



Hellman, Z. and Levy, Y. J. (2020) Dense orbits of the Bayesian updating group action. *Mathematics of Operations Research*, (Accepted for Publication).

This is the author's final accepted version.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/227272/>

Deposited on: 18 December 2020

Enlighten – Research publications by members of the University of Glasgow

<http://eprints.gla.ac.uk>

DENSE ORBITS OF THE BAYESIAN UPDATING GROUP ACTION

ZIV HELLMAN & YEHUDA JOHN LEVY

ABSTRACT. We study dynamic properties of the group action on the simplex that is induced by Bayesian updating. We show that generically the orbits are dense in the simplex, although one must make use of the entire group, hence departing from straightforward Bayesian updating. We demonstrate also the necessity of the genericity of the signalling structure, a relationship to descriptive set theoretical concepts and applications thereof to repeated games of incomplete information, as well as strengthening concerning the group action on itself.

Keywords: Bayesian updating, group actions, descriptive set theory, repeated games.

Classifications: 37B20, 91A26

1. INTRODUCTION

Consider a decision maker implementing Bayesian inference, starting with a discrete prior probability distribution p over a finite set of states, receiving a signal, updating to another distribution q , and repeating this process of inferring posteriors as new signals arrive. Much of contemporary decision theory (as well as other fields) rests squarely on such a scenario, to the extent that many regard it as a major component of the very definition of rational decision making.

We study here aspects of this process from the perspective of group actions and elements from descriptive set theory. To do so, we work with a ‘signalling structure’ model familiar from the literature. Letting K be a finite set of states and M be a finite set of signals, we suppose that for each true state k^* there is a unique distribution of signals $r_{k^*} \in \Delta(M)$. A decision maker begins with a prior $p \in \Delta(K)$ and after observing n i.i.d. generated signals updates to a posterior q by sequential application of Bayes’ Rule.

Date: November 30, 2020.

Ziv Hellman acknowledges research supported by Israel Science Foundation grant 1626/18.

We may identify each signal $m \in M$ with an operator $\phi_m : \Delta(K) \rightarrow \Delta(K)$ that maps a prior p to the posterior inferred from p given signal m . When all r_k have full support, this gives us a set of operators which may be regarded as being the generators of a group G of operators on $\Delta(K)$. If we restrict attention to the induced semi-group (which would mean taking into consideration only ‘forward paths’ of inferences, as it would exclude inverse operators of the form ϕ_m^{-1}) we model the standard process of Bayesian inference. However, this restriction turns out to be too limiting, in a sense we shall make precise.

With the mathematical machineries afforded by the theory of topological dynamics, the theory of diophantine approximations, and some abstract algebra, we deduce here a series of results.

- Preliminarily in Section 3.1, but of great importance for the rest of the paper, the group G generated by the signal operators is described explicitly; it will follow that G is an Abelian group (this is well-known for the semi-group). For a generic signalling structure G is freely generated and acts freely on the simplex (other than the Dirac measures); this implies that the ordering in which signals are received is of no consequence, and if two sequences of signals starting from the same prior lead to the same posterior then those sequences are of the same length and consist of the same signals up to permutation.
- We then in Section 3.2 reach the main result, which states that when there are at least as many signals as states, generically the orbits resulting from the this group operation are dense in the simplex. In particular, it follows that generically the orbits have recurrent points. In fact, strengthening this result in Section 5, we show that when the group acts on itself, the orbits have recurrent points in a strong sense (uniform convergence of the mapping and its derivatives). Examples we present (Section 3.3) will also show the importance of using the entire group of operators – i.e., using the ‘inverse’ operations to Bayesian updating – and not just the semi-group.
- In Section 4, we discuss the relevance of the above results in relation to the concept of *smoothness* of the orbit relation, and through this, to the study of repeated games with (public) incomplete information. Smoothness had been used to study such repeated games in [Hellman and Levy, 2019]. In particular, findings in this paper limit the usefulness of those previous results with respect to repeated games, but also show that generically when the state space is intuitively ‘quite large’ relative to the signal and the action spaces,

stationary equilibria (i.e., equilibria which condition actions solely on the current beliefs) do exist.

2. PRELIMINARIES AND MODEL SETUP

2.1. Mathematical Preliminaries.

2.1.1. *Free Groups, Semigroups, Orbits, and Recurrence.* We provide here a brief, slightly informal, review of the definitions of groups, semi-groups, and orbits; the formal definitions are standard and can be found in any text on abstract algebraic structures. We also define *recurrence*.

Recall that a group (resp. semi-group) consists of a set G with a binary operation, denoted by \cdot , satisfying associativity, identity, and invertibility (resp. just associativity). The free semi-group generated by elements ϕ_1, \dots, ϕ_N consists of all words (which includes the empty word, which is the identity element id) made from these elements. The free group generated by elements ϕ_1, \dots, ϕ_N consists of all words made from the elements $\phi_1, \dots, \phi_N, \phi_1^{-1}, \dots, \phi_N^{-1}$ under identification of $\phi_i \cdot \phi_i^{-1}$ and $\phi_i^{-1} \cdot \phi_i$ with id . A group is *Abelian* if the binary operation is commutative. The free Abelian group (resp. semi-group) results from the free group (resp. semi-group) under identification of $\phi_i^s \cdot \phi_j^t$ with $\phi_j^s \cdot \phi_i^t$ for $s, t \in \{\pm 1\}$ (resp. $\phi_i \cdot \phi_j$ with $\phi_j \cdot \phi_i$).

For a set Ω and collection \mathcal{F} of mappings $\mathcal{F} : \Omega \rightarrow \Omega$, a natural associated semi-group is the smallest semi-group that includes all elements of \mathcal{F} and is closed under composition; this is the semi-group *generated by* \mathcal{F} , where the composition is the binary operation. If the elements of \mathcal{F} are bijections, we can talk about the group *generated by* \mathcal{F} , which is the smallest group containing all elements of \mathcal{F} that is closed under composition and inversion, and contains the identity.

For a set Ω and a group G of bijections of Ω (endowed with the composition operation), the *orbit* of an element $\omega \in \Omega$ is the set $G(x) := \{g(\omega) \mid g \in G\}$. The induced *orbit relation* is the equivalence relation whose classes are the orbits of G . If Ω is endowed with a topological structure, a point x is *recurrent* if it is in the closure of $G(x) \setminus \{x\}$. A subset $A \subseteq \Omega$ is G -invariant if $G(A) = A$.

2.2. Model Setup.

Let K be a finite set of states. We assume that there is one true state $k^* \in K$ and that an agent seeks to ascertain which is that true state. Throughout, for $k \in K$, let δ_k denote the Dirac measure on k , and $\|\cdot\|_\infty$ denote the supremum norm.

The agent begins with a prior $p \in \Delta(K)$ over K . By observation of a signal, the agent then proceeds to update this prior. The space of signals is a finite set M with at least two signals. There is, in addition, a collection of size $|K|$ of probability distributions $r_1, \dots, r_{|K|} \in \Delta(M)$ over M , each with full support. We will call the resulting $r := (r_k[m])_{k,m}$ (which can be thought of as a $|K| \times |M|$ matrix) by the name *signalling structure*.

The interpretation is that if the true state is $k^* \in K$ then a signal from M will be observed in accordance with probability distribution r_{k^*} , that being the k^* -th row in the signalling structure r . That is, $r_k[m]$, the element of the matrix corresponding to the k -th row and m -th column, represents the probability of signal m conditional on state k .

The agent subsequently observing the signal will correspondingly update his or her beliefs in a Bayesian manner. A main purpose of this paper is to study a natural equivalence relation that emerges if one considers two distributions to be in the same equivalence class if one can follow from the other using Bayesian updating by observation of some sequence of signals from M .

Suppose that signal m has been observed by an agent with prior $p \in \Delta(K)$. Then by the standard Bayesian updating procedure, the posterior probability $\phi_m(p)[k]$ that the agent ascribes to state $k \in K$ conditional on signal m and prior p is

$$(2.1) \quad \phi_m(p)[k] := \frac{r_k[m] \cdot p[k]}{\sum_{j \in K} r_j[m] \cdot p[j]}$$

From this, the full posterior distribution $\phi_m(p) \in \Delta(K)$ following receipt of signal m is

$$(2.2) \quad \phi_m(p) = (\phi_m(p)[k])_{k \in K},$$

This may, of course, be repeated sequentially. That is, having received signal m_1 and derived the posterior $\phi_{m_1}(p)$, one may receive the further signal m_2 and update the posterior itself to a new posterior $\phi_{m_2}(\phi_{m_1}(p))$, and so forth. Inductively, define, $B_0, B_1, \dots, B_N, \dots$,

$$(2.3) \quad B_1(p) = \{\phi_m(p) \mid m \in M\}, \quad B_N(p) = \bigcup_{y \in B_{N-1}(p)} B_1(y)$$

and also B_0, B_{-N} by

$$(2.4) \quad B_0(p) = \{p\}, \quad B_{-N}(p) = \{y \mid p \in B_N(y)\}.$$

2.3. The Induced Group and the Induced Orbit Relation.

Consider again $\phi_m(p)$ from Equation (2.2), that is, the move from prior p to a posterior upon receipt of signal m . One may regard the collection ϕ_1, \dots, ϕ_M to be a collection of operators, such that $\phi_j : \Delta(K) \rightarrow \Delta(K)$,

for each $j \in M$. We observe the group G generated by ϕ_1, \dots, ϕ_M under the composition operation. Note that the full group action of G includes both what one might term ‘forward’ and ‘backward’ orbits, as well as their closures on taking further orbits, starting from any prior probability distribution p . By this, we mean that a word in G may include both positive and negative exponents of the operators, hence a word might look, for example, like $\phi_1\phi_2^{-1}\phi_3^{-1}\phi_4$, where ϕ_2^{-1} is the inverse of ϕ_2 , and so forth. The Bayesian inference process by repeated signals in the literature typically involves only forward orbits from a prior p , sequentially moving to a new posterior upon receipt of new signals. If we restrict attention solely to forward orbits one generates not a group but a semi-group. The distinction between the group and semi-group in this context may be substantial, as we will subsequently show.

Remark 2.1. As is well known, the collection of probability distributions $\Delta(K)$ forms a simplex. In this context, what is termed a face of the simplex refers to the collection of probability distributions whose support is concentrated on a particular subset of K . More formally, let $K' \subseteq K$. The face determined by K' is defined as

$$\Delta(K') = \{p \in \Delta(K) \mid p[k] > 0 \implies k \in K'\}.$$

The interior of the face determined by K' is

$$\Delta^0(K') = \{p \in \Delta(K) \mid p[k] > 0 \iff k \in K'\}.$$

When all the elements of $(r_k[m])_{k \in K, m \in M}$ are positive, any two elements $p, q \in \Delta(K)$ that share membership in the same equivalence class of \mathcal{E} also share the same support. From this, we can speak intelligibly about the orbit relations induced on the face, or on the interior of any face of $\Delta(K)$, by \mathcal{E} . Formally, Lemma 3.1 shows that when all the elements of $(r_k[m])_{k \in K, m \in M}$ are positive, it follows that $\phi_i^{\pm 1}(\Delta(K')) = \Delta(K')$ and $\phi_i^{\pm 1}(\Delta^0(K')) = \Delta^0(K')$ for each $i \in M$. \blacklozenge

2.4. Genericity. Recall that a collection $(\alpha_j)_{j=1}^n$ in \mathbb{R} is algebraically independent if there is no non-zero polynomial with rational coefficients $p[x_1, \dots, x_n]$ such that $p[\alpha_1, \dots, \alpha_n] = 0$.

Definition 2.2. *For our purposes in this paper, the term generic signalling structure will mean that $r_k[m] > 0$ for all k, m , and the collection of elements $(\ln(r_k[m]))_{k \in K, m \in M}$ are algebraically independent.*

The collection of real numbers $(\ln(r_k[m]))_{k \in K, m \in M}$ induced by a signalling structure is a manifold of dimension $K \times (M - 1)$ in $\mathbb{R}^{K \times M}$, since $\sum_{m \in M} r_k[m] = 1$ for each k, m . However, we also have the following proposition:

Proposition 2.3. *The set of signalling structures for which the set $(\ln(r_k[m]))_{k \in K, m \in M}$ is not algebraically independent is meagre and of Lebesgue measure 0 in $(\Delta(M))^K$.*

The intuition is that the logarithmic operation is a transcendental function, i.e., does not satisfy any polynomial equation. Showing that this claim of the proposition obtains for the specific relationships which hold among the variables $(\ln(r_k[m]))_{k \in K, m \in M}$ requires some work.¹ Since the arguments below are not used later in the paper, the reader may wish to skip the remainder of this section in a first reading.

We will make use of the concept of a *semi-algebraic function*: that is, a function between Euclidean spaces whose graph can be defined by finitely many polynomial equalities and inequalities (see, e.g., [Bochnak et al., 1998, Ch. 2]).

Recall that a *real analytic function* $f : U \rightarrow \mathbb{R}$ on an open set $U \subseteq \mathbb{R}^K$ is one which is infinitely differentiable in U and satisfies the property that around every point $z \in U$ there is a neighbourhood $V \subseteq U$ throughout which the Taylor expansion at z converges and agrees with f . The following theorem is proven in [Mityagin, 2020]:

Theorem 2.1. *Let $f : U \rightarrow \mathbb{R}$ on a connected open set $U \subseteq \mathbb{R}^K$ be real analytic and not identically 0. Then $\{x \in U \mid f(x) = 0\}$ is of Lebesgue measure 0.*

Remark 2.4. The proof of the statement of Theorem 2.1 in [Mityagin, 2020] shows that $\{x \in U \mid f(x) = 0\}$ is contained in a countable union of sets whose closures have empty interior, and hence is meagre as well.

Remark 2.5. In Proposition 2.3 the natural logarithm \ln can be replaced by any strictly monotonic function $\phi : (0, 1) \rightarrow \mathbb{R}$ with the property that the function $\phi(C - \phi^{-1}(\cdot))$ is not semi-algebraic in any domain for any $C > 0$. Indeed, the last step of the proof of Proposition 2.3 shows that \ln satisfies this condition.

Proposition 2.3 likely generalises further than Remark 2.5 suggests, including allowing the domain to be to more general semi-algebraic sets than the simplex, but we have not pursued this direction.

We now prove Proposition 2.3:

Proof. Since algebraic independence is a condition on countably many polynomials, and countable unions of meagre (resp. Lebesgue measure 0) sets

¹ The authors are grateful to Lior Silberman for sharing a discussion of this point with us.

are also meagre resp. (resp. of Lebesgue measure 0), it suffices to show that for any non-zero polynomial p in $K \times M$ variables, the set of $(r_k[m])$ in $(\Delta(M))^K$ such that $p((\ln(r_k[m]))_{k \in K, m \in M}) = 0$ is meagre and of Lebesgue measure zero. For such p , define

$$F((\alpha_{k,m})_{k \leq |K|, m \leq |M|-1}) := p\left((\alpha_{k,1}, \dots, \alpha_{k,|M|-1}, \ln(1 - \sum_{m=1}^{|M|-1} e^{\alpha_{k,m}}))_{k=1}^K\right)$$

and denote by U the open domain in which the arguments of the logarithms are positive. Observe that the exponential and logarithmic transformations preserve sets of measure 0, and for $\alpha_{k,m} = \ln(r_k[m])$,

$$p((\ln(r_k[m]))_{k \in K, m \in M}) = F((\alpha_{k,m})_{k \leq |K|, m \leq |M|-1}) = 0.$$

Hence, in order to prove the proposition, since F is real analytic in U , it suffices by Theorem 2.1 and Remark 2.1 to show that F is not identically 0.

If F were identically 0, then by holding all but one variable constant, we see that $q(x, \ln(C - e^x)) = 0$ for some non-trivial polynomial q , some $C \in (0, 1)$, in some open interval $I \subseteq \mathbb{R}$; possibly by shrinking I , it follows that the function $f(x) = \ln(C - e^x)$ is semi-algebraic on I . We show that this cannot be.

If it were so, then since differentiable semi-algebraic functions of one variable have semi-algebraic derivatives (e.g., [Bochnak et al., 1998, Prop. 2.9.1]), $\frac{-e^x}{1-e^x}$ is semi-algebraic in some open interval, from which (as semi-algebraic functions are closed under basic algebraic operations, e.g., [Bochnak et al., 1998, Prop. 2.2.5]) e^x is semi-algebraic in some open interval. Hence (e.g., [Neyman and Sorin, 2003, Ch. 6, Sec. 4, Cor. 1]), for some non-trivial polynomial ρ , $\rho(x, e^x) \equiv 0$, which by taking $x \rightarrow \infty$ can be seen to be impossible. ■

3. MAIN RESULTS

3.1. Properties of The Induced Group.

Recall that K is the set of states and M is the set of signals.

Lemma 3.1. *Let r be a positive signalling structure. For any finite integer L , let $m_1, \dots, m_L \in M$ and $s_1, \dots, s_L \in \{\pm 1\}$, denote for each $m \in M$,*

$$u_m = \#\{l \mid m_l = m, s_l = 1\} - \#\{l \mid m_l = m, s_l = -1\}$$

Then, for each $p \in \Delta(K)$, denoting

$$q = (\phi_{m_L}^{s_L} \circ \phi_{m_{L-1}}^{s_{L-1}} \circ \dots \circ \phi_{m_1}^{s_1})(p)$$

we have that p, q have the same support, and for each $k \in K$,

$$(3.1) \quad q[k] = p[k] \frac{\prod_m (r_k[m])^{u_m}}{\sum_{k'} p[k'] \cdot \prod_m (r_{k'}[m])^{u_m}}$$

or, equivalently, for any $k', k'' \in K$ in the common support,

$$(3.2) \quad \frac{q[k'']}{q[k']} = \frac{p[k'']}{p[k']} \cdot \prod_m \left(\frac{r_{k''}[m]}{r_{k'}[m]} \right)^{u_m}$$

Remark 3.2. When $s_1 = \cdots = s_L = 1$, Equation (3.1) has a very intuitive meaning (and an alternative derivation). The right-hand side in that case is the posterior probability associated with a state $k \in K$ being chosen by Nature using distribution p , followed by L signals selected i.i.d. in accordance with distribution r_k . In other words, this is exactly the standard model of Bayesian inference by way of signals data in the i.i.d. setting (and in the even broader context of exchangeability). The conditional probability of receiving a particular sequence with u_m appearances of signal m , given that k' was initially selected by Nature, is $\frac{L!}{\prod_m u_m!} \prod_m (r_{k'}[m])^{u_m}$, and hence using Bayes' Law, the posterior probability that k was selected is

$$p[k] \frac{\frac{L!}{\prod_m u_m!} \prod_m (r_k[m])^{u_m}}{\sum_{k'} p[k'] \cdot \frac{L!}{\prod_m u_m!} \prod_m (r_{k'}[m])^{u_m}}.$$

This is the right-hand side of Equation (3.1).

Proof. Fixing $m \in M$, it follows immediately from the definition of ϕ_m (Equation (see 2.1)) that for any $p \in \Delta(K)$ and any k', k'' for which $p[k'] \neq 0$, it holds that $\phi_m(p)[k'] \neq 0$ and

$$\frac{\phi_m(p)[k'']}{\phi_m(p)[k']} = \frac{r_{k''}[m]}{r_{k'}[m]} \cdot \frac{p[k'']}{p[k']}.$$

It follows that ϕ_m is invertible and for any k', k'' for which $p[k'] \neq 0$, it holds that $\phi_m^{-1}(p)[k'] \neq 0$ and

$$\frac{\phi_m^{-1}(p)[k'']}{\phi_m^{-1}(p)[k']} = \left(\frac{r_{k''}[m]}{r_{k'}[m]} \right)^{-1} \cdot \frac{p[k'']}{p[k']}.$$

Equation (3.2), as well as the claim that p, q have the same support, then follow by induction on L . Equation (3.1) follows by noting that the right-hand side of Equation (3.1) defines the unique $q \in \Delta(K)$ which satisfies Equation (3.2). ■

We deduce immediately the following fairly well-known property of Bayesian updating, applied here to the group which includes the converse operations as well:

Theorem 3.1. *If all the elements of a signalling structure r are positive, then the group generated by ϕ_1, \dots, ϕ_M is an Abelian group.*

Remark 3.3. From the proof of Lemma 3.1 it follows that even if the values of r are not all positive (and hence not all the generators are invertible) we can still establish that the semi-group generated by ϕ_1, \dots, ϕ_M is Abelian; this is well-known.

Theorem 3.2. *For any generic signalling structure r , the Abelian group G generated by ϕ_1, \dots, ϕ_M is freely generated and acts freely on $\Delta(K) \setminus \{\delta_k\}_{k \in K}$.*

Theorem 3.2 implies that if $g_1, g_2 \in G$ and $p \in \Delta(K) \setminus \{\delta_k\}_{k \in K}$ satisfy $g_1(p) = g_2(p)$, then not only does $g_1 = g_2$, but in addition g_1, g_2 are written identically up to permutations as words in ϕ_1, \dots, ϕ_M in the freely generated group. This yields the conclusion that if two sequences of signals starting from the same (non-Dirac) prior lead to the same posterior, then these two sequences must be of the same length and consist of the same signals up to permutation. The only points in the simplex that are fixed under the group action are the corner elements $\{\delta_k\}_{k \in K}$. In particular, it follows for p which is not such a corner element, $B_n(p) \cap B_\ell(p) = \emptyset$ for $n, \ell \in \mathbb{Z}$ with $n \neq \ell$.

Proof. Fix some $u_1, \dots, u_M \in \mathbb{Z}$ and $u'_1, \dots, u'_M \in \mathbb{Z}$, such that for some $p \in \Omega \setminus \{\delta_k\}_{k \in K}$,

$$\phi_1^{u_1} \circ \dots \circ \phi_M^{u_M}(p) = \phi_1^{u'_1} \circ \dots \circ \phi_M^{u'_M}(p).$$

Fix some $k' \neq k''$ with $p[k'] \neq 0$ and $p[k''] \neq 0$; such k', k'' exist since p is not a Dirac measure. From Equation (3.2):

$$\frac{p[k'']}{p[k']} \cdot \prod_m \left(\frac{r_{k''}[m]}{r_{k'}[m]} \right)^{u_m} = \frac{p[k'']}{p[k']} \cdot \prod_m \left(\frac{r_{k''}[m]}{r_{k'}[m]} \right)^{u'_m}$$

which implies

$$\sum_{m \in M} (u_m - u'_m) \cdot (\ln(r_{k'}[m]) - \ln(r_{k''}[m])) = 0.$$

Since the collection $(\ln(r_k[m]))_{k \in K, m \in M}$ is algebraically independent, it follows that $u_m = u'_m$ for all $m \in M$. ■

Recall that $\binom{n}{k}$ denotes the binomial coefficient $\frac{n!}{k!(n-k)!}$. The next corollary follows from Theorem 3.1 and Theorem 3.2, since the number of words, up to permutation, of length N on M symbols is $\binom{N+M-1}{M-1}$:

Corollary 3.4. *For each $p \in \Delta(K)$ and $N \in \mathbb{N}$, $|B_N(p)| \leq \binom{N+M-1}{M-1}$, with equality if p is in $\Delta(K) \setminus \{\delta_k\}_{k \in K}$ and the signaling structure is generic.*

3.2. Density of Orbits.

Example 3.1. Suppose that K , the set of states, and M , the set of signals, are both the two element set $\{1, 2\}$. Let $r_1 = (\alpha, 1 - \alpha)$ and $r_2[1] = (\beta, 1 - \beta)$, where $0 < \alpha \neq \beta < 1$. Then, after beginning with belief $(p, 1 - p)$ and then observing $u_1, u_2 > 0$ times signals 1, 2, respectively, one obtains a belief $(q, 1 - q)$ that satisfies

$$\frac{q}{1 - q} = \frac{p}{1 - p} \left(\frac{r_1[1]}{r_2[1]} \right)^{u_1} \left(\frac{r_1[2]}{r_2[2]} \right)^{u_2} = \frac{p}{1 - p} \left(\frac{\alpha}{\beta} \right)^{u_1} \left(\frac{1 - \alpha}{1 - \beta} \right)^{u_2}$$

i.e.,

$$(3.3) \quad \ln \left(\frac{q}{1 - q} \right) = \ln \left(\frac{p}{1 - p} \right) + u_1 \cdot \ln \left(\frac{\alpha}{\beta} \right) + u_2 \cdot \ln \left(\frac{1 - \alpha}{1 - \beta} \right)$$

As long as $\ln(\frac{\alpha}{\beta}), \ln(\frac{1-\alpha}{1-\beta})$ are linearly independent over \mathbb{Q} (which Proposition 2.3 shows is the case for generic $\alpha, \beta \in (0, 1)$), since $\ln(\frac{\alpha}{\beta}), \ln(\frac{1-\alpha}{1-\beta})$ are of opposite signs, any real number can be approximated arbitrarily well by the right-hand side of Equation (3.3) using $u_1, u_2 \in \mathbb{N}$. This also follows in greater generality from our density results.

Our main theorem is:

Theorem 3.3. *If $|M| \geq |K|$ then for a generic signalling structure r the orbit of every p in the interior $\Delta^0(K)$ of $\Delta(K)$ is dense in $\Delta(K)$.*

As a result, if $|M| \geq |K|$, then for a generic signalling structure r every orbit is dense in the interior of the minimal face in which it is contained. (The minimal face which an orbit is contained is precisely the space of distributions with the same support as the elements in the orbit.)

Remark 3.5. Note that in general it is *not* true under the assumptions of the theorem that the forward orbits, that is, sets of the form $\cup_{n \geq 0} B^n(p)$ for some $p \in \Delta(K)$ (in fact, for any $p \in \Delta^0(K)$) are dense; this is exhibited in an example in Section 3.3. In fact, as we show, in general $\cup_{n \in \mathbb{Z}} B^n(p)$, the forward and backward orbits together, need not be dense even for generic signalling structures. One needs to consider the entire orbit, which closes the forward (respectively, backward) orbit under backward (respectively, forward) operations. \blacklozenge

We postpone the presentation of the proof of Theorem 3.3 slightly to Section 3.4, in favour of some discussion in this section and the next. We begin with some examples which contrast with Theorem 3.3.

Example 3.2. We claim here that if the signalling structure r is such that if there is some $0 < a \in \mathbb{R}_{++}$ such that for each $k, k' \in K$ and $m \in M$, $\frac{r_k[m]}{r_{k'}[m]}$




Figure 1. Non-Dense Orbit, $|M| = |K| = 3$, $t = \frac{1}{2}$
Example 3.2

is of the form a^ℓ for some $\ell \in \mathbb{Z}$ – i.e., if $\ln\left(\frac{r_k[m]}{r_{k'}[m]}\right) \in a\mathbb{Z}$ – then the orbits have no recurrent points. Indeed, by Equation (3.2), it follows that given any $p \in \Delta(K)$ of full support (which we may always assume holds, as we can always restrict K to be the support of p), for any q in the orbit of p , for each $k', k'' \in \mathbb{N}$, $\frac{q[k']}{q[k'']} = a^\ell \cdot \frac{p[k']}{p[k'']}$ for some $(k', k''\text{-dependent}) \ell \in \mathbb{Z}$. \blacklozenge

For a concrete example, if $K = M$ and for some $t \in (0, 1)$,

$$r_m[m'] = \begin{cases} t & \text{if } m = m' \\ \frac{1-t}{|M|-1} & \text{if } m \neq m' \end{cases}$$

then taking $a := \frac{t}{1-t}(|M|-1)$ shows that the orbits have no recurrent points. This is exhibited in Figure 1 for $|M| = |K| = 3$ and $t = \frac{1}{2}$.

In contrast to Theorem 3.3, which assumes $|M| \geq |K|$, we have:

Proposition 3.6. *If $|M| < |K|$ and r is positive, the orbits are not dense; if moreover r is generic, then there are no recurrent points in $\Delta^0(K)$, the interior of the simplex.*

For $|K| > |M| + 1$, the first part of the proposition is just a dimensionality argument; for $|K| = |M| + 1$ the argument is slightly more refined.

Proof. Let $|K| > |M|$. Fix $p \in \Delta^0(K)$. Using Equation (3.2), we see that the set $\cup_{n \in \mathbb{Z}} B_n(p)$ consists of those $q \in \Delta^0(K)$ for which there are

$(u_m)_{m \in M}$ in \mathbb{Z} , such that for $k = 1, \dots, K-1$ (in some enumeration of K),

$$(3.4) \quad \ln\left(\frac{q[k]}{q[k+1]}\right) - \ln\left(\frac{p[k]}{p[k+1]}\right) = \sum_m u_m \cdot \ln\left(\frac{r_k[m]}{r_{k+1}[m]}\right).$$

The image of the mapping $\Delta^0(K) \rightarrow \mathbb{R}^{K-1}$ given by $q \rightarrow (\ln(\frac{q[k]}{q[k+1]}))_{k=1}^{K-1}$ is a homeomorphism, with converse given by $y \rightarrow (C(y) \cdot e^{-\sum_{j < k} y_j})_{k=1}^{K-1}$, where $C(y)$ is defined for the sake of normalisation (and clearly continuous). Letting A be the $(K-1) \times M$ matrix given by

$$(3.5) \quad A = \left(\ln\left(\frac{r_k[m]}{r_{k+1}[m]}\right) \right)_{k \leq K-1, m \leq M}$$

we see that for *any* orbit to be dense in $\Delta^0(K)$, it must be the case that the closure of $\{Au \mid u \in \mathbb{Z}^M\}$ contains an open set in \mathbb{R}^{K-1} . Clearly this cannot be true if $|M| + 1 < |K|$, as the rank of the matrix $A = (\ln(\frac{r_k[m]}{r_{k+1}[m]}))_{k,m}$ is at most $|M| < |K| - 1$. If $|M| + 1 = |K|$ but A is singular, the reasoning is the same. If $|M| + 1 = |K|$ but A is non-singular, then the set $\{Au \mid u \in \mathbb{Z}^M\}$ is closed and consists of isolated points, hence again ruling out that its closure contains an open set.

Next, suppose r is a generic signalling structure. The analysis above shows that to have a recurrent point in $\Delta^0(K)$, 0 must be a recurrent point of $\{Au \mid u \in \mathbb{Z}^M\} \subseteq \mathbb{R}^{K-1}$. However, as $|K| - 1 \geq |M|$, for a generic signalling structure $A = (\ln(r_k[m]) - \ln(r_{k+1}[m]))_{k \leq K-1, m \leq M}$ is of rank M , hence as above, this set is closed and consists of isolated points. ■

3.3. Non-Denseness of Forward Orbits, Etc.

As per Remark 3.5, we show that the forward orbits $\cup_{n \geq 0} B_n(p)$, and even the union of the forward and backward orbits $\cup_{n \in \mathbb{Z}} B_n(p)$, need not be dense, even for a generic signalling structure. If $\cup_{n \geq 0} B_n(p)$ (respectively $\cup_{n \in \mathbb{Z}} B_n(p)$) is dense for some $p \in \Delta^0(K)$, then in particular p must be an accumulation point; fixing some enumeration of K , we would have by Equation (3.2) that for each $\varepsilon > 0$, there are non-negative integers $(u_m)_{m \in M}$ (respectively $(u_m)_{m \in M}$ all weakly of the same sign), not all 0, such that

$$(3.6) \quad \left| \sum_{m \in M} u_m \cdot \ln(r_k[m]/r_{k+1}[m]) \right| < \varepsilon, \quad \forall k = 1, \dots, K-1.$$

This could *not* be the case if there were some non-zero vector $w = (w_1, \dots, w_{K-1})$ such that

$$(3.7) \quad \sum_{k=1}^{K-1} w_k \cdot \ln(r_k[m]/r_{k+1}[m]) > 0, \quad \forall m = 1, \dots, M.$$

Indeed, Equation (3.6) states that 0 is a density point of the positive lattice (with the origin removed) generated by the vectors v_1, \dots, v_M defined by $v_m = (\ln(r_k[m]/r_{k+1}[m]))_{k=1}^{K-1} \in \mathbb{R}^{K-1}$. Equation (3.7) states that the vectors v_1, \dots, v_M can be separated from 0 by a hyperplane, which would show that Equation (3.6) cannot hold for non-zero non-negative u .

Example 3.3. We work with a signalling structure with set of states $K = \{\alpha, \beta, \gamma\}$ and set of signals $M = \{1, 2, 3\}$. By abusive but simplifying notation, denote $\alpha = (\alpha_1, \alpha_2, \alpha_3)$ instead of $r_\alpha = (r_\alpha[1], r_\alpha[2], r_\alpha[3])$, and similarly for β, γ :

$$\alpha = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right), \beta = (0.3, 0.1, 0.6), \gamma = (0.3, 0.6, 0.1).$$

Consider the vector $w = (2, 1)$. We need to show that

$$(3.8) \quad 2 \cdot \ln\left(\frac{\alpha_m}{\beta_m}\right) + \ln\left(\frac{\beta_m}{\gamma_m}\right) > 0, \quad m = 1, 2, 3.$$

Indeed, checking this for $m = 1, 2, 3$,

$$\begin{aligned} 2 \cdot \ln\left(\frac{1/3}{0.3}\right) + \ln\left(\frac{0.3}{0.3}\right) &= 2 \ln\left(\frac{1}{0.9}\right) > 0, \\ 2 \cdot \ln\left(\frac{1/3}{0.1}\right) + \ln\left(\frac{0.1}{0.6}\right) &> 2 \cdot \ln(3) - \ln(6) = \ln(1.5) > 0, \\ 2 \cdot \ln\left(\frac{1/3}{0.6}\right) + \ln\left(\frac{0.6}{0.1}\right) &> 2 \ln(1/2) + \ln(6) = \ln(1.5) > 0. \end{aligned}$$

This suffices to show that the forward orbit, or even the forward and backward orbits together, are not dense. Clearly, the same calculations show that Equation (3.8) would hold if α, β, γ were in some sufficiently small neighbourhood of the values specified above, so that for any signalling structure close to the given structure $(r_k[m])_{k,m}$, there do not exist non-negative integers $(u_m)_{m \in M}$, not all 0, nor integers $(u_m)_{m \in M}$ weakly of the same sign and not all 0, such that Equation (3.6) holds for sufficiently small $\varepsilon > 0$. ♦

Figure 2 illustrates how the forward orbit can fail to intersect large sections of this space. Figure 3, which includes both the forward and the backward orbit, shows that in this way much more of the space can be covered; this latter graph, however, still fails to be dense.




Figure 2. Forward Orbit Only




Figure 3. Forward and Backward Orbit.

3.4. Proof of Theorem 3.3. The following is Kronecker's theorem ([Cassels, 1957, Ch. 3]):

Theorem 3.4. *Let $A = (a_{i,j})$ be an $N \times N$ real matrix, $b = (b_1, \dots, b_N)^T \in \mathbb{R}^N$, and $\varepsilon > 0$. Then the following conditions are equivalent:*

- *There are integers $u_1, \dots, u_N, d \in \mathbb{Z}$, such that*

$$(3.9) \quad \left| \sum_{i=1}^N a_{l,i} \cdot u_i - b_l - d \right| < \varepsilon, \quad \forall l, 1 \leq l \leq N.$$

- *Whenever $r_1, \dots, r_N \in \mathbb{Z}$ satisfy*

$$(3.10) \quad \sum_{i=1}^N a_{l,i} \cdot r_i \in \mathbb{Z}$$

then they also satisfy

$$\sum_{i=1}^N b_i \cdot r_i \in \mathbb{Z}.$$

Remark 3.7. The latter condition in particular clearly holds if the elements $(a_{i,j})$ are algebraically independent, since Equation (3.10) can only hold if $r_1 = \dots = r_N = 0$.

Lemma 3.8. *Suppose that $(r_k[m])_{k,m}$ are all positive, that $(\ln(r_k[m]))_{m \in M, k \in K}$ are algebraically independent, and that $|K| \leq |M|$. Given positive $\gamma_1, \dots, \gamma_k > 0$ and $\varepsilon > 0$, there are integers u_1, \dots, u_m, d such that for all $k \in K$*

$$(3.11) \quad (1 - \varepsilon)\gamma_k \cdot e^d \leq \prod_m (r_k[m])^{u_m} \leq (1 + \varepsilon)\gamma_k \cdot e^d.$$

Proof. Without loss of generality, $|K| = |M|$. Denoting $\delta := \min[-\ln(1 - \varepsilon), \ln(1 + \varepsilon)]$, the inequalities of Equation (3.11) are implied by the existence of integers u_1, \dots, u_M such that

$$\left| \sum_{m=1}^M u_m \ln(r_k[m]) - \ln(\gamma_k) - d \right| < \delta, \quad \forall k \in K.$$

The existence of such integers u_1, \dots, u_M follows from Theorem 3.4 and Remark 3.7. ■

Now we can prove Theorem 3.3:

Proof of Theorem 3.3. We will show that any orbit intersecting $\Delta^0(K)$ (and therefore contained in it, by the positivity of the signalling structure) is dense in $\Delta^0(K)$, hence attaining our objective. Fix $p, q \in \Delta^0(K)$ and an arbitrary $\varepsilon > 0$. We will show that there is a point q' in the orbit of p under G such that $q \frac{1-\varepsilon}{1+\varepsilon} < q' < q \frac{1+\varepsilon}{1-\varepsilon}$, coordinate-wise. This will imply that the orbit of p is dense.

Set $\gamma_k = \frac{q[k]}{p[k]}$. Lemma 3.8 shows there are $u_1, \dots, u_M, d \in \mathbb{Z}$, and $1 - \varepsilon < o_1, \dots, o_m < 1 + \varepsilon$ such that for each $k \in K$,

$$(3.12) \quad \prod_m (r_k[m])^{u_m} = e^d \cdot \gamma_k \cdot o_k.$$

Plugging into with Equation (3.1), we have

$$\begin{aligned} q'[k] &:= (\phi_{m_L}^{s_L} \circ \phi_{m_{L-1}}^{s_{L-1}} \circ \dots \circ \phi_1^{s_1})(p)[k] = p[k] \frac{\prod_m (r_k[m])^{u_m}}{\sum_{k'} p[k'] \cdot \prod_m (r_{k'}[m])^{u_m}} \\ &= \frac{p[k] \cdot o_k \cdot \gamma_k \cdot e^d}{\sum_{k'} p[k'] \cdot o_{k'} \cdot \gamma_{k'} \cdot e^d} = \frac{q[k] o_k}{\sum_{k'} q[k'] \cdot o_{k'}} \end{aligned}$$

and hence $q \frac{1-\varepsilon}{1+\varepsilon} < q' < q \frac{1+\varepsilon}{1-\varepsilon}$. ■

4. SMOOTHNESS & REPEATED GAMES

4.1. Relation to Smoothness. A related concept, used to formulate an array of results in [Hellman and Levy, 2019], is that of *smoothness*. A complete, separable, and metrisable space Ω is called a Polish space. A relation \mathcal{E} on a Polish space Ω is said to be *Borel* if \mathcal{E} is a Borel subset of $\Omega \times \Omega$. An equivalence relation is said to be *countable* if each equivalence class is countable. We will abbreviate *countable Borel equivalence relation* as *CBER*. A *countable Borel equivalence relation* (CBER) \mathcal{E} on a Polish space Ω is *smooth* if it possesses a *Borel cross-section*, also known as a *Borel transversal*: that is, there is a Borel subset $B \subseteq \Omega$ which intersects each class of \mathcal{E} in exactly one point.

Theorem 4.1. *Let G be a countable group of homeomorphisms acting on a Polish space Ω . The equivalence relation induced by G is smooth if and only if there is no recurrent point.*

The implication ‘no recurrent point implies smoothness’ in Theorem 4.1 was proven (and in fact does not require continuity of the group action, that is, it holds for any group of Borel automorphisms²) in Theorem 11.1 of [Hellman and Levy, 2019]. For the converse, Lemma 1.1 of [Sullivan et al., 1986] in particular says that the existence of a dense orbit in a perfect Polish space implies that each G -invariant Borel set is either meagre or co-meagre (a.k.a. *generic ergodicity*), where we recall that a topological space Ω is complete if each point is in the closure of its complement. From there, a standard argument establishes non-smoothness. (A sketch of such a standard argument is as follows: suppose there is a recurrent point and restrict attention to the closure of its orbit, which is a perfect space X satisfying $G(X) = X$. One then shows that a Borel transversal T must be non-meagre, as the space can be written as the countable disjoint union $\cup_{g \in G} g(T)$. One then shows that T can be partitioned into non-meagre subsets T_1, T_2 . As a result $G(T_1), G(T_2)$ are G -invariant, Borel disjoint, non-meagre, and Borel.)

Theorem 4.2. *Let r be a generic signalling structure:*

- (a) *For each open face F of $\Delta(K)$ of dimension $1 < \dim(F) \leq |M|$, the equivalence relation \mathcal{E} restricted to F is non-smooth.*
- (b) *For each open face F of $\Delta(K)$ of dimension $\dim(F) > |M|$, the equivalence relation \mathcal{E} restricted to F is smooth.*

² The implication ‘smoothness implies no recurrent point’ is not true if G is a group of Borel automorphisms, even if G is cyclic. Indeed on \mathbb{R} define a topology preserving the standard Borel structure but in which 0 is identified with $\pm\infty$, i.e., 0 is an accumulation of any unbounded sequence; and observe the \mathbb{Z} -action given by $n(x) = x + n$.

Using Theorem 4.1, we see that Part (a) of Theorem 4.2 follows from Theorem 3.3 (denseness of orbits for lower dimension faces), while Part (b) follows from Proposition 3.6 (non-existence of recurrent points for higher dimension faces). Note that statement (b) of Theorem 4.2 is vacuous if $|K| \leq |M|$. In particular, in Example 3.2, the orbit equivalence relation is smooth.

4.2. Repeated Games. We describe here, slightly heuristically, the games discussed more formally in [Hellman and Levy, 2019] (introduced in [Kohlberg and Zamir, 1974] for the zero-sum case, [Forges, 1982] more generally; see also [Neyman and Sorin, 2003, Ch. 21]). In the model there, there are K many I -player strategic form games, G^1, \dots, G^K , with finite action space A^i for each player $i \in I$. A game G^k for some $k \in K$ is chosen, once and for all, at the start of play according to a prior $p \in \Delta(K)$ that is known to the players; however, which game has thus been chosen is not revealed to the players.

The game is played repeatedly, with each player observing all actions. In addition, following a stage at which a profile $a \in A := \prod_i A^i$ was played with underlying true game $k \in K$, a signal from a finite set M is received according to a distribution $q(\cdot | k, a) \in \Delta(M)$. Given a fixed discount rate λ , the payoff over the course of the entire game is the expected discounted sum of stage payoffs.

The players use Bayesian updating to update their beliefs over the true state in K . The orbit equivalence relation is such that p', p'' are in the same orbit if there is a common belief q can be reached from each p' and p'' following *some sequences* of actions and signals. This equivalence agrees with the orbit relationship induced in our paper if the action is chosen randomly with full support (e.g., uniformly on actions). Hence the appropriate signal space would be $A \times M$ instead of M .

We seek Nash equilibria of the resulting infinitely repeated discounted game, satisfying two properties: (i) Borel measurability, and (ii) stationarity, in the sense that the strategies should depend only on the shared common belief in $\Delta(K)$ at each stage. As discussed in [Hellman and Levy, 2019], it is known that such equilibria exist in the zero-sum case, but it is not known whether they exist in general. We also note that if one drops the measurability requirement then existence follows, as each belief can only lead to countably many others and hence existence results on equilibria in dynamic games with countable state space can be applied along with an application of the axiom of choice.

As mentioned, the existence of stationary equilibria in repeated games, even with positive signalling structure, is an open question. In [Hellman and Levy, 2019], it is shown that:

Theorem 4.3. *If a positive signalling structure q in a repeated game of incomplete information induces a smooth orbit equivalence relation in $\Delta^0(K)$, then the game possesses a measurable stationary equilibrium in $\Delta^0(K)$.*

Statement (a) of Theorem 4.2 indicates that the use of Theorem 4.3 turns out to be somewhat limited. Nonetheless, statement (b) of Theorem 4.2 (which relates to smoothness) does show, when combined with Theorem 4.3, that:

Theorem 4.4. *If $|K| > |M| \cdot |A|$, then for a generic signalling structure, the repeated game possesses a stationary equilibrium in $\Delta^0(K)$.*

By genericity here we mean that the collection $(q(m | k, a))_{K \in K, m \in M, a \in A}$ satisfies the genericity assumptions of this paper, namely, positivity as well as algebraic independence of $(\ln(q(m | k, a)))_{K \in K, m \in M, a \in A}$.

5. STRENGTHENING RECURRENCE OF GROUP ACTION

Theorem 3.3 implies that for generic signalling structures every point of $\Delta \setminus \{\delta_k\}_{k \in K}$ is a recurrent point under the group action induced by the Bayesian updating operators. The next result shows something even stronger: for each generic signalling structure, every element of the group G is a recurrent point when the group acts on itself with respect to uniform convergence of mappings and their derivatives. We first state the theorem in such a way that genericity of the signalling structure is not assumed, then add genericity in Corollary 5.1 below.

Theorem 5.1. *If all $(r_k[m])_{k,m}$ are positive and $|K| \leq |M|$, then for each neighbourhood V of the identity in the space of continuous maps $\Delta(K) \rightarrow \Delta(K)$ with the topology of uniform convergence, $C(\Delta(K), \Delta(K))$, and each neighbourhood U of the identity in the space of linear transformations on the tangent space of the simplex, there are infinitely many words in the Abelian group generated by ϕ_1, \dots, ϕ_m such that, for each such word w , the corresponding action g on $\Delta(K)$ is in V and has derivatives only in U , i.e., $Dg(p) \in U$ for all $p \in \Delta(K)$.*

To formalise Theorem 5.1, let $X \subseteq \mathbb{R}^N$ be compact and convex and let $T(X)$ be the tangent space of X , i.e., $T(X) = \text{span}\{u - w \mid u, w \in X\}$. Consider a C^1 differentiable map $\phi : X \rightarrow X$ (formally, this means the existence of a neighbourhood U of X and an extension $\phi : U \rightarrow U$ to a C^1 map). For each $p \in X$ we have that $D\phi(p) \in \mathcal{L}(T(X), T(X))$, the space of linear maps from $T(X)$ to itself.

What Theorem 5.1 says is that for infinitely many words w the corresponding group element in g is ‘very close’ to being the identity, where

‘very close’ means that not only does the action of $g \in G$ send elements to nearby elements, but also that the linear approximation to the action is ‘close to the identity’. Since we have not required genericity in the statement of the theorem, ‘close to’ could mean ‘equal to’; for recurrence, we would require ‘close to but not equal to’.

Clearly, from Theorem 5.1 and the freeness of the induced group (Theorem 3.2) we have:

Corollary 5.1. *If $|K| \leq |M|$, then for generic signalling structure, for each neighbourhood V of the identity in $C(\Delta(K), \Delta(K))$, and each neighbourhood U of the identity in $\mathcal{L}(T(\Delta(K)), T(\Delta(K)))$, there are infinitely many elements g of the group generated by ϕ_1, \dots, ϕ_m with $g \in V$ and $Dg(p) \in U$ for all $p \in \Delta(K)$.*

We remark that Theorem 3.3 does not follow from Theorem 5.1, as the latter speaks only of recurrence, while the former speaks of density.

Before proving Theorem 5.1, we establish Lemma 5.3, which follows from the following proposition, which in turn follows from a more general theorem of Minkowski; see, e.g., [Cassels, 1957, Appendix B]:

Proposition 5.2. *Let A be an $N \times (N - 1)$ real matrix, and $\varepsilon > 0$. Then there are integers $u_1, \dots, u_N \in \mathbb{Z}$, with at least one of those integers non-zero, such that*

$$\left| \sum_{i=1}^N u_i \cdot a_{i,l} \right| < \varepsilon, \quad 1 \leq l \leq N - 1.$$

Lemma 5.3. *Suppose $K \leq M$, and fix some enumeration of K . For positive $(r_k[m])_{k \in K, m \in M}$ and $\beta > 0$, there are integers $u_1, \dots, u_M \in \mathbb{Z}$, with $\sum_{m=1}^M |u_m|$ arbitrarily large, such that:*

$$(5.1) \quad 1 - \beta < \prod_{j=1}^M \left(\frac{r_l[j]}{r_n[j]} \right)^{u_j} < 1 + \beta, \quad \forall l, n \in K.$$

Proof. Without loss of generality, $K = M$. Applying \ln as an operator to Equation (5.1), the required inequalities are equivalent to

$$(5.2) \quad \ln(1 - \beta) < \sum_{j=1}^M u_j \cdot (\ln(r_l[j]) - \ln(r_n[j])) < \ln(1 + \beta), \quad \forall 1 \leq l, n \leq K.$$

Let $\zeta = \min[-\ln(1 - \beta), \ln(1 + \beta)]$. Let A be a matrix whose $a_{j,l}$ entry is $\ln r_l[j] - \ln r_{l+1}[j]$. By Proposition 5.2, and using $\varepsilon = \frac{\zeta}{K-1}$, there are

integers $u_1, \dots, u_N \in \mathbb{Z}$, not all zero, such that

$$\left| \sum_{j=1}^M u_j \cdot (\ln(r_l[j]) - \ln(r_{l+1}[j])) \right| < \frac{1}{K-1} \zeta, \quad \forall l = 1, \dots, K-1,$$

from which Equation (5.2) follows.

From here, we can make the sum $\sum_{m=1}^M |u_m|$ arbitrarily large: we can replace $\ln(1 \pm \beta)$ in Equation (5.2) with $\frac{1}{T} \ln(1 \pm \beta)$ for an arbitrarily large $T \in \mathbb{N}$, and then note that the desired inequalities obtain with $T \cdot u_1, \dots, T \cdot u_M$ replacing u_1, \dots, u_M . ■

We now prove Theorem 5.1:

Proof of Theorem 5.1. Fix a neighbourhood $V \subseteq C(\Delta(K), \Delta(K))$ of the identity in the space of continuous maps. Observe that for any mapping $f : \Delta(K) \rightarrow \Delta(K)$ of the form

$$(5.3) \quad f_k(p) = \frac{p[k] \alpha_k}{\sum_{k'} p[k'] \cdot \alpha_{k'}} = \frac{p[k]}{\sum_{k'} p[k'] \cdot \frac{\alpha_{k'}}{\alpha_k}},$$

there exists $\beta_1 > 0$ such that if $1 - \beta_1 < \frac{\alpha_l}{\alpha_n} < 1 + \beta_1$ for all l, n , then $f \in V$.

Next, fix a neighbourhood $U \subseteq \mathcal{L}(T(X), T(X))$ of the identity. Observe that for any mapping $f : \Delta(K) \rightarrow \Delta(K)$ of the form of Equation (5.3), for positive $\alpha_1, \dots, \alpha_{k'}$ one has

$$\begin{aligned} \langle (Df_k)(p), e_l - e_n \rangle &= (\delta_{l=k} - \delta_{n=k}) \cdot \frac{\alpha_k}{\sum_{k'} p[k'] \cdot \alpha_{k'}} - p[k] \frac{\alpha_k (\alpha_l - \alpha_n)}{(\sum_{k'} p[k'] \cdot \alpha_{k'})^2} \\ &= (\delta_{l=k} - \delta_{n=k}) \left(\sum_{k'} p[k'] \cdot \frac{\alpha_{k'}}{\alpha_k} \right)^{-1} - p[k] \left(\frac{\alpha_l}{\alpha_k} - \frac{\alpha_n}{\alpha_k} \right) \left(\sum_{k'} p[k'] \cdot \frac{\alpha_{k'}}{\alpha_k} \right)^{-2}. \end{aligned}$$

Fix $\varepsilon > 0$. Clearly there exists $\beta_2 > 0$ such that if $1 - \beta_2 < \frac{\alpha_l}{\alpha_n} < 1 + \beta_2$ for all l, n , then

$$|\langle (Df_k)(p), e_l - e_n \rangle - (\delta_{l=k} - \delta_{n=k})| < \varepsilon$$

i.e.,

$$|\langle (Df_k)(p), e_l - e_n \rangle - \langle e_k, e_l - e_n \rangle| < \varepsilon.$$

It follows that for sufficiently small ε , and correspondingly small β_2 , that $Df(p) \in U$ holds for all p .

Next, recall Equation (3.1); we see that for each word w in the Abelian group generated by ϕ_1, \dots, ϕ_m , $w = \phi_1^{u_1} \cdots \phi_M^{u_M}$ with $u_m \in \mathbb{Z}$ for each $m \in M$, the corresponding mapping $\phi = \phi_1^{u_1} \circ \cdots \circ \phi_M^{u_M}$ is of the form of Equation (5.3) with $\alpha_k = \prod_{j=1}^M (r_k[j])^{u_j}$ for each $k \in K$. Hence to prove

Theorem 5.1, it suffices to show that there are infinitely many M -tuples of integers (u_1, \dots, u_M) for which $1 - \beta < \frac{\prod_{j=1}^M (r_k[j])^{u_j}}{\prod_{j=1}^M (r_l[j])^{u_j}} < 1 + \beta$ for each $k, l \in K$ (with $\beta = \min[\beta_1, \beta_2]$); this indeed follows from Lemma 5.3. ■

REFERENCES

- Bochnak et al., 1998. Bochnak, J., Coste, M., and Roy, M.-F. (1998). *Real Algebraic Geometry*. Springer.
- Cassels, 1957. Cassels, J. (1957). *An Introduction to Diophantine Approximation*. Cambridge tracts in mathematics and mathematical physics. Hafner Publishing Company.
- Forges, 1982. Forges, F. (1982). Infinitely repeated games of incomplete information: Symmetric case with random signals. *International Journal of Game Theory*, 11(3):203–213.
- Hellman and Levy, 2019. Hellman, Z. and Levy, Y. J. (2019). Measurable selection for purely atomic games. 87(2):593–629.
- Kohlberg and Zamir, 1974. Kohlberg, E. and Zamir, S. (1974). Repeated games of incomplete information: The symmetric case. *Ann. Statist.*, 2(5):1040–1041.
- Mityagin, 2020. Mityagin, B. S. (2020). The zero set of a real analytic function. *Mathematical Notes*, 107(3):529–530.
- Neyman and Sorin, 2003. Neyman, A. and Sorin, S. (2003). *Stochastic Games and Applications*. NATO science series: Mathematical and physical sciences. Springer Verlag, Netherlands.
- Sullivan et al., 1986. Sullivan, D., Weiss., B., , and Wright, J. (1986). Generic dynamics and monotone complete c^* -algebras. *Trans. Amer. Math. Soc.*, 295:795–809.

DEPARTMENT OF ECONOMICS, BAR-ILAN UNIVERSITY, RAMAT GAN, ISRAEL.
ZIV.HELLMAN@BIU.AC.IL

ADAM SMITH BUSINESS SCHOOL, UNIVERSITY OF GLASGOW, GLASGOW, UNITED KINGDOM. JOHN.LEVY@GLASGOW.AC.UK