



HAL
open science

Agafonov's Theorem for finite and infinite alphabets and probability distributions different from equidistribution

Thomas Seiller, Jakob Grue Simonsen

► To cite this version:

Thomas Seiller, Jakob Grue Simonsen. Agafonov's Theorem for finite and infinite alphabets and probability distributions different from equidistribution. 2022. hal-02993635v2

HAL Id: hal-02993635

<https://hal.science/hal-02993635v2>

Preprint submitted on 23 May 2022 (v2), last revised 7 Nov 2022 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

AGAFONOV'S THEOREM FOR FINITE AND INFINITE ALPHABETS AND PROBABILITY DISTRIBUTIONS DIFFERENT FROM EQUIDISTRIBUTION

THOMAS SEILLER^a AND JAKOB G. SIMONSEN^b

^a CNRS, LIPN – UMR 7030 Université Sorbonne Paris Nord

Current address: 99 Avenue Jean-Baptiste Clément, 93430 Villetaneuse, FRANCE

e-mail address: seiller@lipn.fr

^b Department of Computer Science, University of Copenhagen

Current address: Universitetsparken 1, Copenhagen Ø, DK-2100, DENMARK

e-mail address: simonsen@di.ku.dk

ABSTRACT. An infinite sequence α over an alphabet Σ is μ -distributed w.r.t. a probability map μ if, for every finite string w , the limiting frequency of w in α exists and equals $\mu(w)$. We prove the following result for any finite or countably infinite alphabet Σ : every finite-state selector over Σ selects a μ -distributed sequence from every μ -distributed sequence *if and only if* μ is induced by a Bernoulli distribution on Σ , that is a probability distribution on the alphabet extended to words by taking the product. The primary – and remarkable – consequence of our main result is a complete characterization of the set of probability maps, on finite and infinite alphabets, for which finite-state selection preserves μ -distributedness. The main positive takeaway is that (the appropriate generalization of) Agafonov's Theorem holds for Bernoulli distributions (rather than just equidistributions) on both finite and countably infinite alphabets. As a further consequence, we obtain a result in the area of symbolic dynamical systems: the shift-invariant measures μ on Σ^ω such that any finite-state selector preserves the property of genericity for μ , are exactly the positive Bernoulli measures.

CONTENTS

1. Introduction	2
1.1. Contribution	3
1.2. Related work	4
1.3. Overview of techniques and the proof of the main theorem	6
2. Preliminaries	7
2.1. Strategies	9
2.2. Finite-State Selectors and selection by DFAs	10
3. Main result	11
4. Non-preservation of μ -distributedness for non-Bernoulli measures	12
5. Finite-state selectors preserve μ -distributedness for Bernoulli measures	14
5.1. Ancillary definitions and results	15
5.2. Preservation of Bernoulli μ_p -distributedness under finite-state selection	17
6. An application in symbolic dynamics: characterizing measures where genericity is preserved by DFAs	25
6.1. Shift spaces and genericity	25
7. Future work	26
References	27
Appendix A. Auxiliary proofs and definitions	30
A.1. Automata and selectors	30
A.2. μ -distribution	30
A.3. Symbolic dynamical systems	35

1. INTRODUCTION

Let $\alpha = x_1x_2\cdots$ be an infinite sequence over a finite alphabet Σ . A string $w \in \Sigma^*$ is said to occur in α with limiting frequency f if, for each $\epsilon > 0$, $\lim_{N \rightarrow \infty} \frac{\#_w(x_1 \cdots x_N)}{N} = f$, where $\#_w(x_1 \cdots x_N)$ is the number of times that w occurs as a contiguous subsequence in $x_1 \cdots x_N$. α is said to be *normal* if every finite string of length n over Σ occurs with limiting frequency $|\Sigma|^{-n}$ in α [12]. By standard results, the fractional part the base- b expansion of almost all real numbers is a normal sequence for $b \geq 2$, so for base 10, almost all real numbers have the digit “0” occurring 1-in-10 times in all sufficiently long finite prefixes of their digit expansion, have “11” occurring 1-in-100 times, “110” occurring 1-in-1000 times, and so on. Concrete examples of normal sequences include Champernowne’s sequence 1234567891011 \cdots [20], the Copeland-Erdős sequence 235711131719 \cdots consisting of concatenating the prime numbers [24], and for any polynomial f with positive integer coefficients the sequence $f(1)f(2)f(3)\cdots$ [25].

A *finite-state selector* is a DFA that selects those symbols x_m from α such that $x_1 \cdots x_{m-1}$ is accepted by the DFA. The sequence of selected symbols may thus be finite or infinite. *Agafonov’s Theorem* states that a sequence α is normal iff any DFA that selects an infinite sequence from α , selects a normal sequence. Colloquially, Agafonov’s Theorem can be stated as: “any constant-space algorithm must preserve normality”.

The purpose of this paper is twofold: (I) we study whether analogues of Agafonov’s Theorem holds if the distribution of finite strings is *different* from equidistribution, i.e.

whether distributions where finite strings s are allowed to occur with frequency distinct from $|\Sigma|^{-|s|}$; and (II) we study extensions of Agafonov's Theorem to infinite alphabets (which in the traditional setup in Agafonov's Theorem is meaningless as there is no equidistributed probability distribution on a countably infinite set).

As an example, consider the (non-normal) sequence $\alpha = 010101 \dots$. Clearly, every finite bit string occurs in α with some well-defined frequency (the simplest way to see this is that for each $n > 0$, there are exactly two distinct substrings of length n in α – one starting with 0 and one starting with 1), and the frequencies thus induce a probability distribution on $\{0, 1\}^n$ for each n . In particular 0 and 1 each occur with limiting frequency 1/2, but any DFA that selects symbols at even positions will select the sequence 111..., and thus the probability distribution on $\{0, 1\}$ is *not* preserved, showing that Agafonov's Theorem in general fails to hold.

In addition to being intrinsically interesting, our study of Agafonov's Theorem is motivated by the fact that constant-space algorithms are usually employed in reactive programming languages used for signal processing (see Section 1.2.2 below), both for transduction and selection, and Agafonov's Theorem is a strong guarantee that such algorithms will always preserve one notion of randomness for infinite strings, namely that a random length- n subsequence is exactly $|\Sigma|^{-n}$ – as the above example shows, selection from sequences where 0 and 1 are known to occur with probability 1/2 is not enough – stronger guarantees such as normality must hold. Conversely, normality is a very strong requirement; in some infinite sequences, certain element may occur with much higher frequency than others, and one tantalizing way of generating new sequences having the same distribution of finite subsequences could be to simply let a DFA select elements from the original sequence, which in general is only possible if (the appropriate analogue) of Agafonov's Theorem holds.

The motivation for studying infinite alphabets is that the study of normality is closely tied to the study of symbolic dynamics and (information-theoretic) coding theory [9, 10, 35, 36], and that both areas have witnessed recent advances using infinite alphabets [13, 29, 11, 62, 37], in particular the techniques of Madritsch and Mance [37] have allowed construction of Champernowne-like sequences for various distributions over infinite alphabets.

1.1. Contribution. The formal statement of the main theorem can be found in Theorem 3.1 below. In plain language, we prove that:

Let Σ be a non-empty finite or countably infinite alphabet, and let $\mu : \Sigma^ \rightarrow [0, 1]$ be a probability map (i.e., for all $n \geq 0$, $1 = \sum_{w \in \Sigma^n} \mu(w)$) such that there exists at least one $\alpha \in \Sigma^\omega$ that is μ -distributed. Then, the following are equivalent:*

- (1) μ is induced by a positive Bernoulli probability distribution p on Σ , i.e. for every $a_1, \dots, a_n \in \Sigma$, $\mu(a_1 \dots a_n) = \prod_{i=1}^n p(a_i)$, and for every $a \in \Sigma$, $p(a) > 0$.
- (2) For every DFA A over Σ and every μ -distributed sequence $\alpha \in \Sigma^\omega$, if A selects an infinite sequence from α , then the selected sequence is μ -distributed.

The above result completely characterizes the probability maps preserved by selection by DFAs, both for finite and infinite alphabets, and Agafonov's Theorem follows immediately as a corollary. We briefly review the roadmap and techniques used for the proof of the main result in Section 1.3.

As the study of distributions associated to limiting frequencies of finite strings in (right-)infinite strings is cryptomorphic to the study of shift-invariant probability measures on the shift space (Σ^ω, s) equipped with the σ -algebra induced by the basis of cylinder sets on Σ , we obtain as a corollary a result in the field of symbolic dynamical systems, namely a complete characterization of the shift-invariant probability measures ν for which any finite-state selector preserves genericity for ν , see Section 6.

1.2. Related work.

1.2.1. *Agafonov's Theorem and its generalizations.* Agafonov's Theorem [45] was one of the end results of multiple efforts grappling with the two notions of (i) *kollektiv* (roughly, $\alpha \in \{0, 1\}^\omega$ is a kollektiv wrt. a set \mathcal{S} of selection strategies if the limiting frequency of 1 is unchanged after applying any strategy in \mathcal{S} to α^1), and (ii) *admissible sequence* and its relation to the notion of normal sequence [22, 54, 55, 48, 47]. Agafonov's Theorem itself had a virtually unknown precursor in a beautiful result by Postnikova [53] that showed, with the terminology of the present paper, that $\alpha \in \{0, 1\}^\omega$ is normal iff the distribution of 1s is preserved by selection strategies depending only on a finite word (see Definition 2.10 for the formal definition of *Postnikova strategies*).

Both Postnikova [53] and Agafonov [46] considered selection functions on sequences in $\{0, 1\}^\omega$ where the limiting distribution of 1 was $0 < p < 1$ (i.e., considered a Bernoulli distribution on $\{0, 1\}$), but considering Bernoulli distributions instead of the special case of equidistributions seems to have disappeared almost completely from all later work. One possible reason for this is that only the short version (without proofs or explanation of techniques) of Agafonov's result [45] appeared in English as [1]; in contrast, the original longer paper in Russian [46] was published in a more obscure journal, and was never translated. We have provided a (very) embellished account of the arguments in [46] on the preprint server *arXiv*² where we expand Agafonov's terse use of existing results of the time in much more detail, including using more basic arguments with modern methods (e.g., using concentration bounds directly instead of appealing to the law of large numbers) and add further embellishments to Agafonov's original arguments. Many of the results in the present paper exist due to insights obtained due to this embellishment, rather than the original proof itself.

For equidistribution, the earliest extension to arbitrary alphabets seems to be by Broglio and Liardet [15], and a number of authors have since re-proved Agafonov's Theorem in the special case of equidistribution using a variety of methods; for example, using predictors defined from finite automata (for $\Sigma = \{0, 1\}$) [44], using compressibility arguments [6, 5, 59], and a combination of automata-theoretic and probabilistic methods similar to Agafonov's original reasoning [16].

Agafonov's Theorem itself has been generalized to treat selectors that are not necessarily (induced by prefix selection by) finite automata [2, 5, 67, 17], and some generalizations consider selectors based on relaxed finiteness criteria of the syntactic monoid of a language selecting prefixes of infinite sequences [31, 69]; conversely, results by Merkle and Reimann

¹The exact definition of kollektiv differs subtly across different authors, compare e.g. [68], [21], and [53]. The original notion of kollektiv introduced by von Mises [68] had no constraints on the set \mathcal{S} , but this turned out to be essentially fruitless [64, 54, 32, 23].

²<https://arxiv.org/abs/2007.03249>.

show that adding just slight computational power to the selection strategies beyond finite automata – e.g. using a Pushdown automaton with unary stack alphabet instead of a DFA³ renders Agafonov's Theorem invalid [40]. Similarly, selection by finite automata has been extended, and analogues for Agafonov's Theorem been proved, in other settings than selection from elements of the set Σ^ω , e.g. for shifts of finite type [16]. All of these results only consider normality rather than more general classes of distributions on finite strings.

Conversely, *construction* of normal sequences (as opposed to selecting normal sequences from other normal ones) has been investigated thoroughly for more than a hundred years [61, 20, 42, 65, 39, 52], including explicit construction of real numbers with normal expansion for any integer base $b \geq 2$ [34, 56, 3], and real numbers with normal expansion in non-integer bases [66, 38]. Among this work, the result of most use to the present paper is the construction by Madritsch and Mance of generic sequences for any shift-invariant probability measure μ [37] – these are essentially sequences that are μ -distributed using the terminology of the present paper (see Definition 2.5).

In very recent work, Carton [16] proves that, for any Markov measure μ on Σ^ω induced by a pair (P, π) of a stochastic $|\Sigma| \times |\Sigma|$ matrix and a stationary distribution π for P , any sequence selected from a μ -distributed sequence by a finite-state selector from a particular subset of μ -compatible selectors, will be μ -distributed. Roughly, a finite-state selector is compatible, if it can only read consecutive symbols of Σ with non-zero transition probability in P and every state has only incoming transitions of at most one symbol from Σ . In contrast, we consider the full set of finite-state selectors. Moreover, Carton's results are restricted to the case of finite alphabets.

1.2.2. *Streams and selection from infinite sequences.* Infinite streams are typically used to model situations where data elements arrive, no upper bound on the length of the stream is known a priori, and the focus is not on resource use as a function of the length of the stream; for example, infinite streams have been studied extensively in event-level differential privacy [26, 33], and in semantics of lazy programming languages such as HASKELL [51].

Selection of (substreams of) elements from infinite streams has been investigated from a practical perspective since the 1960s [63], and is typically performed by specialized stream processing languages, e.g. LUSTRE [18] and ESTEREL [8], typically for use in reactive programming (e.g., for signal processing or circuit design). As they are designed for real-time processing, these languages typically allow only very constrained operations – any program in both LUSTRE and ESTEREL can be compiled to a finite state transducer automaton (and deterministic program selecting a subsequence from its input is hence a finite-state selector as in Agafonov's Theorem).

In typical algorithmic treatments of stream processing, one typically studies unordered, finite sequences of elements from a very large, or infinite, set [41]. The problems considered typically have strong constraints, e.g. that only a single pass over the stream is allowed and that each element can only be observed once, and often involve a *sketch*—a data structure that stores information about the elements seen in the stream and allows to answer predefined queries. A classic example is estimating the frequency moments of the distribution of elements in the stream using sketches with low memory in both alphabet size and stream length

³In fact, one of the strategies considered by Merkle and Riemann, which consists in computing the language $\{ww^R \mid w \in \Sigma^*\}$ where w^R is the reverse of w , can be computed by an arguably less expressive model of computation, namely two-way automata with two heads [28].

[4, 30, 14]. Our work can be seen as a variation of streaming where the alphabet size may be infinite, the stream itself is infinite, and the distribution of element is not limited to the set of elements, but also has requirements on the finite subsequences of elements in the stream; in this setting, our main result is that any constant-space sketch sampling an infinite stream in real-time preserves the distribution of finite subsequences iff the distribution is induced by a Bernoulli distribution on the set of elements.

1.3. Overview of techniques and the proof of the main theorem. The main result has two directions: (I) proving that if μ -distributedness is preserved by selection by any DFA, then μ is necessarily induced by a Bernoulli distribution, and (II) any μ induced by a Bernoulli distribution is preserved across selection by any DFA.

For (I), we prove the more general result that if μ is not induced by a Bernoulli distribution on Σ , selection by a particular *Postnikova strategy* (roughly, a Postnikova strategy selects an element of the sequence if and only if it follows a fixed finite word) will select a non- μ -distributed infinite sequence from a – bespoke – μ -distributed sequence. The Postnikova strategy contains prefixes on the form $u \cdot w \in \Sigma^*$ for a *fixed* w chosen such that $w \cdot a \in \Sigma^{|w|+1}$ is a minimal witness string such that $\mu(w \cdot a) \neq \mu(w) \cdot \mu(a)$. Using basic constructions, we can then prove that the Postnikova strategy can be implemented by a DFA that simulates a sliding fixed-width window.

For (II), most of the modern methods of proving Agafonov’s Theorem (e.g., [6, 5, 59]) are not immediately adaptable because they use methods that are particular to equidistributions on finite alphabets (e.g., lossless finite-state compressors [6] or automatic Kolmogorov complexity [59]) – and we consider both Bernoulli distributions and infinite alphabets. Instead, we work along the general lines of Agafonov’s original proof [45] that more heavily uses probabilistic reasoning.

The key insights in Agafonov’s original proof was (i) that any *strongly connected* finite automaton (containing at least one accepting state) applied to a normal sequence must select (always, not just with probability 1) more than a constant fraction of elements from any sufficiently long finite substring of its input, and (ii) that selecting more than a constant fraction of sufficiently long substrings entails that each element of Σ must be selected with approximately equal probability, by the Law of Large Numbers. In Agafonov’s original approach (for $\Sigma = \{0, 1\}$), an appeal to the Strong Law of Large Numbers was used in conjunction with the product measure on the product topology on $\{0, 1\}^\omega$, thus required reasoning about cylinder sets centered on sets A of finite strings; and to avoid “double-counting” the probabilities, these sets had to be prefix-free. We avoid this difficulty by using concentration bounds to tally the occurrences of elements $a \in \Sigma$ in block decompositions of finite prefixes of α .

The proof that any DFA selects a μ -distributed infinite sequence from a μ -distributed infinite sequence then follows by observing that (i) any run of a DFA on an infinite sequence eventually reaches a strongly connected component C of the DFA that is recurrent (i.e., the run can never exit C), and (ii) that any such component induces an irreducible Markov chain, whence we can apply the Ergodic Theorem for Markov Chains to conclude that accepting states are reached infinitely often and with appropriate frequency.

The extension to infinite alphabets is surprisingly straightforward in most proofs: essentially, instead of using combinatorial estimates for finite sets, we have to ensure that series taken over infinite alphabets converge properly, but almost all instances involve series that (i) have non-negative elements, and (ii) are bounded above, whence the usual reasoning

about absolutely convergent series can be employed. Similarly, the classic results for finite automata that we use need to be re-stated and re-proved in the case of infinite alphabets, but this in general turns out to be doable without too much leg-work (e.g. Lemma A.2). One caveat is that several important ancillary results have standard proofs that use combinatorial arguments on finite sets, and we thus need to provide alternative proofs using different methods.

2. PRELIMINARIES

Definition 2.1. We assume a non-empty, possibly (countably) infinite, alphabet Σ and denote by λ the empty string; the sets of finite and right-infinite sequences of elements of Σ are denoted by Σ^* and Σ^ω , respectively. Elements of Σ^* are ranged over by v, w, \dots , and elements of Σ^ω by α, β, \dots . If $\alpha = a_1 a_2 \dots \in \Sigma^\omega$ and N is a positive integer, we denote by $\alpha|_{\leq N}$ the finite string $a_1 a_2 \dots a_N$.

Given $v \in \Sigma^*$ and $u \in \Sigma^* \cup \Sigma^\omega$, we write $v \cdot w$ for the element of $\Sigma^* \cup \Sigma^\omega$ obtained by concatenation. For words $v \in \Sigma^*$ and $w \in \Sigma^* \cup \Sigma^\omega$, v is said to be a *prefix* of w , written $v \preceq w$, if there exists $u \in \Sigma^* \cup \Sigma^\omega$ such that $w = u \cdot v$. If $v \preceq w$ and $v \neq w$, v is said to be a *proper prefix* of w , written $v \prec w$. For any $w \in \Sigma^*$, the *cylinder set* of w , denoted $[w]$, is the subset of Σ^ω defined by $[w] = \{\alpha \in \Sigma^\omega : \alpha = w \cdot \beta, \beta \in \Sigma^\omega\}$, that is the set of right-infinite sequences that have w as prefix.

Definition 2.2. Let Σ be a non-empty, possibly (countably) infinite, alphabet. A *probability map* (over Σ) is a map $\mu : \Sigma^+ \rightarrow [0, 1]$ such that, for all positive integers n , the series

$$\sum_{a_1 \dots a_n \in \Sigma^n} \mu(a_1 \dots a_n)$$

is convergent with limit 1. Note that convergence implies absolute convergence here.

A probability map μ is said to be:

- *induced by a Bernoulli distribution* $p : \Sigma \rightarrow [0, 1]$ if, for all positive integers n , and all $a_1, \dots, a_n \in \Sigma$, $\mu(a_1 \dots a_n) = \prod_{i=1}^n \mu(a_i) = \prod_{i=1}^n p(a_i)$.
- *invariant* if, for all $w \in \Sigma^*$ the series $\sum_{a \in \Sigma} \mu(w \cdot a)$ and $\sum_{a \in \Sigma} \mu(a \cdot w)$ are convergent with limit $\mu(w)$.
- (when Σ is finite) *equidistributed* if, for any $w \in \Sigma^n$, $\mu(w) = |\Sigma|^{-n}$.

Observe that an equidistributed μ is also Bernoulli. For alphabets $|\Sigma| > 1$, any map $p : \Sigma \rightarrow [0, 1]$ such that the series $\sum_{a \in \Sigma} p(a)$ converges to 1 induces a probability map μ_p by setting $\mu_p(a_1 \dots a_n) \triangleq \prod_{j=1}^n p(a_j)$. For finite alphabets Σ , this map is equidistributed iff $p(a) = |\Sigma|^{-1}$ for every $a \in \Sigma$.

The expression ‘‘induced by a Bernoulli distribution’’ is justified by the fact that Bernoulli probability maps correspond directly to the measure of cylinders in Bernoulli shifts [60]⁴.

Proposition 2.3. *A probability map μ induced by a Bernoulli distribution is invariant.*

⁴In the literature on normal numbers, the word Bernoulli is sometimes used slightly differently, for example Schnorr and Stimm [57] use the term ‘‘Bernoulli sequence’’ for sequences that are equidistributed in our terminology. We also note that μ -distributed sequences (defined on the next page) w.r.t. Bernoulli distributions were first introduced by Postnikov and I. I. Piatetski-Shapiro under the name ‘‘Bernoulli normal sequences’’ [49].

Proof. For any $w \in \Sigma^*$, $\sum_{a \in \Sigma} \mu(aw) = \sum_{a \in \Sigma} \mu(a)\mu(w) = \sum_{a \in \Sigma} \mu(w)\mu(a) = \sum_{a \in \Sigma} \mu(wa)$. And $\sum_{a \in \Sigma} \mu(w)\mu(a) = \mu(w)\sum_{a \in \Sigma} \mu(a) = \mu(w)$. \square

We shall need probability maps to act as “measures” on (possibly infinite) sets of finite strings:

Definition 2.4. Let Σ be a non-empty alphabet, let $W \subseteq \Sigma^*$, and let μ be a probability map over Σ . If $W = \emptyset$, we define $\mu(W) = 0$. If $\sum_{w \in W} \mu_p(w)$ converges, we define $\mu_p(W) = \sum_{w \in W} \mu_p(w)$.

Observe that as $\mu_p(w) \geq 0$ for all $w \in \Sigma^*$, if $\sum_{w \in W} \mu_p(w)$ converges, it is absolutely convergent (hence, we do not need to specify an ordering of W).

We are interested in the probability maps whose values can be realized as the limiting frequencies of finite words in right-infinite sequences over Σ .

Definition 2.5. Let $v = v_1 \cdots v_N$ and $w = w_1 \cdots w_n$ be finite words over Σ . We denote by $\#_w(v)$ the number of occurrences of w in v , that is, the quantity

$$|\{j \leq N + 1 - n : v_j v_{j+1} \cdots v_{j+n-1} = w_1 w_2 \cdots w_n\}|$$

Let μ be a probability map over Σ , and be α is a right-infinite sequence over Σ . If the limit

$$\lim_{N \rightarrow \infty} \frac{\#_w(\alpha|_{\leq N})}{N}$$

exists and is equal to some real number f , we say that w occurs in α with limiting frequency f . If every $w \in \Sigma^+$ occurs in α with limiting frequency $\mu(w)$, we say that α is μ -distributed.

Proposition 2.6. Let μ be a probability map over Σ . If there exists a μ -distributed sequence, then μ is invariant.

Proof. Let μ be a probability map over Σ and $\alpha = a_1 a_2 \dots$ a μ -distributed sequence. We consider $w = w_1 w_2 \dots w_k \in \Sigma^k$ and note that for all $N > 0$:

$$\left| \sum_{a \in \Sigma} \#_{wa}(\alpha|_{\leq N}) - \#_w(\alpha|_{\leq N}) \right| \leq 1.$$

Indeed, every occurrence of w as $a_i a_{i+1} \dots a_{i+k}$ such that $i > 1$ is also an occurrence of $b \cdot w$ for a (unique) $b \in \Sigma$, so the expressions $\#_w(\alpha|_{\leq N})$ and $\sum_{a \in \Sigma} \#_{wa}(\alpha|_{\leq N})$ are equal if and only if $a_1 a_2 \dots a_k \neq w$ and their difference is equal to 1 otherwise.

Thus

$$\left| \frac{\sum_{a \in \Sigma} \#_{wa}(\alpha|_{\leq N})}{N} - \frac{\#_w(\alpha|_{\leq N})}{N} \right| \leq \frac{1}{N}.$$

We therefore obtain that:

$$\left| \frac{\sum_{a \in \Sigma} \#_{wa}(\alpha|_{\leq N})}{N} - \mu(w) \right| \leq \left| \frac{\sum_{a \in \Sigma} \#_{wa}(\alpha|_{\leq N})}{N} - \frac{\#_w(\alpha|_{\leq N})}{N} \right| + \left| \frac{\#_w(\alpha|_{\leq N})}{N} - \mu(w) \right|.$$

Since both expressions on the right converge to 0, the left-hand side converges to zero, showing that $\mu(w) = \lim_{n \rightarrow \infty} \frac{\sum_{a \in \Sigma} \#_{wa}(\alpha|_{\leq n})}{n} = \sum_{a \in \Sigma} \lim_{n \rightarrow \infty} \frac{\#_{wa}(\alpha|_{\leq n})}{n} = \sum_{a \in \Sigma} \mu(wa)$.

Similarly, for all $N > 0$:

$$\left| \sum_{a \in \Sigma} \#_{wa}(\alpha|_{\leq N}) - \#_w(\alpha|_{\leq N}) \right| \leq 1,$$

by a similar argument as the one used above, noting that the number of occurrences is different if and only if $a_{N-k+1}a_{N-k+2}\dots a_N = w$. We then conclude that $\mu(w) = \sum_{a \in \Sigma} \mu(aw)$ in the same way. \square

Observe that an infinite sequence α is normal in the usual sense iff it is μ -distributed for (the unique) equidistributed probability map μ over Σ . Also observe that it is not all probability maps μ for which there exists a μ -distributed sequence.

Example 2.7. An example of a probability map that is *not* Bernoulli, but such that there is at least one μ -distributed right-infinite sequence, is the map μ over $\Sigma = \{0, 1\}$ defined by $\mu(w) = 1/2$ if w does not contain any of the strings 00 or 11 (note that for each positive integer n , there are exactly two such strings of length n , namely 0101010 \dots and 101010 \dots), and $\mu(w) = 0$ otherwise. Observe that the right-infinite sequence 010101 \dots is μ -distributed.

In contrast to all previous work on Agafonov's Theorem, we allow countably infinite alphabets Σ . Alphabets of larger cardinality do not in general have probability measures realizable by considering limiting frequencies of elements of Σ^ω – simply because most elements of Σ cannot occur at all in a single element of Σ^ω .

One reason why previous generalizations of Agafonov's Theorem have not considered infinite alphabets is that there can be no equidistribution on a countably infinite set. However, there are Bernoulli measures μ on countably infinite alphabets Σ and μ -distributed infinite sequences over Σ .

Example 2.8. An example of a countably infinite alphabet with a Bernoulli measure is $\Sigma = \mathbb{N}$ and $p(n) = 6/(\pi n)^2$ (note that we have $\sum_{n \in \Sigma} p(n) = 1$). In general, any convergent series $\sum_{n=1}^{\infty} a_n$ where every a_n is non-negative induces a Bernoulli distribution on \mathbb{N} by setting $p(n) = a_n / \sum_{n=1}^{\infty} a_n$. Each such Bernoulli distribution p induces an invariant probability map μ_p , and by a result of Madritsch and Mance [37], there exists a μ_p -distributed sequence.

Remark 2.9. As we consider possibly infinite alphabets, we often have to consider infinite series instead of finite sums in the proofs. In most cases, these series will have elements that are known to be non-negative, and the sum of all partial sums will be bounded above, whence the series will be absolutely convergent and the order of summation can thus be changed freely. A trivial example of use is to consider some $B \subseteq \Sigma$ and note that $\sum_{a \in B} p(a) = 1 - \sum_{a \in \Sigma \setminus B} p(a)$ (as $\sum_{a \in B} p(a) \leq 1$, $\sum_{a \in \Sigma \setminus B} p(a) \leq 1$, and $p(a) \geq 0$ the two series are absolutely convergent, and $\sum_{a \in B} p(a) + \sum_{a \in \Sigma \setminus B} p(a) = \sum_{a \in \Sigma} p(a) = 1$).

2.1. Strategies.

Definition 2.10. Let Σ be an alphabet. A *strategy* S over Σ is a subset $S \subseteq \Sigma^*$.

Given a strategy S and $\alpha \in \Sigma^\omega$, we define the sequence *selected* by S , denoted $S[\alpha]$, as follows: if $i_1, i_2, \dots, i_k, \dots$ is the (increasing) sequence of indices i_j such that $\alpha|_{<i_j} \in S$, then $S[\alpha]_j = \alpha_{i_j}$. When $w \in \Sigma^*$ is a finite word, we define $S[w]$ *mutatis mutandis*.

A strategy S is a *Postnikova strategy* if there is $w \in \Sigma^*$ such that $S = \Sigma^*w$.

Thus, $S[\alpha]$ is simply the subsequence of symbols from α that are “picked out” by applying S to prefixes of α . Note also that if $w \in S$, then in any word on the form $w \cdot b \cdot v$, S must pick b . Thus, S cannot be made to, for instance, only pick out a single symbol from Σ – it must select “the next symbol” after any $w \in S$. This precludes, for example, constant-memory strategies from selecting only 0s from a normal binary sequence.

Our primary object of study is the case where S is a regular language, described next.

2.2. Finite-State Selectors and selection by DFAs. As we treat both finite and (countably) infinite alphabets, we must consider automata over possibly infinite alphabets. Every automaton has a finite number of states as usual, but as the alphabet is infinite and a deterministic automaton has transitions on all symbols from every state, the underlying graph of the automaton will be infinitely branching. To keep notations simple, we refer to deterministic automata with a finite number of states as “DFA”s as usual, even if the underlying alphabet is infinite.

Definition 2.11. A *finite-state selector* over Σ is a DFA $A = (Q, \delta, q_s, Q_F)$, where Q is the set of states, q_s is the unique start state, Q_F is the set of accepting states, and $\delta : Q \times \Sigma \rightarrow Q$ is the transition relation.

A DFA is strongly connected if its underlying directed graph (states are nodes, transitions are edges) is strongly connected.

Denote by $L(A)$ the language accepted by the automaton. If $\alpha = a_1 a_2 \dots$ is a finite or right-infinite sequence over Σ , the subsequence *selected by* A is the (possibly empty) sequence of letters a_n such that the prefix $a_1 \dots a_{n-1} \in L(A)$, that is, the automaton when started on the finite word $a_1 \dots a_{n-1}$ in state q_s ends in an accepting state after having read the entire word. The *run* of S on input α is the sequence of states visited when S is applied to α from the starting state. For $(q, w) = (q, w_1 \dots w_n) \in Q \times \Sigma^*$, we use the notation $\delta^*(q, w)$ to denote the state $\delta(\dots \delta(\delta(q, w_1), w_2) \dots w_n)$, that is, the state reached by starting from q and following the (unique) path induced by w .

Observe that a DFA may select an empty, finite or infinite sequence when run on a right-infinite word.

Definition 2.12. Let A be a DFA. A strongly connected component C in (the underlying directed graph of) S is said to be *recurrent* if, for every state p in C and every $a \in \Sigma, \delta(p, a)$ is a state in C (i.e., once a run of S on some infinite word reaches a state in C , the run cannot leave C).

Definition 2.13. Let $A = (Q, \Sigma, \delta, q_0, F)$ be a connected DFA. For all $q \in Q$, we denote by A_q the automaton $(Q, \Sigma, \delta, q, F)$, i.e. where the state q is chosen as the initial state.

Definition 2.14. Let $A = (Q, \Sigma, \delta, q_0, F)$ be a connected DFA, and let $q \in Q$. Let α be a right-infinite sequence over Σ . We denote by $A_q[x]$ the subsequence $\bar{\alpha}$ of α *picked out* by A_q , that is, $w_i \in \bar{w}$ if and only if $A_q(w_{<i})$ reaches an accepting state.

We shall use the following fundamental result in automata theory⁵:

Lemma 2.15 (Lemma 2.6 of [57]). *For every DFA $A = (Q, \delta, q_s, Q_F)$ over (the possibly infinite) alphabet Σ , there is a word $w \in \Sigma^*$ such that, for every $q \in Q$, there is a strongly connected recurrent component C of (the underlying directed graph of S) such that $\delta^*(q, w) \in C$.*

Corollary 2.15.1. *Let μ be a probability map induced by a positive Bernoulli distribution on Σ , let A be a DFA over Σ , and let $\alpha \in \Sigma^\omega$ be μ -distributed. Then, the run of S on α eventually reaches a strongly connected recurrent component of A .*

⁵The result in [57] is stated for finite alphabets, but the proof method carries through for infinite alphabets as well. We provide a proof in Appendix A.

Proof. Let w be the word obtained from Lemma 2.15. As α is p -distributed, w appears in α , so write $\alpha = v w \alpha'$, and let q be the state of A reached after $|v|$ transitions in the run of S on α . Then, after at most a further $|w|$ transitions, the run reaches a state in a strongly connected component of (the underlying directed graph of) A . \square

Corollary 2.15.1 ensures that we can assume without loss of generality that the finite-state selectors we treat are strongly connected. Note that the corollary does not imply that the strongly connected recurrent component contains an accepting state (indeed, the automata may have an empty set of accepting states). Thus, some automata do not always select infinite sequences, and additional assumptions are needed if this is desirable (this is discussed in Remark 3.2 below). However, this is not an issue for our main result which states that the output of a selector applied to a normal sequence is again normal *as long as it is infinite*.

3. MAIN RESULT

Theorem 3.1. *Let Σ be a non-empty (finite or infinite) alphabet and μ be a probability map such that there exists at least one $\alpha \in \Sigma^\omega$ that is μ -distributed. Then, the following statements are equivalent:*

- (1) μ is induced by a positive Bernoulli distribution p on Σ , that is, for every $a_1 \cdots a_n \in \Sigma$, $\mu(a_1 \cdots a_n) = \prod_{i=1}^n p(a_i)$, and $p(a) > 0$ for all $a \in \Sigma$;
- (2) (Postnikova property) for every finite word $w \in \Sigma^*$ and μ -distributed sequence $\alpha \in \Sigma^\omega$, if the sequence selected from α by the Postnikova strategy $\mathcal{S}_w = \{u \in \Sigma^* \mid \exists v \text{ s.t. } u = vw\}$ is infinite, then it is μ -distributed;
- (3) (Agafonov property) For every DFA A over Σ and every μ -distributed sequence $\alpha \in \Sigma^\omega$, if the sequence selected from α by A is infinite, then it is μ -distributed.

Proof. For the implication $1 \Rightarrow 3$, 2.15.1 yields that any run of a finite-state selector on a μ -distributed sequence eventually reaches a strongly connected recurrent component; the restriction of any DFA to the state set of one of its recurrent component is also a DFA, and the result now follows by 5.11. The implication $3 \Rightarrow 2$ is clear from the definitions since the considered strategies are computed by finite automata (4.3). Lastly, 4.1 and 4.2 prove that $2 \Rightarrow 1$. \square

Remark 3.2. Theorem 3.1 addresses the case where a DFA or Postnikova strategy selects an infinite sequence from a μ -distributed sequence. If one wants to restrict attention to automata that *always* select an infinite subsequence from any μ -distributed sequence, extra conditions sometimes occur in the literature, e.g. that every cycle in the (underlying graph of the) DFA contains an accepting state [6] ensuring that an infinite subsequence is selected from any (not just μ -distributed sequence). Another condition that ensures that an infinite subsequence is selected from any μ -distributed sequence is to consider only DFAs such that every strongly connected recurrent component contains at least one accepting state. In this case, Corollary 2.15.1 ensures that any run on the automaton on a μ -distributed sequence will reach a strongly recurrent component, and Lemma 5.7 below then ensures that the DFA accepts an infinite subsequence from α .

4. NON-PRESERVATION OF μ -DISTRIBUTEDNESS FOR NON-BERNOULLI MEASURES

We first prove that if μ is a probability map such that any DFA selects a μ -distributed right-infinite sequence from any μ -distributed right-infinite sequence, then μ must be Bernoulli. This is an immediate consequence of a stronger property proved in Lemma 4.1 below.

The idea of the proof is that if μ is *not* Bernoulli, there exists a word $a_1 \cdots a_k$ such that $\mu(a_1 \cdots a_{k-1}) = \prod_{j=1}^{k-1} \mu(a_j)$, but $\mu(a_1 \cdots a_{k-1}a_k) \neq \mu(a_1 \cdots a_{k-1}) \cdot \mu(a_k)$. One can then construct a finite-state selector that acts like a “sliding window” of size $k - 1$, that is, remembers the last $k - 1$ letters scanned and accepts if these are $a_1 \cdots a_k$. This selector will select every letter following $a_1 \cdots a_{k-1}$; after a prefix of length N of a right-infinite sequence has been scanned, approximately $N \cdot \mu(a_1 \cdots a_{k-1})$ have been selected, and approximately $N \cdot \mu(a_1 \cdots a_{k-1}a_k)$ of these will be the symbol a_k . But then the limiting frequency of a_k in the sequence selected will be $\mu(a_1 \cdots a_{k-1}a_k)/\mu(a_1 \cdots a_{k-1}) \neq \mu(a_k)$, and the result follows.

For completeness, we give a fully formal proof after the lemma, but the entirety of the reasoning is essentially as we just described.

Lemma 4.1. *Let $\mu : \Sigma^* \rightarrow [0, 1]$ be a probability map. If μ is not induced by a Bernoulli distribution on Σ , there exists a finite word $w \in \Sigma^*$ such that if $\alpha \in \Sigma^\omega$ is μ -distributed, then the Postnikova strategy $\mathcal{S}_w = \{u \in \Sigma^* \mid \exists v \text{ s.t. } u = vw\}$ selects from α an infinite sequence $\beta \in \Sigma^\omega$ that is not μ -distributed.*

Proof. If no element of Σ^ω is μ -distributed, the lemma is vacuously true. Hence, assume that there is at least one $\alpha \in \Sigma^\omega$ that is μ -distributed. If $|\Sigma| = 1$, then there is exactly one probability map on Σ^* , namely the one that assigns probability 1 to the unique element of Σ^k for every $k \geq 0$, and this probability map is clearly Bernoulli, and the lemma is thus vacuously true. Hence, in the remainder of the proof, assume that $|\Sigma| \geq 2$.

Assume that μ is not induced by a Bernoulli distribution on Σ . Then there are k and a word $a_1 \cdots a_{k-1}a_k \in \Sigma^k$ such that $\mu(a_1 \cdots a_{k-1}a_k) \neq \prod_{j=1}^k \mu(a_j)$. Observe that $k = 1$ is impossible (as $\mu(a_1) = \prod_{j=1}^1 \mu(a_1)$), and thus we must have $k \geq 2$. Assume without loss of generality that k is minimal among such k , and hence that $\mu(a_1 \cdots a_{k-1}) = \prod_{j=1}^{k-1} \mu(a_j)$, and note that this implies $\mu(a_1 \cdots a_{k-1}a_k) \neq \mu(a_1 \cdots a_{k-1}) \cdot \mu(a_k)$.

Assume for contradiction that $\mu(a_1 \cdots a_{k-1}) = 0$. Then $\mu(a_i) = 0$ for at least one a_i and thus $\mu(a_1 \cdots a_{k-1}a_k) = 0$, because the fact that there is at least one μ -distributed right-infinite sequence entails that $\mu(a_1 \cdots a_{k-1}a_k) > 0$ implies $\mu(a_i) \geq \mu(a_1 \cdots a_{k-1}a_k) > 0$. But this is a contradiction as we would then have $\mu(a_1 \cdots a_{k-1}a_k) = 0 = \prod_{j=1}^k \mu(a_j)$. Thus, $\mu(a_1 \cdots a_{k-1}) > 0$.

As $\mu(a_1 \cdots a_{k-1}a_k) \neq \mu(a_1 \cdots a_{k-1}) \cdot \mu(a_k)$, $\mu(a_1 \cdots a_{k-1}) > 0$, and we have $\mu(a_1 \cdots a_{k-1}a_k) \leq \mu(a_1 \cdots a_{k-1})$ (because μ is invariant by Proposition 2.6), there is a real number γ with $0 < \gamma < 1$ such that:

$$\left| \frac{\mu(a_1 \cdots a_{k-1}a_k)}{\mu(a_1 \cdots a_{k-1})} - \mu(a_k) \right| > \gamma$$

We now consider the Postnikova strategy \mathcal{S}_w with $w = a_1 \cdots a_{k-1}$, i.e. the strategy that selects exactly the symbols following the occurrences of $a_1 \cdots a_{k-1}$ in α .

Let $\alpha \in \Sigma^\omega$ be μ -distributed. Then, for every $\epsilon > 0$, there is an $N_\epsilon > 0$ such that for all $n > N_\epsilon$ we have:

$$\left| \frac{\#_{a_1 \cdots a_{k-1}}(\alpha|_{\leq n})}{n} - \mu(a_1 \cdots a_{k-1}) \right| \leq \epsilon$$

Hence

$$n\mu(a_1 \cdots a_{k-1}) - n\epsilon \leq \#_{a_1 \cdots a_{k-1}}(\alpha|_{\leq n}) \leq n\mu(a_1 \cdots a_{k-1}) + n\epsilon \quad (4.1)$$

and

$$n\mu(a_1 \cdots a_{k-1}a_k) - n\epsilon \leq \#_{a_1 \cdots a_{k-1}a_k}(\alpha|_{\leq n}) \leq n\mu(a_1 \cdots a_{k-1}a_k) + n\epsilon \quad (4.2)$$

As $\mu(a_1 \cdots a_{k-1}) > 0$ and \mathcal{S}_w selects the symbol after each occurrence of $a_1 \cdots a_{k-1}$, \mathcal{S}_w selects an infinite sequence β from α . Let $\beta^{(n)} \in \Sigma^*$ be the finite sequence selected by \mathcal{S}_w from $\alpha|_{\leq n}$. Observe that we have $|\beta^{(n)}| = \#_{a_1 \cdots a_{k-1}}(\alpha|_{< n})$, and $\#_{a_k}(\beta^{(n)}) = \#_{a_1 \cdots a_{k-1}a_k}(\alpha|_{< n})$. The fraction of occurrences $\#_{a_k}(\beta^{(n)})/|\beta^{(n)}|$ of a_k in $\beta^{(n)}$ thus satisfies:

$$\frac{\#_{a_k}(\beta^{(n)})}{|\beta^{(n)}|} = \frac{\#_{a_1 \cdots a_{k-1}a_k}(\alpha|_{< n})}{\#_{a_1 \cdots a_{k-1}}(\alpha|_{< n})} = \frac{\#_{a_1 \cdots a_{k-1}a_k}(\alpha|_{< n})}{n} \cdot \frac{n}{\#_{a_1 \cdots a_{k-1}}(\alpha|_{< n})}$$

and hence, by (4.1) and (4.2), for all $n > N$:

$$\frac{\mu(a_1 \cdots a_{k-1}a_k) - \epsilon}{\mu(a_1 \cdots a_{k-1}) + \epsilon} \leq \frac{\#_{a_k}(w_{A,n})}{|w_{A,n}|} \leq \frac{\mu(a_1 \cdots a_{k-1}a_k) + \epsilon}{\mu(a_1 \cdots a_{k-1}) - \epsilon} \quad (4.3)$$

Consider an arbitrary δ with $0 < \delta < \gamma/2$. By (4.3), for all sufficiently small ϵ , we have

$$\left| \frac{\mu(a_1 \cdots a_{k-1}a_k)}{\mu(a_1 \cdots a_{k-1})} - \frac{\#_{a_k}(\beta^{(n)})}{|\beta^{(n)}|} \right| < \delta$$

and thus for all $n > N_\epsilon$:

$$\begin{aligned} \gamma < \left| \frac{\mu(a_1 \cdots a_{k-1}a_k)}{\mu(a_1 \cdots a_{k-1})} - \mu(a_k) \right| &\leq \left| \frac{\mu(a_1 \cdots a_{k-1}a_k)}{\mu(a_1 \cdots a_{k-1})} - \frac{\#_{a_k}(\beta^{(n)})}{|\beta^{(n)}|} \right| + \left| \frac{\#_{a_k}(\beta^{(n)})}{|\beta^{(n)}|} - \mu(a_k) \right| \\ &< \delta + \left| \frac{\#_{a_k}(\beta^{(n)})}{|\beta^{(n)}|} - \mu(a_k) \right| < \frac{\gamma}{2} + \left| \frac{\#_{a_k}(\beta^{(n)})}{|\beta^{(n)}|} - \mu(a_k) \right| \end{aligned}$$

whence:

$$\left| \frac{\#_{a_k}(\beta^{(n)})}{|\beta^{(n)}|} - \mu(a_k) \right| > \frac{\gamma}{2}$$

and as the sequence $(\beta^{(n)})_{n \in \mathbb{N}}$ consists of prefixes of the sequence $\mathcal{S}_w[\alpha]$ selected by \mathcal{S}_w from α , and is eventually increasing, the frequency of occurrences of a_k differs infinitely often from $\mu(a_k)$ by at least $\gamma/2$, $\mathcal{S}_w[\alpha]$ cannot be μ -distributed. \square

Lemma 4.1 shows that if a probability map is not induced by a Bernoulli distribution on Σ , some Postnikova strategy will select a non- μ -distributed sequence from *any* μ -distributed sequence. In case μ is induced by a Bernoulli distribution, but not a positive Bernoulli distribution, we can show the weaker result that there will be a Postnikova strategy that selects a non- μ -distributed sequence from *some* μ -distributed sequences (and this is sufficient for our main Theorem).

Lemma 4.2. *Let $\mu : \Sigma^* \rightarrow [0, 1]$ be a probability map induced by a Bernoulli distribution on Σ that is not positive. Then there exists a finite word $w \in \Sigma^*$ and μ -distributed $\alpha \in \Sigma^\omega$ such that the Postnikova strategy $\mathcal{S}_w = \{u \in \Sigma^* \mid \exists v \text{ s.t. } u = vw\}$ selects from α an infinite sequence $\beta \in \Sigma^\omega$ that is not μ -distributed.*

Proof. As μ is not positive, pick $b \in \Sigma$ such that $\mu(b) = 0$, and let $\Gamma \subseteq \Sigma$ be a maximal subset such that the restriction of μ to Γ is a positive Bernoulli distribution (observe that Γ is non-empty because μ is a probability map and $1 = \sum_{a \in \Sigma} \mu(a)$ thus implies $\mu(a) > 0$ for some $a \in \Sigma$). By [37] there exists a μ -distributed infinite sequence $\beta \in \Gamma^\omega$; notice that β can be assumed w.l.o.g. to not contain any occurrences of b . Let $\alpha \in \Sigma^\omega$ be obtained by inserting the string bb at positions $2, 4, 8, 16, \dots$. Then, α is μ -distributed because (i) every $v \in \Gamma^*$ occurs with the same limiting frequency as in β^6 , and every $v \in \Sigma^*$ that contains an element of $\Sigma \setminus \Gamma$ occurs in α with limiting frequency 0. Set $w = b$; then the Postnikova strategy $\mathcal{S}_w = \{u \in \Sigma^* \mid \exists v \text{ s.t. } u = vw\}$ selects from α a sequence $\beta = \mathcal{S}_w[\alpha]$ such that, for every $n > 0$, $\#_b(\beta|_{\leq n}) \geq n/2 - 1$. Thus, the limiting frequency of b in β is not 0, and hence is not $\mu(b)$, proving that β is not μ -distributed. \square

Lemma 4.3. *Let $w \in \Sigma^*$. The Postnikova strategy $\{u \in \Sigma^* \mid \exists v \text{ s.t. } u = vw\}$ is computable by a strongly connected DFA over Σ .*

Proof. Note that the alphabet can possibly be infinite in the following proof.

We write m the length of the word w , and write w_1, w_2, \dots, w_m the bits of w . We design a finite state selector M_w with exactly 2^m states which will select a bit of the input if and only if it is preceded by the word w . Let $M_w = (\mathbf{2}^m, \delta, q_s, Q_F)$ be defined as follows:

- $\mathbf{2}^m = \{(b_1, b_2, \dots, b_m) : b_i \in \{0, 1\}\}$ is the set of binary sequences of length m ; those will represent a sequence of bits where $b_j = 1$ if and only if the previous j bits of the input coincide with the first j bits of the input;
- q_s the initial state is chosen to be the sequence $(0, 0, \dots, 0) \in \mathbf{2}^m$;
- Q_F the set of accepting states is equal to the set of sequences $\{(b_1, b_2, \dots, b_m) \in \mathbf{2}^m : b_m = 1\}$;
- δ the transition function is defined as $\delta(b_1, b_2, \dots, b_m; a) = (c_1, c_2, \dots, c_m)$ where $c_j = 1$ if and only if $b_{j-1} = 1$ and $a = w_j$ for $j \neq 1$, and $c_1 = 1$ if and only $a = w_1$.

The fact that this automaton computes the Postnikova strategy is clear from the definition. We now show it is strongly connected by showing that any state (b_1, b_2, \dots, b_m) is reachable from an arbitrary state. For this, we consider a word $u_{b_1, b_2, \dots, b_m} = u_1, \dots, u_m$ defined by $u_i = w_i$ if and only if $b_i = 1$ (and thus $u_i \neq w_i$ whenever $b_i = 0$). We then claim that the automaton, starting from any state $c \in \mathbf{2}^m$, reaches the state (b_1, b_2, \dots, b_m) when given the word u_{b_1, b_2, \dots, b_m} as input. \square

5. FINITE-STATE SELECTORS PRESERVE μ -DISTRIBUTEDNESS FOR BERNOULLI MEASURES

The sequence of auxiliary results of this section follows the general lines of Agafonov's original proof in Russian for the case $\Sigma = \{0, 1\}$ [45], but with multiple proofs needing more careful analysis and adapted techniques.

⁶The key observation here is that since the 'bb's are inserted at exponentially increasing positions, the frequency of occurrence of all other strings is decreased by a very small (and quickly decaying) factor.

5.1. Ancillary definitions and results.

Definition 5.1. Let Σ be an alphabet, $\alpha = x_1x_2\cdots x_n\cdots \in \Sigma^\omega$, and let n be a positive integer. The n -block decomposition of α is the sequence $(\alpha_{(n,r)})_{r \geq 1}$ where $\alpha_{(n,r)} = x_{(r-1)n+1} \cdots x_{rn} \in \Sigma^n$.

Thus, $\alpha_{(n,1)}$ is the string of the first n symbols of α , $\alpha_{(n,2)}$ is the string of the next n symbols, and so forth.

Definition 5.2. Let μ be a probability map over Σ and $\alpha = x_1x_2\cdots x_n\cdots \in \Sigma^\omega$. We say that α is μ -block-distributed if, for each $n \geq 1$ and every $w \in \Sigma^n$, the n -block decomposition $(\alpha_{(n,r)})_{r \geq 1}$ of α satisfies:

$$\lim_{k \rightarrow \infty} \frac{|i \leq k : \alpha_{(n,k)} = w|}{k} = \mu(w)$$

For finite alphabets and the special case of p being an equidistribution on Σ , it is straightforward to prove that the properties of being μ_p -distributed and μ_p -block-distributed are equivalent [43, 19, 50]. For the present paper, we only use that μ_p -distributedness implies μ_p -block-distributedness, which follows by tedious, but standard counting arguments on sufficiently large finite prefixes of $\alpha|_{\leq N}$ using the same reasoning as the original proof by Niven and Zuckerman for finite alphabets and normality [43], *mutatis mutandis*:

Proposition 5.3. *Let μ_p be a probability map induced by a Bernoulli distribution p on alphabet Σ . If $\alpha \in \Sigma^\omega$ is μ_p -distributed, it is μ_p -block distributed.*

We now prove that finite-state selectors can be composed appropriately; this will later be a key ingredient in reducing the problem of selecting finite strings $w \in \Sigma^*$ with frequency $\mu_p(w)$ to the problem of selecting single symbols $a \in \Sigma$ with frequency $p(a)$.

Proposition 5.4 (Finite-State selectors are compositional). *Let A and B be DFAs over the same alphabet. Then there is a DFA C such that, for each sequence w , $C[w] = B[A[w]]$. If A and B are both strongly connected and A contains at least one accepting state, C can be chosen to be strongly connected.*

Proof. Let $A = (Q^A, \Sigma, \delta^A, q_0^A, F^A)$ and $B = (Q^B, \Sigma, \delta^B, q_0^B, F^B)$. Define $Q^C = Q^A \times Q^B$, and set $q_0^C = (q_0^A, q_0^B)$ and $F^C = F^A \times F^B$. For each $q^B \in Q^B$, define the set $D_{q^B} = \{(q, q^B) : q \in Q^A\} \subseteq Q^C$. Observe that $Q^C = \bigcup_{q^B \in Q^B} D_{q^B}$ and that for $q^B, r^B \in Q^B$ with $q^B \neq r^B$, we have $D_{q^B} \cap D_{r^B} = \emptyset$, and thus $\{D_{q^B} : q^B \in Q^B\}$ is a partitioning of Q^C . Hence, the transition relation, δ^C , of C may be defined by defining it separately on each subset D_{q^B} :

$$\delta^C((q, q^B), a) = \begin{cases} (r, q^B) & \text{if } q \notin F^A \text{ and } \delta^A(q, a) = r \\ (r, r^B) & \text{if } q \in F^A \text{ and } \delta^A(q, a) = r \text{ and } \delta^B(q^B, a) = r^B \end{cases}$$

Thus, when C processes its input, it freezes the current state q^B of B (the freezing is represented by staying within D_{q^B}) and simulates A until an accepting state of A is reached (i.e. just before A would select the next symbol); on the next transition, C unfreezes the current state of B and moves to the next state r^B of B and then freezes it and continues with a simulation of A .

Observe that a symbol is picked out by C iff the state is an element of $F^C = F^A \times F^B$ iff the symbol is the next symbol read after simulation of A reaches an accepting state of A when the current frozen state of B is an accepting state of B .

By construction, C is strongly connected if both A and B are: for any pair of states (q_1^A, q_1^B) and (q_2^A, a_2^B) in Q^C , strong connectivity of B implies that there is a directed path from q_1^B to q_2^B in B . Let $q_1^B, q_{1,2}^B, q_{1,3}^B, \dots, q_{1,k}^B$ be the states along this path. Strong connectivity of A and the assumption that there is some $q_1^F \in F^A$ imply that there is a directed path from (q_1^A, q_1^B) to (q_1^F, q_1^B) in C , and by definition of δ^C , there is a transition in C from (q_1^F, q_1^B) to $(q_1^B, q_{1,2}^B)$. A straightforward induction on k now completes the proof. \square

The following shows that to prove that the property of being μ_p -distributed is preserved under finite-state selection, it suffices to prove that the limiting frequency of each $a \in \Sigma$ exists and equals $p(a)$.

Lemma 5.5. *Let μ_p be a probability map induced by a Bernoulli distribution p on Σ , and let $\alpha \in \Sigma^\omega$ be μ_p -distributed. The following are equivalent:*

- For all strongly connected DFAs A , if $A[\alpha]$ is infinite, then $A[\alpha]$ is μ_p -distributed.
- For all strongly connected DFAs A and all $a \in \Sigma$, if $A[\alpha]$ is infinite, then the limiting frequency of a in $A[\alpha]$ exists and equals $p(a)$.

Proof. If, for all A such that $A[\alpha]$ is infinite, $A[\alpha]$ is μ_p -distributed, then in particular the limiting frequency of a in $A[\alpha]$ exists and is equal to $p(a)$ for all A .

Conversely, suppose that, for all strongly connected DFAs A and all $a \in \Sigma$, if $A[\alpha]$ is infinite, then the limiting frequency of a in $A[\alpha]$ exists and equals $p(a)$. Let A be such a DFA. If A has no accepting states, there is nothing to prove, so assume that A has at least one accepting state.

We will prove by induction on $k \geq 0$ that the limiting frequency of every $v_1 \cdots v_k v_{k+1} \in \Sigma^{k+1}$ exists and equals $\mu_p(v_1 \cdots v_k v_{k+1})$.

- $k = 0$: This is the supposition.
- $k \geq 1$. Suppose that the result has been proved for $k - 1$. Let $v_1 \cdots v_k \in \Sigma^k$; by the induction hypothesis, the limiting frequency of $v_1 \cdots v_k$ in $A[w]$ is $\mu_p(v_1 \cdots v_k)$. We claim that there is a strongly connected DFA B that, from any sequence, selects the symbol after each occurrence of $v_1 \cdots v_k$. To see that such a DFA exists, let there be a state for each element of Σ^k and think of the state as the current length- k string in a “sliding window” that moves over w one symbol at the time; when the window is moved one step, the DFA transits to the state representing the new length- k string in the window, i.e. for any $a \in \Sigma$, from the state representing the word $w_1 \cdots w_k$, there are transitions to $w_2 \cdots w_k a$; it thus every state is reachable from every other state in at most k transitions. The unique final state of B is the state representing $v_1 \cdots v_k$; the start state of B can be chosen to be any state representing a string $w_1 \cdots w_k$ such that there are exactly k transitions to the final state; for example, let $a \in \Sigma \setminus \{v_1\}$. Then from the state representing a^k , the final state cannot be reached in $k - 1$ or fewer steps, but every state is reachable in k steps.

By 5.4, there is a strongly connected DFA C such that $C[w] = B[A[w]]$ for all $w \in \Sigma^*$.

For any $a \in \Sigma$ and any sufficiently large positive integer N , we have:

$$\frac{\#_a(C[\alpha|_{\leq N}])}{|C[\alpha|_{\leq N}]|} = \frac{\#_a(B[A[\alpha|_{\leq N}]])}{|B[A[\alpha|_{\leq N}]]|} = \frac{\#_{v_1 \cdots v_k a}(A[\alpha|_{\leq N}])}{\#_{v_1 \cdots v_k}(A[\alpha|_{\leq N}])}$$

By the induction hypothesis, for every $\epsilon > 0$, we have, for all sufficiently large N , that $\left| \frac{\#_a(C[\alpha|_{\leq N}])}{|C[\alpha|_{\leq N}]|} - p(a) \right| < \epsilon$, and hence:

$$\left| \frac{\#_{v_1 \cdots v_k a}(A[\alpha|_{\leq N}])}{\#_{v_1 \cdots v_k}(A[\alpha|_{\leq N}])} - p(a) \right| < \epsilon \quad (5.1)$$

But for all sufficiently large N , the induction hypothesis also furnishes that:

$$\left| \frac{\#_{v_1 \dots v_k}(A[\alpha|_{\leq N}])}{|A[\alpha|_{\leq N}]|} - \mu_p(v_1 \dots v_k) \right| < \epsilon \quad (5.2)$$

But as:

$$\frac{\#_{v_1 \dots v_k a}(A[\alpha|_{\leq N}])}{|A[\alpha|_{\leq N}]|} = \frac{\#_{v_1 \dots v_k a}(A[\alpha|_{\leq N}])}{\#_{v_1 \dots v_k}(A[\alpha|_{\leq N}])} \cdot \frac{\#_{v_1 \dots v_k}(A[\alpha|_{\leq N}])}{|A[\alpha|_{\leq N}]|}$$

5.1 and 5.2 thus yield:

$$\begin{aligned} \left| \frac{\#_{v_1 \dots v_k a}(A[\alpha|_{\leq N}])}{|A[\alpha|_{\leq N}]|} - \mu_p(v_1 \dots v_k a) \right| &= \left| \frac{\#_{v_1 \dots v_k a}(A[\alpha|_{\leq N}])}{\#_{v_1 \dots v_k}(A[\alpha|_{\leq N}])} \cdot \frac{\#_{v_1 \dots v_k}(A[\alpha|_{\leq N}])}{|A[\alpha|_{\leq N}]|} - \mu_p(v_1 \dots v_k) p(a) \right| \\ &< \epsilon^2 + \epsilon \left(\frac{\#_{v_1 \dots v_k a}(A[\alpha|_{\leq N}])}{\#_{v_1 \dots v_k}(A[\alpha|_{\leq N}])} + \frac{\#_{v_1 \dots v_k}(A[\alpha|_{\leq N}])}{|A[\alpha|_{\leq N}]|} \right) \\ &\leq \epsilon^2 + 2\epsilon \end{aligned}$$

Hence, for all $a \in \Sigma$, the limiting frequency of $v_1 \dots v_k a$ in $A[\alpha]$ exists and equals $\mu_p(v_1 \dots v_k a)$, as desired. \square

5.2. Preservation of Bernoulli μ_p -distributedness under finite-state selection. By Lemma 5.5 we may restrict our attention to proving that the frequency of single symbols from Σ are preserved under selection by DFAs. The strategy will be to consider an arbitrary strongly connected DFA A , split the set of finite words Σ^* into multiple classes that depend on the selection behaviour of A , and use a combination of concentration bounds and basic Markov chain theory applied to these classes to obtain upper and lower bounds on the frequency with which A selects each symbol from A .

Definition 5.6. Let $A = (Q, \Sigma, \delta, q_0, F)$ be a strongly connected DFA. For any probability distribution $p : \Sigma \rightarrow [0, 1]$, any $b \in [0, 1]$, $n \in \mathbb{N}$, and any $\epsilon > 0$, we define sets $D_n^p(b, \epsilon)$, $E_n(b, q)$ and $G_n(b, \epsilon, q)$ as follows:

$$\begin{aligned} D_n^p(b, \epsilon, q) &= \left\{ w \in \Sigma^n : |A_q[w]| > bn \text{ and } \sup_{a \in \Sigma} \left| \frac{\#_a(A_q[w])}{|A_q[w]|} - p(a) \right| < \epsilon \right\} \\ D_n^p(b, \epsilon) &= \bigcap_{q \in Q} D_n^p(b, \epsilon, q) \\ E_n(b, q) &= \{w \in \Sigma^n : |A_q[w]| \leq bn\} \\ E_n(b) &= \bigcup_{q \in Q} E_n(b, q) \\ G_n(b, \epsilon, q) &= \left\{ w \in \Sigma^n : |A_q[w]| > bn \text{ and } \sup_{a \in \Sigma} \left| \frac{\#_a(A_q[w])}{|A_q[w]|} - p(a) \right| \geq \epsilon \right\} \\ G_n(b, \epsilon) &= \bigcup_{q \in Q} G_n(b, \epsilon, q) \end{aligned}$$

Observe that, for all b, n, ϵ ,

$$\Sigma^n = E_n(b) \cup D_n^p(b, \epsilon) \cup G_n(b, \epsilon)$$

(and also note that $E_n(b)$ and $G_n(b, \epsilon)$ are not necessarily disjoint).

Lemma 5.7. *Let p be a positive Bernoulli distribution on Σ , and let $S = (Q, \Sigma, \delta, q_s, Q_F)$ be a strongly connected finite automaton with $Q_F \neq \emptyset$, and let n a positive integer. Then there exists a real number $c > 0$ such that for all real numbers $\epsilon > 0$ we have $\lim_{n \rightarrow \infty} \mu_p(E_n(c - \epsilon)) = 0$.*

Proof. S induces a stochastic $|Q| \times |Q|$ matrix \mathbf{P} by setting

$$\mathbf{P}_{ij} = \sum_{a \in \Sigma} p(a) \cdot 1_{\delta(i,a)=j}.$$

Observe that if Σ is infinite, the fact that (i) $p(a) \cdot 1_{\delta(i,a)=j} \geq 0$, (ii) $p(a) \cdot 1_{\delta(i,a)=j} \leq p(a)$, and (iii) $\sum_{a \in \Sigma} p(a) = 1$ entails that the series $\sum_{a \in \Sigma} p(a) \cdot 1_{\delta(i,a)=j}$ is absolutely convergent.

Note also that $\mathbf{P}_{ij} = 0$ iff there are no transitions from i to j in Q on a symbol $a \in \Sigma$ with $p(a) > 0$. As S is strongly connected, there exists a path from state i to state j for each $i, j \in Q$. Let v be the word along this path; as $p(a) > 0$ for all $a \in \Sigma$, we have $\mu_p(v) > 0$, whence for each i, j there is an integer n_{ij} such that $\mathbf{P}_{ij}^{n_{ij}} > 0$, that is, \mathbf{P} (and its associated Markov chain) is irreducible. As all states of a finite Markov chain with irreducible transition matrix are positive recurrent, standard results (see, e.g., [58, Thm. 54]) yield that there is a unique positive stationary distribution $\pi : Q \rightarrow [0, 1]$ (s.t., for all $i \in Q$, we have $\pi(i) > 0$ and $\pi(i) = \sum_{j \in Q} \pi(j) \mathbf{P}_{ij}$). Furthermore, the expected return time M_i to state i satisfies $M_i = 1/\pi(i)$.

Let $(X_n)_{n \geq 0} = (X_0, X_1, X_2, \dots)$ be the Markov chain with transition matrix \mathbf{P} and some initial distribution λ on the states. Consider, for each $i \in Q$, the stochastic variable V_i , where

$$V_i(n) = \sum_{k=0}^{n-1} 1_{X_k=i},$$

that is, $V_i(n)$ is the number of times state i is visited in the first n elements of the Markov chain. As \mathbf{P} is irreducible, the Ergodic Theorem for Markov chains (see, e.g., [58, Thm. 75]) yields that, independently of λ , we have for arbitrary $\epsilon > 0$:

$$\lim_{n \rightarrow \infty} \Pr \left(\left| \frac{V_i(n)}{n} - \pi(i) \right| \geq \epsilon \right) = \lim_{n \rightarrow \infty} \Pr \left(\left| \frac{V_i(n)}{n} - \frac{1}{M_i} \right| \geq \epsilon \right) = 0 \quad (5.3)$$

Let n be a positive integer, let $w = w_1 \cdots w_n \in \Sigma^n$, and let $\mathbf{q}_{S_j}(w) = q_0^w q_1^w \cdots q_n^w$ be the sequence of states visited in the run of S_j on w (i.e., $q_0^w = j$). The probability of observing a state sequence $q_0 \cdots q_n$ in the Markov chain is (when the initial distribution λ has $\lambda(q_0) = \lambda(j) = 1$):

$$\Pr(q_0 \cdots q_n) = \prod_{i=0}^{n-1} \sum_{a \in \Sigma} p(a) 1_{\delta(q_i, a) = q_{i+1}} = \sum_{a_1, \dots, a_n \in \Sigma} p(a_1) 1_{\delta(q_0, a_1) = q_1} \cdots p(a_n) 1_{\delta(q_{n-1}, a_n) = q_n}$$

where we have used the fact that the Cauchy product of two absolutely convergent series is convergent.

As for all integers i with $0 \leq i \leq n$ we have $\delta(q_{i-1}^w, w_i) = q_i^w$, we obtain:

$$\sum_{a_1, \dots, a_n \in \Sigma} p(a_1) 1_{\delta(q_0, a_1) = q_1} \cdots p(a_n) 1_{\delta(q_{n-1}, a_n) = q_n} = \mu_p(\{a_1 \cdots a_n : \mathbf{q}_{S_j}(a_1 \cdots a_n) = q_0 \cdots q_n\})$$

and hence

$$\Pr(q_0 \cdots q_n) = \mu_p(\{w : \mathbf{q}_{S_j}(w) = q_0 \cdots q_n\}) \quad (5.4)$$

Thus, as S is deterministic and every $w_1 \cdots w_n \in \Sigma^n$ occurs along exactly one path of states in S , we have:

$$\begin{aligned} \Pr\left(\left|\frac{V_i(n)}{n} - \pi(i)\right| \geq \epsilon\right) &= \sum_{q_0 q_1 \cdots q_n \in Q^n} \Pr(q_0 \cdots q_n) \mathbf{1}_{|V_i(n)/n - \pi(i)| \geq \epsilon} \\ &= \sum_{q_0 q_1 \cdots q_n \in Q^n} \mu_p(\{w_1 \cdots w_n : \mathbf{q}_{S_j}(w_1 \cdots w_n) = q_0 \cdots q_n\}) \\ &= \mu_p\left(w : \left|\frac{V_i(n)}{n} - \pi(i)\right| \geq \epsilon\right) \end{aligned} \quad (5.5)$$

Hence, by 5.3 and 5.5, we have

$$\lim_{n \rightarrow \infty} \mu_p\left(w : \left|\frac{V_i(n)}{n} - \pi(i)\right| \geq \epsilon\right) = 0 \quad (5.6)$$

If $\mathbf{q}_{S_j}(w) = q_0 \cdots q_n$ and $q_k \in Q_F$ for some k with $0 \leq k \leq n-1$, then S_j selects w_{k+1} . Set $c = \min_{q_i \in Q_F} \pi(i)$ (c is well-defined as $Q_F \neq \emptyset$), and let $i \in Q_F$ be such that $\pi(i) = c$. Then, for all $j \in Q$:

$$\begin{aligned} \mu_p(E_n(c - \epsilon, j)) &= \mu_p(\{w \in \Sigma^n : |S_j[w]| \leq (c - \epsilon)n\}) \leq \mu_p(\{w \in \Sigma^n : V_i(n) \leq (c - \epsilon)n\}) \\ &= \mu_p\left(\left\{w \in \Sigma^n : \frac{V_i(n)}{n} - c \leq -\epsilon\right\}\right) = \mu_p\left(\left\{w \in \Sigma^n : \left|\frac{V_i(n)}{n} - c\right| \geq \epsilon\right\}\right) \end{aligned}$$

And hence, by 5.6, we have $\lim_{n \rightarrow \infty} \mu_p(E_n(c - \epsilon, j)) = 0$, and as $j \in Q$ was arbitrary, we obtain

$$\begin{aligned} \lim_{n \rightarrow \infty} \mu_p(E_n(c - \epsilon)) &= \lim_{n \rightarrow \infty} \mu_p(\cup_{j \in Q} \mu_p(E_n(c - \epsilon, j))) \leq \lim_{n \rightarrow \infty} \sum_{j \in Q} \mu_p(E_n(c - \epsilon, j)) \\ &= \sum_{j \in Q} \lim_{n \rightarrow \infty} \mu_p(E_n(c - \epsilon, j)) = 0 \end{aligned}$$

as desired. \square

Lemma 5.8. *Let S be a strategy, $a \in \Sigma$, b, ϵ be real numbers with $0 < b \leq 1$ and $\epsilon > 0$, and $p : \Sigma \rightarrow [0, 1]$ a positive Bernoulli distribution. Define, for all positive integers n :*

$$\begin{aligned} H_n(b, \epsilon) &= \left\{w \in \Sigma^n : |S(w)| > bn \wedge \left|p(a) - \frac{\#_a(S(w))}{|S(w)|}\right| \geq \epsilon\right\} \\ &= \bigcup_{bn < \ell \leq n} \left\{w \in \Sigma^n : S(w) \in \Sigma^\ell \wedge \left|p(a) - \frac{\#_a(S(w))}{\ell}\right| \geq \epsilon\right\} \end{aligned}$$

Then:

$$\lim_{n \rightarrow \infty} \mu_p(H_n(b, \epsilon)) = 0$$

Proof. Define

$$F_n(b, \epsilon) = \bigcup_{bn < \ell \leq n} \left\{y \in \Sigma^\ell : \left|p(a) - \frac{\#_a(y)}{\ell}\right| \geq \epsilon\right\}$$

Observe that $H_n(b, \epsilon) = \{w \in \Sigma^n : S(w) \in F_n(b, \epsilon)\}$. Thus, $\mu_p(H_n(b, \epsilon)) \leq \mu_p(F_n(b, \epsilon))$ for all n , and it thus suffices to prove that $\lim_{n \rightarrow \infty} \mu_p(F_n(b, \epsilon)) = 0$.

Consider the stochastic variable X_a that is 1 when a is picked from Σ with probability $p(a)$, and 0 otherwise. Then, the mean of X_a is $p(a)$ and the variance of X_a is $p(a)(1 - p(a))$. Now consider performing $\ell \geq 1$ independent Bernoulli trials drawn according to X_a . Define $q : \{0, 1\}^+ \rightarrow [0, 1]$ inductively by $q(1) = p(a)$, $q(0) = 1 - p(a)$, and $q(1c) = p(a)q(c)$ and $q(0c) = (1 - p(a))q(c)$ for $c \in \Sigma^+$, and observe that q induces a probability distribution \bar{q} on Σ^ℓ by setting $\bar{q}(w) = q(w)$. Now, for any $v \in \Sigma^\ell$, $\bar{q}(v)$ is the probability of obtaining v by performing ℓ repeated Bernoulli trials as above.

Define the stochastic variable $X_a^\ell = X_a + X_a + \cdots + X_a$ (ℓ times). Then, X^ℓ counts the number of occurrences of a by performing the ℓ repeated Bernoulli trials.

By the Chernoff bound, X_a^ℓ satisfies:

$$\Pr \left(\left| p(a) - \frac{X_a^\ell}{\ell} \right| \geq \epsilon \right) \leq 2e^{-\frac{\ell\epsilon^2}{3p(a)}} \quad (5.7)$$

Define the map $g : \Sigma \rightarrow \{0, 1\}$ by $g(a) = 1$ and $g(b) = 0$ for all $b \in \Sigma \setminus \{a\}$. Clearly, g extends homomorphically to a map $\tilde{g} : \Sigma^\ell \rightarrow \{0, 1\}^\ell$ by setting $\tilde{g}(c_1 c_2 \cdots c_\ell) = g(c_1)g(c_2) \cdots g(c_\ell)$.

Claim: For any $u \in \{0, 1\}^\ell$,

$$\bar{q}(u) = \mu_p(\{y \in \Sigma^\ell : \tilde{g}(y) = u\}) \quad (5.8)$$

Proof of claim: By induction on ℓ .

- If $\ell = 1$, then if $u = 0$, we have $\{y \in \Sigma^\ell : \tilde{g}(y) = u\} = \Sigma \setminus \{a\}$ and thus:

$$\bar{q}(u) = \bar{q}(0) = q(0) = 1 - p(a) = \sum_{b \in \Sigma \setminus \{a\}} p(b) = \mu_p(\Sigma \setminus \{a\})$$

Similarly, if $u = 1$, we have $\{y \in \Sigma^\ell : \tilde{g}(y) = u\} = \{a\}$, and thus $\bar{q}(u) = \bar{q}(1) = q(1) = p(a) = \mu_p(\{a\})$, as desired.

- If $\ell > 1$, write $u = b_1 \cdots b_{\ell-1} b_\ell$; by the induction hypothesis:

$$\bar{q}(b_1 \cdots b_{\ell-1}) = \mu_p(\{y \in \Sigma^\ell : \tilde{g}(y') = b_1 \cdots b_{\ell-1}\}) = \sum_{\substack{y' \in \Sigma^{\ell-1} \\ \tilde{g}(y') = b_1 \cdots b_{\ell-1}}} \mu_p(y')$$

If $b_\ell = 0$, then:

$$\bar{q}(b_1 \cdots b_{\ell-1} b_\ell) = \bar{q}(b_1 \cdots b_{\ell-1})q(0) = \bar{q}(b_1 \cdots b_{\ell-1})(1 - p(a)) = \sum_{\substack{y' \in \Sigma^{\ell-1} \\ \tilde{g}(y') = b_1 \cdots b_{\ell-1}}} \mu_p(y')(1 - p(a))$$

$$= \sum_{\substack{y' \in \Sigma^{\ell-1} \\ \tilde{g}(y') = b_1 \cdots b_{\ell-1}}} \left(\mu_p(y') \sum_{c \in \Sigma \setminus \{a\}} p(c) \right) = \sum_{\substack{y' \in \Sigma^{\ell-1} \\ \tilde{g}(y') = b_1 \cdots b_{\ell-1}}} \sum_{c \in \Sigma \setminus \{a\}} \mu_p(y')p(c)$$

$$\stackrel{(\dagger)}{=} \sum_{\substack{y' \in \Sigma^{\ell-1} \\ \tilde{g}(y') = b_1 \cdots b_{\ell-1} \\ c \in \Sigma \setminus \{a\}}} \mu_p(y'c) = \sum_{\substack{y \in \Sigma^\ell \\ \tilde{g}(y) = b_1 \cdots b_{\ell-1} b_\ell}} \mu_p(y)$$

$$= \mu_p(\{y \in \Sigma^\ell : \tilde{g}(y) = b_1 \cdots b_{\ell-1} b_\ell\})$$

where (\dagger) follows as both series on the left- and right-hand sides of the equality are absolutely convergent. The proof for the case $b_\ell = 1$ is symmetric, mutatis mutandis.

(End of proof of claim.)

Observe that, for any $y \in \Sigma^\ell$, we have:

$$|p(a) - \#_1(\tilde{g}(y))/\ell| \geq \epsilon \quad \text{iff} \quad |p(a) - \#_a(y)/\ell| \geq \epsilon \quad (5.9)$$

Hence, by Equation (5.8), for any event $\mathcal{U} \subseteq \{0, 1\}^\ell$, we have:

$$\begin{aligned} \Pr(\mathcal{U}) &= \sum_{u \in \mathcal{U}} \bar{q}(u) = \sum_{u \in \mathcal{U}} \mu_p(\{y \in \Sigma^\ell : \tilde{g}(y) = u\}) \\ &= \mu_p(\{y \in \Sigma^\ell : \tilde{g}(y) \in \mathcal{U}\}) \end{aligned} \quad (5.10)$$

The event $|p(a) - X_a^\ell/\ell| \geq \epsilon$ is shorthand for the set

$$\left\{ u \in \{0, 1\}^\ell : \left| p(a) - \frac{\sum_{j=1}^\ell u_j}{\ell} \right| \geq \epsilon \right\} = \left\{ u \in \{0, 1\}^\ell : \left| p(a) - \frac{\#_1(u)}{\ell} \right| \geq \epsilon \right\}$$

We thus obtain:

$$\begin{aligned} \Pr\left(\left| p(a) - \frac{X_a^\ell}{\ell} \right| \geq \epsilon\right) &= \Pr\left(\left\{ u \in \{0, 1\}^\ell : \left| p(a) - \frac{\#_1(u)}{\ell} \right| \geq \epsilon \right\}\right) \\ &= \mu_p\left(\left\{ y \in \Sigma^\ell : \left| p(a) - \frac{\#_1(\tilde{g}(y))}{\ell} \right| \geq \epsilon \right\}\right) \quad \text{by (5.10)} \\ &= \mu_p\left(\left\{ y \in \Sigma^\ell : \left| p(a) - \frac{\#_a(y)}{\ell} \right| \geq \epsilon \right\}\right) \quad \text{by (5.9)} \end{aligned} \quad (5.11)$$

Observe that:

$$\begin{aligned} \mu_p(F_n(b, \epsilon)) &= \mu_p\left(\bigcup_{bn < \ell \leq n} \left\{ y \in \Sigma^\ell \cap F_n(b, \epsilon) : \left| p(a) - \frac{\#_a(y)}{\ell} \right| \geq \epsilon \right\}\right) \\ &= \sum_{bn < \ell \leq n} \mu_p\left(\left\{ y \in \Sigma^\ell \cap F_n(b, \epsilon) : \left| p(a) - \frac{\#_a(y)}{\ell} \right| \geq \epsilon \right\}\right) \\ &\leq \sum_{bn < \ell \leq n} \mu_p\left(\left\{ y \in \Sigma^\ell : \left| p(a) - \frac{\#_a(y)}{\ell} \right| \geq \epsilon \right\}\right) \\ &= \sum_{bn < \ell \leq n} \Pr\left(\left| p(a) - \frac{X_a^\ell}{\ell} \right| \geq \epsilon\right) \quad \text{by 5.11} \\ &\leq \sum_{bn < \ell \leq n} 2e^{-\frac{\ell \epsilon^2}{3p(a)}} \quad \text{by 5.7} \\ &\leq (1-b)n2e^{-\frac{bn\epsilon^2}{3p(a)}} \end{aligned}$$

And thus $\lim_{n \rightarrow \infty} \mu_p(F_n(b, \epsilon)) = 0$, as desired. □

Corollary 5.8.1. *Let b, ϵ be real numbers with $0 < b \leq 1$ and $\epsilon > 0$. Then,*

$$\lim_{n \rightarrow \infty} \mu_p(G_n(b, \epsilon)) = 0$$

Proof. By Lemma 5.8 with S the strategy defined by the automaton A_q , we obtain that:

$$\lim_{n \rightarrow \infty} \mu_p(G_n(b, \epsilon, q)) = 0$$

and as $G_n(b, \epsilon) = \bigcup_{q \in Q} G_n(b, \epsilon, q)$, we have:

$$\mu_p(G_n(b, \epsilon)) \leq \sum_{q \in Q} \mu_p(G_n(b, \epsilon, q))$$

As Q is finite, we hence obtain $\lim_{n \rightarrow \infty} \mu_p(G_n(b, \epsilon)) = 0$. \square

Lemma 5.9. *There is a real number b with $0 < b \leq 1$ such that for all $\epsilon > 0$:*

$$\lim_{n \rightarrow \infty} \mu_p(D_n^p(b, \epsilon)) = 1.$$

Proof. Observe that, for all b with $0 < b \leq 1$:

$$\begin{aligned} \Sigma^n \setminus D_n^p(b, \epsilon) &= \{w \in \Sigma^n : \exists q \in Q. |A_q[w]| \leq bn\} \\ &\cup \left\{ w \in \Sigma^n : \exists q \in Q. |A_q[w]| > bn \wedge \sup_{a \in \Sigma} \left| \frac{\#_a(A_q[w])}{|A_q[w]|} - p(a) \right| \geq \epsilon \right\} \\ &= \left(\bigcup_{q \in Q} E_n(b, q) \right) \cup \left(\bigcup_{q \in Q} G_n(b, \epsilon, q) \right) \end{aligned}$$

and thus:

$$\begin{aligned} \mu_p(\Sigma^n \setminus D_n^p(b, \epsilon)) &\leq \mu_p \left(\bigcup_{q \in Q} E_n(b, q) \right) + \mu_p \left(\bigcup_{q \in Q} G_n(b, \epsilon, q) \right) \\ &= \mu_p(G_n(b, \epsilon)) + \mu_p(E_n(b)) \end{aligned}$$

By 5.7, choose a real number $c > 0$ such that $\lim_{n \rightarrow \infty} \mu_p(E_n(c - \epsilon)) = 0$, and set $b = c - \epsilon$.

By 5.8.1, we obtain that $\lim_{n \rightarrow \infty} G_n(b, \epsilon) = 0$, and thus $\lim_{n \rightarrow \infty} \mu_p(\Sigma^n \setminus D_n^p(b, \epsilon)) = 0$. The result now follows by $\mu_p(D_n^p(b, \epsilon)) = 1 - \mu_p(\Sigma^n \setminus D_n^p(b, \epsilon))$. \square

Lemma 5.10. *Let $p : \Sigma \rightarrow [0, 1]$ be a probability distribution, let $\alpha \in \Sigma^\omega$ be μ_p -block-distributed, and A a strongly connected DFA over Σ . Then, for all $a \in \Sigma$, the limiting frequency of a in the sequence $\beta = A[\alpha]$ exists and equals $p(a)$.*

Proof. For each n, r , let $\beta_{(n,r)}$ be the sequence of symbols picked out from block $\alpha_{(n,r)}$ when A is applied to α ; note that each $\beta_{(n,r)}$ has length between 0 and n .

For each positive integer m , define:

$$L_m = \sum_{i=1}^m |\beta_{(n,i)}|$$

And for each $a \in \Sigma$, define ρ_a^m by:

$$\rho_a^m = \frac{\sum_{i=1}^m \#_a(\beta_{(n,i)})}{L_m}$$

To prove the lemma, it suffices to show that, for any real number $\epsilon > 0$, then for all sufficiently large m , we have $|\rho_a^m - p(a)| < \epsilon$.

Define:

$$I_m = \left\{ i \leq m : \alpha_{(n,i)} \notin D_n^p \left(b, \frac{\epsilon}{2} \right) \right\}$$

where $b < 1$ is a constant to be fixed later in the proof.

And define:

$$\ell_m = \sum_{i \in I_m} |\beta_{(n,i)}|$$

Now, define θ_a^m by:

$$\theta_a^m = \frac{\sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)})}{\sum_{i \in \{1, \dots, m\} \setminus I_m} |y_{(n,i)}|} = \frac{\sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)})}{L_m - \ell_m}$$

That is, θ_a^m is the frequency of occurrences of a when the blocks $\beta_{(n,i)}$ picked out from blocks $\alpha_{(i,r)} \in D_n^p(b, \frac{\epsilon}{2})$ with $i \leq m$ are all concatenated. Observe that, by definition of D_n^p , we have $|\theta_a^m - p(a)| < \frac{\epsilon}{2}$.

We have:

$$\begin{aligned} \rho_a^m - \theta_a^m &= \frac{\sum_{i=1}^m \#_a(\beta_{(n,i)})}{L_m} - \frac{\sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)})}{L_m - \ell_m} \\ &= \left(\frac{\sum_{i \in I_m} \#_a(\beta_{(n,i)})}{L_m} + \frac{\sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)})}{L_m} \right) - \frac{\sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)})}{L_m - \ell_m} \\ &\stackrel{(\dagger)}{=} \frac{\sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)})}{L_m} - \frac{\sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)})}{L_m - \ell_m} + \frac{\sum_{i \in I_m} \#_a(\beta_{(n,i)})}{L_m} \\ &\leq \frac{\sum_{i \in I_m} \#_a(\beta_{(n,i)})}{L_m} \leq \frac{\sum_{i \in I_m} |\beta_{(n,i)}|}{L_m} = \frac{\ell_m}{L_m} \end{aligned} \quad (5.12)$$

where the penultimate inequalities in the last line above follows because $L_m \geq L_m - \ell_m$ implies $\frac{\sum_{i \in \{1, \dots, m\} \setminus I} \#_a(\beta_{(n,i)})}{L} - \frac{\sum_{i \in \{1, \dots, m\} \setminus I} \#_a(\beta_{(n,i)})}{L - \ell} \leq 0$, and the final inequality follows because $\sum_{i \in I_m} \#_a(\beta_{(n,i)}) \leq \sum_{i \in I} |\beta_{(n,i)}| = \ell_m$.

By basic algebra, we have:

$$\frac{\sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)})}{L_m} - \frac{\sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)})}{L_m - \ell_m} = \frac{-\ell_m \sum_{i \in \{1, \dots, m\} \setminus I} \#_a(\beta_{(n,i)})}{L_m(L_m - \ell_m)}$$

and as

$$\sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)}) \leq \sum_{i \in \{1, \dots, m\} \setminus I_m} |\beta_{(n,i)}| = L_m - \ell_m$$

we conclude that:

$$\frac{-\ell_m \sum_{i \in \{1, \dots, m\} \setminus I_m} \#_a(\beta_{(n,i)})}{L_m(L_m - \ell_m)} \geq -\frac{\ell_m}{L_m}$$

and thus by (\dagger) above that:

$$\rho_a^m - \theta_a^m - \frac{\sum_{i \in I_m} \#_a(\beta_{(n,i)})}{L_m} \geq -\frac{\ell_m}{L_m}$$

whence $-\ell_m/L_m \leq \rho_a^m - \theta_a^m$, which combined with (5.12) yields $|\rho_a^m - \theta_a^m| \leq \ell_m/L_m$.

By 5.9 pick a b such that for all $\epsilon > 0$, we have $\lim_{n \rightarrow \infty} \mu_p(D_n^p(b, \epsilon)) = 1$. Choose $\delta > 0$ with $\delta < \frac{b\epsilon}{8}$, and pick $n \in \mathbb{N}$ such that $\mu_p(D_n^p(b, \epsilon)) > 1 - \delta$. Now, pick $\gamma < \frac{b\epsilon}{8}$. Because α is μ_p -block-distributed, there exists $M \in \mathbb{N}$ such that for all $k \geq M$ and all $B \subseteq \Sigma^n$, the prefix $\alpha|_{\leq kn}$ satisfies:

$$\left| \frac{|\{i \leq k : \alpha_{(n,i)} \in B\}|}{k} - \mu_p(B) \right| < \gamma$$

In the particular case $B = D_n^p(b, \epsilon/2)$, we thus have:

$$\left| \frac{|\{i \leq k : \alpha_{(n,i)} \in D_n^p(b, \frac{\epsilon}{2})\}|}{k} - \mu_p \left(D_n^p \left(b, \frac{\epsilon}{2} \right) \right) \right| < \gamma$$

and thus

$$1 - \delta - \frac{|\{i \leq k : \alpha_{(n,i)} \in D_n^p(b, \frac{\epsilon}{2})\}|}{k} \leq \mu_p \left(D_n^p \left(b, \frac{\epsilon}{2} \right) \right) - \frac{|\{i \leq k : \alpha_{(n,i)} \in D_n^p(b, \frac{\epsilon}{2})\}|}{k} < \gamma$$

whence we conclude:

$$\left| \left\{ i \leq k : \alpha_{(n,i)} \in D_n^p \left(b, \frac{\epsilon}{2} \right) \right\} \right| > k(1 - \delta - \gamma) \quad (5.13)$$

By definition of $D_n^p(b, \frac{\epsilon}{2})$, every $\alpha_{(n,i)} \in D_n^p(b, \frac{\epsilon}{2})$ satisfies $|A[\alpha_{(n,i)}]| > bn$, and we thus have:

$$L_m = \sum_{i=1}^m |y_{(n,i)}| = \sum_{i=1}^m |A[\alpha_{(n,i)}]| \geq \left| \left\{ i \leq m : \alpha_{(n,i)} \in D_n^p \left(b, \frac{\epsilon}{2} \right) \right\} \right| bn > m(1 - \delta - \gamma)bn \quad (5.14)$$

Furthermore, by the definition of I_m and (5.13):

$$\begin{aligned} |I_m| &= \left| \left\{ i \leq m : \alpha_{(n,i)} \notin D_n^p \left(b, \frac{\epsilon}{2} \right) \right\} \right| = m - \left| \left\{ i \leq m : \alpha_{(n,i)} \in D_n^p \left(b, \frac{\epsilon}{2} \right) \right\} \right| \\ &< m - m(1 - \delta - \gamma) = m(\delta + \gamma) \end{aligned}$$

But then,

$$\ell_m = \sum_{i \in I_m} |y_{(i,n)}| \leq |I_m|n < mn(\delta + \gamma) \quad (5.15)$$

and thus by 5.14 and 5.15:

$$\frac{\ell_m}{L_m} < \frac{mn(\delta + \gamma)}{m(1 - \delta - \gamma)bn} = \frac{\delta + \gamma}{b(1 - \delta - \gamma)} < \frac{\frac{b\epsilon}{8} + \frac{b\epsilon}{8}}{b \left(1 - \frac{b\epsilon}{8} - \frac{b\epsilon}{8} \right)} < \frac{\frac{\epsilon}{4}}{1 - \frac{1}{4}} < \frac{\epsilon}{2}$$

where we have used that $b\epsilon < 1$ in the penultimate inequality.

We now finally have

$$|\rho_a - p(a)| \leq |\rho_a^m - \theta_a^m| + |\theta_a - p(a)| < \frac{\ell_m}{L_m} + \frac{\epsilon}{2} < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$$

concluding the proof. \square

Lemma 5.11. *Let Σ be an alphabet, p a positive Bernoulli distribution on Σ , let $\alpha \in \Sigma^\omega$ be μ_p -distributed, and let A be a strongly connected DFA over Σ . Then, $A[\alpha]$ is μ_p -distributed.*

Proof. By Lemma 5.5 it suffices to show for every $a \in \Sigma$ and every strongly connected A that the limiting frequency of a in $A[\alpha]$ exists and equals $p(a)$. As α is μ_p -distributed, it follows from Proposition 5.3 that it is μ_p -block-distributed, and the result then immediately follows by Lemma 5.10. \square

6. AN APPLICATION IN SYMBOLIC DYNAMICS: CHARACTERIZING MEASURES WHERE
GENERICITY IS PRESERVED BY DFAS

We now show an application of the main result to the area of symbolic dynamical systems. The following section recalls basic facts about symbolic dynamical systems, including establishing the correspondence between probability maps on Σ^* and probability measures on full shifts.

6.1. Shift spaces and genericity. We briefly introduce basic notions; full accounts can be found in standard textbooks, e.g. [35].

Definition 6.1. Let Σ be a non-empty alphabet. The (one-sided) *shift* $s : \Sigma^\omega \rightarrow \Sigma^\omega$ is the map defined by $s(a_1 a_2 a_3 \dots) = a_2 a_3 \dots$. A *shift space* is a pair (X, s) where $X \subseteq \Sigma^\omega$ is a closed (in the product topology on Σ^ω when Σ is endowed with the discrete topology) subset such that $s(X) = X^7$, and s is the restriction of the shift to X .

As usual, we consider the σ -algebra \mathcal{C} on Σ^ω having the set of cylinders $\{[w] : w \in \Sigma^*\}$ as basis. All measures μ in the remainder of the paper are understood to be measures on $(\Sigma^\omega, \mathcal{C})$.

The standard example of probability measures on shift spaces is the set of Bernoulli measures [60]:

Definition 6.2. A probability measure on the shift space (Σ^ω, s) is a probability measure on Σ^ω with the σ -algebra generated by the cylinder sets $\{[v] : v \in \Sigma^*\}$. A probability measure $\bar{\mu}$ on the full shift is a *Bernoulli measure* if there is a probability distribution $p : \Sigma \rightarrow [0, 1]$ such that the measure of each cylinder satisfies $\bar{\mu}([a_1 \dots a_n]) = \prod_{i=1}^n p(a_i)$. In this case, we say that $\bar{\mu}$ is *induced* by p .

Definition 6.3. Let (X, s) be a shift space. A probability measure $\bar{\mu}$ on X is said to be *shift invariant* if $\bar{\mu}(S^{-1}(A)) = \bar{\mu}(A)$ for all $A \subseteq X$. A finite word $w \in \Sigma^k$ is said to be *admissible* for $\bar{\mu}$ if $\bar{\mu}([w]) > 0$.

A right-infinite sequence $\alpha \in \Sigma^\omega$ is said to be *generic* for $\bar{\mu}$ if, for all words w admissible for $\bar{\mu}$, we have:

$$\lim_{n \rightarrow \infty} \frac{\#_w(\alpha|_{\leq n})}{n} = \bar{\mu}([w])$$

That is, w occurs in α with limiting frequency $\bar{\mu}([w])$.

The study of probability measures on the full shift is cryptomorphic to the study of invariant probability maps; this folklore result is contained in the following two propositions (proofs can be found in Appendix A).

Proposition 6.4. *Every invariant probability map $\mu : \Sigma^* \rightarrow [0, 1]$ induces a shift-invariant probability measure $\bar{\mu} : \Sigma^\omega \rightarrow [0, 1]$ by setting $\bar{\mu}([w]) = \mu(w)$. Conversely, every probability measure $\nu : \Sigma^\omega \rightarrow [0, 1]$ induces a probability map $\underline{\nu} : \Sigma^* \rightarrow [0, 1]$ by defining $\underline{\nu}(w) = \nu([w])$; if ν is shift-invariant, then $\underline{\nu}$ is invariant. Furthermore, $\mu = \underline{\bar{\mu}}$, and $\nu = \bar{\underline{\nu}}$.*

Proposition 6.5. *Let $\mu : \Sigma^* \rightarrow [0, 1]$ be a probability map. The following are equivalent:*

- (1) *There exists a μ -distributed $\alpha \in \Sigma^\omega$.*
- (2) *μ is invariant.*
- (3) *There exists a shift-invariant probability measure ν on Σ^ω such that $\bar{\mu} = \nu$.*

⁷For one-sided shifts, some authors require only $s(X) \subseteq X$; we shall not do so here.

Conversely, let ν be a probability measure on Σ . The following are equivalent:

- (1) There exists $\alpha \in \Sigma^\omega$ that is generic for ν .
- (2) ν is shift-invariant.
- (3) There exists an invariant probability map $\mu : \Sigma^* \rightarrow [0, 1]$ such that $\underline{\nu} = \mu$.

It follows that the shift-invariant probability measures ν on the full shift such that genericity is preserved by finite-state selection, are exactly the Bernoulli measures:

Theorem 6.6. *Let Σ be a non-empty alphabet, and let ν be a shift-invariant measure on the full shift (Σ^ω, s) such that there exists at least one $\alpha \in \Sigma^\omega$ generic for ν . Then, every finite-state selector preserves genericity iff ν is a Bernoulli measure such that all words in Σ^* are admissible.*

Proof. Observe that for a Bernoulli measure $\bar{\mu}$ on the full shift on Σ , all words are admissible iff $\bar{\mu}(a) > 0$ for all $a \in \Sigma$. The Theorem now follows from Theorem 3.1 and Proposition 6.5. \square

7. FUTURE WORK

The most obvious extension of our main results is to attempt to relax the requirement that selection is done by a DFA by using methods similar to Kamae and Weiss [31], and Kamae and Wang [69] where reasoning using a combination of density arguments and relaxed finiteness conditions on the syntactic monoid of the strategy (using our terminology) have been used for normal sequences over binary alphabets. We conjecture that some of these techniques can be adapted to positive Bernoulli distributions on arbitrary finite alphabets.

A different possible thrust is to consider generalizations of Agafonov's Theorem on domains different from infinite sequence over alphabets. However, some results in the – sparse – literature on selection from normal sequence-like objects in other contexts are negative; for example normality is not preserved by arithmetic progressions (so, probably not by finite-state selectors in any reasonable sense) for continued fraction expansions [27]. On the other hand, very recent work by Bergelson et al. has successfully adapted the classical techniques of Kamae and Weiss [31] to show that certain Følner sequences preserve (the appropriate analogue of) normality in cancellative amenable semigroups [7].

REFERENCES

- [1] V. N. Agafonov. Normal sequences and finite automata. *Sov. Math., Dokl.*, 9:324–325, 1968. Originally published in Russian (vol. 179:2, p. 255-266).
- [2] D. Airey and B. Mance. Normality preserving operations for Cantor series expansions and associated fractals, i. *Illinois J. Math.*, 59(3):531–543, 2015.
- [3] C. Aistleitner, V. Becher, A.-M. Scheerer, and T. Slaman. On the construction of absolutely normal numbers. *Acta Arithmetica*, 180:333–346, 2017.
- [4] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *Journal of Computer and System Sciences*, 58(1):137 – 147, 1999.
- [5] V. Becher, O. Carton, and P. A. Heiber. Normality and automata. *Journal of Computer and System Sciences*, 81(8):1592 – 1613, 2015.
- [6] V. Becher and P. A. Heiber. Normal numbers and finite automata. *Theoretical Computer Science*, 477:109–116, 2013.
- [7] V. Bergelson, T. Downarowicz, and J. Vandehey. Deterministic functions on amenable semigroups and a generalization of the kamae-weiss theorem on normality preservation, 2020.
- [8] G. Berry and G. Gonthier. The estereel synchronous programming language: design, semantics, implementation. *Science of Computer Programming*, 19(2):87 – 152, 1992.
- [9] F. Blanchard. Non literal transducers and some problems of normality. *Journal de Théorie des Nombres de Bordeaux*, 5(2):303–321, 1993.
- [10] F. Blanchard, J. M. Dumont, and A. Thomas. Generic sequences, transducers and multiplication of normal numbers. *Israel Journal of Mathematics*, 80:257–287, 1992.
- [11] D. Bontemps, S. Boucheron, and E. Gassiat. About adaptive coding on countable alphabets. *IEEE Transactions on Information Theory*, 60(2):808–821, 2014.
- [12] E. Borel. Les probabilités dénombrables et leurs applications arithmétiques. *Rend. Circ. Matem. Palermo*, 27:247–271, 1909.
- [13] S. Boucheron, A. Garivier, and E. Gassiat. Coding on countably infinite alphabets. *IEEE Trans. Inf. Theory*, 55(1):358–373, 2009.
- [14] V. Braverman and R. Ostrovsky. Generalizing the layering method of Indyk and Woodruff: Recursive sketches for frequency-based vectors on streams. In P. Raghavendra, S. Raskhodnikova, K. Jansen, and J. D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 58–70, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [15] A. Broglio and P. Liardet. Predictions with automata. symbolic dynamics and its applications. *Contemporary Mathematics*, 135:111–124, 1992. Also appeared in Proceedings of the AMS Conference in honor of R. L. Adler. New Haven CT - USA 1991.
- [16] O. Carton. A direct proof of Agafonov's theorem and an extension to shifts of finite type. *Preprint*, 2020.
- [17] O. Carton and J. Vandehey. Preservation of normality by non-oblivious group selection. *Theory of Computing Systems*, 2020.
- [18] P. Caspi, D. Pilaud, N. Halbwachs, and J. Plaice. Lustre: A declarative language for programming synchronous systems. In *Conference Record of the Fourteenth Annual ACM Symposium on Principles of Programming Languages, Munich, Germany, January 21-23, 1987*, pages 178–188, 1987.
- [19] J. W. S. Cassels. On a paper of Niven and Zuckerman. *Pacific J. Math.*, 2(4):555–557, 1952.
- [20] D. G. Champernowne. The construction of decimals normal in the scale of ten. *Journal of the London Mathematical Society*, s1-8(4):254–260, 1933.
- [21] A. Church. On the concept of a random sequence. *Bulletin of the American Mathematical Society*, 46(2):130–135, 1940.
- [22] A. H. Copeland. Admissible numbers in the theory of probability. *American Journal of Mathematics*, 50(4):535–552, 1928.
- [23] A. H. Copeland. Point set theory applied to the random selection of the digits of an admissible number. *American Journal of Mathematics*, 58(1):181–192, 1936.
- [24] A. H. Copeland and P. Erdős. Note on normal numbers. *Bull. Amer. Math. Soc.*, 52(10):857–860, 10 1946.
- [25] H. Davenport and P. Erdős. Note on normal decimals. *Canadian J. Math*, pages 58–63, 1952.
- [26] C. Dwork. Differential privacy in new settings. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 174–183, 2010.

- [27] B. Heersink and J. Vandehey. Continued fraction normality is not preserved along arithmetic progressions. *Archiv der Mathematik*, 106, 09 2015.
- [28] M. Holzer, M. Kutrib, and A. Malcher. Multi-Head Finite Automata: Characterizations, Concepts and Open Problems. In *CSP*, pages 93–107, 2008.
- [29] M. Hosseini and N. Santhanam. On redundancy of memoryless sources over countable alphabets. In *2014 International Symposium on Information Theory and its Applications*, pages 299–303, 2014.
- [30] P. Indyk and D. P. Woodruff. Optimal approximations of the frequency moments of data streams. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 202–208, 2005.
- [31] T. Kamae and B. Weiss. Normal numbers and selection rules. *Israel Journal of Mathematics*, pages 101–110, 1975.
- [32] E. Kamke. Über neuere begründungen der Wahrscheinlichkeitsrechnung. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 42:14–27, 1933.
- [33] G. Kellaris, S. Papadopoulos, X. Xiao, and D. Papadias. Differentially private event sequences over infinite streams. *Proc. VLDB Endow.*, 7(12):1155–1166, 2014.
- [34] M. Levin. Absolutely normal numbers. *Moscow Univ. Math. Bull.*, 34(1):32–39, 1979.
- [35] D. Lind and B. Marcus. *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, 1995.
- [36] M. Madritsch. *Normal Numbers and Symbolic Dynamics*, pages 271–329. Springer International Publishing, Cham, 2018.
- [37] M. Madritsch and B. Mance. Construction of μ -normal sequences. *Monatshefte für Mathematik*, 179:259–280, 2016.
- [38] M. Madritsch, A.-M. Scheerer, and R. Tichy. Computable absolutely Pisot normal numbers. *Acta Arithmetica*, 184:7–29, 2018.
- [39] B. Mance. Cantor series constructions of sets of normal numbers. *Acta Arithmetica*, 156:223–245, 2012.
- [40] W. Merkle and J. Reimann. Selection functions that do not preserve normality. *Theory Comput. Syst.*, 39(5):685–697, 2006.
- [41] S. Muthukrishnan. *Data Streams: Algorithms and Applications*. Now Publishers Inc., 2005.
- [42] Y. Nakai and I. Shiokawa. Discrepancy estimates for a class of normal numbers. *Acta Arithmetica*, 62(3):271–284, 1992.
- [43] I. Niven and H. S. Zuckerman. On the definition of normal numbers. *Pacific J. Math.*, 1(1):103–109, 1951.
- [44] M. G. O’Connor. An unpredictability approach to finite-state randomness. *Journal of Computer and System Sciences*, 37(3):324 – 336, 1988.
- [45] В. Н. Агафонов. Нормальные последовательности и конечные автоматы. *Докл. АН СССР*, 179(2):255–256, 1968.
- [46] В. Н. Агафонов. *Нормальные последовательности и конечные автоматы*, volume 20, pages 123–129. Наука, Академии наук СССР, 1968.
- [47] А. Г. Постников. Арифметическое моделирование случайных процессов. *Тр. МИАН СССР*, 57:3–84, 1960.
- [48] А. Г. Постников and И. И. Пятецкий. Нормальные по Бернулли последовательности знаков. *Изв. АН СССР. Сер. матем.*, 21(4):501–514, 1957.
- [49] А. Г. Постников and И. И. Пятецкий. Нормальные по Бернулли последовательности знаков. *Изв. АН СССР. Сер. матем.*, 21:501–514, 1957.
- [50] Л. П. Постникова. О связи понятий коллектива Мизеса–Черча и нормальной по Бернулли последовательности знаков. *Теория вероятн. и ее примен.*, 6(2):232–234, 1961.
- [51] S. Peyton-Jones. *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, 2003.
- [52] P. Pollack and J. Vandehey. Some normal numbers generated by arithmetic functions. *Canadian Mathematical Bulletin*, 58, 09 2013.
- [53] L. Postnikova. On the connection between the concepts of collectives of Mises–Church and normal Bernoulli sequences of symbols. *Theory of Probability & Its Applications*, 6(2):211–213, 1961. translation of [50] by Eizo Nishiura.
- [54] H. Reichenbach. Axiomatik der wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 34(1):568–619, 1932.

- [55] H. Reichenbach. Les fondements logiques du calcul des probabilités. In *Annales de l'institut Henri Poincaré*, volume 7, pages 267–348, 1937.
- [56] A.-M. Scheerer. Computable absolutely normal numbers and discrepancies. *Mathematics of Computation*, 86, 11 2015.
- [57] C. Schnorr and H. Stimm. Endliche Automaten und Zufallsfolgen. *Acta Informatica*, 1:345–359, 1972.
- [58] R. Serfozo. *Basics of Applied Stochastic Processes*. Probability and Its Applications. Springer-Verlag, 2009.
- [59] A. Shen. Automatic Kolmogorov complexity and normality revisited. In R. Klasing and M. Zeitoun, editors, *Fundamentals of Computation Theory - 21st International Symposium, FCT 2017, Bordeaux, France, September 11-13, 2017, Proceedings*, volume 10472 of *Lecture Notes in Computer Science*, pages 418–430. Springer, 2017.
- [60] P. Shields. *The Theory of Bernoulli Shifts*. Univ. Chicago Press, 1973.
- [61] W. Sierpinski. Démonstration élémentaire du théorème de M. Borel sur les nombres absolument normaux et détermination effective d'une tel nombre. *Bulletin de la Société Mathématique de France*, 45:125–132, 1917.
- [62] J. F. Silva and P. Piantanida. Almost lossless variable-length source coding on countably infinite alphabets. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pages 1–5, 2016.
- [63] R. Stephens. A survey of stream processing. *Acta Informatica*, 34(7):491–541, 1997.
- [64] E. Tornier. Wahrscheinlichkeitsrechnung und Zahlentheorie. erste Mitteilung. *Journal für die reine und angewandte Mathematik*, 1929(160):177–198, 1929.
- [65] J. Vandehey. The normality of digits in almost constant additive functions. *Monatshefte für Mathematik*, 171, 06 2012.
- [66] J. Vandehey. New normality constructions for continued fraction expansions. *Journal of Number Theory*, 166:424 – 451, 2016.
- [67] J. Vandehey. Uncanny subsequence selections that generate normal numbers. *Uniform Distribution Theory*, 12:65–75, 2017.
- [68] R. Von Mises. Grundlagen der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 5(191):52–99, 1919.
- [69] X. Wang and T. Kamae. Selection rules preserving normality. *Israel Journal of Mathematics*, 232:427–442, 2019.

APPENDIX A. AUXILIARY PROOFS AND DEFINITIONS

A.1. Automata and selectors. The following is a proof of the extension of Lemma 2.6 of [57]. The proof follows the original in most details.

Definition A.1. Let $G = (V, E)$ be a directed multigraph, and denote by $\sim \subseteq V \times V$ the equivalence relation such that $v \sim w$ iff there is a directed path from v to w and a directed path from w to v . For every $v \in V$, denote by $[v]_{\sim}$ the equivalence class containing v . Define the partial order $<$ on V/\sim by $\mathcal{V} < \mathcal{W}$ iff there are $v \in \mathcal{V}$ and $w \in \mathcal{W}$ such that there is a directed path from w to v' .

If G has a finite number of nodes, $<$ is clearly well-founded. As $<$ is clearly also transitive, every $\mathcal{W} \in V/\sim$ satisfies $\mathcal{W} > \mathcal{V}$ for some $<$ -minimal $\mathcal{V} \in V/\sim$.

Also observe that every $<$ -minimal \mathcal{V} is a recurrent strongly connected component, because (i) it is strongly connected by definition, and $<$ -minimality implies that no directed path from any node in \mathcal{V} can reach a node in a strongly connected component distinct from \mathcal{V} .

Lemma A.2. *Let $S = (Q, \delta, q_s, Q_F)$ be a finite automaton over a (possibly infinite) alphabet Σ . Then there is a word $w \in \Sigma^*$ such that, for all states $q \in Q$, $\delta^*(q, w)$ is a state in a $<$ -minimal element of Q/\sim .*

Proof. Write $Q = \{q_1, \dots, q_m\}$. We prove by induction on $i \leq m$ that there is a word $w_i \in \Sigma^*$ such that for all $j \leq i$, $\delta^*(s_j, w_i)$ is a state in a $<$ -minimal element of Q/\sim .

$i = 1$:: Let \mathcal{V} be a $<$ -minimal element of Q/\sim such that $[q_1]_{\sim} > \mathcal{V}$. Choose $q \in Q$ such that $[q]_{\sim} = \mathcal{V}$. Then there is a directed path from s_1 to q . Let w_1 be the word along that path, and observe that $\delta^*(q_1, w_1) = q$.

$i > 1$:: Let \mathcal{V} be a $<$ -minimal element of Q/\sim such that $\delta^*(q_{i+1}, w_i) \in \mathcal{V}$, and let $q \in \mathcal{V}$, whence there is a directed path from $\delta^*(q_{i+1}, w_i)$ to q . Let $w' \in \Sigma^*$ be the word along that path, whence $\delta^*(\delta^*(q_{i+1}, w_i), w') = q$. Define $w_{i+1} = w_i \cdot w'$, and observe that $\delta^*(q_{i+1}, w_{i+1}) = q$.

For $j \leq i$, we claim that $\delta^*(q_j, w_{i+1})$ is a state in a $<$ -minimal element of Q/\sim . For, by the Induction Hypothesis, $\delta(q_j, w_i)$ is in a $<$ -minimal element \mathcal{V}_j of Q/\sim , and as $<$ -minimal element are recurrent strongly connected components, no directed path from $\delta(q_j, w_i)$ can end in a state outside \mathcal{V}_j .

□

Proof of Lemma 2.15. Any $<$ -minimal element of Q/\sim is recurrent. By Lemma A.2, there is a word w such that from any state $q \in Q$, $\delta^*(q, w)$ is a state in a recurrent strongly connected component of the automaton. As $p(a) > 0$ for all $a \in \Sigma$, $\mu_p(w) > 0$, and as α is p -distributed, w thus occurs (infinitely often) in α . After the first occurrence of w , the run of A on α has entered a strongly recurrent connected component. □

A.2. μ -distribution.

Proof of Proposition 5.3. We use exactly the same arguments as in the proof by Niven and Zuckerman [43], but using the notation of the present paper. Almost the entirety of the proof in [43] is devoted to counting arguments on finite prefixes of α , and involves neither the size of the alphabet Σ , nor the particular distribution on it; indeed any consideration of

those matters is isolated to a few observations in the beginning of the proof that are then used repeatedly when taking limits later on. We have clearly indicated those observations below, but give the entirety of the proof in the interest of completeness.

Let $w = w_1 \cdots w_v \in \Sigma^v$ be arbitrary. We introduce the following notation:

- For any $t \geq 0$, $w\Sigma^t w$ is the set $\{wuw : u \in \Sigma^t\}$.
- $\#_w^i(n)$ is the number of times that w occurs in $\alpha|_{\leq n}$ at a position congruent to $i \pmod{n}$.
- $\#_w^{i,j}(n) = \#_w^i(n) - \#_w^j(n)$.
- $g : \mathbb{N} \rightarrow \mathbb{N}$ is the function defined by: $g(n) = \sum_{i=1}^{n-1} \#_w^i(n)$.
- $\theta_t(n)$ is the number of occurrences of any element from $w\Sigma^t w$ in $\alpha|_{\leq n}$.
- w' is shorthand for any string of length between $v+1$ and $2v-1$ whose first v digits are w and whose last digits are w , i.e. an “overlap of w with itself”. Such a string does not necessarily exist.

We now treat the part of the proof depending on the cardinality of Σ and μ_p -distributedness (as opposed to finiteness of Σ and equidistribution).

As α is μ_p -distributed, we have

$$\lim_{n \rightarrow \infty} \frac{g(n)}{n} = \mu_p(w) \quad (\text{A.1})$$

and for each fixed $t \geq 0$, we also have:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{\theta_t(n)}{n} &= \lim_{n \rightarrow \infty} \frac{\sum_{a_1 \cdots a_t \in \Sigma^t} \#_{wa_1 \cdots a_t w}(\alpha|_{\leq n})}{n} \\ &= \lim_{n \rightarrow \infty} \sum_{a_1 \cdots a_t \in \Sigma^t} \frac{\#_{wa_1 \cdots a_t w}(\alpha|_{\leq n})}{n} \\ &= \sum_{a_1 \cdots a_t \in \Sigma^t} \lim_{n \rightarrow \infty} \frac{\#_{wa_1 \cdots a_t w}(\alpha|_{\leq n})}{n} \quad (\text{By the Dominated Convergence Theorem}) \\ &= \sum_{a_1 \cdots a_t \in \Sigma^t} \mu_p(wa_1 \cdots a_t w) \\ &= \sum_{a_1 \cdots a_t \in \Sigma^t} \mu_p(w)^2 \prod_{i=1}^t p(a_i) \quad (\text{As } \mu_p \text{ is induced by a Bernoulli distribution}) \\ &= \mu_p(w)^2 \sum_{a_1 \cdots a_t \in \Sigma^t} \prod_{i=1}^t p(a_i) \\ &= \mu_p(w)^2 \prod_{i=1}^t \sum_{a \in \Sigma} p(a) \quad (\text{By monotone convergence}) \\ &= \mu_p(w)^2 \quad (\text{As } 1 = \sum_{a \in \Sigma} p(a)) \end{aligned} \quad (\text{A.2})$$

We shall prove that:

$$\lim_{n \rightarrow \infty} \frac{\#_w^{i,j}(n)}{n} = 0 \quad (\text{A.3})$$

By A.1 and A.3, it follows for any i with $0 \leq i < v$ that:

$$\lim_{n \rightarrow \infty} \frac{\#_w^i(n)}{n} = \frac{\mu_p(w)}{v}$$

and as v and $w \in \Sigma^v$ were arbitrary, that α is μ_p -block-distributed.

The remainder of the proof is devoted to prove A.3 and is only concerned with counting arguments on finite prefixes of α . All arguments from hereon are, modulo notation and use of A.2, completely identical to the proof in [43].

Let $s \geq 0$ be an integer. Observe that $\#_w^i(n+s) - \#_w^i(n)$ is the number of occurrences of w that (1) are in $\alpha|_{\leq n+s}$ at a position congruent to $i \pmod v$, but (2) are not entirely contained in $\alpha|_{\leq n}$. Thus,

$$\sum_{\substack{i < j \\ i \in \{0, \dots, v-2\} \\ j \in \{1, \dots, v-1\}}} (\#_w^i(n+s) - \#_w^i(n)) (\#_w^j(n+s) - \#_w^j(n))$$

is the number of words on the form w' or wuw (for $0 \leq |u| \leq s-v-1$ and $\neg(|u| \equiv 0 \pmod v)$) that occur in $\alpha|_{n+s}$, but such that the initial $|w| = v$ symbols are not entirely contained in $\alpha|_{\leq n}$.

For $n > s$, we define:

$$\sigma = \sum_{m=0}^{s-s} \sum_{\substack{i < j \\ i \in \{0, \dots, v-2\} \\ j \in \{1, \dots, v-1\}}} (\#_w^i(m+s) - \#_w^i(m)) (\#_w^j(m+s) - \#_w^j(m)) \quad (\text{A.4})$$

Consider $\alpha|_{\leq n}$ and any single occurrence of an element wuw with $t = |u| \leq s-v-1$ in $\alpha|_{\leq n}$. The occurrence of wuw is counted a number of times in σ , and it does not occur too close to either end of $\alpha|_{\leq n}$, it is counted at most $s-t-v$ times in σ . If the occurrence of wuw is preceded by at least $s-t-2v$ symbols and is followed by at least $s-t-v-1$ symbols, wuw is counted exactly $s-t-v$ times. Thus, we have:

$$\sigma \geq \sum_{\substack{t=0 \\ \neg(t \equiv 0 \pmod v)}}^{s-v-1} (s-t-v)(\theta_t(n-s) - \theta_t(s)) \quad (\text{A.5})$$

Thus, by A.2 and A.5, we obtain, for *fixed* s :

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{\sigma}{n} &\geq \lim_{n \rightarrow \infty} \left(\sum_{\substack{t=0 \\ \neg(t \equiv 0 \pmod v)}}^{s-v-1} (s-t-v)(\theta_t(n-s) - \theta_t(s)) \right) \\ &= \sum_{\substack{t=0 \\ \neg(t \equiv 0 \pmod v)}}^{s-v-1} \lim_{n \rightarrow \infty} \left(\frac{\sigma}{n} (s-t-v)(\theta_t(n-s) - \theta_t(s)) \right) \\ &= \sum_{\substack{t=0 \\ \neg(t \equiv 0 \pmod v)}}^{s-v-1} (s-t-v) \mu_p(w)^2 \quad (\text{By A.2}) \quad (\text{A.6}) \end{aligned}$$

Choose s such that $s \equiv 0 \pmod v$. Then, A.6 becomes:

$$\lim_{n \rightarrow \infty} \frac{\sigma}{n} \geq \frac{(v-1)(s-v)^2}{2v} \mu_p(w)^2 \quad (\text{A.7})$$

Similarly, we count the number of occurrences of words on the form wuw where $t = |u| \equiv 0 \pmod{v}$ and proceed as above. This yields:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{m=0}^{n-s} \sum_{i=0}^{v-1} \frac{(\#_w^i(m+s) - \#_w^i(m))(\#_w^i(m+s) - \#_w^i(m) - 1)}{2} \quad (\text{A.8})$$

$$= \sum_{\substack{t=0 \\ -(t \equiv 0 \pmod{v})}} (s-t-v)\mu_p(w)^2 = \frac{s(s-v)}{2v}\mu_p(w)^2 \quad (\text{A.9})$$

We have, for fixed s such that $s \equiv 0 \pmod{v}$, that:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{2n} \sum_{m=0}^{n-s} \sum_{i=0}^{v-1} (\#_w^i(m+s) - \#_w^i(m)) &= \lim_{n \rightarrow \infty} \frac{1}{2n} \sum_{m=0}^{n-s} (g(m+s) - g(m)) \\ &= \lim_{n \rightarrow \infty} \left(\frac{1}{2n} \sum_{m=n-s+1}^n g(m) - \frac{1}{2n} \sum_{m=0}^{s-1} g(m) \right) \\ &= \lim_{n \rightarrow \infty} \frac{1}{2n} \sum_{m=n-s+1}^n g(m) \\ &= \frac{s\mu_p(w)}{2} \quad (\text{By A.1}) \end{aligned} \quad (\text{A.10})$$

Thus, by A.7 and A.10, we obtain:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{m=0}^{n-s} \sum_{i=0}^{v-1} (\#_w^i(m+s) - \#_w^i(m))^2 = s\mu_p(w) + \frac{s(s-v)}{v}\mu_p(w)^2 \quad (\text{A.11})$$

For fixed s with $s \equiv 0 \pmod{v}$, A.4, A.7, and A.11 now yield:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{m=0}^{n-s} \sum_{\substack{i < j \\ i \in \{0, \dots, v-2\} \\ j \in \{1, \dots, v-1\}}} (\#_w^i(m+s) - \#_w^i(m) - (\#_w^j(m+s) - \#_w^j(m))) \\ \leq (v-1)s\mu_p(w) + (v-q)(s-v)\mu_p(w)^2 \end{aligned} \quad (\text{A.12})$$

Noting that $\sum_{i=1}^n x_i^2 \geq \frac{1}{n}(\sum_{i=1}^n x_i)^2$, we obtain:

$$\begin{aligned}
& \sum_{m=0}^{n-s} (\#_w^i(m+s) - \#_w^i(m) - (\#_w^j(m+s) - \#_w^j(m))) \\
& \geq \frac{1}{n-s+1} \left(\sum_{m=0}^{n-s} \#_w^i(m+s) - \#_w^i(m) - \#_w^j(m+s) + \#_w^j(m) \right)^2 \\
& = \frac{1}{n-s+1} \left(\sum_{m=0}^{n-s} (\#_w^{i,j}(m+s) - \#_w^{i,j}(m)) \right)^2 \\
& = \frac{1}{n-s+1} \left(\sum_{m=0}^{s-1} \#_w^{i,j}(n-m) - \sum_{m=0}^{s-1} \#_w^{i,j}(m) \right)^2 \tag{A.13}
\end{aligned}$$

Now, A.12 and A.13 imply:

$$\limsup_{n \rightarrow \infty} \frac{1}{n(n-s+1)} \sum_{\substack{i < j \\ i \in \{0, \dots, v-2\} \\ j \in \{1, \dots, v-1\}}} \left(\sum_{m=0}^{s-1} \#_w^{i,j}(n-m) - \sum_{m=0}^{s-1} \#_w^{i,j}(m) \right)^2 \tag{A.14}$$

$$\leq (v-1)s\mu_p(w) + (v-1)(s-v)\mu_p(w)^2 \tag{A.15}$$

By the definition of $\#_w^{i,j}$, we have $|\#_w^{i,j}(m)| < m$, and thus, for fixed s , we have:

$$\lim_{n \rightarrow \infty} \frac{1}{n(n-s+1)} \left(\sum_{m=0}^{s-1} \#_w^{i,j}(m) \right)^2 = 0$$

and:

$$\lim_{n \rightarrow \infty} \frac{1}{n(n-s+1)} \sum_{m=0}^{s-1} \#_w^{i,j}(n-m) \sum_{m=0}^{s-1} \#_w^{i,j}(m) = 0$$

which in turn imply, by A.15, that:

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \frac{1}{n(n-s+1)} \sum_{\substack{i < j \\ i \in \{0, \dots, v-2\} \\ j \in \{1, \dots, v-1\}}} \left(\sum_{m=0}^{s-1} \#_w^{i,j}(n-m) \right)^2 \\
& = \limsup_{n \rightarrow \infty} \frac{1}{n(n-s+1)} \sum_{\substack{i < j \\ i \in \{0, \dots, v-2\} \\ j \in \{1, \dots, v-1\}}} \left(s\#_w^{i,j}(n) + \sum_{m=0}^{s-1} (\#_w^{i,j}(n-m) - \#_w^{i,j}(n)) \right)^2 \\
& \leq (v-1)s\mu_p(w) + (v-1)(s-v)\mu_p(w)^2 \tag{A.16}
\end{aligned}$$

But $|\#_w^{i,j}(n-m) - \#_w^{i,j}(n)| < 2m$, so from A.16 we obtain:

$$\limsup_{n \rightarrow \infty} \frac{1}{n(n-s+1)} \sum_{\substack{i < j \\ i \in \{0, \dots, v-2\} \\ j \in \{1, \dots, v-1\}}} s^2 (\#_w^{i,j}(n))^2 \leq (v-1)s\mu_p(w) + (v-1)(s-v)\mu_p(w)^2 \quad (\text{A.17})$$

which in turn implies, for fixed $s \equiv 0 \pmod{v}$, that:

$$\limsup_{n \rightarrow \infty} \frac{(\#_w^{i,j}(n))^2}{n^2} = \limsup_{n \rightarrow \infty} \frac{(\#_w^{i,j}(n))^2}{n(n-s+1)} \leq \frac{(v-1)\mu_p(w)}{s} + \frac{(v-1)(s-v)\mu_p(w)^2}{s^2} \quad (\text{A.18})$$

As the expression on the right-hand side of A.18 can be made arbitrarily small by choosing s large enough, we obtain:

$$\lim_{n \rightarrow \infty} \frac{|\#_w^{i,j}(n)|}{n} = 0$$

and hence:

$$\lim_{n \rightarrow \infty} \frac{\#_w^i(n)}{n} = \lim_{n \rightarrow \infty} \frac{\#_w^j(n)}{n}$$

as desired. \square

A.3. Symbolic dynamical systems.

Proof of Proposition 6.4. The two identities $\mu = \bar{\mu}$, and $\nu = \underline{\nu}$ follow directly from the definitions. If μ is invariant, then for every cylinder $[w]$, we have $\bar{\mu}([w]) = \mu(w) = \sum_{a \in \Sigma} \mu(w \cdot a) = \sum_{a \in \Sigma} \bar{\mu}([w \cdot a])$; from this, and the observation that $1 = \mu(\epsilon) = \bar{\mu}([\epsilon]) = \bar{\mu}(\Sigma^\omega)$, that $\bar{\mu}$ is a probability measure on Σ with the sigma algebra generated by the cylinder sets. In addition, as μ is invariant, we have for any cylinder $[w]'$ that:

$$\bar{\mu}(S^{-1}([w])) = \bar{\mu} \left(\bigcup_{a \in \Sigma} [a \cdot w] \right) = \sum_{a \in \Sigma} \bar{\mu}([a \cdot w]) = \sum_{a \in \Sigma} \mu(a \cdot w) = \mu(w) = \bar{\mu}([w])$$

whence $\bar{\mu}$ is shift-invariant.

Conversely, if ν is a shift-invariant probability measure on Σ^ω , we have for any w that:

$$\underline{\nu}(w) = \nu([w]) = \nu \left(\bigcup_{a \in \Sigma} [w \cdot a] \right) = \sum_{a \in \Sigma} \nu([w \cdot a]) = \sum_{a \in \Sigma} \underline{\nu}(w \cdot a)$$

and

$$\underline{\nu}(w) = \nu([w]) = \nu(S^{-1}([w])) = \nu \left(\bigcup_{a \in \Sigma} [a \cdot w] \right) = \sum_{a \in \Sigma} \nu([a \cdot w]) = \sum_{a \in \Sigma} \underline{\nu}(a \cdot w)$$

showing that $\underline{\nu}$ is invariant. \square

Proof of Proposition 6.5. For the first part, we prove $1 \Rightarrow 2 \Rightarrow 3 \Rightarrow 1$.

If there is a μ -distributed $\alpha \in \Sigma^\omega$, then for any $w \in \Sigma^*$ and any $\epsilon > 0$, for all sufficiently large n we have $\sup_{b \in \Sigma \cup \{\lambda\}} |\#_{b \cdot w}(\alpha|_{\leq n})/n - \mu(w)| < \epsilon$. Observe that every occurrence of a word on the form $a \cdot w$ in α contains an occurrence of w , and hence $\#_w(\alpha|_{\leq n}) \geq \sum_{a \in \Sigma} \#_{a \cdot w}(\alpha|_{\leq n})$. Conversely, for every occurrence of w starting at some

position $i \geq 2$ in α , there is exactly one $a \in \Sigma$ such that the word $a \cdot w$ occurs at position $i - 1$, whence $\#_w(\alpha|_{\leq n}) \leq 1 + \sum_{a \in \Sigma} \#_{a \cdot w}(\alpha|_{\leq n})$, and hence:

$$\begin{aligned}
\left| \mu(w) - \sum_{a \in \Sigma} \mu(a \cdot w) \right| &= \left| \mu(w) - \frac{\#_w(\alpha|_{\leq n})}{n} + \frac{\#_w(\alpha|_{\leq n})}{n} - \sum_{a \in \Sigma} \mu(a \cdot w) \right| \\
&\leq \left| \mu(w) - \frac{\#_w(\alpha|_{\leq n})}{n} \right| + \left| \frac{\#_w(\alpha|_{\leq n})}{n} - \sum_{a \in \Sigma} \mu(a \cdot w) \right| \\
&< \epsilon + \frac{1}{n} + \left| \frac{\sum_{a \in \Sigma} \#_{a \cdot w}(\alpha|_{\leq n})}{n} - \sum_{a \in \Sigma} \mu(a \cdot w) \right| \\
&\leq \epsilon + \frac{1}{n} + \sum_{a \in \Sigma} \left| \frac{\#_{a \cdot w}(\alpha|_{\leq n})}{n} - \mu(a \cdot w) \right| \\
&< \epsilon + \frac{1}{n} + |\Sigma|\epsilon
\end{aligned}$$

and as ϵ was arbitrary, we thus have $\mu(w) = \sum_{a \in \Sigma} \mu(a \cdot w)$. The case for $\mu(w) = \sum_{a \in \Sigma} \mu(w \cdot a)$ is symmetric, *mutatis mutandis*, and hence μ is invariant. If μ is invariant, then by Proposition 6.4, $\bar{\mu}$ is a shift-invariant probability measure on Σ^ω . If ν is a shift-invariant probability measure on Σ^ω such that $\bar{\mu} = \nu$, then by [37, Main Thm. 2.1], there exists $\alpha \in \Sigma^\omega$ generic for $\bar{\mu}$, and thus for any admissible $w \in \Sigma^*$ $\lim_{n \rightarrow \infty} \frac{\#_w(\alpha|_{\leq n})}{n} = \nu(w) = \bar{\mu}([w]) = \mu(w)$. Observe that any inadmissible word $w = a_1 \cdots a_n$ has $\prod_{i=1}^n \mu(a_i) = \mu(w) = 0$, whence $\mu(a_i) = 0$ for some i , and hence $\lim_{n \rightarrow \infty} \frac{\#_w(\alpha|_{\leq n})}{n} \leq \lim_{n \rightarrow \infty} \frac{\#_{a_i}(\alpha|_{\leq n})}{n} = 0$. Hence, α is μ -distributed.

For the second part, we prove $1 \Rightarrow 3 \Rightarrow 2 \Rightarrow 1$. Assume that α is generic for ν . By construction, $\underline{\nu}$ is a probability map such that α is $\underline{\nu}$ -distributed, and by the first part of the proposition, $\underline{\nu}$ is invariant, as desired. If $\underline{\nu}$ is an invariant probability map, then as any measurable A can be written as a disjoint union of cylinder sets, and as we for any cylinder $[w]$ have $S^{-1}([w]) = \cup_{a \in \Sigma} [a \cdot w]$, we obtain

$$\nu(S^{-1}([w])) = \underline{\nu}(\cup_{a \in \Sigma} [a \cdot w]) = \sum_{a \in \Sigma} \underline{\nu}(a \cdot w) = \underline{\nu}(w) = \nu([w])$$

showing that ν is shift-invariant. Finally, if ν is shift-invariant, it follows from [37, Main Thm. 2.1], there exists $\alpha \in \Sigma^\omega$ generic for ν , as desired. \square