty, k}$ in Lemma 4.13. To this end, we use the additional property established in [27, Proposition 3.9]:

$$\mathbf{P}[(X_k)_{k \leq 0} \in \cdot | \mathcal{Y}_{-\infty, 0}] \sim \mathbf{P}[(X_k)_{k \leq 0} \in \cdot | \mathcal{Y}_{-\infty, \infty}] \quad \text{a.s.}$$

Theorem 3.14 implies that the past tail σ -field $\bigcap_n \mathcal{X}_{-\infty, n}$ is $\mathbf{P}[\cdot | \mathcal{Y}_{-\infty, 0}]$ -trivial a.s. (cf. [33]). Thus a standard argument [21, Theorem III.14.10] yields

$$\left\| \mathbf{P}[(X_k)_{k \leq n} \in \cdot | \mathcal{Y}_{-\infty, 0}] - \mathbf{P}[(X_k)_{k \leq n} \in \cdot | \mathcal{Y}_{-\infty, \infty}] \right\|_{\text{TV}} \xrightarrow{n \rightarrow -\infty} 0 \quad \text{in } L^1.$$

Therefore, by stationarity and a simple truncation argument, we have

$$\operatorname{ess\,sup}_{u \in \mathbb{U}_0} \left| \mathbf{E}[l(u, X_0) | \mathcal{Y}_{-\infty, n}] - \mathbf{E}[l(u, X_0) | \mathcal{Y}_{-\infty, \infty}] \right| \xrightarrow{n \rightarrow -\infty} 0 \quad \text{in } L^1.$$

This yields the following consequence.

Corollary 4.14. *Under the assumptions of Theorem 3.15, we have*

$$\frac{1}{T} \sum_{k=1}^T \{l(u_k, X_k) - \mathbf{E}[l(u_k, X_k) | \mathcal{Y}_{-\infty, k}]\} \xrightarrow{T \rightarrow \infty} 0 \quad \text{a.s.}$$

for any admissible strategy \mathbf{u} . In particular,

$$\liminf_{T \rightarrow \infty} L_T(\mathbf{u}) \geq \mathbf{E} \left[\operatorname{ess\,inf}_{u \in \mathbb{U}_0} \mathbf{E}[\ell_0(u) | \mathcal{Y}_{-\infty, 0}] \right] = L^* \quad \text{a.s.}$$

Proof (Sketch). Following almost verbatim the proof of Lemma 4.1, one can prove

$$\frac{1}{T} \sum_{k=1}^T \{ \mathbf{E}[l(u_k, X_k) | \mathcal{Y}_{-\infty, k}] - \mathbf{E}[l(u_k, X_k) | \mathcal{Y}_{-\infty, k+r}] \} \xrightarrow{T \rightarrow \infty} 0 \quad \text{a.s.}$$

for any $r \in \mathbb{N}$. On the other hand, we have

$$\begin{aligned} & \limsup_{T \rightarrow \infty} \left| \frac{1}{T} \sum_{k=1}^T \{ \mathbf{E}[l(u_k, X_k) | \mathcal{Y}_{-\infty, k+r}] - \mathbf{E}[l(u_k, X_k) | \mathcal{Y}_{-\infty, \infty}] \} \right| \\ & \leq \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^T \operatorname{ess\,sup}_{u \in \mathbb{U}_k} \left| \mathbf{E}[l(u, X_k) | \mathcal{Y}_{-\infty, k+r}] - \mathbf{E}[l(u, X_k) | \mathcal{Y}_{-\infty, \infty}] \right| \\ & = \mathbf{E} \left[\operatorname{ess\,sup}_{u \in \mathbb{U}_0} \left| \mathbf{E}[l(u, X_0) | \mathcal{Y}_{-\infty, r}] - \mathbf{E}[l(u, X_0) | \mathcal{Y}_{-\infty, \infty}] \right| \right] \quad \text{a.s.} \end{aligned}$$

by the ergodic theorem. It was shown above that the latter quantity converges to zero as $r \rightarrow \infty$, and the result now follows using Lemma 4.13. \square

The remainder of the proof of Theorem 3.15 is identical to that of Theorem 2.6 modulo trivial modifications, and is therefore omitted.

Acknowledgment

This work was partially supported by NSF grant DMS-1005575.

References

1. Algoet, P.H.: The strong law of large numbers for sequential decisions under uncertainty. *IEEE Trans. Inform. Theory* **40**(3), 609–633 (1994)
2. Bellow, A., Losert, V.: The weighted pointwise ergodic theorem and the individual ergodic theorem along subsequences. *Trans. Amer. Math. Soc.* **288**(1), 307–345 (1985)
3. Berend, D., Bergelson, V.: Mixing sequences in Hilbert spaces. *Proc. Amer. Math. Soc.* **98**(2), 239–246 (1986)
4. Cappé, O., Moulines, E., Rydén, T.: *Inference in hidden Markov models*. Springer, New York (2005)
5. Chigansky, P., van Handel, R.: A complete solution to Blackwell’s unique ergodicity problem for hidden Markov chains. *Ann. Appl. Probab.* **20**(6), 2318–2345 (2010)
6. Conze, J.P.: Convergence des moyennes ergodiques pour des sous-suites. In: *Contributions au calcul des probabilités*, pp. 7–15. *Bull. Soc. Math. France, Mém. No. 35*. Soc. Math. France, Paris (1973)
7. Crisan, D., Rozovskiĭ, B. (eds.): *The Oxford handbook of nonlinear filtering*. Oxford University Press, Oxford (2011)
8. Del Moral, P., Ledoux, M.: Convergence of empirical processes for interacting particle systems with applications to nonlinear filtering. *J. Theoret. Probab.* **13**(1), 225–257 (2000)
9. Dellacherie, C., Meyer, P.A.: *Probabilities and potential*. C. North-Holland, Amsterdam (1988)
10. Dudley, R.M.: *Uniform central limit theorems*. Cambridge University Press, Cambridge (1999)
11. Etemadi, N.: An elementary proof of the strong law of large numbers. *Z. Wahrsch. Verw. Gebiete* **55**(1), 119–122 (1981)
12. Grothendieck, A.: Produits tensoriels topologiques et espaces nucléaires. *Mem. Amer. Math. Soc.* **1955**(16), 140 (1955)
13. Halmos, P.R.: In general a measure preserving transformation is mixing. *Ann. of Math. (2)* **45**, 786–792 (1944)
14. van Handel, R.: The stability of conditional Markov processes and Markov chains in random environments. *Ann. Probab.* **37**(5), 1876–1925 (2009)
15. van Handel, R.: Uniform time average consistency of Monte Carlo particle filters. *Stochastic Process. Appl.* **119**(11), 3835–3861 (2009)
16. van Handel, R.: On the exchange of intersection and supremum of σ -fields in filtering theory. *Israel J. Math.* (2012). To appear
17. van Handel, R.: The universal Glivenko-Cantelli property. *Probab. Th. Rel. Fields* (2012). To appear
18. Hoffmann-Jørgensen, J.: Uniform convergence of martingales. In: *Probability in Banach spaces, 7* (Oberwolfach, 1988), *Progr. Probab.*, vol. 21, pp. 127–137. Birkhäuser Boston, Boston, MA (1990)
19. Kallenberg, O.: *Foundations of modern probability*, second edn. Springer-Verlag, New York (2002)
20. Kunita, H.: Asymptotic behavior of the nonlinear filtering errors of Markov processes. *J. Multivariate Anal.* **1**, 365–393 (1971)
21. Lindvall, T.: *Lectures on the coupling method*. Dover Publications Inc., Mineola, NY (2002). Corrected reprint of the 1992 original

22. Meyn, S., Tweedie, R.L.: Markov chains and stochastic stability, second edn. Cambridge University Press, Cambridge (2009)
23. Neveu, J.: Discrete-parameter martingales. North-Holland, Amsterdam (1975)
24. Nobel, A.B.: On optimal sequential prediction for general processes. *IEEE Trans. Inform. Theory* **49**(1), 83–98 (2003)
25. Pollard, D.: A user's guide to measure theoretic probability. Cambridge University Press, Cambridge (2002)
26. Rudolph, D.J.: Pointwise and L^1 mixing relative to a sub-sigma algebra. *Illinois J. Math.* **48**(2), 505–517 (2004)
27. Tong, X.T., van Handel, R.: Conditional ergodicity in infinite dimension (2012). Preprint
28. Totoki, H.: On a class of special flows. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete* **15**, 157–167 (1970)
29. van der Vaart, A.W., Wellner, J.A.: Weak convergence and empirical processes. Springer-Verlag, New York (1996)
30. Volkonskiĭ, V.A., Rozanov, Y.A.: Some limit theorems for random functions. I. *Theor. Probability Appl.* **4**, 178–197 (1959)
31. Walters, P.: An introduction to ergodic theory. Springer-Verlag, New York (1982)
32. Weissman, T., Merhav, N.: Universal prediction of random binary sequences in a noisy environment. *Ann. Appl. Probab.* **14**(1), 54–89 (2004)
33. von Weizsäcker, H.: Exchanging the order of taking suprema and countable intersections of σ -algebras. *Ann. Inst. H. Poincaré Sect. B (N.S.)* **19**(1), 91–100 (1983)