

Complexity and Mixed Strategy Equilibria

Tai-Wei Hu*

Northwestern University

This version: Oct 2012

Abstract

Unpredictable behavior is central to optimal play in many strategic situations because predictable patterns leave players vulnerable to exploitation. A theory of unpredictable behavior based on differential complexity constraints is presented in the context of repeated two-person zero-sum games. Each player's complexity constraint is represented by an endowed oracle and a strategy is feasible iff it can be implemented with an oracle machine using that oracle. When one player's oracle is more complicated than the other player's, no equilibrium exists without one player fully exploiting the other. If each player has an incompressible sequence (relative to the opponent's oracle) according to the Kolmogorov complexity, an equilibrium exists in which equilibrium payoffs are equal to those of the stage game and all equilibrium strategies are unpredictable. A full characterization of history-independent equilibrium strategies is also obtained.

JEL classification: D01, D80

Keywords: Kolmogorov complexity; objective probability; frequency theory of probability; mixed strategy; zero-sum game; algorithmic randomness

*E-mail: t-hu@kellogg.northwestern.edu; Corresponding address: 2001 Sheridan Road, Jacobs Center 548, Evanston, IL 60208-0001, United States; Phone number: 847-467-1867.

1 Introduction

Unpredictable behavior is central to optimal play in many strategic situations, especially in social interactions with conflicts of interest. There are many illustrative examples from competitive sports, such as the direction of tennis serves and penalty kicks in soccer. Other relevant examples include secrecy in military affairs, bluffing behavior in poker, and tax auditing. A prototypical example is the matching pennies game: two players simultaneously present a coin and one player wins if the sides of the coins coincide while the other wins if the sides differ. In these situations, it is optimal for players to aim at being unpredictable to avoid detectable patterns that leave them vulnerable to exploitation, as Von Neumann and Morgenstern [16] point out. However, this intuition has not been formalized. The classical model using mixed strategies does not provide a foundation for unpredictable behavior based on pattern-exploitation; rather, it simply assumes that mixed strategies are unpredictable, either in one-shot games or repeated situations. Such a foundation requires an explicit model of the source of unpredictability and a formal derivation that shows unpredictable behavior avoids pattern-exploitation in equilibrium. Because patterns exist only in repeated plays, such a model can only be found in the context of repeated plays.

This paper formalizes equilibrium unpredictable behavior based on pattern-exploitation in the context of (infinitely) repeated two-person zero-sum games. Pattern exploitation is modeled by differential *complexity constraints*, which give boundaries for strategies that players can implement and also restrict patterns that players can exploit. Depending on their complexity constraints, a pattern can be simple to one player but can be complicated and not exploitable to another. A salient feature of this framework is the possibility for both players to implement complicated strategies relative to the other player's complexity constraint.

The complexity constraints are borrowed from the computability theory originated by Turing [23]. The core concept in this literature is *relative computability*, and the building block for this concept is a model of computation called *oracle machines*. An oracle

machine formalizes the intuitive notion of an algorithm or a finite procedure. It differs from a Turing machine in that it allows the algorithm to use bits of information from an external source called an *oracle*, an infinite binary sequence. An oracle θ represents certain pieces of information or knowledge that cannot be obtained from a mechanical procedure; the set of θ -computable functions (i.e., functions that can be computed by an oracle machine with θ as the oracle) captures the algorithmic content of the information contained in θ .

Player i 's complexity constraint is represented by his endowed oracle θ^i , which captures player i 's strategic insight that may not be obtainable through mechanical procedures. The complexity constraint captures limitations on implementation, and each player i has to implement his actions through an algorithm modeled by an oracle machine (using θ^i as the oracle). Because the set of θ^i -computable functions is countable, player i can only implement countably many strategies and hence most strategies in the standard model are not feasible for him. The computability relation is then a measure of complexity: an oracle θ^j is more complicated than another oracle θ^i if there is a oracle machine that computes θ^i using θ^j as the oracle. This relation is reflexive and transitive but not complete. If θ^i is computable from θ^j , then any of player i 's strategies is simple in the sense that it can be simulated and exploited by player j . On the other hand, if θ^i is not computable from θ^j and *vice versa*, both players can implement some strategies that are complicated relative to the other player's complexity constraint.

The goal of this paper is to find appropriate complexity constraints which would lead to equilibrium behavior that may be called unpredictable and characterize such behavior. Interpreting the complexity constraints as constraints on strategy-implementation instead of on players' rationality, Nash equilibrium is applicable. My first result analyzes the case where one player's oracle is stronger than the other's. I show that a Nash equilibrium exists if the stronger oracle is sufficiently complex, and, in this equilibrium, the stronger player fully exploits the other. When one player's oracle is stronger, any strategy of the weaker player can be simulated and hence the full exploitation result follows. This result shows that if an equilibrium exists without full exploitation (note that full exploitation

implies perfect predictability of one's strategy by the other and is the opposite to unpredictable behavior), then the two players' oracles have to be incomparable in terms of the computability relation.

The main result is the existence of an equilibrium in which unpredictable behavior emerges. I give a sufficient condition, called *mutual complexity*, for such existence. I use Kolmogorov complexity [11], which measures how compressible a sequence is, to formulate mutual complexity. Sequences computable from an oracle have low Kolmogorov complexity relative to that oracle. However, sequences with high Kolmogorov complexity relative to an oracle display unpredictable features relative to that oracle. Mutual complexity is formulated through the notion of *incompressible sequences* relative to an oracle, adopted from the algorithmic randomness literature (see the survey paper Downey et al. [6]). It states that each player can, with the aid of his own oracle, compute an incompressible sequence relative to the other player's oracle. Mutual complexity holds for uncountably many pairs of oracles.

It is also shown that, under mutual complexity, the equilibrium payoffs are the same as those of the stage game. This implies, together with the first result, that equilibrium strategies are not computable from the other player's oracle. Moreover, I obtain a full characterization of equilibrium history-independent strategies when the stage game has a unique mixed equilibrium: a sequence of actions is an equilibrium strategy if and only if the limit frequencies in any of its subsequences that can be selected in a computable way relative to the opponent's oracle are consistent with the mixed equilibrium strategy of the stage-game. A similar characterization result also holds for other zero-sum stage games with necessary modifications. These results give a precise criterion for unpredictability in repeated zero-sum games.

The characterization result can be used to devise tests (for example, that compare the average payoffs across different actions from the opponent) that are employed in the empirical literature about repeated zero-sum games. It has been shown (Walker and Wooders [24], Palacios-Huerta [19], and Palacios-Huerta and Volij [20]) that those tests are sufficiently powerful to distinguish between plays from amateur subjects, which generally fail

to exhibit equilibrium unpredictable behavior, and plays from professional player, which may not be rejected as equilibrium plays. Moreover, that empirical literature generally assumes that the equilibrium hypothesis can only be consistent with observed i.i.d. plays. However, the above characterization result implies existence of equilibrium strategies that are inconsistent with any i.i.d. process in terms of statistical regularities. Indeed, in my framework, unpredictability is a *relative* notion; a strategy is unpredictable or not depends on the opponent’s complexity constraint. This result formalizes the comment in Palacios-Huerta and Volij [20] which argues that “as long as players behavior is largely unpredictable to other players, [...] we may safely say that the minimax theory does well in explaining our soccer players choices.”

Now I turn to related literature. My model is closely related to the machine game framework (Aumann [1]) that employs finite automata to model strategic complexity. Ben-Porath [2] adopts that approach and shows, assuming randomization over automata, that the player with a sufficiently bigger automaton fully exploits the other in an infinitely repeated zero-sum game. The use of mixed strategies is essential for equilibrium existence there, while players in my model can only use pure strategies. The crucial difference is that the complexity measure in that literature is a linear order—any two automata are comparable in size, while two oracles can be incomputable relative to each other. Another difference is the use of incomputable strategies in my model. Those strategies are appropriate for a model of unpredictable behavior. First, any computable strategy can be implemented with an algorithm and hence cannot be genuinely unpredictable.¹ Second, the use of incomputable oracles does not contradict the Church-Turing thesis; that thesis only asserts that all finite procedures can be captured by Turing machines but does not imply that all human insights are bounded by computable ones.²

The rest of the paper is organized as follows. Section 2 formulates repeated games with

¹This argument is consistent with the requirement proposed by McKelvey [15] that a theory of strategic interaction should be ‘publication-proof,’ that is, it should survive its own publication.

²See Golding and Wegner [7] for an argument that refutes the “strong Church-Turing thesis” that asserts the equivalence between Turing computability and all forms of intelligence. See also Soare [22] for the historical remarks on the thesis.

complexity constraints and gives a characterization of equilibrium unpredictable behavior. Section 3 gives the existence result and then discusses some patterns in equilibrium strategies that are not consistent with an i.i.d. process. Section 4 gives some concluding remarks. All proofs are in Section 5.

2 Strategic complexity in repeated zero-sum games

2.1 Preliminaries on relative computability

A function f with arguments and values in natural numbers is *computable* if there exists a computer program that *computes* f , i.e., if n is in the domain of f then the program halts on input n and produces output $f(n)$, and if n is not in the domain of f then the algorithm does not halt on input n and runs forever. The formal definition is based on a model of idealized computations using *Turing machines* (see Odifreddi [18] for details), which are programs run on an idealized computer without memory or time restrictions. The celebrated Church-Turing hypothesis states that Turing-computability captures our intuition of a finite procedure or an algorithm; that is, a function can be computed by an algorithm if and only if it can be computed by a Turing machine.

Two remarks on computable functions are in order. First, the domain of a computable function can be a strict subset of \mathbb{N} : an algorithm may run into an infinite loop and never produce an output for some inputs. The notation $f : \subset \mathbb{N} \rightarrow \mathbb{N}$ is sometimes used in the following to emphasize that the domain of f , denoted by $\text{dom}(f)$, is a subset of \mathbb{N} . When $f(n)$ is defined for every natural number n , f is said to be *total*. Second, the notion of computability can be extended to other sets of mathematical objects such as \mathbb{N}^k or strings over a finite set. These sets can be *effectively identified* with \mathbb{N} : there are computable ways to *encode* elements of those sets as natural numbers, i.e. there exist computable one-to-one correspondences, called codings, between these sets and the set \mathbb{N} . Consider \mathbb{N}^2 as an example: Every pair (m, n) of natural numbers can be encoded as the number $(n + m)(n + m + 1)/2 + n$, which is a computable function.

For later purposes some discussion about finite strings is useful. For any finite set X let X^* be the set of finite strings over X , that is, $X^* = \bigcup_{n \in \mathbb{N}} X^n$, where $X^0 = \{\epsilon\}$ and ϵ is the empty string. X^* can be effectively identified with \mathbb{N} . Take the set $\{0, 1\}^*$ as an example: Every finite binary sequence $\sigma = (\sigma_0, \dots, \sigma_{n-1})$ can be encoded as the number $\sum_{t=0}^{n-1} \sigma_t 2^{n-1-t} + 2^n - 1$, which is computable. Codes for other finite sets X can be similarly constructed. Given the coding, computable functions to and from X^* are well-defined. In what follows, the notion of computability will be applied to sets that can be effectively identified with \mathbb{N} , assuming a fixed coding but without constructing specific codes.

Most functions that can be described explicitly are computable almost by definition. However, the existence of incomputable functions can be easily shown by a counting argument. Because the set of computable functions, which has the same cardinality as the set of computer programs (which are finite sequences of symbols), is countable, most functions are not computable. The celebrated result in computability theory, the Enumeration Theorem ([18], Theorem II.1.5), gives an effective enumeration of all computable functions: It states the existence of a binary computable function $U : \mathbb{C} \times \mathbb{N}^2 \rightarrow \mathbb{N}$, called the *universal Turing machine*, such that for every computable function f there is an m such that $f(\cdot) \cong U(m, \cdot)$; that is, such that $f(n)$ is defined if and only if $U(m, n)$ is defined, and when both are defined their values coincide. The Enumeration Theorem also provides an explicit example of an incomputable function: the characteristic function for $\text{dom}(U)$, the *halting problem*, which consists of all pairs (m, n) such that $U(m, n)$ is defined.

Now I turn to the notion of an *oracle machine*, which is a Turing machine with access to a black box, called an *oracle*. The only difference from Turing machines is that oracle machines may “call” to the oracle during computation. Calling to an oracle is similar to calling another program in programming languages: a program calls to another function f which returns values that are not directly computed by the program. Oracle machines generalize this idea and allow the oracle to be any infinite binary sequence $\theta = (\theta_0, \theta_1, \dots) \in \{0, 1\}^{\mathbb{N}}$. If a function f can be computed with an oracle machine P using θ as the oracle, then the function f is said to be computable relative to θ , or θ -computable. θ -computability captures what can be computed with an algorithm using

information contained in θ which may not be obtainable by mechanical processes.

Although the oracle may contain infinite bits of information, any actual computation only uses finitely many bits through a specific machine P . Thus, if P halts at input n using θ as the oracle, then there is a number k such that for any oracle θ' with $\theta'_i = \theta_i$ for all $i = 0, \dots, k-1$, P halts at n for θ' with the same output. This uniform feature leads to a useful technique called *relativization* that extends many results in Turing-computability to computability relative to an oracle without changing the proofs substantially. One such extension is the Enumeration Theorem relative to an oracle θ : There exists a binary θ -computable function $U^\theta : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that for every θ -computable function f there is an m such that $f(\cdot) \cong U^\theta(m, \cdot)$. Thus, only countably many functions are θ -computable. Another result is that the characteristic function for $\text{dom}(U^\theta)$, the *halting problem relative to θ* , is not θ -computable. That function is denoted by θ^H . However, θ is θ^H -computable.

An oracle can be regarded as a function over \mathbb{N} and hence, given two oracles θ and η , it is legitimate to ask whether θ is η -computable or not. It is easy to show that if θ is η -computable and if a function f is θ -computable, then f is also η -computable. The computability relation then gives a partition over oracles called *Turing degrees*: θ and η belong to the same degree if and only if θ is η -computable and η is θ -computable. For any oracle θ , θ and θ^H do not belong to the same degree.

Turing degrees give one way to classify the oracles in terms of their complexity. However, when η is not θ -computable, it does not tell how incomputable η is in any sense. Another concept, Kolmogorov complexity [11], is more useful; it measures the complexity of a finite object in terms of its minimum description length for a given language. A language is a mapping from finite strings over $\{0, 1\}$ to those strings, i.e., a function $L : \mathbb{N} \rightarrow \mathbb{N}$ (notice that a language may not be a total function). Its domain consists of legitimate descriptions and its range consists of objects to be described. I consider prefix-free languages only, which have nice connections to unpredictability (see Li and Vitányi [12] for a useful discussion on different languages). A language L is prefix-free if for any descriptions $\sigma, \tau \in \text{dom}(L)$, σ is not an initial segment of τ ; that is, either τ is shorter than σ or $\sigma = (\sigma_0, \dots, \sigma_{k-1}) \neq (\tau_0, \dots, \tau_{k-1})$. Notice that if a language L is

prefix-free, then L is not total. The Kolmogorov complexity of a string σ is defined as

$$K_L(\sigma) = \min\{|\tau| : \tau \in \{0, 1\}^*, L(\tau) = \sigma\},$$

and $K_L(\sigma) = \infty$ if there is no $\tau \in \text{dom}(L)$ such that $L(\tau) = \sigma$.

For a given oracle θ , because $\{0, 1\}^*$ can be effectively identified with \mathbb{N} , it is legitimate to speak of θ -computable prefix-free languages. Using the relativized Enumeration Theorem, it can be shown ([17], Proposition 2.2.7) that there exists a θ -computable *universal prefix-free language* for θ , denoted by L_θ ,³ that satisfies the following property: for any θ -computable prefix-free language L , there is a constant $c \in \mathbb{N}$ (which depends on L) such that $L(\sigma) \cong L_\theta(0^c 1 \sigma)$ for all $\sigma \in \{0, 1\}^*$. The complexity measure K_{L_θ} is denoted by K_θ ; L_θ is universal in the sense that for any θ -computable prefix-free language L , there is a constant c such that $K_\theta(\sigma) \leq K_L(\sigma) + c$ for all $\sigma \in \{0, 1\}^*$. There is a uniform upper bound for the Kolmogorov complexity: for any oracle θ , there is a constant $d \in \mathbb{N}$ such that $K_\theta(\sigma) \leq |\sigma| + 2 \log_2 |\sigma| + d$ for all $\sigma \in \{0, 1\}^*$, where $|\sigma|$ is the length of the string.

Given a universal prefix-free machine, a string $\sigma \in \{0, 1\}^*$ is said to be *d-incompressible relative to θ* for some $d \in \mathbb{N}$ if $K_\theta(\sigma) > |\sigma| - d$. An oracle η is said to be *incompressible relative to θ* if its initial segments are all *d-incompressible relative to θ* for some $d \in \mathbb{N}$, that is, $K_\theta(\eta[n]) > n - d$ for all $n \in \mathbb{N}$, where $\eta[n]$ is the initial segment of η with length n . It is known that the set of incompressible sequences relative to a fixed oracle θ has measure 1 under the Lebesgue measure over $\{0, 1\}^{\mathbb{N}}$ ([5], Corollary 6.2.6). Moreover, incompressible oracles relative to θ are not θ -computable: if an oracle η is θ -computable, then for some constant $d \in \mathbb{N}$, $K_\theta(\eta[n]) \leq 2 \log_2 n + d$, and hence is not incompressible relative to θ . However, an incompressible sequence η relative to θ does not have any statistically “rare” properties relative to θ and hence incompressibility connects complexity to unpredictability as well.

³Universal prefix-free languages relative a fixed oracle θ are not unique; however, the complexity measure derived from any such language only differs by a constant. In what follows I fix a particular universal language.

2.2 Strategic complexity

Here I propose a model of repeated two-person zero-sum games with complexity constraints. The stage game is a finite two-person zero-sum game $g = \langle X_1, X_2, h_1, h_2 \rangle$, where for $i = 1, 2$, X_i is the set of player i 's actions and $h_i : X_1 \times X_2 \rightarrow \mathbb{Q}$ is the von Neumann-Morgenstern utility function for player i , with \mathbb{Q} being the set of rational numbers.⁴ g is zero-sum and hence $h_1(x_1, x_2) + h_2(x_1, x_2) = 0$ for all $(x_1, x_2) \in X_1 \times X_2$. My analysis focuses on repeated games whose stage games have no pure equilibrium.⁵ However, the main result extends to general n -person (non-zero sum) games.

Each player i is endowed with an oracle θ^i and is restricted to use an oracle machine to implement his strategy; hence, a strategy is feasible for player i if and only if it is θ^i -computable. The set of all θ^i -computable functions that are total is denoted by $\mathcal{C}(\theta^i)$. Notice that, by the relative Enumeration Theorem, the set $\mathcal{C}(\theta^i)$ is countable and hence the set of feasible strategies for each player is only countable. Therefore, for player i , any strategy more powerful than $(\theta^i)^H$, the halting problem relative to θ^i , is not feasible.

Definition 2.1. Let $g = \langle X_1, X_2, h_1, h_2 \rangle$ be a finite zero-sum game and let θ^1, θ^2 be two oracles. The *repeated game with oracles* θ^1, θ^2 based on the stage game g , denoted by $RG(g, \theta^1, \theta^2)$, is a triple $\langle \mathcal{A}_1, \mathcal{A}_2, u_1, u_2 \rangle$ such that

- (a) $\mathcal{A}_i = \{\alpha_i : X_{-i}^* \rightarrow X_i : \alpha_i \in \mathcal{C}(\theta^i)\}$ is the set of player i 's strategies;
- (b) $u_i : \mathcal{A}_1 \times \mathcal{A}_2 \rightarrow \mathbb{R}$ is player i 's payoff function defined as

$$u_i(\alpha_1, \alpha_2) = \liminf_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{h_i(\xi_t^{\alpha,1}, \xi_t^{\alpha,2})}{T}, \quad (1)$$

where $(\xi_t^{\alpha,1}, \xi_t^{\alpha,2})$ is the outcome of period t for the strategy profile $\alpha = (\alpha_1, \alpha_2)$ defined by $\xi_0^{\alpha,j} = \alpha_j(\epsilon)$ and for any $t \geq 0$, $\xi_{t+1}^{\alpha,j} = \alpha_j(\xi_0^{\alpha,-j}, \xi_1^{\alpha,-j}, \dots, \xi_t^{\alpha,-j})$, for both $j = 1, 2$.

Because the average payoffs may not converge, taking limit inferior or other modification is necessary. Here I choose limit inferior and make the criterion symmetric between

⁴I assume that payoffs are rational-valued because of computability issues.

⁵If the stage game has a pure equilibrium, there are obvious equilibria in the repeated game as well where unpredictability plays no role.

the players. This choice is useful to obtain a full characterization of equilibrium strategies but not essential for equilibrium existence. Notice also that the game $RG(g, \theta^1, \theta^2)$ is not zero-sum because limit inferior is only super-additive but not additive. Hence, the value for $RG(g, \theta^1, \theta^2)$ (the unique equilibrium payoff) may not exist. However, unique equilibrium payoffs can still be obtained by imposing the maximin criterion. For this reason I introduce the notion of *secured equilibrium*. Here $u_i(\alpha_i; \alpha_{-i})$ stands for $u_i(\alpha_1, \alpha_2)$.

Definition 2.2. (1) The *security level* for a strategy $\alpha_i \in \mathcal{A}_i$, denoted by $s_i(\alpha_i)$, is given by $s_i(\alpha_i) = \inf_{\alpha_{-i} \in \mathcal{A}_{-i}} u_i(\alpha_i; \alpha_{-i})$.

(2) A *secured equilibrium* in the repeated game $RG(g, \theta^1, \theta^2)$ is a pair of strategies $(\alpha_1^*, \alpha_2^*) \in \mathcal{A}_1 \times \mathcal{A}_2$ such that (a) for $i = 1, 2$, $\alpha_i^* \in \arg \max_{\alpha_i \in \mathcal{A}_i} s_i(\alpha_i)$; (b) $s_1(\alpha_1^*) + s_2(\alpha_2^*) = 0$.

(3) When a secured equilibrium (α_1^*, α_2^*) exists, the *value* for player i of the repeated game $RG(g, \theta^1, \theta^2)$, denoted by $V_i(g, \theta^1, \theta^2)$, is defined to be $s_i(\alpha_i^*)$.

The following lemma shows that secured equilibrium refines Nash equilibrium in $RG(g, \theta^1, \theta^2)$. This holds because g is zero-sum. All proofs are in Section 5.

Lemma 2.1. *Suppose that (α_1^*, α_2^*) is a secured equilibrium in $RG(g, \theta^1, \theta^2)$. Then, (α_1^*, α_2^*) is also a Nash equilibrium in $RG(g, \theta^1, \theta^2)$.*

Now I turn to equilibrium analysis of $RG(g, \theta^1, \theta^2)$. The goal is to investigate conditions on (θ^1, θ^2) under which unpredictable behavior would emerge in equilibrium and to study the structure of unpredictable behavior when it emerges. I begin with a negative result where unpredictable behavior does not happen. Intuitively, if one player can fully capture the behavioral pattern in any strategy of his opponent, then in equilibrium the opponent's behavior would be perfectly predictable and fully exploited; this would not be called "unpredictable behavior." I formalize this intuition with the following proposition, which deals with the case where θ^2 is θ^1 -computable, and show that player 1 fully exploits player 2 in this case. Indeed, when θ^2 is θ^1 -computable, any strategy of player 2 can be simulated by player 1 in the sense that it is also θ^1 -computable and hence its behavioral pattern can be fully captured by player 1. Recall that for any oracle θ , θ^H is the characteristic function for the halting problem relative to θ ; θ is θ^H -computable but θ^H is not θ -computable.

Proposition 2.1. *Let $g = \langle X_1, X_2, h_1, h_2 \rangle$ be a two-person zero-sum game without any pure equilibrium.*

(a) *If $(\theta^2)^H$ is θ^1 -computable, then a secured equilibrium exists in $RG(g, \theta^1, \theta^2)$.*

(b) *Suppose that θ^2 is θ^1 -computable. If a secured equilibrium exists, then*

$$V_1(g, \theta^1, \theta^2) = \min_{x_2 \in X_2} \max_{x_1 \in X_1} h_1(x_1, x_2).$$

The intuition behind the proof for part (b) is straightforward. Consider the matching pennies game $g^{MP} = \langle \{x_1, x_2\}, \{y_1, y_2\}, h_1, h_2 \rangle$ with

$$h_1(x_1, y_1) = 1 = h_1(x_2, y_2) \text{ and } h_1(x_1, y_2) = -1 = h_1(x_2, y_1).$$

Suppose that θ^2 is θ^1 -computable, and suppose that there is a secured equilibrium, and hence a Nash equilibrium, in $RG(g, \theta^1, \theta^2)$. For any player 2's equilibrium strategy, there is a θ^1 -computable function that simulates that strategy and player 1 can devise a θ^1 -computable strategy to exactly match that strategy. Hence, the equilibrium payoff for player 1 has to be 1. Notice that this result also implies that there is no equilibrium when $\mathcal{C}(\theta^1) = \mathcal{C}(\theta^2)$, i.e., when the two oracles belong to the same Turing degree.

The proof for part (a) is more involved. By part (b) one has to construct a strategy α_1^* for player 1 that fully exploits player 2 to prove equilibrium existence. The idea for the construction in the proof is the following. First enumerate player 2's strategies as $\alpha_2^0, \alpha_2^1, \dots, \alpha_2^k, \dots$; against α_1^* , any two strategies α_2^k and α_2^l either give the same outcome across all periods or there is a finite period when the two strategies give different actions. In period 0, α_1^* assumes that player 2 is playing α_2^0 and chooses an optimal action against α_2^0 ; if that hypothesis is proved wrong at some period, α_1^* finds the minimum index in the list such that the associated strategy is consistent with the observed history and then uses that strategy as the working hypothesis to choose an optimal action. At any point of time α_1^* always has a working hypothesis about player 2's strategy and it looks for the next possible strategy when proved wrong. Because every strategy of player 2 appears in the list, player 2's strategy will be found out by α_1^* at a finite time and hence α_1^* fully exploits any strategy of player 2.

To ensure that α_1^* is θ^1 -computable, θ^1 has to be sufficiently powerful relative to

θ^2 . Indeed, although the Enumeration Theorem gives a θ^2 -effective list of θ^2 -computable functions, it does not tell which ones are *total* and hence qualify as strategies. The assumption that $(\theta^2)^H$ is θ^1 -computable ensures that when α_1^* finds a potential hypothesis α_2^k for player 2's strategy, player 1 can check whether α_2^k is a valid strategy in the sense that it gives an action for the next period.

As a corollary of Proposition 2.1, if a secured equilibrium exists and neither player fully exploits the other in equilibrium, then each player's equilibrium strategy is not computable relative to the opponent's oracle and hence the two players' oracles are incomparable in terms of the computability relation. This result gives a necessary condition for the existence of a secured equilibrium where neither player fully exploits the other: $\mathcal{C}(\theta^i) - \mathcal{C}(\theta^{-i}) \neq \emptyset$ for both $i = 1, 2$.

To discuss unpredictability I focus on the case where $V_i(g, \theta^1, \theta^2)$ is equal to the value v_i^* of the stage game for both i 's. Conditions for existence of such an equilibrium will be established in the next section, which requires other structures than the computability relation. Here I give a characterization of equilibrium history-independent strategies in such an equilibrium. A strategy $\alpha_i : X_{-i}^* \rightarrow X_i$ is history-independent if α_i gives the same action for all histories $\sigma \in X_{-i}^*$ with the same length; such a strategy can be identified with an infinite sequence $\xi^i = (\xi_0^i, \xi_1^i, \dots) \in X_i^{\mathbb{N}}$.

A concept called selection functions is necessary for the characterization. Given a finite set X , a *selection function for X* is a total function $r : X^* \rightarrow \{0, 1\}$. r can be used to choose a subsequence ξ^r from a sequence $\xi = (\xi_0, \xi_1, \xi_2, \dots) \in X^{\mathbb{N}}$ as follows: for all $t \in \mathbb{N}$, $\xi_t^r = \xi_{g(t)}$, where $g(0) = \min\{t : r(\xi[t]) = 1\}$, and $g(t) = \min\{s : r(\xi[s]) = 1, s > g(t-1)\}$ for $t > 0$.⁶ $\xi[t] = (\xi_0, \xi_1, \dots, \xi_{t-1})$ is the initial segment of ξ with length t . Such a selection function might not produce an infinite subsequence, but only a finite initial segment. I use $\Delta(X)$ to denote the set of probability distributions over X with rational probability values, i.e., $\Delta(X) = \{p \in ([0, 1] \cap \mathbb{Q})^X : \sum_{x \in X} p[x] = 1\}$.

Theorem 2.1 (Equilibrium characterization). *Suppose that a secured equilibrium exists in $RG(g, \theta^1, \theta^2)$ with $V_i(g, \theta^1, \theta^2) = v_i^*$ for both $i = 1, 2$, where v_i^* is the value of g for i .*

⁶If at some t , there is no $s > g(t-1)$ such that $r(\xi[s]) = 1$, then $g(t)$ is undefined.

(a) A θ^i -computable sequence $\xi \in X_i^{\mathbb{N}}$ is an equilibrium strategy in $RG(g, \theta^1, \theta^2)$ if for any θ^{-i} -computable selection function r for X_i such that ξ^r is an infinite sequence, there is an equilibrium mixed strategy $p^i \in \Delta(X_i)$ such that $c_x(x) = 1$ and $c_x(y) = 0$ if $x \neq y$)

$$\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{c_{x_i}(\xi_t^r)}{T} = p^i[x_i] \text{ for all } x_i \in X_i. \quad (2)$$

(b) Let $\xi \in X_i^{\mathbb{N}}$ be an equilibrium strategy in $RG(g, \theta^1, \theta^2)$. For any θ^{-i} -computable selection function r for X_i such that $\liminf_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{r(\xi[t])}{T} > 0$, the limit of any convergent subsequence of $\{(\sum_{t=0}^{T-1} \frac{c_{x_i}(\xi_t^r)}{T})_{x_i \in X_i}\}_{T \in \mathbb{N}}$ is an equilibrium mixed strategy of g .

(c) Suppose that g has a unique equilibrium (p^1, p^2) . If $\xi \in X_i^{\mathbb{N}}$ is an equilibrium strategy in $RG(g, \theta^1, \theta^2)$, then for any θ^{-i} -computable selection function r for X_i such that $\liminf_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{r(\xi[t])}{T} > 0$, $\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{c_{x_i}(\xi_t^r)}{T} = p^i[x_i]$ for all $x_i \in X_i$.

The proof of Theorem 2.1 (a) is based on the following observation. For a candidate equilibrium (history-independent) strategy ξ^i for player i , its payoff against a strategy $\alpha_j \in \mathcal{A}_j$ depends on the collection of selection functions $\{r^y\}_{y \in X_j}$ such that $r^y(\sigma) = 1$ if $\alpha_j(\sigma) = y$ and $r^y(\sigma) = 0$ otherwise for each $\sigma \in X_i^*$. Notice that each r^y is θ_j -computable because α_j is. Indeed, for any T , the payoff $\sum_{t=0}^{T-1} \frac{h_i(\xi_t^i; \alpha_j(\xi^i[t]))}{T}$ is a weighted average of $\{\sum_{t=0}^{T_y-1} \frac{h_i((\xi^i)^{r^y}; y)}{T_y}\}_{y \in X_j}$, where $T_y = \sum_{t=0}^{T-1} r^y(\xi^i[t])$. Condition (2) ensures that for any y , the average payoff along the subsequence $(\xi^i)^{r^y}$ against y is greater or equal to the value v_i^* of the stage game and hence the overall payoff for ξ^i against α_j is also at least v_i^* .

The necessary conditions in (b) and (c) are valid only for equilibrium strategies in a secured equilibrium. Their reasoning is also based on the above observation. If a convergent subsequence of $\{(\sum_{t=0}^{T-1} \frac{c_{x_i}(\xi_t^r)}{T})_{x_i \in X_i}\}_{T \in \mathbb{N}}$ fails to converge to a mixed equilibrium of the stage game, then r can be used to devise a strategy α_j that gives player i a payoff that is strictly lower than v_i^* . Notice that the statement (c) is a directly corollary of (b) for repeated games whose stage games have unique equilibria.

Assuming that players adopt history-independent strategies, Theorem 2.1 implies that the limit frequencies in equilibrium strategies are consistent with equilibrium mixed strategies of the stage game. This frequency implication extends also to subsequences of the

equilibrium strategies that can be effectively selected by the opponent. One such subsequence is to select places where the opponent chooses a particular action, and this implies that, in equilibrium, the average payoffs are the same along each subsequence of plays where the opponent uses a particular action. This test has been used in the empirical literature to test the equilibrium hypothesis in the context of repeated zero-sum games. While that literature typically employs other tests, such as those testing serial independence, our characterization result suggests that tests along this line are the only relevant ones for the equilibrium hypothesis.

3 Complexity and Unpredictability

3.1 Kolmogorov complexity and existence

Here I show that a secured equilibrium exists with the values equal to those of the stage game when both players' oracles are sufficiently complex relative to each other. The sufficient condition is based on incompressible sequences introduced in Section 2.1. Recall that an oracle η is incompressible relative to another oracle θ (called θ -incompressible hereafter) if for some d , $K_\theta(\eta[n]) > n - d$ for all $n \in \mathbb{N}$.

Definition 3.1. Two oracles θ^1 and θ^2 are *mutually complex* if there are oracles $\nu^1 \in \mathcal{C}(\theta^1)$ and $\nu^2 \in \mathcal{C}(\theta^2)$ such that for both $i = 1, 2$, ν^i is θ^{-i} -incompressible.

As mentioned in Section 2.1, the set of incompressible sequences relative to a fixed oracle η has measure 1. Van Lambalgen's Theorem ([5], Theorem 6.9.1) shows that if θ is incompressible, then θ^1 is θ^2 -incompressible and θ^2 is θ^1 -incompressible, where $\theta_t^1 = \theta_{2t}$ and $\theta_t^2 = \theta_{2t+1}$ for all $t \in \mathbb{N}$; hence, θ^1 and θ^2 are mutually complex. As a result, this implies existence of mutually complex oracles; moreover, it shows that there are uncountably many pairs of oracles satisfying mutual complexity, as stated in the following proposition. Its proof is omitted as it directly follows from the above discussions.

Proposition 3.1. *There are uncountably many pairs of oracles that satisfy mutual complexity.*

The following theorem states that mutual complexity is a sufficient condition for the existence of a secured equilibrium with value equal to that of the stage game.

Theorem 3.1 (Existence). *Let g be a finite zero-sum game. Suppose that θ^1 and θ^2 are mutually complex. Then there exists a secured equilibrium in $RG(g, \theta^1, \theta^2)$ with $V_i(g, \theta^1, \theta^2) = v_i^*$ for both $i = 1, 2$, where v_i^* is the value of g for i .*

This existence result does not depend on the limit inferior criterion. In fact, the theorem remains the same if limit superior is adopted or anything that lies between limit inferior and limit superior. Moreover, a similar existence result also holds for any finite N -person stage games (see Section 4).

Recall that by Theorem 2.1 (a), a θ^i -computable sequence satisfying the condition (2) for any θ^{-i} selection function r gives a minimal payoff that is at least equal to the equilibrium payoff v_i^* of the stage game g . It is known that (see also Lemma 3.1 and related discussions below) if ν^i is θ^{-i} -incompressible for both $i = 1, 2$, then ν^i satisfies $\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{(\nu^i)_t^r}{T} = \frac{1}{2}$ for any θ^{-i} -computable selection function r for $\{0, 1\}$ such that $(\nu^i)^r$ is an infinite sequence. Therefore, if g is the matching pennies game, then (ν^1, ν^2) is a secured equilibrium by Theorem 2.1 (a). However, (ν^1, ν^2) is not an equilibrium for other games such as $g^0 = \langle \{x_1, x_2\}, \{y_1, y_2\}, h_1, h_2 \rangle$ given by $h_1(x_1, y_1) = 2 = h_1(x_2, y_2)$, $h_1(x_1, y_2) = -1$, and $h_1(x_2, y_1) = -4$. The unique equilibrium in g^0 is (p^1, p^2) such that $p^1[x_1] = \frac{1}{3} = p^2[y_1]$ and $p^1[x_2] = \frac{2}{3} = p^2[y_2]$, and hence, by Theorem 2.1, an equilibrium strategy ξ for player 1 has limit frequency described by p^1 but ν^1 has frequency $(\frac{1}{2}, \frac{1}{2})$.

Nevertheless, ν^1 can still be used to compute a strategy ξ with appropriate limit frequencies; in the following I illustrate the construction by considering $X_1 = \{x_1, x_2\}$ and $p^1[x_1] = \frac{1}{3}$ as in the above example g^0 . Beginning from the θ^2 -incompressible sequence $\nu^1 \in \{0, 1\}^{\mathbb{N}}$, the construction takes three steps.

(1) Convert ν^1 into ζ over $\{w_1, w_2, w_3, w_4\}$ as follows: $\zeta_n = w_1$ if $(\nu_{2n}^1, \nu_{2n+1}^1) = (0, 0)$, $\zeta_n = w_2$ if $(\nu_{2n}^1, \nu_{2n+1}^1) = (0, 1)$, $\zeta_n = w_3$ if $(\nu_{2n}^1, \nu_{2n+1}^1) = (1, 0)$, and $\zeta_n = w_4$ if $(\nu_{2n}^1, \nu_{2n+1}^1) = (1, 1)$. The limit frequency in ζ is described by $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$

(2) Convert ζ into η over $\{w_1, w_2, w_3\}$ by dropping the action w_4 in ζ . The limit frequency

in η is then described by $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

(3) Convert η into ξ over $\{x_1, x_2\}$ by replacing action w_1 in η by x_1 in ξ and by combining actions w_2 and w_3 in η into a single action x_2 in ξ ; that is, $\xi_n = x_1$ if $\eta_n = w_1$ and $\xi_n = x_2$ otherwise. The limit frequency in ξ is then described by $(\frac{1}{3}, \frac{2}{3})$.

In fact, for any finite set X and any $p \in \Delta(X)$, ν^1 can be used to compute a sequence with limit frequency p following a similar three-step procedure as above. It is rather straightforward to check the constructed sequence has the desired limit frequencies, but it is less obvious to see whether the limit frequencies are constant over all subsequences selected by any θ^2 -computable selection functions, as required by condition (2) in Theorem 2.1. Indeed, the main task in proving Theorem 3.1 is to demonstrate that this condition holds for the constructed sequence. The main tool for the proof is a concept of (algorithmic) randomness based on *betting functions* that is connected to incompressibility.

Given a finite set X and a distribution $p \in \Delta(X)$, a *betting function* over X for p is a nonnegative function $B : X^* \rightarrow \mathbb{R}_+$ such that $B(\sigma) = \sum_{x \in X} p[x]B(\sigma x)$ for all $\sigma \in X^*$. Consider $X = \{0, 1\}$ and $p = (\frac{1}{2}, \frac{1}{2})$. Then a betting function B describes the stakes owned by the gambler by betting $B(\sigma 0) - B(\sigma)$ on 0 after observing the partial history σ in a fair gamble when the underlying process is i.i.d. $(\frac{1}{2}, \frac{1}{2})$. Other betting functions have similar interpretations. Intuitively, if a sequence $\xi = (\xi_0, \xi_1, \dots, \xi_t, \dots) \in X^{\mathbb{N}}$ is random, then there should be no betting function that generates unbounded payoffs by playing against ξ . Formally, a betting function B for p is said to *succeed* over an infinite sequence $\xi \in X^{\mathbb{N}}$ if B reaches unbounded payoffs by betting against ξ , that is, $\limsup_{n \rightarrow \infty} B(\xi[n]) = \infty$. The set of infinite sequences over which B succeeds is denoted by $\text{succ}(B)$.

Of course, because the sequence ξ is deterministic, there always exists a betting function B that succeeds over ξ . However, this concept becomes useful when considering betting functions that are complexity-constrained by a fixed oracle. Given an oracle θ , a betting function B for $p \in \Delta(X)$ is said to be θ -*effective* if it can be computably approximated from below relative to θ , that is, if there is a θ -computable function $C : \mathbb{N} \times X^* \rightarrow \mathbb{Q}$ such that $\lim_{s \rightarrow \infty} C(s, \sigma) = B(\sigma)$, $C(s+1, \sigma) \geq C(s, \sigma)$ for all $(s, \sigma) \in \mathbb{N} \times X^*$, and $C(s, \cdot)$

is a betting function for p for each s . If for a sequence $\xi \in X^{\mathbb{N}}$, there is no θ -effective betting function for $p \in \Delta(X)$ that succeeds over ξ , then we may say that ξ is θ -random for p . This terminology is borrowed from the algorithmic randomness literature (see Downey *et al.* [6] or Nies [17] for a survey). This notion of randomness is closely connected to incompressibility, as the following lemma shows. This lemma is a direct implication of the relativized versions of Theorem 3.2.9 and Proposition 7.2.6 in Nies [17] and hence its proof is omitted. It shows that incompressibility can be a substitute for randomness.

Lemma 3.1. *Let θ be an oracle. $\nu \in \{0, 1\}^{\mathbb{N}}$ is θ -incompressible if and only if ν is θ -random for $(\frac{1}{2}, \frac{1}{2})$.*

To prove Theorem 3.1, first I show, in Lemma 5.1, that for any finite set X and any $p \in \Delta(X)$, a θ -incompressible sequence ν can be used to compute a θ -random sequence $\xi \in X^{\mathbb{N}}$ for p , following the three-step construction mentioned previously. Moreover, Lemma 5.2 shows that if $\xi \in X^{\mathbb{N}}$ is a θ -random sequence for p , then it satisfies the condition (2) in Theorem 2.1; that is, $\lim_{T \rightarrow \infty} \frac{\sum_{t=0}^{T-1} c_x(\xi_t)}{T} = p[x]$ for each $x \in X$ for any θ -computable selection function r . Therefore, given the equilibrium (p^1, p^2) of the stage game g , for both $i = 1, 2$, the θ^{-i} -incompressible sequence ν^i can be used to compute a θ^{-i} -random sequence ξ^i for p^i , and hence (ξ^1, ξ^2) is a secured equilibrium.

3.2 Criterion for unpredictability

In the last section I show that under mutual complexity, each player i has a θ^{-i} -random history-independent strategy whose limit frequency is consistent with an equilibrium mixed strategy of the stage game. However, the characterization result, Theorem 2.1, shows that the condition for being an equilibrium strategy for player i in $R(g, \theta^1, \theta^2)$ is weaker than θ^{-i} -randomness. The following theorem shows that there are always equilibrium strategies for player i that do not satisfy θ^{-i} -randomness.

Theorem 3.2. *Suppose that there exists a secured equilibrium in $RG(g, \theta^1, \theta^2)$ with values $V_i(g, \theta^1, \theta^2) = v_i^*$ for both i 's, where v_i^* is the value of g for i . For each i , an equilibrium strategy ξ^i exists such that ξ^i is not θ^{-i} -random for any non-degenerate $p \in \Delta(X_i)$.*

Intuitively, a sequence is not random for $p \in \Delta(X)$ if it has a pattern that allows a betting function to exploit it; such a pattern may be regarded as a statistical irregularity that can be used to reject the hypothesis that the sequence is generated by an i.i.d. process. Thus, even when unpredictable behavior emerges in equilibrium, an equilibrium strategy does not have to be like an i.i.d. sequence to be sufficiently unpredictable. Notice that in Theorem 3.2 mutual complexity is not assumed.

The proof of Theorem 3.2 goes as follows. Consider a fixed player i . Assuming that a secured equilibrium exists with the same equilibrium payoffs as the stage game, there are two possible situations: (1) for any non-degenerate distribution $p \in \Delta(X_i)$, any θ^{-i} -random sequence ξ for p is not θ^i -computable; (2) for some non-degenerate distribution $p \in \Delta(X_i)$, there is a θ^i -computable sequence ξ that is also θ^{-i} -random sequence ξ for p . Situation (1) is not excluded by Theorem 3.1 because mutual complexity is only a sufficient condition, but Theorem 3.2 trivially holds if the oracles are in this situation. On the other hand, if the oracles are in situation (2), it can be shown that the oracle θ^i can be used to compute a θ^{-i} -incompressible sequence. In this case, there are equilibrium strategies that violate a particular statistical regularity for i.i.d. sequences called the Law of the Iterated Logarithm (LIL). Consider a finite set X and a distribution $p \in \Delta(X)$. LIL states that the following condition holds for almost all ξ in $X^{\mathbb{N}}$ (with respect to μ_p):

$$\limsup_{T \rightarrow \infty} \frac{|\sum_{t=0}^T c_x(\xi_t) - Tp[x]|}{\sqrt{2p[x](1-p[x])T \log \log T}} = 1. \quad (3)$$

This law gives the exact convergence rate of the frequency in an i.i.d. sequence, but the following proposition shows that equilibrium strategies in the repeated game under mutual complexity can converge to the limit much more slowly than an i.i.d. sequence.

Proposition 3.2. *Suppose that θ^1 and θ^2 are mutually complex. For any non-degenerate equilibrium mixed strategy p^i of g , there exists an equilibrium strategy ξ^i in $RG(g, \theta^1, \theta^2)$ such that*

- (a) *it satisfies (2) for $X = X_i$, $p = p^i$, and $\theta = \theta^{-i}$;*
- (b) *for some $x \in X_i$ with $p^i[x] \in (0, 1)$,*

$$\lim_{T \rightarrow \infty} \frac{\sum_{n=0}^T c_x(\xi_n^i) - Tp^i[x]}{\sqrt{2p^i[x](1-p^i[x])T \log \log T}} = \infty. \quad (4)$$

Theorem 3.2 shows that, although certain patterns are detectable, it may not be feasible to transform them into a strategy that exploits them. Moreover, this result shows that, in repeated zero-sum games failure of certain statistical regularities does not entail the rejection of the equilibrium hypothesis. Indeed, Theorem 2.1 and Theorem 3.2 imply that only the limit frequencies along various subsequences are relevant for the equilibrium hypothesis; other patterns may be misleading to test the equilibrium hypothesis.

4 Concluding remarks

As mentioned in the introduction, here the complexity constraints are interpreted as constraints on implementation of strategies for players than as constraints on players' rationality. This interpretation is consistent with the machine game literature (see Chatterjee and Sabourian [4]). An alternative interpretation would understand equilibrium as the result of a learning process, and it seems possible to regard the complexity constraints as constraints on rationality under that interpretation. However, it requires a full model of learning process to formalize that interpretation.

Theorem 3.1 can be extended to arbitrary N -person finite games. The notion of mutual complexity can be extended to the N -person case: N oracles $(\theta^1, \dots, \theta^N)$ are *mutually complex* if for each i , θ^i is incompressible relative to the oracle that combines all the other oracles $(\theta^j)_{j \neq i}$.⁷ If the repeated game with complexity constraints satisfies mutual complexity, then for any mixed equilibrium of the stage game, there exists a Nash equilibrium consisting of history-independent strategies whose payoffs are the same as those of the stage game equilibrium.

Many results in the paper can be extended to finitely repeated games. One approach to accomplish this is to consider asymptotic properties of unpredictable behavior in long but finitely repeated games. To model complexity constraints in finitely repeated games, it is then necessary to impose resource restrictions on players' oracle machines but maintain the basic structure of mutual complexity. Under these assumptions, it is possible to obtain

⁷For a precise formulation and the formal result, see Section 5 in the supplemental material [10].

ε -equilibrium with values arbitrarily close to those of the stage game in a long but finitely repeated game with ε vanishing as the length of the game approaches infinity. Notice that if ε -equilibrium can be obtained in a finitely repeated game, then ε -equilibrium with discounting is not hard to get by manipulating the ε 's.

5 Proofs

Proof of Lemma 2.1: First notice that for any (α_1, α_2) , $u_1(\alpha_1, \alpha_2) + u_2(\alpha_1, \alpha_2) \leq 0$:

$$\begin{aligned} u_1(\alpha_1, \alpha_2) &= \liminf_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{h_1(\xi_t^{\alpha_1, 1}, \xi_t^{\alpha_1, 2})}{T} \leq \limsup_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{h_1(\xi_t^{\alpha_1, 1}, \xi_t^{\alpha_1, 2})}{T} \\ &= -\liminf_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{-h_2(\xi_t^{\alpha_1, 1}, \xi_t^{\alpha_1, 2})}{T} = -u_2(\alpha_1, \alpha_2). \end{aligned}$$

Now, for each $\alpha_i \in \mathcal{A}_i$, $u_j(\alpha_j^*; \alpha_i) \geq s_j(\alpha_j^*) = v_j^*$. Hence, $u_i(\alpha_i^*; \alpha_j^*) \geq v_i^* = -v_j^* \geq -u_j(\alpha_j^*; \alpha_i) \geq u_i(\alpha_i; \alpha_j^*)$. So (α_1^*, α_2^*) is a Nash equilibrium. \square

Proof of Proposition 2.1: (a) Let $\{\phi_0, \phi_1, \dots, \phi_k, \dots\}$ be an θ^2 -computable enumeration of (partial) functions from X_1^* to X_2 which are computable relative to θ^2 . Such an enumeration exists by the relativized Enumeration Theorem. For any strategy $\alpha_2 \in \mathcal{A}$, there exists some k so that $\alpha_2 = \phi_k$; for any k , $\phi_k \in \mathcal{A}_2$ if and only if ϕ_k is total.

Define functions $g_1 : \bigcup_{n=0}^{\infty} (X_1^n \times X_2^n) \rightarrow \mathbb{N}$ and $g_2 : \bigcup_{n=0}^{\infty} (X_1^n \times X_2^n) \rightarrow X_1$ as follows.

$$g_1(\epsilon, \epsilon) = 0; \tag{5}$$

$$g_1(\sigma^1, \sigma^2) = \min\{k : \phi_k(\sigma^1) \text{ is defined and } \phi_k(\sigma^1[t]) = \sigma_t^2 \text{ for } t = 0, \dots, n-1\}$$

$$\text{if } (\sigma^1, \sigma^2) \in X_1^n \times X_2^n \text{ for some } n > 0;$$

$$g_2(\epsilon, \epsilon) = \arg \max_{x \in X_1} h_1(x, \phi_0(\epsilon)); \tag{6}$$

$$g_2(\sigma^1, \sigma^2) = \arg \max_{x \in X_1} h_1(x, \phi_{g_1(\sigma^1, \sigma^2)}(\sigma^1)) \text{ if } (\sigma^1, \sigma^2) \in X_1^n \times X_2^n \text{ for some } n > 0.$$

Here it is assumed that the maximizer for x is unique; if not, the x with the minimum index can be used. The functions g_1 and g_2 are computable relative to $(\theta^2)^H$ and hence are computable relative to θ^1 .

Define the strategy α_1^* as follows:

$$\alpha_1^*(\epsilon) = g_2(\epsilon, \epsilon); \quad \alpha_1^*(\sigma^2) = g_2(\sigma^1, \sigma^2) \text{ with } \sigma_t^1 = \alpha_1^*(\sigma^2[t]) \text{ for } t = 0, \dots, |\sigma| - 1.$$

Notice that α_1^* is defined inductively. Now I show that for any strategy $\alpha_2 \in \mathcal{A}_2$, $u_1(\alpha_1^*, \alpha_2) \geq v_1 = \min_{x_2 \in X_2} \max_{x_1 \in X_1} h_1(x_1, x_2)$.

Consider an arbitrary strategy α_2 , and let $\alpha_2 = \phi_k$. Let $(\xi^{\alpha,1}, \xi^{\alpha,2})$ be the sequence of actions induced by (α_1^*, α_2) as defined in (1). First I show that there exist $\bar{T} \in \mathbb{N}$ and $l \leq k$ such that for all $T \geq \bar{T}$, $g_1(\xi^{\alpha,1}[T], \xi^{\alpha,2}[T]) = l$ and $\alpha_2(\xi^{\alpha,1}[t]) = \phi_l(\xi^{\alpha,1}[t])$ for all $t \in \mathbb{N}$. Let l be the smallest index such that $\phi_l(\xi^{\alpha,1}[t]) = \xi_t^{\alpha,2}$ for all $t \in \mathbb{N}$. Such l exists because $\alpha_2 = \phi_k$; hence, $l \leq k$. Moreover, for each $l' < l$, there exists some $T_{l'}$ such that either $\phi_{l'}(\xi^{\alpha,1}[T_{l'}])$ is undefined or $\phi_{l'}(\xi^{\alpha,1}[T_{l'}]) \neq \xi_{T_{l'}}^{\alpha,2}$. So for $T \geq \bar{T} = \max\{T_{l'} : l' < l\} + 1$, $g_1(\xi^{\alpha,1}[T], \xi^{\alpha,2}[T]) = l$. Moreover, for all $t \geq \bar{T}$, $h_1(\xi_t^{\alpha,1}, \xi_t^{\alpha,2}) \geq v_1$ by construction. Hence, $u_1(\alpha_1^*, \alpha_2) \geq v_1$. Finally, let α_2^* be such that $\alpha_2^*(\sigma) \in \arg \max_{x_2 \in X_2} \min_{x_1 \in X_1} h_2(x_1, x_2)$. Then, for any $\alpha_1 \in \mathcal{A}_1$, $u_2(\alpha_1, \alpha_2^*) \geq -v_1$.

(b) Suppose that there is a secured equilibrium. I first show that for any $\alpha_2 \in \mathcal{A}_2$, there is strategy $\alpha_1^* \in \mathcal{A}_1$ such that $u_1(\alpha_1^*, \alpha_2) \geq v_1 = \min_{x_2 \in X_2} \max_{x_1 \in X_1} h_1(x_1, x_2)$. Given α_2 , α_1^* is constructed as follows: for all $\sigma \in X_2^*$ with length t , $\alpha_1^*(\sigma) = \zeta_t$, with ζ defined as

$$\zeta_0 = \arg \max_{x \in X_1} h_1(x, \alpha_2(\epsilon)), \text{ and for } t > 0, \zeta_t = \arg \max_{x \in X_1} h_1(x, \alpha_2(\zeta[t])).$$

Here it is assumed that the maximizer for x is unique; if not, the x with the minimum index can be used. ζ is θ^2 -computable because α_2 is and hence α_1^* is θ^1 -computable because θ^2 is θ^1 -computable. By construction, for all $t \in \mathbb{N}$,

$$h_2(\zeta_t, \alpha_2(\zeta[t])) = -h_1(\zeta_t, \alpha_2(\zeta[t])) \leq -\min_{x_2 \in X_2} \max_{x_1 \in X_1} h_2(x_1, x_2) = -v_1.$$

Hence, $u_2(\alpha_1^*, \alpha_2) \leq \liminf_{T \rightarrow \infty} -v_1 = -v_1$. It follows that $s_2(\alpha_2) \leq -v_1$ for any $\alpha_2 \in \mathcal{A}_2$. Now, let $y^* \in \arg \max_{x_2 \in X_2} (\min_{x_1 \in X_1} h_2(x_1, x_2))$. Let $\alpha_2^* \in \mathcal{A}_2$ be such that $\alpha_2^*(\tau) = y^*$ for all $\tau \in X_1^*$. Then $s_2(\alpha_2^*) \geq -v_1$ and hence $\max_{\alpha_2 \in \mathcal{A}_2} s_2(\alpha_2) = -v_1$. \square

Proof of Theorem 2.1: (a) Suppose that ξ satisfies (2) for $i = 1$. It is sufficient to show that $s_1(\xi) \geq v_1^*$, that is, for all $\alpha_2 \in \mathcal{A}_2$,

$$\liminf_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{h_1(\xi_t, \alpha_2(\xi[t]))}{T} \geq v_1^*. \quad (7)$$

Let $\alpha_2 \in \mathcal{A}_2$ be given. For each $y \in X_2$, let $r^y : X_1^* \rightarrow \{0, 1\}$ be the selection function for X_1 such that $r^y(\sigma) = 1$ if $\alpha_2(\sigma) = y$, and $r^y(\sigma) = 0$ otherwise.

Define $L_y(T) = \#\{t \in \mathbb{N} : 0 \leq t \leq T-1, r^y(\xi[t]) = 1\}$ and $\xi^y = (\xi)^{r^y}$. It is easy to see that r^y is θ^2 -computable because α_2 is. Let

$$\mathcal{E}^1 = \{y \in X_2 : \lim_{T \rightarrow \infty} L_y(T) = \infty\} \text{ and } \mathcal{E}^2 = \{y \in X_2 : \lim_{T \rightarrow \infty} L_y(T) < \infty\}.$$

For each $y \in \mathcal{E}^2$, let $B_y = \lim_{T \rightarrow \infty} L_y(T)$ and let $C_y = \sum_{t=0}^{B_y-1} h_1(\xi_t^y, y)$. On the other hand, for any $y \in \mathcal{E}^1$, because ξ satisfies (2) and because r^y is a θ^2 -computable selection function, there is a mixed equilibrium strategy $p^1 \in \Delta(X_1)$ such that for any $x \in X_1$, $\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{c_x(\xi_t^y)}{T} = p^1[x]$. Therefore,

$$\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{h_1(\xi_t^y, y)}{T} = \lim_{T \rightarrow \infty} \sum_{x \in X_1} \sum_{t=0}^{T-1} \frac{c_x(\xi_t^y) h_1(x, y)}{T} = \sum_{x \in X_1} p^1[x] h_1(x, y) \geq v_1^*.$$

The last inequality comes from the fact that p^1 is an equilibrium mixed strategy of g .

I claim that for any $\varepsilon > 0$, there is some T' such that $T > T'$ implies that

$$\sum_{t=0}^{T-1} \frac{h_1(\xi_t, \alpha_2(\xi[t]))}{T} \geq v_1^* - \varepsilon. \quad (8)$$

Fix some $\varepsilon > 0$. Let T_1 be so large that $T > T_1$ implies that, for all $y \in \mathcal{E}^1$,

$$\sum_{t=0}^{T-1} \frac{h_1(\xi_t^y, y)}{T} \geq v_1^* - \frac{\varepsilon}{2|X_2|}, \quad (9)$$

and, for all $y \in \mathcal{E}^2$, $\frac{C_y}{T} > -\frac{\varepsilon}{2|X_2|}$. Let T' be so large that, for all $y \in \mathcal{E}_1$, $L_y(T') > T_1$ and $v^* \sum_{y \in \mathcal{E}_1} \frac{L_y(T)}{T} \geq v^* - \frac{\varepsilon}{2}$ for all $T > T'$. If $T > T'$, then

$$\begin{aligned} \sum_{t=0}^{T-1} \frac{h_1(\xi_t, \alpha_2(\xi[t]))}{T} &= \sum_{y \in \mathcal{E}_1} \frac{L_y(T)}{T} \sum_{t=0}^{L_y(T)-1} \frac{h_1(\xi_t^y, y)}{L_y(T)} + \sum_{y \in \mathcal{E}_2} \sum_{t=0}^{L_y(T)-1} \frac{h_1(\xi_t^y, y)}{T} \\ &\geq \sum_{y \in \mathcal{E}_1} \frac{L_y(T)}{T} \left(v_1^* - \frac{\varepsilon}{2|X_2|} \right) - \sum_{y \in \mathcal{E}_2} \frac{\varepsilon}{2|X_2|} \geq v_1^* - \varepsilon. \end{aligned}$$

Notice that L_y is weakly increasing, and $L_y(T) \leq T$ for all T . Thus, $T > T'$ implies that $L_y(T) \geq L_y(T') > T_1$, and so $T > T_1$. This proves (8), which implies (7).

(b) Let ξ be an equilibrium history-independent strategy for player $i = 1$. Because $\liminf_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{r(\xi[t])}{T} > 0$, there exist some $\varepsilon_0 > 0$ and some $\bar{T} \in \mathbb{N}$ such that for all

$T \geq \bar{T}$, $\sum_{t=0}^{T-1} \frac{r(\xi[t])}{T} > 2\varepsilon_0$. I will prove by contradiction. Suppose that for some infinite sequence $\{T_0^0 < T_1^0 < \dots < T_j^0 < \dots\}$ there exists $p^1 \in \Delta(X_1)$ that is not an equilibrium mixed strategy of g such that for all $x \in X_1$, $\lim_{j \rightarrow \infty} \frac{\sum_{t=0}^{T_j^0-1} r(\xi[t])c_x(\xi_t)}{\sum_{t=0}^{T_j^0-1} r(\xi[t])} = p^1[x]$. Consider two cases:

(b.1) $\lim_{T_j \rightarrow \infty} \sum_{t=0}^{T_j-1} \frac{r(\xi[t])}{T_j} = 1$. Define a strategy $\alpha_2 : X_1^* \rightarrow X_2$ as follows: $\alpha_2(\sigma) = y^1 \in \arg \min_{y \in X_2} h_1(p^1, y)$ if $r(\sigma) = 1$ and $\alpha_2(\sigma) = y^2$ for some arbitrary $y^2 \in X_2$ otherwise. Then, $\lim_{j \rightarrow \infty} \sum_{t=0}^{T_j-1} \frac{h_1(\xi_t, \alpha_2(\xi[t]))}{T_j} = h_1(p^1, y^1) < v_1^* - \delta < v_1^*$.

(b.2) For some $\varepsilon_2 > 0$, $\liminf_{T_j \rightarrow \infty} \sum_{t=0}^{T_j-1} \frac{r(\xi[t])}{T_j} < 1 - \varepsilon_2$. Now, let $\{T_j^1\}_{j=0}^\infty$ be a subsequence of $\{T_j\}_{j=0}^\infty$ such that

$$(1) \lim_{j \rightarrow \infty} \sum_{t=0}^{T_j^1-1} \frac{r(\xi[t])}{T_j^1} = a \in (0, 1);$$

$$(2) \text{ for each } x \in X_1, \lim_{j \rightarrow \infty} \frac{\sum_{t=0}^{T_j^1-1} (1-r(\xi[t]))c_x(\xi_t)}{\sum_{t=0}^{T_j^1-1} (1-r(\xi[t]))} = p^2[x] \text{ for some } p^2 \in \Delta(X_1).$$

Because p^1 is not an equilibrium mixed strategy of g , there exists some $\delta > 0$ so that $\min_{x_2 \in X_2} h_1(p^1, x_2) < v_1^* - \delta$. Define a strategy $\alpha_2 : X_1^* \rightarrow X_2$ as follows: $\alpha_2(\sigma) \in \arg \min_{y \in X_2} h_1(p^1, y)$ if $r(\sigma) = 1$ and $\alpha_2(\sigma) \in \arg \min_{y \in X_2} h_1(p^2, y)$ otherwise. In particular, let y^1 and y^2 be the actions chosen for $r(\sigma) = 1$ and $r(\sigma) = 0$, respectively. Then,

$$\begin{aligned} \lim_{j \rightarrow \infty} \sum_{t=0}^{T_j^1-1} \frac{h_1(\xi_t, \alpha_2(\xi[t]))}{T_j^1} &= \lim_{j \rightarrow \infty} \left[\frac{1}{T_j^1} \sum_{t=0}^{T_j^1-1} r(\xi[t]) \right] \sum_{x \in X_1} \left[\frac{\sum_{t=0}^{T_j^1-1} r(\xi[t])c_x(\xi_t)}{\sum_{t=0}^{T_j^1-1} r(\xi[t])} \right] h_1(x, y^1) \\ &+ \lim_{j \rightarrow \infty} \left[\frac{1}{T_j^1} \sum_{t=0}^{T_j^1-1} (1-r(\xi[t])) \right] \sum_{x \in X_1} \left[\frac{\sum_{t=0}^{T_j^1-1} (1-r(\xi[t]))c_x(\xi_t)}{\sum_{t=0}^{T_j^1-1} (1-r(\xi[t]))} \right] h_1(x, y^2) \\ &= ah_1(p^1, y^1) + (1-a)h_1(p^2, y^2) < a(v_1^* - \delta) + (1-a)v_1^* = v_1^* - \alpha\delta < v_1^*. \end{aligned}$$

Thus, in either case, the strategy α_2 is such that $u_1(\xi, \alpha_2) < v_1^*$, a contradiction to the optimality of ξ .

(c) Let ξ be an equilibrium history-independent strategy for player i . Because g has a unique mixed equilibrium $p = (p^1, p^2)$, any convergent subsequence of $\left\{ \frac{\sum_{t=0}^{T-1} r(\xi[t])c_{\bar{x}}(\xi_t)}{\sum_{t=0}^{T-1} r(\xi[t])} \right\}_{T \in \mathbb{N}}$ has limit p^i . Let $\varepsilon > 0$ be given. Fix some $\bar{x} \in X_i$. If for there is a subsequence $\{T_s\}_{s \in \mathbb{N}}$

such that

$$\left| \frac{\sum_{t=0}^{T_s-1} r(\xi[t])c_{\bar{x}}(\xi_t)}{\sum_{t=0}^{T_s-1} r(\xi[t])} - p^i[\bar{x}] \right| > \varepsilon$$

for all $s \in \mathbb{N}$, then there is another subsequence $\{T_v\}_{v \in \mathbb{N}}$ of $\{T_s\}_{s \in \mathbb{N}}$ such that

$$\lim_{v \rightarrow \infty} \frac{\sum_{t=0}^{T_v-1} r(\xi[t])c_x(\xi_t)}{\sum_{t=0}^{T_v-1} r(\xi[t])} = q[x] \text{ for each } x \in X$$

for some $q \in \Delta(X)$. But $|q[\bar{x}] - p[\bar{x}]| \geq \varepsilon$ and hence $q \neq p^i$, a contradiction to part (b).

Thus, for some \bar{T} , $T \geq \bar{T}$ implies that

$$\left| \frac{\sum_{t=0}^{T-1} r(\xi[t])c_{\bar{x}}(\xi_t)}{\sum_{t=0}^{T-1} r(\xi[t])} - p^i[\bar{x}] \right| \leq \varepsilon.$$

□

Proof of Theorem 3.1: Suppose, by mutual complexity, that $\nu^i \in \{0, 1\}^{\mathbb{N}}$ is θ^i -computable and is θ^{-i} -incompressible, for $i = 1, 2$. By Lemma 3.1, ν^i is also θ^{-i} -random for $(\frac{1}{2}, \frac{1}{2})$. Given an equilibrium (p^1, p^2) of the stage game g , I first show, in the following lemma, that ν^i can compute a sequence ξ^i that is θ^{-i} -random for p^i .

Lemma 5.1. *Suppose that $\nu \in \{0, 1\}^{\mathbb{N}}$ is θ -random for $(\frac{1}{2}, \frac{1}{2})$ for some oracle θ . For any finite set X and any $p \in \Delta(X)$, there is a $\xi \in \mathcal{C}(\nu)$ that is θ -random for p .*

Proof. Let $X = \{x_1, \dots, x_K\}$ and let $p[x_i] = \frac{l_i}{L}$, where $l_1, \dots, l_K, L \in \mathbb{N}$. Let m be such that $2^{m-1} < L \leq 2^m$. The construction takes three steps: first step transforms ν into ζ , which is θ -random over a set of 2^m actions for the uniform distribution; the second step transforms ζ into η by dropping $2^m - L$ actions from ζ and makes η a θ -random sequence over a set of L actions for the uniform distribution; finally, the third step transforms η into ξ by combining l_i actions in η into a single action x_i in ξ and makes ξ a θ -random sequence over X for p .

(Step 1). Let $W = \{w_1, \dots, w_{2^m}\}$. Enumerate the set $\{0, 1\}^m$ as $\{\rho^1, \dots, \rho^{2^m}\}$. Construct $\zeta \in W^{\mathbb{N}}$ from ν as follows:

$$\text{For each } n \in \mathbb{N}, \zeta_n = w_i \text{ if } (\xi_{nm}^0, \xi_{nm+1}^0, \dots, \xi_{(n+1)m-1}^0) = \rho^i.$$

I show that ζ is θ -random for $(2^{-m}, 2^{-m}, \dots, 2^{-m})$ by contradiction. Suppose that there exists a θ -effective betting function B for $(2^{-m}, 2^{-m}, \dots, 2^{-m})$ that succeeds over ζ . I will device a betting function C that succeeds over ν .

Define the mapping $\Gamma^1 : \bigcup_{N=0}^{\infty} \{0, 1\}^{Nm} \rightarrow W^*$ by setting $\Gamma^1(\sigma) = \tau$ with $\tau_n = w_i$ if $|\sigma| = Nm$ and $(\sigma_{nm}, \sigma_{nm+1}, \dots, \sigma_{(n+1)m-1}) = \rho^i$ for each $n = 0, 1, \dots, N-1$. Construct C as follows.

(1.1) For all σ with $|\sigma| = Nm$ for some $N \in \mathbb{N}$, $C(\sigma) = B(\Gamma^1(\sigma))$.

(1.2) Suppose that C is defined over all strings σ 's with length $Nm - (k-1)$, $1 \leq k < m-1$. Consider a string σ with $|\sigma| = Nm - k$. Take $C(\sigma) = \frac{1}{2}C(\sigma 0) + \frac{1}{2}C(\sigma 1)$.

C is a betting function. First, by construction, $C(\sigma) = \frac{1}{2}C(\sigma 0) + \frac{1}{2}C(\sigma 1)$ holds for all σ with $|\sigma| \neq Nm$ for any N . Moreover, if $|\sigma| = Nm - k$ for some k between 1 and $m-1$, then $C(\sigma) = \sum_{\rho \in \{0,1\}^k} 2^{-k} C(\sigma \rho)$. Thus, if $|\sigma| = Nm$, then $C(\sigma 0) = \sum_{\rho \in \{0,1\}^{m-1}} \frac{1}{2^{m-1}} C(\sigma 0 \rho)$ and $C(\sigma 1) = \sum_{\rho \in \{0,1\}^{m-1}} \frac{1}{2^{m-1}} C(\sigma 1 \rho)$, and hence

$$\frac{1}{2}C(\sigma 0) + \frac{1}{2}C(\sigma 1) = \sum_{\rho \in \{0,1\}^m} \frac{1}{2^m} C(\sigma \rho) = \sum_{w \in W} \frac{1}{2^m} B(\Gamma^1(\sigma)w) = B(\Gamma^1(\sigma)) = C(\sigma).$$

Finally, C is θ -effective because B is and Γ is computable. Moreover, for any n , $B(\zeta[n]) = C(\Gamma^{-1}(\zeta[n])) = C(\nu[nm])$. C succeeds over ν because B succeeds over ζ , a contradiction.

(Step 2). Because ζ is θ -random for the uniform distribution, there are infinitely many k 's such that $\zeta_k = w_1$. I show this by contradiction; suppose that $\zeta_n \neq w_1$ for all $n > N$. Let B be a computable betting function for $(2^{-m}, 2^{-m}, \dots, 2^{-m})$ defined as follows: (a) $B(\epsilon) = 1 = B(\sigma)$ if $|\sigma| \leq N$; (b) $B(\sigma w_1) = 0$ and $B(\sigma w_i) = \frac{2^m}{2^m-1} B(\sigma)$ for all $i > 1$ if $|\sigma| > N$. $B(\zeta[N+n]) = (\frac{2^m}{2^m-1})^n \rightarrow \infty$ as $n \rightarrow \infty$, a contradiction.

Let $Z = \{w_1, \dots, w_L\} \subset W$. Construct $\eta \in Z^{\mathbb{N}}$ from ζ as follows. First define $g : \mathbb{N} \rightarrow \mathbb{N}$ by (a) $g(0) = \min\{k \in \mathbb{N} : \xi_k^1 \in Z\}$; (b) for $n \geq 0$, $g(n+1) = \min\{k > g(n) : \xi_k^1 \in Z\}$. g is total. Then define η by setting $\eta_n = \zeta_{g(n)}$ for all $n \in \mathbb{N}$.

η is θ -random for the uniform distribution $(\frac{1}{L}, \dots, \frac{1}{L})$ over Z . To show this, suppose, by contradiction, that a θ -effective betting function B for $(\frac{1}{L}, \dots, \frac{1}{L})$ succeeds over η . Construct a θ -effective betting function C for $(2^{-m}, 2^{-m}, \dots, 2^{-m})$ as follows.

(2.1) Define $\Gamma^2 : W^* \rightarrow Z^*$ by setting $\Gamma^2(\sigma) = \tau$, where τ is obtained from σ by eliminating all the occurrences of w_{L+1}, \dots, w_{2^m} in σ .

(2.2) Define C by setting $C(\sigma) = B(\Gamma^2(\sigma))$ for all $\sigma \in W^*$.

By construction, for any $\sigma \in W^*$, $C(\sigma w_i) = B(\Gamma^2(\sigma)w_i)$ if $i \leq L$ and $C(\sigma w_i) = B(\Gamma^2(\sigma))$ if $i > L$. C is a betting function for $(2^{-m}, \dots, 2^{-m})$: let $\sigma \in W^*$,

$$\begin{aligned} \sum_{i=1}^{2^m} 2^{-m} C(\sigma w_i) &= \left[\sum_{i \leq L} L^{-1} B(\Gamma^2(\sigma)w_i) \right] \frac{L}{2^m} + \left[\sum_{j > L} 2^{-m} B(\Gamma^2(\sigma)) \right] \\ &= \frac{L}{2^m} B(\Gamma^2(\sigma)) + \frac{2^m - L}{2^m} B(\Gamma^2(\sigma)) = B(\Gamma^2(\sigma)) = C(\sigma). \end{aligned}$$

C is θ -effective because B is. Finally, $C(\zeta[g(n)]) = B(\eta[n])$ for all $n \in \mathbb{N}$. So C succeeds over ζ because B succeeds over η , a contradiction.

(Step 3). Let $X = \{x_1, \dots, x_K\}$. Construct $\xi \in X^{\mathbb{N}}$ from η as follows. Let $L_0 = 0$; for $k \geq 0$, let $L_{k+1} = L_k + l_{k+1}$. Define ξ by setting $\xi_n = x_k$ if $\zeta_n \in \{w_{L_{k-1}+1}, \dots, w_{L_k}\}$ for all $n \in \mathbb{N}$. Now I show that ξ is θ -random for $p = (\frac{l_1}{L}, \dots, \frac{l_K}{L})$.

Suppose, by contradiction, that a θ -effective betting function B over X for p succeeds over ξ . Construct a θ -effective betting function C over Z as follows.

(3.1) Define $\Gamma^3 : Z^* \rightarrow X^*$ by setting $\Gamma^3(\sigma) = \tau$ with $\tau_n = x_k$ if $\sigma_n \in \{w_{L_{k-1}+1}, \dots, w_{L_k}\}$ for all $n = 0, \dots, |\sigma| - 1$.

(3.2) Define C by setting $C(\sigma) = B(\Gamma^3(\sigma))$ for all $\sigma \in Z^*$.

Thus, for any $\sigma \in Z^*$, $C(\sigma w_i) = B(\Gamma^3(\sigma)x_k)$ if $i \in \{L_{k-1} + 1, \dots, L_k\}$. C is a betting function for $(\frac{1}{L}, \dots, \frac{1}{L})$: for any $\sigma \in Z^*$,

$$\sum_{i=1}^L \frac{1}{L} C(\sigma w_i) = \sum_{k=1}^K \frac{l_k}{L} B(\Gamma^3(\sigma)x_k) = \sum_{k=1}^K p[x_k] B(\Gamma^3(\sigma)x_k) = B(\Gamma^3(\sigma)) = C(\sigma).$$

C is θ -effective because B is. Finally, $C(\eta[n]) = B(\xi[n])$ for all $n \in \mathbb{N}$. So C succeeds over η because B succeeds over ξ , a contradiction. Thus, ξ is θ -random for p . ξ is ν -computable because both ζ and η are. \square

Then, I give a lemma that shows θ -random sequences for p satisfies condition (2) in Theorem 2.1 for any θ -computable selection function r .

Lemma 5.2. *Let X be a finite set and let $p \in \Delta(X)$. If $\xi \in X^{\mathbb{N}}$ is θ -random for p , then*

$$\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{c_x(\xi_t^r)}{T} = p[x] \text{ for all } x \in X \quad (10)$$

for any θ -computable selection function r for X such that ξ^r is an infinite sequence.

Proof. First I show that if ξ is θ -random for p , and if r is a θ -computable selection function for X such that ξ^r is infinite, then ξ^r is also θ -random for p . Suppose, by contradiction, that a θ -effective betting function B defeats ξ^r . Construct a betting function C for p as follows: $C(\sigma) = B(\sigma^r)$ for all $\sigma \in X^*$.

C is also a betting function for p : let $\sigma \in X^*$ be given; if $r(\sigma) = 0$, then $(\sigma x)^r = \sigma^r$ and hence $C(\sigma x) = C(\sigma)$ for all x , which implies that $\sum_{x \in X} p[x]C(\sigma x) = C(\sigma)$; if $r(\sigma) = 1$, then $(\sigma x)^r = \sigma^r x$ and hence $C(\sigma x) = B(\sigma^r x)$ for all x , which implies that $\sum_{x \in X} p[x]C(\sigma x) = \sum_{x \in X} p[x]B(\sigma^r x) = B(\sigma^r) = C(\sigma)$.

C is θ -effective because B is θ -effective and r is θ -computable. Finally, for any $n \in \mathbb{N}$, there exists $m_n \geq n$ such that $(\xi[m_n])^r = \xi^r[n]$. Thus, for any n , $C(\xi[m_n]) = B(\xi^r[n])$. Thus, C succeeds over ξ because B succeeds over ξ^r , a contradiction.

Now I show that if ξ is θ -random for p , then $\lim_{T \rightarrow \infty} \frac{\sum_{t=0}^{T-1} c_x(\xi_t)}{T} = p[x]$ for each $x \in X$. Suppose, by contradiction, that for some $y \in X$ with $p[y] \in (0, 1)$ and some $\varepsilon > 0$, there is an infinite sequence $T_0 < T_1 < \dots < T_j < \dots$ such that $\lim_{T_j \rightarrow \infty} \frac{\sum_{t=0}^{T_j-1} c_y(\xi_t)}{T_j} = q_y < p[y] - 2\varepsilon$. Let $d = \frac{1}{2(1-p[y])}$; construct a computable betting function B for p as follows:

Take $B(\epsilon) = 1$; $B(\sigma y) = (1 - d(1 - p[y]))B(\sigma)$ and $B(\sigma x) = (1 + dp[y])B(\sigma)$ for all $x \neq y$.

B is a betting function for p :

$$\sum_{x \in X} p[x]B(\sigma x) = p[y](1 - d(1 - p[y]))B(\sigma) + \sum_{x \neq y} p[x](1 + dp[y])B(\sigma) = B(\sigma).$$

Let $C \equiv (1 - d(1 - p[y]))^{p[y]-\varepsilon} (1 + dp[y])^{1-p[y]+\varepsilon}$. It is straightforward to check that $C > 1$. Let J be so large that $j > J$ implies that $\frac{\sum_{t=0}^{T_j-1} c_y(\xi_t)}{T_j} < p[y] - \varepsilon$. For all $j > J$,

$$B(\xi[T_j]) = \left[(1 - d(1 - p[y]))^{\frac{\sum_{t=0}^{T_j-1} c_y(\xi_t)}{T_j}} (1 + dp[y])^{1 - \frac{\sum_{t=0}^{T_j-1} c_y(\xi_t)}{T_j}} \right]^{T_j}$$

$$\geq [(1 - d(1 - p[y]))^{p[y]-\varepsilon} (1 + dp[y])^{1-p[y]+\varepsilon}]^{T_j} = C^{T_j}$$

Thus, $\lim_{j \rightarrow \infty} B(\xi[T_j]) \geq \lim_{j \rightarrow \infty} C^{T_j} = \infty$, that is, B succeeds over ξ , a contradiction. \square

Proof of Theorem 3.1: Suppose that $\nu^i \in \{0, 1\}^{\mathbb{N}}$ is both θ^i -computable and is θ^i -incompressible. Let p^i be an equilibrium mixed strategy of g and let v_i^* be the value of g for player i . By Lemma 3.1 ν^i is also θ^i -random for $(\frac{1}{2}, \frac{1}{2})$. By Lemma 5.1, there is a ν^i -computable sequence ξ^i that is θ^i -random for p^i . Finally, by Lemma 5.2, for any θ^i -computable selection function r for X_i , the convergence condition (2) in Theorem 2.1 holds if $(\xi^i)^r$ is an infinite sequence. Thus, by the arguments in the proof of Theorem 2.1 (a), $s_i(\xi^i) \geq v_i^*$. Thus, (ξ^1, ξ^2) is a secured equilibrium. \square

Proof of Theorem 3.2: (Sketch) Suppose that there exists some θ^i -computable sequence ζ^i that is θ^i -random for some non-degenerate $p \in \Delta(X_i)$. Then, a θ^i -random sequence $\eta^i \in \{0, 1\}^{\mathbb{N}}$ for $(\frac{1}{2}, \frac{1}{2})$ can be constructed from ζ^i so that η^i is ζ^i -computable. The construction works as follows. Let $y \in X$ be such that $p[y] \in (0, 1)$. Define $\eta^0 \in \{0, 1, 2\}^{\mathbb{N}}$ by $\eta_n^0 = 0$ if $\zeta_{2n} = y$ and $\zeta_{2n+1} \neq y$, $\eta_n^0 = 1$ if $\zeta_{2n} \neq y$ and $\zeta_{2n+1} = y$, $\eta_n^0 = 2$ otherwise. Then η^0 is θ^i -random for $(p[y](1 - p[y]), p[y](1 - p[y]), 1 - 2p[y](1 - p[y]))$. η^i is obtained from η^0 by deleting all occurrences of 2 in η^0 . Then, η^i is θ^i -random for $(\frac{1}{2}, \frac{1}{2})$. Thus, an equilibrium strategy ξ^i exists which is computable from η^i and which fails LIL as in (4) by Proposition 3.2. \square

Proof of Proposition 3.2: (Sketch) The notion of betting functions can be extended to other measures; in particular, consider a sequence of distributions $\mathbf{p} = \{p^t \in \Delta(X) : t \in \mathbb{N}\}$ and let $\mu_{\mathbf{p}}$ be the measure over $X^{\mathbb{N}}$ that is independently generated by \mathbf{p} . A θ^i -incompressible sequence η^i can be used to compute a θ^i -random sequence ξ^i w.r.t. $\mu_{\mathbf{p}}$ that is η^i -computable. If $p^t \rightarrow p$ but slowly (say $\sum_{x \in X} |p^t[x] - p[x]|$ converges to 0 with a rate lower than $\frac{1}{t^{0.5}}$), then ξ^i satisfies (4). See the supplemental material [10], Theorem 4.2, for the full proof. \square

References

- [1] Aumann, R.J. (1981). "Survey of Repeated Games," in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern, Vol 4 of Gesellschaft, Recht, Wirtschaft, Wissenschaftsverlag*, edited by V. Bohm, Bibliographisches Institut, Mannheim, pp. 11-42.
- [2] Ben-Porath, E. (1993), "Repeated Games with Finite Automata," *Journal of Economic Theory*. vol. 59, pp. 17-32.
- [3] Chaitin, G. J. (1975). "A Theory of Program Size Formally Identical to Information Theory," *Journal of the ACM*. vol. 22 pp. 329-340.
- [4] Chatterjee, K. and H. Sabourian. (2009). "Complexity and Game Theory." in *Springer Encyclopedia of Complexity and Systems Science*. Springer, 2009.
- [5] Downey, R., D. Hirschfeldt. (2010). *Algorithmic Randomness and Complexity*. Springer.
- [6] Downey, R., D. Hirschfeldt, A. Nies, and S. Terwijn. (2006). "Calibrating Randomness," *Bulletin Symbolic Logic*. vol. 12, pp. 411-491.
- [7] Goldin, D. and Wegner, P. (2008). "The Interactive Nature of Computing: Refuting the Strong ChurchTuring Thesis," *Minds and Machines*. vol. 18, pp. 17-38.
- [8] Harsanyi, J. C. (1973). "Games with Randomly Disturbed Payoffs: a New Rationale for Mixed-strategy Equilibrium Points," *Int. J. Game Theory*. vol. 2, pp. 123.
- [9] Hu, T-W. (2009). "Expected Utility Theory from the Frequentist Perspective," *Economic Theory*, forthcoming.
- [10] Hu, T-W. (2012). "Complexity and Mixed Strategy Equilibria: Supplemental Material," *Working Paper*.
- [11] Kolmogorov, A. N. (1965). "Three Approaches to the Quantitative Definition of Information," in *Problems of Information Transmission (Problemy Peredachi Informatsii)*. vol. 1, pp. 1-7.

- [12] Li, Ming and Paul Vitányi, (1997) *An introduction to Kolmogorov complexity and its applications*. Springer.
- [13] Luce, R. Duncan and Howard Raiffa. (1957). *Games and Decisions: Introduction and Critical Survey*. Dover.
- [14] Martin-Löf, P. (1966). “The Definition of Random Sequences.” *Information and Control*, vol. 9, pp. 602-619.
- [15] McKelvey, R. D. (1999). “The *hard* sciences.” *Proc. Natl. Acad. Sci. USA*, vol. 96, pp. 10549.
- [16] von Neumann, J. and O. Morgenstern (1944). “*Theory of Games and Economic Behavior*.” Princeton University.
- [17] Nies, A. (2009). *Computability and Randomness*. Cambridge.
- [18] Odifreddi, P. G. (1989). *Classical Recursion Theory*. (Vol. I). North-Holland Publishing Co., Amsterdam.
- [19] Palacios-Huerta, I. (2003). “Professionals Play Minimax,” *Review of Economic Studies*. vol. 70, pp. 395-415.
- [20] Palacios-Huerta, I. and O. Volij (2008). “Experientia Docet: Professionals Play Minimax in Laboratory Experiments,” *Econometrica*. vol. 76, pp. 71-115.
- [21] Rubinstein, A. (1991). “Comments on the Interpretation of Game Theory,” *Econometrica*, vol. 59, pp. 909-924.
- [22] Soare, R. (2009). “Turing Oracle Machines, Online Computing, and Three Displacements in Computability Theory,” *Annals of Pure and Applied Logic*. vol. 160, pp. 368-399.
- [23] Turing, A. (1939). “Systems of Logic Based on Ordinals,” *Proc. London math. soc.*, vol. 45.

- [24] Walker, M. and Wooders, J. (2001). “Minimax Play at Wimbledon,” *American Economic Review*. vol. 91, pp. 1521-1538.