Note

on

Wide Range No-Regret[1]

Yuichi Noguchi

Graduate School of Economics

Hitotsubashi University

2-1 Naka, Kunitachi, Tokyo 186-8601

Japan

E-mail: ynoguchi@econ.hit-u.ac.jp

Phone: +81-42-580-8598

Fax: +81-42-580-8882

[1] I am very grateful to Drew Fudenberg for his encouragement and detailed comments. I also thank Tatsuro Ichiishi for his helpful comment.

Abstract

Fudenberg and Levine's no-regret approach can show that conditional smooth fictitious play is (universally) as good as any countably many alternatives. *Journal of Economic Literature* Classification Numbers: C72, C73, D83.

Key Words: conditional smooth fictitious play, universal conditional consistency, wide range no-regret.

1 Introduction

The purpose of this note is to clarify the implication of Fudenberg and Levine's approach from a "no-regret test" point of view. Fudenberg and Levine (1999) propose the concept of "conditional consistency" as an *optimality* criterion, taking into consideration regularities of opponent behaviors, and they show "universal conditional consistency theorem" (abbreviated to UCC theorem): for any learning rule (under plausible assumptions), "conditional smooth fictitious play" on the rule passes conditional consistency against *all* opposing strategies. Noguchi (1999) develops their approach by constructing optimal learning rules, on which conditional smooth fictitious play passes stronger optimality criteria against *many* opposing strategies. Those works evaluate a player's strategy from an optimality point of view: they *directly* check whether a player's strategy performs well against opponent strategies.

On the other hand, there is an intimately related research line (Lehrer (2001), (2003), and Sandroni et al. (2000)) in which the performance of a player's strategy or prediction is *indirectly* evaluated by carrying out *tests*. They consider that a tester obtains a realized *infinite* sequence of actions after an infinitely repeated game ends, and then checks whether the infinite sequence passes a given set of tests.

Especially Lehrer (2003) introduces *no-regret test* about time-average payoff: a no-regret test defines periods and alternative actions such that in the defined periods, realized payoffs are compared with alternative payoffs that would have been obtained if the alternative actions had been played. The test is passed at a realized infinite sequence of actions if averages of the realized payoffs are eventually at least as good as averages of the alternative payoffs.

He insists that Fudenberg and Levine's approach is quite restrictive from the no-regret test point of view: their strategy (i.e., conditional smooth fictitious play) can only pass a *finite* set of "time-dependent no-regret tests." The main result in Lehrer (2003) is a generalization of their theorem in the no-regret test direction: he shows that there exists a strategy which passes a *wide range* of no-regret tests in the following sense: a tester gives any probability distribution over the set of all no-regret tests. Then, after playing an infinitely repeated game, the tester obtains an infinite sequence of realized actions, and then stochastically chooses a no-regret test according to the given probability distribution, so that the sequence of actions always passes the stochastically chosen test. That is, for the infinite sequence there exists a support of the distribution such that the infinite sequence passes all tests in the support.

This result does *not* assure that the strategy almost surely passes all tests in the support of the distribution against any opposing strategy. An important fact, however, is that in the case that the distribution over no-regret tests has a *countable* support, the result has a strong implication: the strategy passes any countably many no-regret tests simultaneously against all opposing strategies with probability one. Therefore, his result is the most significant in the countable support case.

In this note we shall show that conditional smooth fictitious play may be extended to obtain the significant result by using a method in Noguchi (1999). It may be meaningful because a simple strategy such as conditional smooth fictitious play has the same property as the strategy in Lehrer (2003).

Finally, we remark that a general class of adaptive strategies in Hart and Mas-Colell (2001) may be extended to *conditional* strategies just in the same way as smooth fictitious play, and those conditional adaptive strategies also have (generalized) universal conditional consistency (Noguchi (2002, unpublished)). Thus, all results in this paper also hold for the large class of adaptive strategies.

The paper is organized as follows. In Section 2, we give a model and basic notions. In Section 3, we show that Fudenberg and Levine approach works for any countably many no-regret tests. Section 4 concludes.

2 The Model

2.1 The basic model

We focus on one player who plays an infinitely repeated game against one opponent; the opponent might be a machine, Nature, or multiple players. A player's payoff at a stage game is denoted by $u(a, y)$, where a is a player's action in a finite set A and y is an opponent action in a finite set Y. The set of all mixed actions over S is given by $\Delta(S)$. Let $u(\lambda, \pi)$ denote the player's expected payoff of mixed actions $\lambda \in \Delta(A)$ and $\pi \in \Delta(Y)$: $u(\lambda, \pi) := \sum_{a,y} u(a, y)\lambda[a]\pi[y]$. A finite history (up to time T) in the repeated game is a sequence of actions, denoted by $h_T := (a_1, y_1, \cdots, a_T, y_T)$. An infinite history is denoted by $h_\infty := (a_1, y_1, a_2, y_2, \cdots)$. We write H for the set of all finite histories including the null history $h_0 := \emptyset$, and H^∞ is the set of all infinite histories. We denote a behavior strategy of a player by $\sigma : H \to \Delta(A)$ and that of an opponent by $\rho : H \to \Delta(Y)$ respectively. Let $\mu_{(\sigma,\rho)}$ designate the stochastic process over H^∞ induced by playing σ and ρ.

2.2 No-regret test

Following Lehrer (2003), we define a *replacing scheme* as a function $g : H \times A \to A$ and an *activeness function* as a function $I : H \times A \to \{0, 1\}$. A replacing scheme prescribes alternative actions, and an active function indicates periods when realized payoffs are compared with alternative payoffs; if $I(h_{T-1}, a_T) = 1$ for a realized history (h_{T-1}, a_T), we say that I is *active at time* T. A pair (g, I) of replacing scheme and activeness function is called an *alternative* or a *no-regret test.* Let $\bar{I}(h_{T-1}, a_T)$ designate the number of times that I has been active (up to time T): $\bar{I}(h_{T-1}, a_T) := \sum_{t=1}^{T} I(h_{t-1}, a_t)$. We say that a strategy σ is *universally as good as* $\{(g_\lambda, I_\lambda)\}_{\lambda \in \Lambda}$, or that σ *universally passes* $\{(g_\lambda, I_\lambda)\}_{\lambda \in \Lambda}$, if for all ρ and all (g_λ, I_λ), if $\bar{I}_\lambda(h_{T-1}, a_T) \to \infty$,

$$\liminf_{T \to \infty} \frac{\sum_{t=1}^{T} I_\lambda(h_{t-1}, a_t)[u(a_t, y_t) - u(g_\lambda(h_{t-1}, a_t), y_t)]}{\bar{I}_\lambda(h_{T-1}, a_T)} \geq 0, \ \mu_{(\sigma,\rho)} - a.s.$$

This inequality means that conditional averages of realized payoffs on I_λ are eventually at least as good as conditional averages of alternative payoffs on I_λ (with probability one) if I_λ is active infinitely often. It is not difficult to show that there is *no* player's strategy that is universally as good as an *uncountable* set of alternatives in general.[1]

2.3 Conditional smooth fictitious play

We consider conditional (categorical) smooth fictitious play. A *classification rule* $\mathcal{R}$, which is a player's learning rule, is a partition of $H \times A$.[2] An element in $\mathcal{R}$ is called a *category*, denoted by γ; note that γ may be considered as a subset of $H \times A$. If a realized history $(h_{t-1}, a_t) \in \gamma$, we say that time t is a $\gamma-$*effective period*, or that γ is *effective at time* t. Each category has prior samples, which is represented by a vector d_0^γ; the prior sample size for γ is defined by $n_0^\gamma := \sum_y d_0^\gamma[y]$.

When time T is $\gamma-$effective, a player collects observed opposing actions in past $\gamma-$effective periods, which is represented by a vector d_{T-1}^γ: each component $d_{T-1}^\gamma[y]$ is the number of times that y has occurred in $\gamma-$effective periods up to time $T-1$. Let $n_{T-1}^\gamma := \sum_y d_{T-1}^\gamma[y]$. Then, the player obtains the *augmented empirical distribution* $\tilde{D}_{T-1}^\gamma$ of observed and prior samples: $\tilde{D}_{T-1}^\gamma := \tilde{d}_{T-1}^\gamma / \tilde{n}_{T-1}^\gamma$ where $\tilde{d}_{T-1}^\gamma := d_0^\gamma + d_{T-1}^\gamma$ and $\tilde{n}_{T-1}^\gamma := n_0^\gamma + n_{T-1}^\gamma$.

An effective category may be endogenous in the sense that which category is effective in a current period may depend on which player's action is realized in the current period while which action is realized may depend on which category is effective. Then, following Fudenberg and Levine (1999), we define *conditional categorical smooth fictitious play σ on $\mathcal{R}$* as follows: suppose h_{T-1} is a realized past history up to the last period. Let γ_a be the category that is effective at time T if a player's action a is realized at time T: $(h_{T-1}, a) \in \gamma_a$. Then, for each a, a player obtains a smooth approximate best response $BR^v(\tilde{D}_{T-1}^{\gamma_a})$ to the augmented empirical distribution $\tilde{D}_{T-1}^{\gamma_a}$.[3] Let Z be a matrix whose columns consist of the best responses: $Z := [BR^v(\tilde{D}_{T-1}^{\gamma_a})]_{a\in A}$. It follows from Perron-Frobenius Theorem that there exists a unique mixed action $\lambda \in \Delta(A)$ such that $Z\lambda = \lambda$. Then, let $\sigma(h_{T-1}) := \lambda$.

2.4 Class

In this paper a subset of $H \times A$ will be called a *class*, generically denoted by β. When a realized history $(h_{t-1}, a_t) \in \beta$, we say that time t is a $\beta-$*active period*, or that β is *active at time t*. Especially, if β being active does not depend on player's current actions at all, i.e., $\beta = X \times A$ for some subset X of H, then it will be called an *uncalibrated class*. Given a history h_T, let n_T^β denote the number of times that β has been active (up to time T) and D_T^β denote the empirical distribution of opposing actions observed in $\beta-$active periods (up to time T).

3 Implication of Universal Conditional Consistency

3.1 Passing infinitely many replacing schemes

Lehrer (2003) characterizes Fudenberg and Levine's UCC theorem (1999) (Fudenberg and Levine no-regret theorem in his terminology) by the following statement: for any *finite* partition on the set of all positive integers, say $\{B_i\}_{i=1}^{i=k}$, there exists a strategy σ (precisely, conditional smooth fictitious play) such that σ is as good as $((a^* \mid B_i), \mathbf{1})$ for all $a \in A$ and all $1 \leq i \leq k$. A replacing scheme $(a^* \mid B_i)$ means that a is taken as an alternative action at periods in B_i, i.e., $(a^* \mid B_i)(h_{T-1}, a_T) = a$ if $T \in B_i$, and $(a^* \mid B_i)(h_{T-1}, a_T) = a_T$, otherwise. An activeness function $\mathbf{1}$ means checking all times, that is, $\mathbf{1}(h_{s-1}, a_s) = 1$ for all $(h_{s-1}, a_s) \in H \times A$. But Fudenberg and Levine's universal conditional consistency has a much stronger implication. We shall first state Fudenberg and Levine's UCC theorem precisely. Given a history h_T, let K_T be the number of categories that have been effective (up to time T), n_T^γ denote the number of $\gamma-$effective periods (up to time T), and D_T^γ denote the empirical distribution of opponent actions observed in $\gamma-$effective periods (up to time T). The maximum payoff against π is given by $V(\pi) := \max_a u(a, \pi)$.

Proposition 1 *(Fudenberg and Levine (1999)) Suppose that a classification rule* $(\mathcal{R}, (d_0^\gamma)_\gamma)$ *satisfies the following two assumptions:*

$$(A1)\ \lim_{T\to\infty} \frac{K_T}{T} = 0 \text{ for all } h_\infty \in H^\infty,\ (A2)\ \sup_\gamma n_0^\gamma < \infty.$$

Then, conditional categorical smooth fictitious play σ *on* $\mathcal{R}$ *passes conditional consistency for all opposing strategies: for all* ρ

$$\limsup_{T\to\infty} \sum_{\gamma\in\mathcal{R}} \frac{n_T^\gamma}{T} V(D_T^\gamma) - \frac{1}{T}\sum_{t=1}^{T} u(a_t, y_t) \leq 0,\ \mu_{(\sigma,\rho)} - a.s.$$

Proof. See Fudenberg and Levine (1999).[4] ■

In UCC theorem, averages are taken over payoffs in all current and past periods, so that an activeness function must be $\mathbf{1}$. However, we never have to restrict replacing schemes to a type such as finite time-dependent replacements $\{(a^* \mid B_i)\}_{a\in A, 1\leq i\leq k}$. Precisely, conditional smooth fictitious play universally passes any countably many replacing schemes. We shall show this as our first main result.

Theorem 1 *Let* $\{g_i\}_{i=1}^{i=\infty}$ *be any countable family of replacing schemes. Then, there exists a classification rule* $(\mathcal{R}_0, (d_0^\gamma)_\gamma)$ *such that conditional categorical smooth fictitious play on* $\mathcal{R}_0$ *is universally as good as* $\{(g_i, \mathbf{1})\}_i$*: for all* ρ *and all* g_i*,*

$$\liminf_{T\to\infty} \frac{\sum_{t=1}^{T}[u(a_t, y_t) - u(g_i(h_{t-1}, a_t), y_t)]}{T} \geq 0,\ a.s.$$

Proof. Our proof is based on that of time-average optimality in Noguchi (1999). Fix any countable family $\{g_i\}_i$ of replacing schemes.

<u>Step 1</u>: Define a finite partition $\mathcal{P}_i$ of $H \times A$ for each replacing scheme g_i by the following equivalence relation on $H \times A$:

$$(h_{s-1}, a_s) \sim_{\mathcal{P}_i} (h'_{s'-1}, a'_{s'}) \overset{def.}{\Leftrightarrow} g_i(h_{s-1}, a_s) = g_i(h'_{s'-1}, a'_{s'}).$$

Thus, any class in $\mathcal{P}_i$ is characterized by $b \in A$, so that we denote a class in $\mathcal{P}_i$ by $\beta_i(b)$: $\beta_i(b) = \{(h_{s-1}, a_s) \in H \times A \mid g_i(h_{s-1}, a_s) = b\}$ and $\mathcal{P}_i = \{\beta_i(b) \mid b \in A\}$. Further, we define a finite partition $\mathcal{D}_i$ of $H \times A$ as follows:

$$(h_{s-1}, a_s) \sim_{\mathcal{D}_i} (h'_{s'-1}, a'_{s'}) \overset{def.}{\Leftrightarrow} (h_{s-1}, a_s) \sim_{\mathcal{P}_j} (h'_{s'-1}, a'_{s'}) \text{ for all } 1 \leq j \leq i.$$

$\mathcal{D}_i$ is finer than $\mathcal{P}_j$ for all $1 \leq j \leq i$, and $\mathcal{D}_{i+1}$ is finer than $\mathcal{D}_i$ for all i.[5]

Step 2: Based on $\{\mathcal{D}_i\}_i$, we construct a classification rule $\mathcal{R}_0$, which represents the following player's behavior: the player takes conditional smooth fictitious play on $\mathcal{D}_1$ first, after some periods he switches to a finer partition $\mathcal{D}_2$ and plays conditional fictitious play on $\mathcal{D}_2$ for a while, and again switches to a finer one $\mathcal{D}_3$, and so on. That is, he switches to finer and finer partitions. We introduce an index function $i : H \times A \to \mathbb{N}$ which prescribes the timing of switching those partitions:[6] let $i(h_0, a_1) := 1$, and

$$i(h_T, a_{T+1}) := i(h_{T-1}, a_T) + 1, \text{ if } \frac{\sum_{j=1}^{i(h_{T-1}, a_T)+1} \#\mathcal{D}_j}{T} < \frac{1}{i(h_{T-1}, a_T)},$$

$$i(h_T, a_{T+1}) := i(h_{T-1}, a_T), \text{ otherwise},$$

where $\#$ denotes the cardinality of a set. Then, a classification rule $\mathcal{R}_0$ is defined as follows:

$$(h_{s-1}, a_s) \sim_{\mathcal{R}_0} (h'_{s'-1}, a'_{s'}) \overset{def.}{\Leftrightarrow} i(h_{s-1}, a_s) = i(h'_{s'-1}, a'_{s'}) \text{ and } (h_{s-1}, a_s) \sim_{\mathcal{D}_{i_0}} (h'_{s'-1}, a'_{s'})$$

where $i_0 := i(h_{s-1}, a_s) = i(h'_{s'-1}, a'_{s'})$. Take prior samples $\{d_0^\gamma\}_\gamma$ arbitrarily such that $\sup_\gamma n_0^\gamma < \infty$, where $n_0^\gamma := \sum_y d_0^\gamma[y]$.

Step 3: It is easy to show that the classification rule $(\mathcal{R}_0, (d_0^\gamma)_\gamma)$ satisfies Assumptions $(A1)$ and $(A2)$ in Proposition 1; its proof is given in Appendix A. Thus, by Proposition

1, conditional smooth fictitious play on $\mathcal{R}_0$ has the universal property of conditional consistency on $\mathcal{R}_0$: for all ρ

$$\limsup_{T\to\infty} \sum_{\gamma\in\mathcal{R}_0} \frac{n_T^\gamma}{T} V(D_T^\gamma) - \frac{1}{T}\sum_{t=1}^{T} u(a_t, y_t) \leq 0,\ a.s.$$

Step 4: By the definition of $\mathcal{R}_0$, $\mathcal{R}_0$ is *eventually finer* than any $\mathcal{P}_i$ for all $h_\infty \in H^\infty$: for all $\mathcal{P}_i$ there exists T_i such that for all $\gamma \in \mathcal{R}_0$ there exists $\beta \in \mathcal{P}_i$ such that for all $T \geq T_i$, if $(h_{T-1}, a_T) \in \gamma$, then $(h_{T-1}, a_T) \in \beta$. This, together with the convexity of $V(\cdot)$, implies that for all i and all $h_\infty \in H^\infty$,

$$\limsup_{T\to\infty} \sum_{\beta\in\mathcal{P}_i} \frac{n_T^\beta}{T} V(D_T^\beta) - \sum_{\gamma\in\mathcal{R}_0} \frac{n_T^\gamma}{T} V(D_T^\gamma) \leq 0.$$

Step 5: Finally, we prove that conditional categorical smooth fictitious play on $\mathcal{R}_0$ is universally as good as $\{(g_i, \mathbf{1})\}_i$. Take any $(g_i, \mathbf{1})$. As defined in Step 1, $\beta_i(b) = \{(h_{s-1}, a_s) \in H \times A \mid g_i(h_{s-1}, a_s) = b\}$ and $\mathcal{P}_i = \{\beta_i(b) \mid b \in A\}$. Thus, for all $h_\infty \in H^\infty$ and all T,

$$\begin{aligned}
&\frac{\sum_{t=1}^{T}[u(a_t, y_t) - u(g_i(h_{t-1}, a_t), y_t)]}{T} \\
= \;& -\frac{1}{T}\Big[\sum_{b\in A} \sum_{\substack{(h_{t-1}, a_t)\in\beta_i(b)\\ 1\leq t\leq T}} u(b, y_t) - \sum_{t=1}^{T} u(a_t, y_t)\Big] \\
= \;& -\Big[\sum_{b\in A} \frac{n_T^{\beta_i(b)}}{T} u(b, D_T^{\beta_i(b)}) - \frac{1}{T}\sum_{t=1}^{T} u(a_t, y_t)\Big] \\
\geq \;& -\Big[\sum_{b\in A} \frac{n_T^{\beta_i(b)}}{T} V(D_T^{\beta_i(b)}) - \frac{1}{T}\sum_{t=1}^{T} u(a_t, y_t)\Big] \\
= \;& -\Big[\sum_{\beta\in\mathcal{P}_i} \frac{n_T^{\beta}}{T} V(D_T^{\beta}) - \frac{1}{T}\sum_{t=1}^{T} u(a_t, y_t)\Big].
\end{aligned}$$

Combining this inequality with universal conditional consistency in Step 3 and the inequality in Step 4, we obtain the desired result. Indeed, for all ρ,

$$\begin{aligned}
& \liminf_{T\to\infty} \frac{\sum_{t=1}^{T}[u(a_t, y_t) - u(g_i(h_{t-1}, a_t), y_t)]}{T} \\
\geq \; & -\limsup_{T\to\infty} \sum_{\beta\in\mathcal{P}_i} \frac{n_T^\beta}{T} V(D_T^\beta) - \frac{1}{T}\sum_{t=1}^{T} u(a_t, y_t) \\
\geq \; & -\limsup_{T\to\infty} \sum_{\gamma\in\mathcal{R}_0} \frac{n_T^\gamma}{T} V(D_T^\gamma) - \frac{1}{T}\sum_{t=1}^{T} u(a_t, y_t) \\
\geq \; & 0, \; a.s.
\end{aligned}$$

■

3.2 Passing infinitely many no-regret tests

In the last subsection, an activeness function must be $\mathbf{1}$: payoff comparison must take place *all times.* Noguchi (1999) gives a *classwise generalization* of universal conditional consistency, which may remove the restriction: we may freely select periods when realized payoffs are compared with alternative payoffs. In other words, we may arbitrarily choose activeness functions as well as replacing schemes. To do so, we extend categorical smooth fictitious play to "weighted smooth fictitious play."

For any countable set Ω of classes, we obtain a weight function $w_\Omega : H \to \mathbb{R}_+^A$, which will be defined precisely in Appendix B.[7] Then, *conditional weighted smooth fictitious play σ on $\mathcal{R}$* is defined as follows: suppose h_{T-1} is a realized past history up to the last period. Let γ_a be the category that is effective at time T if a is realized at time T: $(h_{T-1}, a) \in \gamma_a$. Then, for each a, we obtain a *weighted* smooth approximate best response $w_\Omega(h_{T-1})[a] \cdot BR^v(\tilde{D}_{T-1}^{\gamma_a})$. Let Z be a matrix whose columns consist of the weighted best responses: $Z := [w_\Omega(h_{T-1})[a] \cdot BR^v(\tilde{D}_{T-1}^{\gamma_a})]_a$. Moreover, let J be a matrix whose diagonal elements are the weights $w_\Omega(h_{T-1})[a]$'s, and whose off-diagonal elements are all zero. Then, there always exists a mixed action $\lambda^* \in \Delta(A)$ such that $Z\lambda^* = J\lambda^*$.[8] Finally, let $\sigma(h_{T-1}) := \lambda^*$. Note that conditional categorical smooth fictitious play is a special case that all weights are equal at every period.

To obtain a classwise generalization of UCC theorem, we impose several assumptions on a classification rule $(\mathcal{R}, (d_0^\gamma)_\gamma)$ and a countable set Ω of classes. The first assumption requires that a classification rule be eventually finer than any class.

Assumption $(B1)$ *For all* $h_\infty \in H^\infty$ *and all* $\beta \in \Omega$, *there exists* T_0 *such that for all* $\gamma \in \mathcal{R}$, *either*

$$\begin{aligned} \text{for all } T &\geq T_0, \text{ if } (h_{T-1}, a_T) \in \gamma, \text{ then } (h_{T-1}, a_T) \in \beta, \\ \text{or for all } T &\geq T_0, \text{ if } (h_{T-1}, a_T) \in \gamma, \text{ then } (h_{T-1}, a_T) \notin \beta. \end{aligned}$$

The second assumption is that the number of effective categories grows quite slowly in active periods of any class. Given a history h_T, let K_T^β denote the number of categories that have been effective in $\beta-$active periods (up to time T).

Assumption $(B2)$ *For all* $h_\infty \in H^\infty$ *and all* $\beta \in \Omega$, *if* $n_T^\beta \to \infty$ *as* $T \to \infty$, *then*

$$\lim_{T\to\infty} \frac{K_T^\beta}{T} = 0.$$

The third one is that prior samples are uniformly bounded.

Assumption $(B3)$ $\sup_\gamma n_0^\gamma < \infty$.

Under the assumptions above, a classwise generalization of UCC theorem obtains for conditional weighted smooth fictitious play.[9] *Classwise* conditional consistency for a countable set Ω of classes requires that conditional consistency hold in active periods of any class in Ω (if that class is active infinitely many times). To describe it precisely, we say that time t is a $\beta\gamma-$*effective period*, if a realized history $(h_{t-1}, a_t) \in \beta \bigcap \gamma$, i.e., time t is both $\beta-$active and $\gamma-$effective. Given a history h_T, let $n_T^{\beta\gamma}$ denote the number of $\beta\gamma-$effective periods (up to time T), and $D_T^{\beta\gamma}$ denote the empirical distribution of opposing actions observed in $\beta\gamma-$effective periods (up to time T).

We shall state a classwise generalization of UCC theorem as follows.

Proposition 2 *Suppose that a classification rule $(\mathcal{R}, (d_0^\gamma)_\gamma)$ and a countable set Ω of classes satisfy Assumptions (B1), (B2) and (B3). Then, conditional weighted smooth fictitious play σ on $\mathcal{R}$ passes classwise conditional consistency for Ω against all opposing strategies: for all ρ and all $\beta \in \Omega$, if $n_T^\beta \to \infty$ as $T \to \infty$, then*

$$\limsup_{T\to\infty} \sum_{\gamma \in \mathcal{R}} \frac{n_T^{\beta\gamma}}{n_T^\beta} V(D_T^{\beta\gamma}) - \frac{1}{n_T^\beta} \sum_{\substack{(h_{t-1}, a_t) \in \beta \\ 1 \le t \le T}} u(a_t, y_t) \le 0, \ \mu_{(\sigma,\rho)} - a.s.$$

Proof. See Appendix B. ■

Remark 1 *Let Ω be a countable set of uncalibrated classes; recall that β is an uncalibrated class if $\beta = X \times A$ for some subset X of H. Then, Proposition 2 holds for conditional categorical smooth fictitious play because all weights are equal in that case and conditional weighted smooth fictitious play is reduced to conditional categorical smooth fictitious play. (See Appendix B for details.)*

Making use of the classwise generalization of UCC theorem, we show our second main result: conditional smooth fictitious play is universally as good as any countably many alternatives.

Theorem 2 *Let* $\{(g_i, I_i)\}_i$ *be any countable family of alternatives. Then, there exists a classification rule* $(\mathcal{R}_1, (d_0^\gamma)_\gamma)$ *such that conditional weighted smooth fictitious play on* $\mathcal{R}_1$ *is universally as good as* $\{(g_i, I_i)\}_i$*.*

Proof. Our proof is based on that of conditioning-class optimality in Noguchi (1999). Let $\{(g_i, I_i)\}_i$ be any countable family of alternatives.

Step 1: For each (g_i, I_i), we define a finite partition $\mathcal{P}_i$ of $H \times A$ by the following equivalence relation:

$$(h_{s-1}, a_s) \sim_{\mathcal{P}_i} (h'_{s'-1}, a'_{s'}) \overset{def.}{\Leftrightarrow} g_i(h_{s-1}, a_s) = g_i(h'_{s'-1}, a'_{s'}) \text{ and } I_i(h_{s-1}, a_s) = I_i(h'_{s'-1}, a'_{s'}).$$

Thus, any class in $\mathcal{P}_i$ is characterized by $(b, \tau) \in A \times \{0, 1\}$, so that we denote a class in $\mathcal{P}_i$ by $\beta_i(b, \tau)$: $\beta_i(b, \tau) = \{(h_{s-1}, a_s) \in H \times A \mid g_i(h_{s-1}, a_s) = b, I_i(h_{s-1}, a_s) = \tau\}$ and $\mathcal{P}_i = \{\beta_i(b, \tau) \mid b \in A, \tau \in \{0, 1\}\}$. Define partitions $\{\mathcal{D}_i\}_i$ in the same way as in Step 1 in the proof of Theorem 1.

Step 2: We construct a classification rule $\mathcal{R}_1$ based on $\{\mathcal{D}_i\}_i$; in $\mathcal{R}_1$ the player switches *classes*, instead of partitions. To define $\mathcal{R}_1$ precisely, we introduce functions $j : H \times A \to \mathbb{N}$ and $\beta : H \times A \to \bigcup_i \mathcal{D}_i$, where $\bigcup_i \mathcal{D}_i$ is the set of all classes in $\{\mathcal{D}_i\}_i$. Those functions describe a player's selection of classes in $\{\mathcal{D}_i\}_i$. To define $j(\cdot)$ and $\beta(\cdot)$, we introduce two other functions $N : H \times A \to \mathbb{N}$, and $M : H \times A \to \mathbb{N}$. They are all defined recursively as follows:

- $j(h_0, a_1) := 1$ and $\beta(h_0, a_1) := \beta$, where $(h_0, a_1) \in \beta$ and $\beta \in \mathcal{D}_1$. Further, let $N(h_0, a_1) := 0$ and $M(h_0, a_1) := 0$.

• Suppose that $j(h_{s-1}, a_s)$, $\beta(h_{s-1}, a_s)$, $N(h_{s-1}, a_s)$, and $M(h_{s-1}, a_s)$ are defined for $1 \leq s \leq T-1$. Then, let

$$N(h_{T-1}, a_T) := \#\{(h_{s-1}, a_s) \mid (h_{T-1}, a_T) \in \beta(h_{s-1}, a_s), (h_{s-1}, a_s) \leq h_{T-1}\}$$

where $(h_{s-1}, a_s) \leq h_{T-1}$ means that (h_{s-1}, a_s) is an initial subhistory of h_{T-1}.

If $N(h_{T-1}, a_T) = 0$, define $M(h_{T-1}, a_T) := 0$. Otherwise, i.e., $N(h_{T-1}, a_T) \geq 1$, let

$$M(h_{T-1}, a_T) := \max\{j(h_{s-1}, a_s) \mid (h_{T-1}, a_T) \in \beta(h_{s-1}, a_s), (h_{s-1}, a_s) \leq h_{T-1}\}.$$

Define $j(h_{T-1}, a_T)$ and $\beta(h_{T-1}, a_T)$ as follows: if $N(h_{T-1}, a_T) = 0$, let $j(h_{T-1}, a_T) := 1$, and

$$j(h_{T-1}, a_T) := M(h_{T-1}, a_T) + 1, \text{ if } \frac{\sum_{i=1}^{M(h_{T-1}, a_T)+1} \#\mathcal{D}_i}{N(h_{T-1}, a_T)} < \frac{1}{M(h_{T-1}, a_T)},$$

$$j(h_{T-1}, a_T) := M(h_{T-1}, a_T), \text{ otherwise.}$$

Finally, let

$$\beta(h_{T-1}, a_T) := \beta, \text{ where } (h_{T-1}, a_T) \in \beta \text{ and } \beta \in \mathcal{D}_{j(h_{T-1}, a_T)}.$$

Given $j(\cdot)$ and $\beta(\cdot)$, we define a classification rule $\mathcal{R}_1$ as follows:

$$(h_{s-1}, a_s) \sim_{\mathcal{R}_1} (h'_{s'-1}, a'_{s'}) \overset{def.}{\Leftrightarrow} j(h_{s-1}, a_s) = j(h'_{s'-1}, a'_{s'}) \text{ and } \beta(h_{s-1}, a_s) = \beta(h'_{s'-1}, a'_{s'}).$$

Note that each category $\gamma \in \mathcal{R}_1$ has its corresponding index $i(\gamma)$ and class $\beta(\gamma)$: $i(\gamma) := j(h_{s-1}, a_s)$ and $\beta(\gamma) := \beta(h_{s-1}, a_s)$ for all $(h_{s-1}, a_s) \in \gamma$; note that $\gamma \subset \beta(\gamma)$ and $\beta(\gamma) \in \mathcal{D}_{i(\gamma)}$. Further, the correspondence is one to one: $\gamma \neq \gamma' \Rightarrow i(\gamma) \neq i(\gamma')$ or $\beta(\gamma) \neq \beta(\gamma')$. (Take $\{d_0^\gamma\}_\gamma$ arbitrarily such that $\sup_\gamma n_0^\gamma < \infty$.)

Step 3: Define a class $\beta(i) := \{(h_{s-1}, a_s) \in H \times A \mid I_i(h_{s-1}, a_s) = 1\}$; note that $\beta(i) = \bigcup_{b \in A} \beta_i(b, 1)$. Let $\Omega_1 := \{\beta(i) \mid i = 1, 2, \cdots\}$. Then, it is not difficult to show that $\mathcal{R}_1$ satisfies Assumptions $(B1)$, $(B2)$ and $(B3)$ with Ω_1; its proof is given in Appendix A. Thus, by Proposition 2, conditional weighted smooth fictitious play on $\mathcal{R}_1$ universally passes classwise conditional consistency for Ω_1: for all ρ and all $\beta \in \Omega_1$, if $n_T^\beta \to \infty$, then

$$\limsup_{T\to\infty} \sum_{\gamma \in \mathcal{R}_1} \frac{n_T^{\beta\gamma}}{n_T^\beta} V(D_T^{\beta\gamma}) - \frac{1}{n_T^\beta} \sum_{\substack{(h_{t-1}, a_t) \in \beta \\ 1 \le t \le T}} u(a_t, y_t) \le 0, \ a.s.$$

Step 4: $\mathcal{R}_1$ also satisfies Assumption $(B1)$ with the set $\bigcup_i \mathcal{P}_i$ of all classes in $\{\mathcal{P}_i\}_i$; see Appendix A. Recall that $\beta(i) = \bigcup_{b \in A} \beta_i(b, 1)$ for all i. These, together with the convexity of $V(\cdot)$, induce that for all i and all $h_\infty \in H^\infty$, if $n_T^{\beta(i)} \to \infty$,

$$\begin{aligned}
& \limsup_{T\to\infty} \sum_{\gamma \in \mathcal{R}_1} \frac{n_T^{\beta(i)\gamma}}{n_T^{\beta(i)}} V(D_T^{\beta(i)\gamma}) - \frac{1}{n_T^{\beta(i)}} \sum_{\substack{(h_{t-1}, a_t) \in \beta(i) \\ 1 \le t \le T}} u(a_t, y_t) \\
= & \limsup_{T\to\infty} \sum_{\gamma \in \mathcal{R}_{1,T}^{\beta(i)}} \frac{n_T^{\gamma}}{n_T^{\beta(i)}} V(D_T^{\gamma}) - \frac{1}{n_T^{\beta(i)}} \sum_{\substack{(h_{t-1}, a_t) \in \beta(i) \\ 1 \le t \le T}} u(a_t, y_t) \\
\ge & \limsup_{T\to\infty} \sum_{b \in A} \frac{n_T^{\beta_i(b,1)}}{n_T^{\beta(i)}} \{V(D_T^{\beta_i(b,1)}) - \frac{1}{n_T^{\beta_i(b,1)}} \sum_{\substack{(h_{t-1}, a_t) \in \beta_i(b,1) \\ 1 \le t \le T}} u(a_t, y_t)\}
\end{aligned}$$

where $\mathcal{R}_{1,T}^{\beta(i)}$ is the set of all categories that have been effective in $\beta(i)-$active periods (up to time T). The first equality holds because $\mathcal{R}_1$ satisfies Assumption $(B1)$ with Ω_1. The second inequality holds because $\beta(i) = \bigcup_{b \in A} \beta_i(b, 1)$ and $\mathcal{R}_1$ satisfies Assumption $(B1)$ with $\bigcup_i \mathcal{P}_i$.

Step 5: Take any (g_i, I_i). By the definition of $\beta_i(b, \tau)$ in Step 1, $g_i(h_{s-1}, a_s) = b$ and $I_i(h_{s-1}, a_s) = 1 \Leftrightarrow (h_{s-1}, a_s) \in \beta_i(b, 1)$. From this it follows that for all $h_\infty \in H^\infty$ and all

T,

$$\begin{aligned}
&\frac{\sum_{t=1}^{T} I_i(h_{t-1},a_t)[u(a_t,y_t)-u(g_i(h_{t-1},a_t),y_t)]}{\bar{I}_i(h_{T-1},a_T)} \\
= \;& -\sum_{b\in A}\frac{n_T^{\beta_i(b,1)}}{\bar{I}_i(h_{T-1},a_T)}\{\frac{1}{n_T^{\beta_i(b,1)}}\sum_{\substack{(h_{t-1},a_t)\in\beta_i(b,1)\\ 1\le t\le T}}[u(b,y_t)-u(a_t,y_t)]\} \\
= \;& -\sum_{b\in A}\frac{n_T^{\beta_i(b,1)}}{\bar{I}_i(h_{T-1},a_T)}\{u(b,D_T^{\beta_i(b,1)})-\frac{1}{n_T^{\beta_i(b,1)}}\sum_{\substack{(h_{t-1},a_t)\in\beta_i(b,1)\\ 1\le t\le T}}u(a_t,y_t)\} \\
\ge \;& -\sum_{b\in A}\frac{n_T^{\beta_i(b,1)}}{\bar{I}_i(h_{T-1},a_T)}\{V(D_T^{\beta_i(b,1)})-\frac{1}{n_T^{\beta_i(b,1)}}\sum_{\substack{(h_{t-1},a_t)\in\beta_i(b,1)\\ 1\le t\le T}}u(a_t,y_t)\}.
\end{aligned}$$

<u>Step 6</u>: Finally we show that conditional weighted smooth fictitious play on $\mathcal{R}_1$ is universally as good as $\{(g_i,I_i)\}_i$. Fix any (g_i,I_i). Note that $\bar{I}_i(h_{T-1},a_T)=n_T^{\beta(i)}$. Then, from Steps 3, 4 and 5 it follows that for all ρ, if $\bar{I}_i(h_{T-1},a_T)\to\infty$, then

$$\begin{aligned}
&\liminf_{T\to\infty}\frac{\sum_{t=1}^{T} I_i(h_{t-1},a_t)[u(a_t,y_t)-u(g_i(h_{t-1},a_t),y_t)]}{\bar{I}_i(h_{T-1},a_T)} \\
\ge \;& -\limsup_{T\to\infty}\sum_{b\in A}\frac{n_T^{\beta_i(b,1)}}{\bar{I}_i(h_{T-1},a_T)}\{V(D_T^{\beta_i(b,1)})-\frac{1}{n_T^{\beta_i(b,1)}}\sum_{\substack{(h_{t-1},a_t)\in\beta_i(b,1)\\ 1\le t\le T}}u(a_t,y_t)\} \\
\ge \;& -\limsup_{T\to\infty}\sum_{\gamma\in\mathcal{R}_1}\frac{n_T^{\beta(i)\gamma}}{n_T^{\beta(i)}}V(D_T^{\beta(i)\gamma})-\frac{1}{n_T^{\beta(i)}}\sum_{\substack{(h_{t-1},a_t)\in\beta(i)\\ 1\le t\le T}}u(a_t,y_t) \\
\ge \;& 0,\ a.s.
\end{aligned}$$

■

If an activeness function I does not depend on player's current actions, i.e., $I(h_s,a)=I(h_s,a')$ for all $a,a'\in A$ and all $h_s\in H$, we say that I is *uncalibrated*. Then, we obtain the following corollary of Theorem 2.

Corollary 1 *Let $\{(g_i, I_i)\}_i$ be any countable family of replacing schemes and uncalibrated activeness functions. Then, there exists a classification rule $(\mathcal{R}_1, (d_0^\gamma)_\gamma)$ such that conditional categorical smooth fictitious play on $\mathcal{R}_1$ is universally as good as $\{(g_i, I_i)\}_i$.*

Proof. Define $\mathcal{R}_1$ and Ω_1 in the same way as in the proof of Theorem 2. Since I_i is uncalibrated, $\beta(i)$ is an uncalibrated class. Thus, Ω_1 is a countable set of uncalibrated classes. Then, as mentioned in Remark 1, conditional weighted smooth fictitious play on $\mathcal{R}_1$ is degenerated to conditional categorical smooth fictitious play on $\mathcal{R}_1$. The remaining may do just in the same way as the proof of Theorem 2. ■

4 Conclusion

By using Fudenberg and Levine's no-regret approach we have shown that conditional categorical smooth fictitious play universally passes any infinitely many no-regret tests as far as activeness functions are uncalibrated. Furthermore, even in the case that activeness functions may depend on player's current actions, we have extended categorical smooth fictitious play to weighted smooth fictitious play, so that conditional weighted smooth fictitious play is universally as good as any countably many alternatives. Finally, we remark that we can extend conditional smooth fictitious play further to obtain exactly the main result in Lehrer (2003).

Appendix A

• $\mathcal{R}_0$ satisfies Assumptions $(A1)$ and $(A2)$.

$(A2)$ is obvious. To show $(A1)$, we first show the following claim.

Claim A1 *$i : H \times A \to \mathbb{N}$ is a non-decreasing function in time such that for all $h_\infty \in H^\infty$, (1) $i(h_0, a_1) = 1$, (2) $0 \leq i(h_T, a_{T+1}) - i(h_{T-1}, a_T) \leq 1$ for all T, and (3) $\lim_{T\to\infty} i(h_{T-1}, a_T) = \infty$.*

Proof. All are obvious except that $\lim_{T\to\infty} i(h_{T-1}, a_T) = \infty$. Suppose $\lim_{T\to\infty} i(h_{T-1}, a_T) < \infty$ for some h_∞. It means that there exists T_0 such that $i(h_{T_0-1}, a_{T_0}) = i(h_{T-1}, a_T)$ for all $T \geq T_0$. It, in turn, implies that

$$\frac{\sum_{j=1}^{i(h_{T_0-1}, a_{T_0})+1} \#\mathcal{D}_j}{T} \geq \frac{1}{i(h_{T_0-1}, a_{T_0})} \text{ for all } T \geq T_0.$$

But then, $\frac{1}{i(h_{T_0-1}, a_{T_0})}$ and $\sum_{j=1}^{i(h_{T_0-1}, a_{T_0})+1} \#\mathcal{D}_j$ are positive constants. Thus, for large T

$$\frac{\sum_{j=1}^{i(h_{T_0-1}, a_{T_0})+1} \#\mathcal{D}_j}{T} < \frac{1}{i(h_{T_0-1}, a_{T_0})}.$$

This is a contradiction. ■

Lemma A1 shows that $\mathcal{R}_0$ satisfies Assumption $(A1)$.

Lemma A1 *For all $h_\infty \in H^\infty$,*

$$\lim_{T\to\infty} \frac{K_T}{T} = 0.$$

Proof. Note that $K_T \leq \sum_{j=1}^{i(h_{T-1}, a_T)} \#\mathcal{D}_j$. Let $m(T) := \max\{t \mid i(h_{t-1}, a_t) = i(h_{T-1}, a_T) - 1\}$. Then, for all T, switching occurs at time $m(T)+1$: $i(h_{m(T)}, a_{m(T)+1}) =$

$i(h_{m(T)-1}, a_{m(T)}) + 1$. Thus,

$$\frac{K_T}{T} \leq \frac{\sum_{j=1}^{i(h_{T-1},a_T)} \#\mathcal{D}_j}{m(T)} = \frac{\sum_{j=1}^{i(h_{m(T)-1},a_{m(T)})+1} \#\mathcal{D}_j}{m(T)} < \frac{1}{i(h_{m(T)-1}, a_{m(T)})} = \frac{1}{i(h_{T-1}, a_T) - 1}.$$

Therefore, $\lim_{T\to\infty} \frac{K_T}{T} = 0$ because $\lim_{T\to\infty} i(h_{T-1}, a_T) = \infty$ by Claim A1. ∎

- $\mathcal{R}_1$ satisfies Assumptions $(B1)$, $(B2)$ and $(B3)$ with both Ω_1 and $\bigcup_i \mathcal{P}_i$.

$(B3)$ is obvious. Further, it suffices to prove that $\mathcal{R}_1$ satisfies Assumptions $(B1)$ and $(B2)$ with $\bigcup_i \mathcal{D}_i$. (From this it immediately follows that $\mathcal{R}_1$ also satisfies $(B1)$ and $(B2)$ with Ω_1 and $\bigcup_i \mathcal{P}_i$.) We first show the following claim.

Claim A2 (1) *For all* $h_\infty \in H^\infty$, $\lim_{T\to\infty} N(h_{T-1}, a_T) = \infty$. (2) *For all* $h_\infty \in H^\infty$, $\lim_{T\to\infty} j(h_{T-1}, a_T) = \infty$.

Proof. (1) Suppose not: $\liminf_{T\to\infty} N(h_{T-1}, a_T) = n_0 < \infty$ for some h_∞. Then, there exist infinitely many T_m such that $N(h_{T_m-1}, a_{T_m}) = n_0$. By the definition of $j(\cdot)$ and $N(\cdot)$, $j(h_{T-1}, a_T) \leq N(h_{T-1}, a_T) + 1$ for all T. Hence, $j(h_{T_m-1}, a_{T_m}) \leq n_0 + 1$ for all T_m. It means that $\beta(h_{T_m-1}, a_{T_m}) \in \bigcup_{i=1}^{n_0+1} \mathcal{D}_i$ for all T_m, where $\bigcup_{i=1}^{n_0+1} \mathcal{D}_i$ is the set of all classes in $\{\mathcal{D}_i\}_{i=1}^{i=n_0+1}$. Since the number of classes in $\bigcup_{i=1}^{n_0+1} \mathcal{D}_i$ is finite, $j(h_{T_m-1}, a_{T_m}) = j(h_{T_{m'}-1}, a_{T_{m'}})$ and $\beta(h_{T_m-1}, a_{T_m}) = \beta(h_{T_{m'}-1}, a_{T_{m'}})$ for some T_m, $T_{m'}$ with $T_m < T_{m'}$. But then, by the definition of $N(\cdot)$ this means that $N(h_{T_m-1}, a_{T_m}) < N(h_{T_{m'}-1}, a_{T_{m'}})$. It contradicts that $N(h_{T_m-1}, a_{T_m}) = N(h_{T_{m'}-1}, a_{T_{m'}}) = n_0$.

(2) Suppose not: $\liminf_{T\to\infty} j(h_{T-1}, a_T) = i_0 < \infty$ for some h_∞. It implies that $j(h_{T_m-1}, a_{T_m}) = i_0$ for infinitely many T_m. Since $\mathcal{D}_{i_0}$ only has finite classes, there exists $\beta_0 \in \mathcal{D}_{i_0}$ such that $\beta(h_{T_l-1}, a_{T_l}) = \beta_0$ for some infinite subsequence $\{T_l\}_l$ of $\{T_m\}_m$; clearly $j(h_{T_l-1}, a_{T_l}) = i_0$

for all T_l. It means that $M(h_{T_l-1}, a_{T_l}) = i_0$ for all T_l. But then, $N(h_{T_l-1}, a_{T_l}) \to \infty$ from (1). Therefore, for large T_l

$$\frac{\sum_{i=1}^{M(h_{T_l-1},a_{T_l})+1} \#\mathcal{D}_i}{N(h_{T_l-1}, a_{T_l})} < \frac{1}{M(h_{T_l-1}, a_{T_l})}.$$

This, together with the definition of $j(\cdot)$, implies that $j(h_{T_l-1}, a_{T_l}) = M(h_{T_l-1}, a_{T_l}) + 1 = i_0 + 1$. This is a contradiction. ■

Lemmas A2 and A3 show that $\mathcal{R}_1$ satisfies Assumptions $(B1)$ and $(B2)$ with $\bigcup_i \mathcal{D}_i$.

Lemma A2 *For all $h_\infty \in H^\infty$ and all $\beta \in \bigcup_i \mathcal{D}_i$, there exists T_0 such that for all $\gamma \in \mathcal{R}_1$, either*

$$\begin{aligned} \textit{for all } T &\geq T_0, \textit{ if } (h_{T-1}, a_T) \in \gamma, \textit{ then } (h_{T-1}, a_T) \in \beta, \\ \textit{or for all } T &\geq T_0, \textit{ if } (h_{T-1}, a_T) \in \gamma, \textit{ then } (h_{T-1}, a_T) \notin \beta. \end{aligned}$$

Proof. Take any infinite history h_∞ and any class β. Then, $\beta \in \mathcal{D}_{i_0}$ for some i_0. Since $\lim_{T\to\infty} j(h_{T-1}, a_T) = \infty$ from Claim A2 (2), there exists T_0 such that $j(h_{T-1}, a_T) \geq i_0$ for all $T \geq T_0$. Recall $i(\gamma)$ and $\beta(\gamma)$ are the corresponding index and class to γ (see Step 2 in the proof of Theorem 2). Therefore, for any effective category γ from time T_0 on, $i(\gamma) \geq i_0$. Further, $\gamma \subset \beta(\gamma)$ and $\beta(\gamma) \in \mathcal{D}_{i(\gamma)}$. Since $\mathcal{D}_{i(\gamma)}$ is finer than $\mathcal{D}_{i_0}$, either $\beta(\gamma) \subset \beta$, or $\beta(\gamma) \bigcap \beta = \emptyset$. Hence, either $\gamma \subset \beta$ when $\beta(\gamma) \subset \beta$, or $\gamma \bigcap \beta = \emptyset$ when $\beta(\gamma) \bigcap \beta = \emptyset$. ■

Lemma A3 *For all $h_\infty \in H^\infty$ and all $\beta \in \bigcup_i \mathcal{D}_i$, if $n_T^\beta \to \infty$ as $T \to \infty$, then*

$$\lim_{T\to\infty} \frac{K_T^\beta}{n_T^\beta} = 0.$$

Proof. Fix any infinite history h_∞ and any class β with $n_T^\beta \to \infty$. Then, $\beta \in \mathcal{D}_{i_0}$ for some i_0. Let $\{T_s\}_s$ be the calendar of $\beta-$active periods: $(h_{T_s-1}, a_{T_s}) \in \beta$ for all $s = 1, 2, \cdots$. Define $j^\beta(h_{T-1}, a_T) := \max_{T_s \leq T} j(h_{T_s-1}, a_{T_s})$. Clearly, $j^\beta(h_{T-1}, a_T) \to \infty$ as $T \to \infty$ because of Claim A2 (2). Furthermore, let $N^\beta(h_{T-1}, a_T) := N(h_{n(T)-1}, a_{n(T)})$ where $n(T) := \min\{T_s \mid j(h_{T_s-1}, a_{T_s}) = j^\beta(h_{T-1}, a_T)\}$. Obviously, $n(T) \to \infty$ and $N^\beta(h_{T-1}, a_T) \to \infty$ as $T \to \infty$ because of Claim A2. Let T_0 be a calendar time for h_∞ and β in Lemma A2. Then, $n_T^\beta \geq N(h_{n(T)-1}, a_{n(T)}) - T_0$ for all T. Moreover, there exists T_1 such that $j^\beta(h_{T-1}, a_T) \geq i_0 + 1$ for all $T \geq T_1$. Thus, for all $T \geq T_1$, switching to a finer class occurs at time $n(T)$: $j(h_{n(T)-1}, a_{n(T)}) = M(h_{n(T)-1}, a_{n(T)}) + 1$. By the definition of $n(T)$, $j(h_{n(T)-1}, a_{n(T)}) = j^\beta(h_{T-1}, a_T)$ for all T. Thus $j^\beta(h_{T-1}, a_T) = M(h_{n(T)-1}, a_{n(T)}) + 1$ for all $T \geq T_1$. Recall that categories have the following properties: $\gamma \subset \beta(\gamma)$ and $\beta(\gamma) \in \mathcal{D}_{i(\gamma)}$, and $\gamma \neq \gamma' \Rightarrow i(\gamma) \neq i(\gamma')$ or $\beta(\gamma) \neq \beta(\gamma')$. Further, for any effective category γ in $\beta-$active periods up to time T, $i(\gamma) \leq j^\beta(h_{T-1}, a_T)$. These induce that $K_T^\beta \leq \sum_{i=1}^{j^\beta(h_{T-1}, a_T)} \#\mathcal{D}_i = \sum_{i=1}^{M(h_{n(T)-1}, a_{n(T)})+1} \#\mathcal{D}_i$ for all $T \geq T_1$. Thus, for all $T \geq T_1$

$$\begin{aligned}
\frac{K_T^\beta}{n_T^\beta} &\leq \frac{\sum_{i=1}^{M(h_{n(T)-1}, a_{n(T)})+1} \#\mathcal{D}_i}{N(h_{n(T)-1}, a_{n(T)})} \frac{N(h_{n(T)-1}, a_{n(T)})}{N(h_{n(T)-1}, a_{n(T)}) - T_0} \\
&\leq \frac{1}{M(h_{n(T)-1}, a_{n(T)})} \frac{N^\beta(h_{T-1}, a_T)}{N^\beta(h_{T-1}, a_T) - T_0} \\
&= \frac{1}{j^\beta(h_{T-1}, a_T) - 1} \frac{N^\beta(h_{T-1}, a_T)}{N^\beta(h_{T-1}, a_T) - T_0}.
\end{aligned}$$

The first inequality and the third equality are obvious. The second inequality holds because switching occurs at time $n(T)$. Since $j^\beta(h_{T-1}, a_T) \to \infty$ and $N^\beta(h_{T-1}, a_T) \to \infty$, the desired result follows. ■

Appendix B

The proof of Proposition 2 is based on that of universal classwise conditional consistency in Noguchi (1999). Before proving Proposition 2, we shall describe several definitions. Let $V^v(\pi) := \max_{\lambda \in \Delta(A)} u(\lambda, \pi) + v(\lambda)$. Then, we make payoff perturbations as follows: First of all, let a family $\{v_m\}_m$ of payoff perturbations be such that (1) $\|v_m\| \to 0$ as $m \to \infty$, (2) $m\|v_m - v_{m-1}\| \to 0$ as $m \to \infty$, and (3) $\max_{\pi,\pi',\pi'' \in \Delta(Y)} \|(\pi' - \pi'')\partial^2 V^{v_m}(\pi)(\pi' - \pi'')\| \leq C_0 \cdot m^{\frac{1}{2}}$, where C_0 is a positive constant. (For example, we can do so with the logistic function $v(\lambda; \kappa) = -\frac{1}{\kappa}\sum_{a \in A} \lambda(a) \log \lambda(a)$ by changing κ.) Then, let $v_{n_0^\gamma + n - 1}$ be the payoff perturbation in the nth $\gamma-$effective period. In other words, payoff perturbations depend only on *augmented sample sizes* $\tilde{n}_{T-1}^\gamma = n_{T-1}^\gamma + n_0^\gamma$.

Next, we define a random variable $X_T(h_\infty)[\beta]$ as

$$X_T(h_\infty)[\beta] := \tilde{n}_T^\gamma V^v(\tilde{D}_T^\gamma) - \tilde{n}_{T-1}^\gamma V^v(\tilde{D}_{T-1}^\gamma) - u(a_T, y_T) - \|v_{\tilde{n}_{T-1}^\gamma}\| - \frac{C_0}{(\tilde{n}_T^\gamma)^{\frac{1}{2}}}, \text{ if } (h_{T-1}, a_T) \in \beta,$$

$$X_T(h_\infty)[\beta] := 0, \text{ otherwise,}$$

where $v = v_{\tilde{n}_{T-1}^\gamma}$ and γ is the effective category at time T: $(h_{T-1}, a_T) \in \gamma$. Let $\bar{X}_T[\beta] := \frac{1}{n_T^\beta}\sum_{t=1}^T X_t[\beta]$ and $[\bar{X}_T]_+(h_\infty)[\beta] := \max\{0, \bar{X}_T(h_\infty)[\beta]\}$. Take any probability distribution $\mathbf{p} = (p_\beta)_\beta$ on Ω such that $p_\beta > 0$ for all $\beta \in \Omega$. Then, we define a weight function $w_\Omega : H \to \mathbb{R}_+^A$ as follows:

$$w_\Omega(h_T)[a] := \sum_{\beta \ni (h_T, a)} p_\beta \cdot [\bar{X}_T]_+[\beta] \cdot \frac{1}{n_{T+1}^\beta}, \text{ if } (h_T, a) \in \beta \text{ for some } \beta \in \Omega,$$

$$w_\Omega(h_T)[a] := 0, \text{ otherwise.}$$

Let $\langle , \rangle$ denote an inner product on L^2: $\langle X, Y \rangle := \sum_\beta p_\beta \cdot X[\beta] \cdot Y[\beta]$ and $L^2 := \{Y \in \mathbb{R}^\infty \mid \sum_\beta p_\beta \cdot (Y[\beta])^2 < \infty\}$. Define $\| \cdot \| := \sqrt{\langle \cdot, \cdot \rangle}$. Let μ be the stochastic process on

H^∞ induced by conditional weighted smooth fictitious play σ and any opposing strategy ρ, and the product measure of μ and $\mathbf{p}$ is denoted by $\mu \times \mathbf{p}$. Let e_y be the unit vector in which all coordinates are zero except that the $y-$coordinate is one.

Proof of Proposition 2.

Step 1: By Envelope Theorem, $\partial V^v(\tilde{D}_T^\gamma) = [u(BR^v(\tilde{D}_T^\gamma), y)]_y$, where $v = v_{\tilde{n}_T^\gamma}$. When $(h_T, a_{T+1}) \in \beta$, this and Taylor Theorem induce that

$$\begin{aligned} X_{T+1}(h_\infty)[\beta] &\leq \tilde{n}_{T+1}^\gamma \partial V^v(\tilde{D}_T^\gamma)(\frac{1}{\tilde{n}_{T+1}^\gamma}(e_{y_{T+1}} - \tilde{D}_T^\gamma)) + \frac{C_0 \cdot (\tilde{n}_T^\gamma)^{\frac{1}{2}}}{\tilde{n}_{T+1}^\gamma} + V^v(\tilde{D}_T^\gamma) \\ &\quad -u(a_{T+1}, y_{T+1}) - \|v_{\tilde{n}_T^\gamma}\| - \frac{C_0}{(\tilde{n}_{T+1}^\gamma)^{\frac{1}{2}}} \\ &\leq u(BR^v(\tilde{D}_T^\gamma), y_{T+1}) - u(a_{T+1}, y_{T+1}) \end{aligned}$$

where γ is the effective category at time $T+1$: $(h_T, a_{T+1}) \in \gamma$. Let $\delta_T^\beta := n_T^\beta - n_{T-1}^\beta$; $\delta_T^\beta = 1$ if β is active at time T, and $\delta_T^\beta = 0$ otherwise. Let $E_\mu[\cdot \mid h_T]$ be conditional expectation on h_T with respect to μ. Then, it follows from the inequality above that for all $h_T \in H$,

$$\begin{aligned} & E_\mu[\langle [\bar{X}_T]_+, \frac{\delta_{T+1}^\beta}{n_{T+1}^\beta} X_{T+1} \rangle \mid h_T] \\ &= E_\mu[\sum_\beta p_\beta \cdot [\bar{X}_T]_+[\beta] \cdot \frac{\delta_{T+1}^\beta}{n_{T+1}^\beta} X_{T+1}[\beta] \mid h_T] \\ &\leq \sum_a \sigma(h_T)[a](\sum_{\beta \ni (h_T, a)} p_\beta \cdot [\bar{X}_T]_+[\beta] \cdot \frac{1}{n_{T+1}^\beta}\{u(BR^v(\tilde{D}_T^{\gamma_a}), \rho(h_T)) - u(a, \rho(h_T))\}) \\ &= u(Z\sigma(h_T), \rho(h_T)) - u(J\sigma(h_T), \rho(h_T)) \\ &= 0. \end{aligned}$$

Step 2: Define $[\bar{X}_T]_- := \bar{X}_T - [\bar{X}_T]_+$ and $proj_{L^2_-}(\bar{X}_T)(h_\infty) := \arg\min_{Y \in L^2_-} \|\bar{X}_T(h_\infty) - Y\|$, where $L^2_- := \{Y \in L^2 \mid Y[\beta] \leq 0 \text{ for all } \beta \in \Omega\}$. Then, $[\bar{X}_T]_-(h_\infty) = proj_{L^2_-}(\bar{X}_T)(h_\infty)$

and $\langle [\bar{X}_T]_+(h_\infty), [\bar{X}_T]_-(h_\infty)\rangle = 0$. Thus, letting $E_{\mu\times\mathbf{p}}[\cdot]$ be expectation with respect to $\mu \times \mathbf{p}$, the second inequality in Step 1 implies that

$$\begin{aligned}
& \sum_{t=1}^{\infty} E_{\mu\times\mathbf{p}}[(\bar{X}_t - proj_{L^2_-}(\bar{X}_t)) \cdot (\frac{\delta^\beta_{t+1}}{n^\beta_{t+1}}(X_{t+1} - proj_{L^2_-}(\bar{X}_t)))] \\
= & \sum_{t=1}^{\infty} E_\mu[\langle \bar{X}_t - proj_{L^2_-}(\bar{X}_t), \frac{\delta^\beta_{t+1}}{n^\beta_{t+1}}(X_{t+1} - proj_{L^2_-}(\bar{X}_t))\rangle] \\
= & \sum_{t=1}^{\infty} E_\mu[\langle [\bar{X}_t]_+, \frac{\delta^\beta_{t+1}}{n^\beta_{t+1}} X_{t+1}\rangle] \\
\leq & \ 0.
\end{aligned}$$

Therefore, we can apply Theorem 4 (and Corollary 1) in Lehrer (2002), so that for all $\beta \in \Omega$, if $n^\beta_T \to \infty$, $[\bar{X}_T]_+(\beta) \to 0$, $\mu - a.s.$

Step 3: Given h_∞ and β, let T_0 be a calendar time for h_∞ and β in Assumption $(B1)$. For notational simplicity, without loss of generality we may assume $T_0 = 1$. Then, given h_∞, let $\mathcal{R}^\beta_T$ denote the set of all categories that have been effective in $\beta-$active periods up to time T, and let $\mathcal{R}^\beta_T(n) := \{\gamma \in \mathcal{R}^\beta_T \mid n^{\beta\gamma}_T \geq n\}$. Assumption $(B2)$ implies the following: $(*)$ if $n^\beta_T \to \infty$, then for all $\varepsilon > 0$ and all n, there exists $T_{\varepsilon,n}$ such that for all $T \geq T_{\varepsilon,n}$, $\sum_{\gamma\in\mathcal{R}^\beta_T(n)} \frac{n^{\beta\gamma}_T}{n^\beta_T} \geq 1 - \varepsilon$. Furthermore, by Assumption $(B1)$, for all T and all $\gamma \in \mathcal{R}^\beta_T$, $n^\gamma_T = n^{\beta\gamma}_T$ (recall $T_0 = 1$). Therefore, $(*)$ implies that for all $h_\infty \in H^\infty$ and all $\beta \in \Omega$, if $n^\beta_T \to \infty$, then

$$\lim_{T\to\infty} \sum_{\gamma\in\mathcal{R}^\beta_T} \frac{n^\gamma_T}{n^\beta_T}[\frac{1}{n^\gamma_T}(\frac{1}{(n^\gamma_0+1)^{\frac{1}{2}}} + \cdots + \frac{1}{(\tilde{n}^\gamma_T)^{\frac{1}{2}}})] = 0.$$

Since $\|v_m\| \to 0$ as $m \to \infty$, $(*)$ also induces that for all $h_\infty \in H^\infty$ and all $\beta \in \Omega$, if $n^\beta_T \to \infty$,

$$\lim_{T\to\infty} \sum_{\gamma\in\mathcal{R}^\beta_T} \frac{n^\gamma_T}{n^\beta_T}[\frac{1}{n^\gamma_T}(\|v_{n^\gamma_0}\| + \cdots + \|v_{\tilde{n}^\gamma_{T-1}}\|)] = 0.$$

Note that $\tilde{n}_t^\gamma \| V^{v_{\tilde{n}_t^\gamma}}(\tilde{D}_t^\gamma) - V^{v_{\tilde{n}_{t-1}^\gamma}}(\tilde{D}_t^\gamma)\| \leq \tilde{n}_t^\gamma \| v_{\tilde{n}_t^\gamma} - v_{\tilde{n}_{t-1}^\gamma}\|$, and that $m\|v_m - v_{m-1}\| \to 0$ as $m \to \infty$. These, together with $(*)$, imply that for all $h_\infty \in H^\infty$ and all $\beta \in \Omega$, if $n_T^\beta \to \infty$, then

$$\begin{aligned} & \lim_{T\to\infty} \sum_{\gamma \in \mathcal{R}_T^\beta} \frac{n_T^\gamma}{n_T^\beta}[\frac{1}{n_T^\gamma}\sum_{t=1}^{T} \tilde{n}_t^\gamma (V^{v_{\tilde{n}_t^\gamma}}(\tilde{D}_t^\gamma) - V^{v_{\tilde{n}_{t-1}^\gamma}}(\tilde{D}_t^\gamma))] \\ \leq \; & \lim_{T\to\infty} \sum_{\gamma \in \mathcal{R}_T^\beta} \frac{n_T^\gamma}{n_T^\beta}[\frac{1}{n_T^\gamma}\sum_{m=1}^{n_T^\gamma} (n_0^\gamma + m)\|v_{n_0^\gamma+m} - v_{n_0^\gamma+m-1}\|] \\ = \; & 0. \end{aligned}$$

Furthermore, $(*)$ and $(B3)$ imply that for all $h_\infty \in H^\infty$ and all $\beta \in \Omega$, if $n_T^\beta \to \infty$,

$$\lim_{T\to\infty} \sum_{\gamma \in \mathcal{R}} \frac{n_T^{\beta\gamma}}{n_T^\beta} V(D_T^{\beta\gamma}) - \sum_{\gamma \in \mathcal{R}_T^\beta} \frac{\tilde{n}_T^\gamma}{n_T^\beta} V(\tilde{D}_T^\gamma) = 0,$$

and since payoff perturbations decrease as time proceeds, $(*)$ induces that for all $h_\infty \in H^\infty$ and all $\beta \in \Omega$, if $n_T^\beta \to \infty$, then

$$\lim_{T\to\infty} \sum_{\gamma \in \mathcal{R}_T^\beta} \frac{\tilde{n}_T^\gamma}{n_T^\beta} V(\tilde{D}_T^\gamma) - \sum_{\gamma \in \mathcal{R}_T^\beta} \frac{\tilde{n}_T^\gamma}{n_T^\beta} V^{v_{\tilde{n}_T^\gamma}}(\tilde{D}_T^\gamma) = 0.$$

<u>Step 4</u>: Let $Z_T^\beta := n_T^\beta S_T^\beta - n_{T-1}^\beta S_{T-1}^\beta$, where $S_T^\beta := \sum_{\gamma \in \mathcal{R}_T^\beta} \frac{\tilde{n}_T^\gamma}{n_T^\beta} V^{v_{\tilde{n}_T^\gamma}}(\tilde{D}_T^\gamma) - \frac{1}{n_T^\beta} \sum_{\substack{(h_{t-1},a_t)\in\beta \\ 1\leq t\leq T}} u(a_t, y_t)$. If β is active at time T, then $Z_T^\beta = \tilde{n}_T^\gamma V^{v_{\tilde{n}_T^\gamma}}(\tilde{D}_T^\gamma) - \tilde{n}_{T-1}^\gamma V^{v_{\tilde{n}_{T-1}^\gamma}}(\tilde{D}_{T-1}^\gamma) - u(a_T, y_T)$, where γ is the effective category at time T. By Assumption $(B1)$, without loss of generality we may assume that if β is not active at time T, then $Z_T^\beta = 0$. Finally, from this and Steps

2 and 3 it follows that for all $\beta \in \Omega$, if $n_T^\beta \to \infty$,

$$
\begin{aligned}
& \limsup_{T\to\infty} \sum_{\gamma\in\mathcal{R}} \frac{n_T^{\beta\gamma}}{n_T^\beta} V(D_T^{\beta\gamma}) - \frac{1}{n_T^\beta} \sum_{\substack{(h_{t-1},a_t)\in\beta \\ 1\le t\le T}} u(a_t, y_t) \\
= & \limsup_{T\to\infty} \sum_{\gamma\in\mathcal{R}_T^\beta} \frac{\tilde{n}_T^\gamma}{n_T^\beta} V^{v_{\tilde{n}_T^\gamma}}(\tilde{D}_T^\gamma) - \frac{1}{n_T^\beta} \sum_{\substack{(h_{t-1},a_t)\in\beta \\ 1\le t\le T}} u(a_t, y_t) \\
= & \limsup_{T\to\infty} \frac{1}{n_T^\beta} \sum_{t=1}^{T} Z_t^\beta \\
= & \limsup_{T\to\infty} \{ \frac{1}{n_T^\beta} \sum_{t=1}^{T} Z_t^\beta - \sum_{\gamma\in\mathcal{R}_T^\beta} \frac{n_T^\gamma}{n_T^\beta} [\frac{1}{n_T^\gamma} \sum_{t=1}^{T} \tilde{n}_t^\gamma (V^{v_{\tilde{n}_t^\gamma}}(\tilde{D}_t^\gamma) - V^{v_{\tilde{n}_{t-1}^\gamma}}(\tilde{D}_t^\gamma))] \\
& - \sum_{\gamma\in\mathcal{R}_T^\beta} \frac{n_T^\gamma}{n_T^\beta} [\frac{1}{n_T^\gamma} (\|v_{n_0^\gamma}\| + \cdots + \|v_{\tilde{n}_{T-1}^\gamma}\|)] - C_0 \sum_{\gamma\in\mathcal{R}_T^\beta} \frac{n_T^\gamma}{n_T^\beta} [\frac{1}{n_T^\gamma} (\frac{1}{(n_0^\gamma+1)^{\frac{1}{2}}} + \cdots + \frac{1}{(\tilde{n}_T^\gamma)^{\frac{1}{2}}})]\} \\
= & \limsup_{T\to\infty} \bar{X}_T(\beta) \\
\le & \limsup_{T\to\infty} [\bar{X}_T]_+(\beta) \\
= & 0, \ a.s.
\end{aligned}
$$

■

References

Fudenberg, D., D. K. Levine, 1999. Conditional Universal Consistency. Games and Econ. Behav. 29, 104-130.

Hart, S., A. Mas-Colell, 2001. A General Class of Adaptive Strategies. J. Econ. Theory 98, 26-54.

Lehrer, E., 2001. Any Inspection is Manipulable. Econometrica 69, 1333-1347.

Lehrer, E., 2002. Approachability in Infinite Dimentional Spaces. Int. J. Game Theory 31, 253-268.

Lehrer, E., 2003. A Wide Range No-Regret Theorem. Games and Econ. Behav. 42, 101-115.

Noguchi, Y., 1999. Optimal Properties of Conditional Fictitious Play, in Ph.D. Thesis at Harvard Univ.

Noguchi, Y., 2002. Conditional Adaptive Strategies and Generalized UCC. Mimeo., Hitotsubashi Univ.

Sandroni, A., R. Smorodinsky, R.V. Vohra, 2000. Calibration with Many Checking Rules. Forthcoming in Math. Oper. Res.

Footnotes

1. Lehrer (2003) obtains the main result for an uncountable case because he uses a weaker concept of goodness than universal goodness.

2. Fudenberg and Levine (1999) define a classification rule as a function from $H \times A$ to a countable set of categories. But their definition is equivalent to ours.

3. Let $v : int(\Delta(A)) \to \mathbb{R}$ be a payoff perturbation such that v is smooth and strictly concave, and $\|\partial v(\lambda)\| \to \infty$ as λ approaches to the boundary of $\Delta(A)$. Define $\|v\| := \sup_\lambda \mid v(\lambda) \mid$. Then, a smooth approximate best response to π is defined by $BR^v(\pi) := \arg\max_\lambda u(\lambda, \pi) + v(\lambda)$, where $\|v\|$ is very small.

4. Strictly speaking, Fudenberg and Levine (1999) show universal $\varepsilon-$conditional consistency for all $\varepsilon > 0$. But we can make $\varepsilon = 0$ by changing payoff perturbations as time proceeds. See Appendix B.

5. We say that a partition $\mathcal{P}$ is finer than a partition $\mathcal{Q}$ if for all $\beta \in \mathcal{P}$ there exists $\hat{\beta} \in \mathcal{Q}$ such that $\beta \subset \hat{\beta}$. Further, when $\beta \subset \hat{\beta}$ we say that β is a finer class than $\hat{\beta}$.

6. The set of all natural numbers is denoted by $\mathbb{N}$.

7. A also denotes the cardinality of itself. $\mathbb{R}^A$ is an $A-$dimensional Euclidean space. $\mathbb{R}^A_+ := \{x \in \mathbb{R}^A \mid x[a] \geq 0 \text{ for all } a \in A\}$.

8. Let $f(\lambda) := \frac{J^{-1}Z\lambda}{\sum_{b\in A}(J^{-1}Z\lambda)[b]}$. It is a continuous function from $\Delta(A)$ to $\Delta(A)$. By the fixed point theorem, there exists a fixed point λ^* of f. Then, $Z\lambda^* = \alpha J\lambda^*$, where $\alpha =$

$\sum_{b\in A}(J^{-1}Z\lambda^*)[b]$. Note that $\sum_{b\in A}(Z\lambda^*)[b] = \sum_{b\in A}\sum_{a\in A} w_\Omega(h_{T-1})[a]\cdot BR^v(\tilde{D}^{\gamma_a}_{T-1})[b]\cdot \lambda^*[a] = \sum_{a\in A} w_\Omega(h_{T-1})[a]\cdot\lambda^*[a] = \sum_{b\in A}(J\lambda^*)[b]$. Thus, $\alpha = 1$. Therefore, $Z\lambda^* = J\lambda^*$.

9. The uniform boundedness of prior sample sizes is imposed to make our argument simple. Indeed, we only need the following weaker assumption instead of Assumptions $(B2)$ and $(B3)$: for all $h_\infty \in H^\infty$ and all $\beta\in\Omega$, if $n_T^\beta \to \infty$ as $T\to\infty$, then

$$\lim_{T\to\infty}\frac{\sum_{\gamma\in\mathcal{R}_T^\beta} n_0^\gamma}{n_T^\beta} = 0$$

where $\mathcal{R}_T^\beta$ is the set of all categories that have been effective in $\beta-$active periods up to time T.