Priced Learning

Sanjay Jain *[1], Junqi Ma[2] and Frank Stephan **[3]

[1] School of Computing, National University of Singapore, Singapore 117417.
Email: sanjay@comp.nus.edu.sg
[2] School of Computing, National University of Singapore, Singapore 117417.
Email: ma.junqi@nus.edu.sg
[3] Department of Mathematics and Department of Computer Science, National University of Singapore, Singapore 119076.
Email: fstephan@comp.nus.edu.sg

Abstract. In iterative learning the memory of the learner can only be updated when the hypothesis changes; this results in only finitely many updates of memory during the overall learning history. Priced learning relaxes this constraint on the update of memory by imposing some price on the updates of the memory – depending on the current datum – and requiring that the overall sum of the costs incurred has to be finite. There are priced-learnable classes which are not iteratively learnable. The current work introduces the basic definitions and results for priced learning. This work also introduces various variants of priced learning.

1 Introduction

Learning from positive data in Gold's model [6] has the following basic scenario: Given a class of r.e. languages and an unknown language L from this class, the learner observes an infinite list containing all and only the elements of L. The order and multiplicity of the elements of L in the list may be arbitrary. As the learner is observing the members of the list, it outputs a sequence of hypotheses about what the input language might be. For learning the language, this sequence of hypotheses is required to converge to a grammar for L (for every input list of elements of L as above).

In general, in the above learning process, the learner can memorise all data observed so far and do comprehensive calculations without restrictions to the memory amount and usage. Several approaches have been formalised in order to restrict the amount of memory used. One of them is the strict separation between short-term and long-term memory where, whenever a datum is read, some computations are done in an unlimited short-term memory and then the data for the next round of learning is archived in a size-constrained long-term memory [4, 7]. The other approaches to limit the memory usage do not count bits and bytes, but rather allow only fixed number of elements of the input to be memorized. In such models, if the learner does not update its memory/hypothesis on a receipt of a datum, then at later stages it is not able

* Supported in part by NUS grant numbers R146-000-181-112 and C252-000-087-001.
** Supported in part by NUS grant number R146-000-181-112.

to know whether the corresponding datum was observed or not. In its most restrictive form, this type of learning is called incremental or iterative learning [3, 8, 13]. An iterative learner can memorise data only when it revises the hypothesis and not at any other point of time; thus it can, through the overall learning history, only finitely often revise its long-term memory. This restriction is quite severe and, for example, the class consisting of the set $\{1, 2, \ldots\}$ and all sets of the form $\{0, 1, \ldots, x\}$, where x is a natural number, is not iteratively learnable. To see this, note that if the learner first sees data from the infinite set $\{1, 2, \ldots\}$, then, after some time, the learner will need to converge to a grammar for $\{1, 2, \ldots\}$. Thereafter the learner cannot archive any data, in particular it cannot archive the maximum datum x seen so far. If later the datum 0 shows up, the learner knows that the input set must be of the form $\{0, 1, \ldots, x\}$; however, the learner can no longer recover the exact value of x from its memory.

The idea of priced learning is to relax the severe constraints placed on iterative learning. In priced learning, a price is charged for each update of the memory and it is required that during the learning process, the overall costs incurred is finite. In the case that this price charged is always the same constant c for each memory update, the corresponding notion would be exactly that of iterative learning. However, by allowing the price to depend on the datum x read when the memory is updated, one can give discounts for various types of data and permit, under certain circumstances, to update the memory infinitely often during the learning process. This concept of priced learning, however, still does not permit, in general, to do every possible update of the memory. Priced learning is therefore between the two extremes of iterative learning and the original unrestricted model of learning by Gold where the memory could be updated as often as desired. For the priced learning model $\mathbf{Priced}_f\mathbf{Ex}$ introduced in this paper, the costs incurred are guided by a price function f which is a parameter of the learning criterion. A $\mathbf{Priced}_f\mathbf{Ex}$-learner incurs, at every update of the memory on input datum x, the cost $1/(f(x)+1)$ (where f is a recursive function). We refer the reader to Section 3 for the formal definitions. On reading a datum x, the learner has to decide whether it wants to update the memory so that on one hand, all important information are preserved in the long term memory and, on the other hand, the overall price of these updates does not go to infinity.

The present work investigates when, for different price functions f and g, $\mathbf{Priced}_f\mathbf{Ex}$-learnability implies $\mathbf{Priced}_g\mathbf{Ex}$-learnability. This depends on the existence of sets S such that $\sum_{x \in S} \frac{1}{f(x)+1}$ is finite but $\sum_{x \in S} \frac{1}{g(x)+1}$ is infinite (see Theorem 9 and Theorem 11). Theorem 15 gives a characterisation of classes which are $\mathbf{Priced}_f\mathbf{Ex}$-learnable for some recursive function f: these are exactly the classes which are learnable by a set-driven learner. Here a set-driven learner [9] is a learner whose output conjecture depends only on the set of elements it has observed in the input and not on their order or amount of repetitions.

Further results are based on whether a class can be learnt when the price functions are from a natural class $\mathcal{C}$ of price functions, irrespective of which price function from the class is actually chosen. Here, one may allow the learner to be chosen depending on the price function, uniformly or non-uniformly based on the program for the price function, or may even require it to be the same learner for all the price functions from the class $\mathcal{C}$. Section 6 onwards explore such questions. Let $\mathbf{MF} = \{f : f \text{ is recursive and unbounded and } \forall x\, [f(x) \leq f(x+1)]\}$ and $\mathbf{FF} = \{f : f$ is recursive and $\forall y\, [\mathrm{card}(f^{-1}(y))$ is finite$]\}$. Note that **MF** is a proper subset of **FF**. It is shown in Theorem 17 that there exists a class of languages which can be $\mathbf{Priced}_f$-learnt for every $f \in \mathbf{FF}$, but which cannot be iteratively learnt. On the other hand, Theorem 18 shows advantages of price function being from **MF** compared to it being from **FF**. Results in Section 7 (along with some results in Section 6) deal with the question of when and what kind of uniformity constraints can be placed on learners which learn a class with respect to a price function from a class of price functions such as **MF** or **FF**. Finally, Theorem 27 deals with the case when one can get a single learner for all price functions from an r.e. class $\mathcal{C}$ of price functions, when the class is learnable with respect to every price function in $\mathcal{C}$.

2 Notations and Preliminaries

Any unexplained recursion theoretic notation is from Rogers' book [11]. The symbol $\mathbb{N}$ denotes the set of natural numbers $\{0, 1, 2, \ldots\}$. Subsets of natural numbers are referred to as languages. We let $\emptyset$, $\subseteq$, $\subset$, $\supseteq$ and $\supset$ respectively denote empty set, subset, proper subset, superset and proper superset. Cardinality of a set S is denoted by $\mathrm{card}(S)$. The maximum and minimum of a set S are denoted by $\max(S)$ and $\min(S)$, respectively, where $\max(\emptyset) = 0$ and $\min(\emptyset) = \infty$. $A \oplus B$ denotes the symmetric difference of A and B, that is $A \oplus B = (A - B) \cup (B - A)$. For a natural number n, we say that $A =^n B$ iff $\mathrm{card}(A \oplus B) \leq n$. We say that $A =^* B$ iff $\mathrm{card}(A \oplus B) < \infty$. $\forall^\infty$ and $\exists^\infty$ respectively denote 'for all but finitely many' and 'there exist infinitely many'.

Let $\langle \cdot, \cdot \rangle$ denote a recursive pairing function which is a bijection from $\mathbb{N} \times \mathbb{N}$ to $\mathbb{N}$. Without loss of generality, we assume that pairing function is monotonic in both its arguments; in particular $\langle 0, 0 \rangle = 0$. The pairing function can be extended to pairing of n-tuples by using $\langle x_1, x_2, \ldots, x_n \rangle = \langle x_1, \langle x_2, \ldots, x_n \rangle \rangle$. We define projection functions π_1, π_2 as $\pi_1(\langle x_1, x_2 \rangle) = x_1$ and $\pi_2(\langle x_1, x_2 \rangle) = x_2$. Projection functions for pairs of n-tuples can be defined similarly as $\pi_i^n(\langle x_1, x_2, \ldots, x_n \rangle) = x_i$.

For a set L, let $\mathrm{cyl}(L) = \{\langle x, i \rangle : x \in L, i \in \mathbb{N}\}$.

By φ we denote a fixed *acceptable* programming system for the partial computable functions mapping $\mathbb{N}$ to $\mathbb{N}$ [11]. By φ_i we denote the partial function computed by the i-th program in the φ-system. Thus, i is also called a program/index for φ_i. Let $\mathcal{R}$ denote the class of all total recursive functions. For any partial function η, we let $\eta(x)\downarrow$ denote that $\eta(x)$ is defined and $\eta(x)\uparrow$ denote that $\eta(x)$ is undefined. By Φ we denote an arbitrary fixed Blum complexity measure [2] for the φ-system. Let

$$\varphi_{i,s}(x) = \begin{cases} \varphi_i(x), & \text{if } x < s \text{ and } \Phi_i(x) < s; \\ \uparrow, & \text{otherwise.} \end{cases}$$

Let $W_i = \mathrm{domain}(\varphi_i)$. W_i can be considered as the language accepted by the i-the φ-program φ_i. We also say that i is a grammar/index for W_i. Thus, $W_0, W_1, \ldots$, is an acceptable programming system for recursively enumerable languages. The symbol $\mathcal{E}$ denotes the set of all recursively enumerable (r.e.) languages. The symbol L ranges over $\mathcal{E}$; the symbol $\mathcal{L}$ ranges over subsets of $\mathcal{E}$. By $\overline{L}$, we denote the complement of L, that is $\overline{L} = \mathbb{N} - L$. By $L(x)$ we denote the value of the characteristic function of L at x; that is, $L(x) = 1$ if $x \in L$ and $L(x) = 0$ if $x \notin L$. By $W_{i,s}$ we denote the set $\{x < s : \Phi_i(x) < s\}$.

3 Models of Learning

We now present some concepts from language learning theory. A *text* is a mapping from $\mathbb{N}$ to $\mathbb{N} \cup \{\#\}$. The *content* of a text T, denoted content(T) is the set of natural numbers in the range of T, that is, $\{T(m) : m \in \mathbb{N}\} - \{\#\}$. We say that T is a *text for L* if content$(T) = L$. Intuitively, $\#$ can be considered as pauses in the presentation of data to a learner. We let T range over texts. $T[n]$ denotes the initial segment of a text T with length n, that is, $T(0), T(1), \ldots, T(n-1)$.

Initial segments of texts are called finite *sequences* (or just sequences). We let σ and τ range over finite sequences. We let content(σ) denote the set of natural numbers in the range of σ. Let SEQ denote the set of all finite sequences. We assume some fixed computable ordering of members SEQ (so that we can talk about the least sequence, etc.). We let $|\sigma|$ denote the length of σ. We let Λ denote the empty sequence and $\sigma[n]$ denote the initial segment of σ of length n (where $\sigma[n] = \sigma$, if $n \geq |\sigma|$). Let $\sigma \diamond \tau$ denote the concatenation of two finite sequence σ and τ. Similarly, let $\sigma \diamond T$ denote the concatenation of finite sequence σ and text T. For $x \in \mathbb{N} \cup \{\#\}$, we let $\sigma \diamond x$ denote the concatenation of σ with the finite sequence containing just one element x. Let $\sigma \preceq T$ and $\sigma \preceq \tau$ denote that σ is an initial segment of T and τ respectively.

Definition 1 (Based on Gold [6]).

(a) A *learner* $\mathbf{M}$ is a (possibly partial) recursive mapping from $(\mathbb{N} \cup \{?\}) \times (\mathbb{N} \cup \{\#\})$ to $(\mathbb{N} \cup \{?\}) \times (\mathbb{N} \cup \{?\})$. A learner has initial memory mem_0 and initial hypothesis hyp_0.

(b) Suppose a learner $\mathbf{M}$ with initial memory mem_0 and initial hypothesis hyp_0 is given. Let $mem_0^{\mathbf{M},T} = mem_0$ and $hyp_0^{\mathbf{M},T} = hyp_0$. For $k \geq 0$, let $(mem_{k+1}^{\mathbf{M},T}, hyp_{k+1}^{\mathbf{M},T}) = \mathbf{M}(mem_k^{\mathbf{M},T}, T(k))$. Extend the definition of $\mathbf{M}$ to initial segments of texts by letting $\mathbf{M}(T[k]) = (mem_k^{\mathbf{M},T}, hyp_k^{\mathbf{M},T})$.
Intuitively, $mem_k^{\mathbf{M},T}$ is considered as the memory of the learner after having seen data $T[k]$ and $hyp_k^{\mathbf{M},T}$ is considered as the hypothesis of the learner after having seen data $T[k]$. Without loss of generality we assume that the memory is revised whenever the hypothesis is revised. Sometimes, we just say $\mathbf{M}(T[k]) = hyp_k^{\mathbf{M},T}$, where the memory of the learner is implicit.

(c) We say that $\mathbf{M}$ converges on a text T to a hypothesis hyp if the sequence $hyp_0^{\mathbf{M},T}, hyp_1^{\mathbf{M},T}, hyp_2^{\mathbf{M},T}, \ldots$ converges syntactically to hyp.

(d) We say that $\mathbf{M}$ **Ex**-learns a language L on text T if $\mathbf{M}$ is defined on all initial segments of the text T and $\mathbf{M}$ converges on T to a hypothesis hyp such that $W_{hyp} = L$.

(e) We say that $\mathbf{M}$ **Ex**-learns a language L (written: $L \in \mathbf{Ex}(\mathbf{M})$), if $\mathbf{M}$ **Ex**-learns L from each text T for L.

(f) We say that $\mathbf{M}$ **Ex**-learns a class $\mathcal{L}$ of languages (written: $\mathcal{L} \subseteq \mathbf{Ex}(\mathbf{M})$) iff $\mathbf{M}$ **Ex**-learns each $L \in \mathcal{L}$.

(g) $\mathbf{Ex} = \{\mathcal{L} : \exists \mathbf{M}\, [\mathcal{L} \subseteq \mathbf{Ex}(\mathbf{M})]\}$.

We say that $\mathbf{M}$ changes its mind (hypothesis, conjecture) at $T[n+1]$, if $hyp_n^{\mathbf{M},T} \neq hyp_{n+1}^{\mathbf{M},T}$. We say that $\mathbf{M}$ changes its memory at $T[n+1]$, if $mem_n^{\mathbf{M},T} \neq mem_{n+1}^{\mathbf{M},T}$ or $hyp_n^{\mathbf{M},T} \neq hyp_{n+1}^{\mathbf{M},T}$; that is, a learner cannot bring down the costs of mind changes by only changing the hypothesis and not doing adequate bookkeeping in the memory. Blum and Blum [1] gave a useful lemma for learnability of languages by learners based on topological considerations.

Definition 2 (Fulk [5]). Suppose $\mathbf{M}(\sigma) = (mem, hyp)$. Sequence σ is said to be a *stabilising sequence* for $\mathbf{M}$ on a language L if

(i) $\mathrm{content}(\sigma) \subseteq L$ and
(ii) for all τ such that $\sigma \subseteq \tau$ and $\mathrm{content}(\tau) \subseteq L$, $\mathbf{M}(\tau) = (\cdot, hyp)$. That is, $\mathbf{M}$ does not change its hypothesis beyond σ on any text T for L.

Definition 3 (Blum and Blum [1]). A sequence σ is said to be a *locking sequence* for $\mathbf{M}$ on a language L if

(i) σ is a stabilising sequence for $\mathbf{M}$ on L and
(ii) $W_{hyp} = L$, where $\mathbf{M}(\sigma) = (mem, hyp)$.

Lemma 4 (Blum and Blum [1]). *Suppose* $\mathbf{M}$ $\mathbf{Ex}$*-learns* L*. Then the following statements hold:*

(i) *There exists a locking sequence for* $\mathbf{M}$ *on* L*;*
(ii) *Every stabilising sequence for* $\mathbf{M}$ *on* L *is a locking sequence for* $\mathbf{M}$ *on* L*.*

A lemma similar to above can be shown for all of the criteria of learning considered in this paper.

For some of the results in this paper it is useful to consider set-driven learning, where the output of the learner depends just on the set of inputs seen.

Definition 5 (Osherson, Stob and Weinstein [9]).

(a) A learner $\mathbf{M}$ is said to be *set-driven* iff for all σ, τ, if $\mathrm{content}(\sigma) = \mathrm{content}(\tau)$, and $\mathbf{M}(\sigma) = (mem, hyp)$, then $\mathbf{M}(\tau) = (mem', hyp)$, for some memory mem'. That is, the hypothesis of the learner depends just on the content of the input and not on the exact order or the number of repetitions of the elements presented.
(b) $\mathbf{SD} = \{\mathcal{L} : \exists \mathbf{M}\, [\mathbf{M} \text{ is set-driven and } \mathcal{L} \subseteq \mathbf{Ex}(\mathbf{M})]\}$.

A related notion of rearrangement independence, where the learner is able to base its conjectures on the content and length of the input was considered by [12]. We now consider iterative learning.

Definition 6 (Based on Wexler and Culicover [13] and Wiehagen [14]).

(a) A learner $\mathbf{M}$ is said to be *iterative* if for any text T, $mem_k^{\mathbf{M},T} = hyp_k^{\mathbf{M},T}$ (where it is possible that both are undefined).
(b) $\mathbf{M}$ $\mathbf{ItEx}$-learns a language L (class of languages $\mathcal{L}$) iff $\mathbf{M}$ is iterative and $\mathbf{M}$ $\mathbf{Ex}$-learns L (class of languages $\mathcal{L}$).
(c) $\mathbf{ItEx} = \{\mathcal{L} : \exists \mathbf{M}\, [\mathbf{M} \text{ is iterative and } \mathcal{L} \subseteq \mathbf{Ex}(\mathbf{M})]\}$.

It can be shown that iterative learning is restrictive: $\mathbf{ItEx} \subset \mathbf{Ex}$ (see [14]).

There are several other variations of iterative learning considered in the literature where the memory is constrained in some way. For example, a learner $\mathbf{M}$ has k-bounded example memory [8, 10] iff for all texts T and n, $mem_n^{\mathbf{M},T}$ represents a set of size at most k, where $mem_0^{\mathbf{M},T} = \emptyset$ and $mem_{n+1}^{\mathbf{M},T} \subseteq mem_n^{\mathbf{M},T} \cup \{T(n)\}$.

As iterative learning (and most of its variations considered until now) put severe constraints on what can be memorised by a learner, we consider a variation where we allow the memory of grow arbitrarily large. However, there is a cost associated with each change of memory: if the learner changes its memory (or hypothesis) after seeing datum x, then it is charged some cost $f(x)$. Now the learner is said to identify the input language (for any given text) iff besides **Ex**-learning the language, the total cost charged to the learner on the input text is finite. Formally, this is defined as follows.

A *price function* is a recursive function mapping $\mathbb{N} \cup \{\#\}$ to $\mathbb{N}$.

Definition 7. Suppose f is a price function and $\mathcal{C}$ is a class of price functions.

(a) $\mathbf{M}$ $\mathbf{Priced}_f\mathbf{Ex}$-learns a language L (written: $L \in \mathbf{Priced}_f\mathbf{Ex}(\mathbf{M})$) iff $\mathbf{M}$ **Ex**-learns L and for all texts T for L, $\sum_{n:\mathbf{M}(T[n+1]) \neq \mathbf{M}(T[n])} \frac{1}{f(T(n))+1} < \infty$.
(b) $\mathbf{Priced}_f\mathbf{Ex} = \{\mathcal{L} : \exists \mathbf{M}\, [\mathcal{L} \subseteq \mathbf{Priced}_f\mathbf{Ex}(\mathbf{M})]\}$.
(c) $\mathbf{Priced}_{\mathcal{C}}\mathbf{Ex} = \{\mathcal{L} : \forall f \in \mathcal{C}\, [\mathcal{L} \in \mathbf{Priced}_f\mathbf{Ex}]\}$.

Note that in the above definition $\mathbf{M}(T[n])$ is taken as the pair $(mem_n^{\mathbf{M},T}, hyp_n^{\mathbf{M},T})$.

For a learner $\mathbf{M}$ and price function g, $\mathbf{cost}_{\mathbf{M},g}(\sigma)$ is defined as follows. Let $S = \{m < |\sigma| : \mathbf{M}$ makes a memory change or a mind change at $\sigma[m+1]\}$. Let $\mathbf{cost}_{\mathbf{M},g}(\sigma) = \sum_{m \in S} 1/(g(\sigma(m))+1)$ and $\mathbf{cost}_{\mathbf{M},g}(T) = \sup_{\sigma \preceq T} \mathbf{cost}_{\mathbf{M},g}(\sigma)$. Here $\mathbf{cost}_{\mathbf{M},g}(\sigma)$ ($\mathbf{cost}_{\mathbf{M},g}(T)$) is called the cost of $\mathbf{M}$ on input σ (input T) with respect to the cost function g.

Let $price_f(S) = \sum_{x \in S} \frac{1}{f(x)+1}$. Intuitively, $price_f(S)$ denotes the cost with respect to the price function f if a mind change is made by the learner exactly once with respect to each member of S.

Let $\mathbf{M}_0, \mathbf{M}_1, \ldots$ denote a recursive enumeration of all the learning machines. The following definition considers when a learner can be effectively obtained from the index of cost functions from a class of cost functions.

Definition 8. Suppose $\mathcal{C} \subseteq \mathcal{R}$. $\mathbf{EffPriced}_{\mathcal{C}}\mathbf{Ex} = \{\mathcal{L} : \exists f \in \mathcal{R}\, \forall i \in \mathbb{N}$ [if $\varphi_i \in \mathcal{C}$ then $\mathbf{M}_{f(i)}$ $\mathbf{Priced}_{\varphi_i}\mathbf{Ex}$-learns $\mathcal{L}$]$\}$.

4 General Cost Functions

The theorems in this section explore when $\mathbf{Priced}_f\mathbf{Ex} \subseteq \mathbf{Priced}_g\mathbf{Ex}$. Theorem 9 shows when $\mathcal{L} \in \mathbf{Priced}_f\mathbf{Ex}$ implies $\mathcal{L} \in \mathbf{Priced}_g\mathbf{Ex}$, and Theorem 11 shows when such an implication does not hold.

Theorem 9. *Suppose f and g are price functions for which no set S satisfies that $price_f(S)$ is finite and $price_g(S)$ is infinite. Then* $\mathbf{Priced}_f\mathbf{Ex} \subseteq \mathbf{Priced}_g\mathbf{Ex}$.

Proof. Suppose $\mathbf{M}$ $\mathbf{Priced}_f\mathbf{Ex}$-learns $\mathcal{L}$. Then, consider the following learner $\mathbf{M}'$.

Suppose T is the input text. Then, $\mathbf{M}'(T[n]) = (mem', hyp)$, where $(mem, hyp) = \mathbf{M}(T'[n])$ and $mem' = (mem, S)$, where T' is a modification of T (as described below) and S is a set formed by taking the elements in $T'[n]$ on which $\mathbf{M}$ made a memory / mind change.

We now describe T' (along with the set S of elements memorised by $\mathbf{M}$ as mentioned above). $T'(0) = T(0)$. Inductively, define $T'(n)$ as follows. If there exists an $n' < n$ such that $T'(n') = T(n)$ and $\mathbf{M}(T'[n']) \neq \mathbf{M}(T'[n'+1])$, then let $T'(n) = \#$; otherwise $T'(n) = T(n)$. Note that this transformation of T to T' can be done by $\mathbf{M}'$ by just remembering the elements S on which $\mathbf{M}$ made a memory / mind change.

Note that $\mathrm{content}(T') = \mathrm{content}(T)$, and M' on T converges to the same hypothesis as M converges to on T' (if any). Now suppose $\mathbf{cost}_{\mathbf{M},f}(T')$ is finite. Let $X = \{i : \mathbf{M}(T'[i]) \neq \mathbf{M}(T'[i+1])\} = \{i : \mathbf{M}'(T[i]) \neq \mathbf{M}'(T[i+1])\}$. Note that, by construction, for each $x \in \mathbb{N}$, there exists at most one i in X such that $T'(i) = x$. As $\mathbf{cost}_{M,f}(T')$ is finite, there are only finitely many $i \in X$ such that $T'(i) = \#$ (otherwise, $\mathbf{cost}_{M,f}(T')$ will be unbounded). Let $\mathrm{card}(\{i \in X : T'(i) = \#\}) = n_{\#}$. It follows that for each $x \in \mathbb{N} \cup \{\#\}$, there are at most $n_{\#} + 1$ many $i \in X$ such that $T(i) = x$. Also, by the hypothesis of the theorem we have that $price_g(\{T(i) : i \in X\})$ is finite. As $\mathbf{cost}_{\mathbf{M}',g}(T) \leq (n_{\#} + 1)(price_g(\{T(i) : i \in X\}))$, it follows that $\mathbf{cost}_{\mathbf{M}',g}(T)$ is finite. Hence, $\mathbf{M}'$ $\mathbf{Priced}_g\mathbf{Ex}$-learns any language which is $\mathbf{Priced}_f\mathbf{Ex}$-learnt by $\mathbf{M}$. □

Proposition 10. *Let S be an infinite r.e. set, g be a price function such that $price_g(S)$ is infinite. Then, $\mathcal{L} = \{S - \{x\} : x \in S - \min(S)\}$ is not* $\mathbf{Priced}_g\mathbf{Ex}$*-learnable.*

Proof. Suppose by way of contradiction that $\mathbf{M}$ $\mathbf{Priced}_g\mathbf{Ex}$-learns $\mathcal{L}$.

Suppose there are distinct $x, y \in S - \{\min(S)\}$ such that for some σ satisfying $\{x, y\} \subseteq \mathrm{content}(\sigma) \subseteq S$, $\mathbf{M}$ did not change its memory and hypothesis whenever it received x, y in the input σ. Let T be a text for $S - \{x, y\}$. Let σ' be obtained from σ by deleting the occurrences of x and σ'' be obtained from σ by deleting the occurrences of y. Now, $\mathbf{M}$ converges to the same hypothesis on both $\sigma' \diamond T$ and $\sigma'' \diamond T$, however $\sigma' \diamond T$ and $\sigma'' \diamond T$ are respectively texts for two different languages $S - \{x\}$ and $S - \{y\}$ in $\mathcal{L}$. Thus, $\mathbf{M}$ does not $\mathbf{Priced}_g\mathbf{Ex}$-learn at least one of these languages.

On the other hand, if such distinct x, y and corresponding σ as above do not exist, then on any text T for a language $L \in \mathcal{L}$, the cost incurred by $\mathbf{M}$ on T is infinite.

Thus, $\mathbf{M}$ does not $\mathbf{Priced}_g\mathbf{Ex}$-learn $\mathcal{L}$. □

Theorem 11. *Suppose f and g are price functions such that there exists a set S satisfying that $price_f(S)$ is finite but $price_g(S)$ is infinite. Then* $\mathbf{Priced}_f\mathbf{Ex} \not\subseteq \mathbf{Priced}_g\mathbf{Ex}$.

Proof. Suppose there is a set S such that $price_f(S)$ is finite and $price_g(S)$ is infinite. Then, for each c, there exists a set S_c such that $price_f(S_c) < 1/(c+1)$ and $price_g(S_c) = \infty$. Such sets can be found by considering tails of the sum of the members in S. Hence one can, by effective search, find disjoint finite sets $\tilde{S}_0, \tilde{S}_1, \ldots$ such that $price_f(\tilde{S}_c) < 1/(c+1)$ and $price_g(\tilde{S}_c) \geq 1$. To see this, define $\tilde{S}_c$ by induction as the first finite set found (in some ordering of finite sets) such that $price_f(\tilde{S}_c) < 1/(w+1)$ and $price_g(\tilde{S}_c) \geq 1$, where $w > c + \max(\{f(x) : x \in \bigcup_{c'<c} \tilde{S}_{c'}\})$; note that the constraint on w implies that $\tilde{S}_c \cap (\bigcup_{c'<c} \tilde{S}_{c'}) = \emptyset$. Now let $\tilde{S}$ be the union of all sets $\tilde{S}_{2^k}$ for $k \in \mathbb{N}$. The set $\tilde{S}$ is an r.e. set.

One can let $\mathcal{L} = \{\tilde{S} - \{x\} : x \in \tilde{S}\}$. It is easy to see that $\mathcal{L} \in$ **Priced**$_f$**Ex**, as on any input text T for $L \in \mathcal{L}$, the learner can memorise all data seen, and, in the limit, output a grammar for $\tilde{S} - \{x\}$, where x is the minimal element, if any, in $\tilde{S} - \text{content}(T)$. Note that such a learner makes a memory / mind change on any x at most once.

By Proposition 10, $\mathcal{L} \notin$ **Priced**$_g$**Ex**. The theorem follows. □

Theorem 9 and Theorem 11 imply that **Priced**$_f$**Ex** = **ItEx** iff $price_f(S)$ is infinite for every infinite subset S of $\mathbb{N}$.

Theorem 11 can be generalised to find classes which have one learner working with all price functions which permit to learn this class. These classes can be chosen such that they are priced learnable for some but not all price functions and hence they do not coincide with iteratively learnable classes.

Theorem 12. *Given any price function f, there is a class $\mathcal{C}_f$ and a learner* **M** *such that for all price functions g, the following conditions are equivalent:*

- **Priced**$_f$**Ex** $\subseteq$ **Priced**$_g$**Ex**;
- **M Priced**$_g$**Ex**-*learns $\mathcal{C}_f$;*
- *There is no set S such that $price_f(S)$ is finite but $price_g(S)$ is infinite.*

Proof. Let $\mathcal{C}_f = \{S - \{x\} : S \text{ is infinite}, x \in S - \{\min(S)\}, W_{\min(S)} = S \text{ and } price_f(S) \text{ is finite }\}$.

Let e be a recursive function such that $W_{e(d,x)} = W_d - \{x\}$. The learner **M** memorises each new datum it observes. After having seen input σ, the hypothesis of the learner is $e(d,x)$, where $d = \min(\text{content}(\sigma))$ and $x = \min(W_{d,\max(\text{content}(\sigma))} - \text{content}(\sigma))$; if no such x exists, then the learner outputs $e(d,d+1)$. Thus, **M** on a text T for a language $L \in \mathcal{C}_f$ converges to the index $e(d,x)$ such that d is the minimum element observed in the text T and x is the minimum element of $W_d - \text{content}(T)$

Note that the learner **M** only updates its memory (and possibly its hypothesis $e(d,x)$) only when it sees a new datum not seen earlier. Thus, for all input texts for $L \in \mathcal{C}_f$, the total cost of memory / mind changes is finite. Note that the learner might incur infinite cost on sets outside $\mathcal{C}_f$.

Now let g be any further price function. First assume the case that for all infinite sets L, if $price_f(L)$ is finite then $price_g(L)$ is finite. It immediately follows that **M** is also **Priced**$_g$**Ex**-learning $\mathcal{C}_f$, as whenever the cost charged to **M** on a text is finite with respect to the price function f, the cost is also finite with respect to price function g.

Now assume that a further learner **N** is **Priced**$_f$**Ex**-learning some class $\mathcal{L}$. Without loss of generality, one can assume the following on **N**:

- Whenever **N** updates its memory on a datum x then it also memorises in some data base that it has observed x;

– Whenever **N** updates the memory on a datum x which it has memorised before then **N** does the same with # in place of x (otherwise **N** could abstain from doing the update, as it also has to learn on texts where some repetitions of the elements are replaced by #).

Now assume that **N** $\mathbf{Priced}_f\mathbf{Ex}$-learns from a text T some language in the class $\mathcal{L}$. Then, without loss of generality, **N** must also $\mathbf{Priced}_f\mathbf{Ex}$-learn from a text T' where all occurrences of x which it has already memorised are replaced by occurrences of # and therefore there can only be finitely many such occurrences – as otherwise the constant cost of # would make the sum go to infinity. Thus, the cost of **N** on T with respect to the price function g would only diverge in the following case: The set S of x occurring in T on which **N** updates the memory satisfies that $\sum_f S$ is finite and $\mathit{price}_g(S)$ is infinite. However, by assumption, such a set does not exist and therefore **N** $\mathbf{Priced}_g\mathbf{Ex}$-learns every set which it is $\mathbf{Priced}_f\mathbf{Ex}$-learning. It follows that $\mathbf{Priced}_f\mathbf{Ex} \subseteq \mathbf{Priced}_g\mathbf{Ex}$.

Second assume the case that there is some infinite set on which the costs of f converge and the costs of g diverge. Thus, as in the proof of Theorem 11, there is also an r.e. such set S; that is, S satisfies $\mathit{price}_f(S)$ is finite and $\mathit{price}_g(S)$ is infinite. Furthermore, without loss of generality, one can assume that $\min(S)$ is an index for S (by a finite modification of the set). Now, by Proposition 10, no learner can $\mathbf{Priced}_g\mathbf{Ex}$-learn $\{S - \{x\} : x > \min(S)\}$. As $\{S - \{x\} : x > \min(S)\} \subseteq \mathcal{C}_f$, it follows that $\mathcal{C}_f \in \mathbf{Priced}_f\mathbf{Ex} - \mathbf{Priced}_g\mathbf{Ex}$. □

5 Priced Learning and Set-Drivenness

Note that a $\mathbf{Priced}_f\mathbf{Ex}$-learner cannot remember how many times some x appears in the input, as it can change its mind on x at most finitely often. This constrains the power of $\mathbf{Priced}_f\mathbf{Ex}$-learning compared to **Ex**-learning. In this section we show that classes which are $\mathbf{Priced}_f\mathbf{Ex}$-learnable with respect to some price function f are learnable by a set-driven learner and vice versa, classes which are set-driven-learnable are $\mathbf{Priced}_g\mathbf{Ex}$-learnable with respect to some price function g.

To this end we first give a modification of the definition of stabilising/locking sequences with respect to a cost function.

Definition 13. A sequence σ is a *g-cost-stabilising sequence* (*g-cost-locking sequence*) for **M** on L, if the following conditions hold:

(a) σ is a stabilising sequence (locking sequence) for **M** on L and
(b) for all τ such that $\sigma \preceq \tau$ and $\mathrm{content}(\tau) \subseteq L$, $\mathbf{cost}_{\mathbf{M},g}(\tau) \le \mathbf{cost}_{\mathbf{M},g}(\sigma) + 1$.

The following can be proved in a way similar to locking sequence lemma in [1].

Lemma 14 (Based on Blum and Blum [1]). *If* **M** $\mathbf{Priced}_g\mathbf{Ex}$*-learns L then there exists a g-cost-locking sequence for* **M** *on L; Furthermore, all g-cost-stabilising sequences for* **M** *on L are g-cost-locking sequences for* **M** *on L.*

Furthermore, note that for any finite set S contained in some set learned by $\mathbf{M}$, and any price function g one can check if σ is a g-cost-stabilising sequence for $\mathbf{M}$ on S. This holds as the number of memory / mind changes of $\mathbf{M}$ on any text extending σ for S will be at most $\max(\{g(x)+1 : x \in S \cup \{\#\}\})$ if σ is indeed a g-cost-stabilising sequence for $\mathbf{M}$ on S. Thus, we can check if the tree of memory / mind changes for $\mathbf{M}$ on S above σ is bounded in depth by $\max(\{g(x)+1 : x \in S \cup \{\#\}\})$.

Theorem 15. $\bigcup_{g\in\mathcal{R}} \mathbf{Priced}_g\mathbf{Ex} = \mathbf{SD}$.

Proof. Suppose $\mathcal{L} \in \mathbf{SD}$ as witnessed by $\mathbf{M}$. Let $g(x) = 1/2^x$. Let $\mathbf{N}(\sigma) = (\text{content}(\sigma), hyp)$, where hyp is the hypothesis of $\mathbf{M}$ on input σ. Note that, as $\mathbf{M}$ is set-driven, $\mathbf{N}$ can easily obtain the hypothesis of $\mathbf{M}$ on input σ, using its memory content(σ), by running $\mathbf{M}$ on any τ such that content$(\tau) =$ content(σ). It is easy to verify that $\mathbf{N}$ $\mathbf{Priced}_g\mathbf{Ex}$-learns $\mathcal{L}$.

Now suppose $\mathcal{L} \in \mathbf{Priced}_g\mathbf{Ex}$ as witnessed by $\mathbf{M}$. Then, consider the following set-driven learner $\mathbf{N}$.

The aim of the set-driven learner $\mathbf{N}$ is to find the least g-cost-stabilising sequence for $\mathbf{M}$ on the input language. $\mathbf{N}$ on input σ searches for the least g-cost-stabilising sequence τ for $\mathbf{M}$ on content(σ), if any, of length at most $|\text{content}(\sigma)|$. Note that one can check if such a stabilising sequence exists. Now, $\mathbf{N}(\sigma) = (\text{content}(\sigma), hyp)$, where hyp is the hypothesis of $\mathbf{M}$ on input τ if a τ as above exists; otherwise hyp is a canonical grammar for content(σ).

For a finite language $L \in \mathcal{L}$, whether there exists a g-cost-stabilising sequence for $\mathbf{M}$ on L of length at most card(L) or not, the learner $\mathbf{N}$ will output a grammar for L after it has received all the elements of L. Thus, $\mathbf{N}$ would $\mathbf{Ex}$-learn L.

For an infinite language $L \in \mathcal{L}$, there exists a g-cost-stabilising sequence for $\mathbf{M}$ on L of finite length, and thus for large enough initial segment of any text T for L, $\mathbf{N}$ will find the least g-cost-stabilising sequence τ. Then, $\mathbf{N}$ will output the hypothesis of $\mathbf{M}$ on τ as its hypothesis. Thus, $\mathbf{N}$ would $\mathbf{Ex}$-learn L.

As $\mathbf{N}$ defined above is set-driven, we have that $\mathcal{L} \in \mathbf{SD}$. □

Theorem 9 and Theorem 11 along with $\mathbf{SD} = \mathbf{Priced}_g\mathbf{Ex}$, for the g used in the proof of Theorem 15, imply that $\mathbf{Priced}_f\mathbf{Ex} = \mathbf{SD}$ iff $price_f(\mathbb{N})$ is finite.

6 Classes of Cost Functions

In this section we will consider priced learning when the price function might be any of the functions in a class of price functions. In particular we look at the classes of monotonic functions ($\mathbf{MF}$) and finite-one functions ($\mathbf{FF}$) which map only finitely many inputs to the same output:

$$\mathbf{MF} = \{f \in \mathcal{R} : \forall x\, [f(x) \leq f(x+1)] \text{ and } f \text{ is unbounded}\};$$
$$\mathbf{FF} = \{f \in \mathcal{R} : \forall y\, [\text{card}(f^{-1}(y)) \text{ is finite}]\}.$$

$\mathbf{FF}$ is natural in some sense as it contains functions which give small costs (below any given value) for large enough x. $\mathbf{MF}$ additionally requires monotonicity in the costs.

The following lemma is useful for proving some results in this paper. It shows that if $\mathcal{L}$ is **Ex**-learnable, then the cylindrification of $\mathcal{L}$ is $\mathbf{Priced_{FF}Ex}$-learnable.

Lemma 16. *Suppose $\mathcal{L} \in \mathbf{Ex}$ and $\mathcal{L}' = \{\text{cyl}(L) : L \in \mathcal{L}\}$. Then, $\mathcal{L}' \in \mathbf{EffPriced_{FF}Ex}$.*

Proof. Suppose $\mathbf{M}$ is a $\mathbf{Ex}$ learner for $\mathcal{L}$. Given a function $g \in \mathbf{FF}$, learner $\mathbf{M}'$ behaves as follows. Initially, it outputs a conjecture for $\emptyset$ until it sees some element $\langle x, i\rangle$ in the input. At which point its memory becomes the sequence $\langle x, i\rangle$. From then onwards it appends new element $\langle y, j\rangle$ in the input to its memory if and only if (i) $g(\langle y, j\rangle) \geq 2^y 2^p$, where p is the number of times elements of the form $\langle y, j'\rangle$ appear in the sequence in the memory. The conjecture of the learner $\mathbf{M}'$ is the cylindrification of the conjecture of $\mathbf{M}$ on the sequence in the memory with the second components of each pair dropped (that is if σ is the sequence in the memory, then conjecture of $\mathbf{M}'$ is index for $\text{cyl}(W_{\mathbf{M}(\tau)})$, where $\tau(j) = \pi_1(\sigma(j))$). It is easy to verify that if $\mathbf{M}$ $\mathbf{Ex}$ learns L, then $\mathbf{M}'$ $\mathbf{Priced}_g\mathbf{Ex}$-learns $\text{cyl}(L)$. Note that the learner $\mathbf{M}'$ can be found effectively from an index for function g. The lemma follows. ☐

The next theorem gives a class $\mathcal{L}$ that can be $\mathbf{Priced}_f\mathbf{Ex}$-learnt with respect to every price function f in $\mathbf{FF}$, even though $\mathcal{L}$ is not iteratively learnable. Moreover, the $\mathbf{Priced}_f\mathbf{Ex}$-learner can be obtained effectively from an index for f.

Theorem 17. $\mathbf{EffPriced_{FF}Ex} \not\subseteq \mathbf{ItEx}$.

Proof. The class $\mathcal{L}_1 \cup \mathcal{L}_2$ with

$$\mathcal{L}_1 = \{W_e : e = \min(W_e) \text{ and } \forall x\,[\{2x, 2x+1\} \not\subseteq W_e]\} \text{ and}$$
$$\mathcal{L}_2 = \{L : \text{card}(L) < \infty \text{ and } \exists x\,[\{2x, 2x+1\} \subseteq L]\}$$

is **Ex**-learnable: The learner, on input σ, checks whether there is an x with $2x, 2x+1 \in \text{content}(\sigma)$. If so then the learner conjectures a canonical index for $\text{content}(\sigma)$ else the learner conjectures $\min(\text{content}(\sigma))$. So, on a text T with both $2x, 2x+1 \in \text{content}(T)$, the learner converges to a canonical index of $\text{content}(T)$ whenever the latter is finite and thus the learner **Ex**-learns $\mathcal{L}_2$; on a text T which contains at most one of $2x, 2x+1$ for any x, the learner converges to the minimum of $\text{content}(T)$ and thus the learner **Ex**-learns $\mathcal{L}_1$. Let

$$\mathcal{L} = \{\text{cyl}(L) : L \in \mathcal{L}_1 \cup \mathcal{L}_2\}.$$

By Lemma 16, $\mathcal{L} \in \mathbf{EffPriced_{FF}Ex}$.

Now it remains to show that $\mathcal{L} \notin \mathbf{ItEx}$. Suppose by way of contradiction that $\mathbf{M}$ $\mathbf{ItEx}$-learns $\mathcal{L}$. Then by Kleene's recursion theorem there exists an e such that W_e is the set of all x for which $\langle x, 0\rangle$ occurs in some σ_s which are defined below. Here, initially, $\sigma_0 = \langle e, 0\rangle$ and in stage s below, σ_{s+1} is constructed.

Stage s.

1. Search for an extension τ of σ_s satisfying the following conditions:

(i) $\mathbf{M}(\sigma_s) \neq \mathbf{M}(\tau)$;

(ii) for all x and i, if $\langle 2x, i\rangle \in \text{content}(\tau)$, then for all j, $\langle 2x+1, j\rangle \notin \text{content}(\tau)$;

(iii) for all i, for all $x < e$, $\langle x, i\rangle \notin \text{content}(\tau)$.

2. If and when such a τ is found, let σ_{s+1} be an extension of τ satisfying the following conditions:

(i) $\forall x, i\,[[\langle x, i\rangle \in \text{content}(\tau)] \Rightarrow [\forall j \leq s\,[\langle x, j\rangle \in \text{content}(\sigma_{s+1})]]]$;

(ii) $\forall x, i\,[[\langle x, i\rangle \in \text{content}(\sigma_{s+1})] \Rightarrow [\exists j\,[\langle x, j\rangle \in \text{content}(\tau)]]]$.

3. Go to stage $s+1$.

End Stage s.

Now if there are infinitely many stages in the above construction, then $T = \bigcup_{s\in\mathbb{N}} \sigma_s$ is a text for $\text{cyl}(W_e)$, where $W_e \in \mathcal{L}_1$. However, $\mathbf{M}$ makes infinitely many mind changes on T. On the other hand if stage s starts but does not finish, then, for $y > \max(\{x : \langle x, 0\rangle \in \text{content}(\sigma_s)\})$, let $L_y = \{x : \langle x, 0\rangle \in \text{content}(\sigma_s)\} \cup \{2y, 2y+1\}$. Then for all texts T extending σ_s for any $\text{cyl}(L_y)$ with $y > \max(\{x : \langle x, 0\rangle \in \text{content}(\sigma_s)\})$, $\mathbf{M}(T)$ converges to $\mathbf{M}(\sigma_s)$, and thus it fails to **ItEx**-learn $\mathcal{L}$, as all such L_y belong to $\mathcal{L}_2$. □

The next theorem shows that some class $\mathcal{L}$ can be $\mathbf{Priced}_f\mathbf{Ex}$-learnt with respect to every price function in **MF**, but not with respect to some price function in **FF**. Furthermore, the $\mathbf{Priced}_f\mathbf{Ex}$ learner for $f \in \mathbf{MF}$ can be obtained effectively from an index for f.

Theorem 18. $\mathbf{EffPriced_{MF}Ex} \not\subseteq \mathbf{Priced_{FF}Ex}$.

Proof. Let $F_0, F_1, \ldots$ be an enumeration of disjoint finite sets such that for all i, if φ_i is total, then for all $x > i$, there exists a $y > \varphi_i(x)$ in F_x. Note that such an enumeration can be easily constructed.

Let $L_0 = \bigcup_{x>0} F_x$. For $e \in \mathbb{N}$, let $L_{e+1} = \bigcup_{x\in D_e\cup\{0\}} F_x$. Let $\mathcal{L} = \{L_e : e \in \mathbb{N}\}$.

Claim. $\mathcal{L} \in \mathbf{EffPriced_{MF}Ex}$.

Let h be a recursive function such that for all i, $\varphi_{h(i)}(x)$ is the first z found in a search such that $\varphi_i(z) > 2^x$. Now, if $\varphi_i \in \mathbf{MF}$, then for all $x > h(i)$, there exists a $y \in F_x$ such that $y > \varphi_{h(i)}(x)$ and thus $\varphi_i(y) > 2^x$.

Given an index i for a cost function $g \in \mathbf{MF}$, consider the following learner $\mathbf{M}$ obtained effectively from i. Learner $\mathbf{M}$ will have a finite set as its memory. Initially, $\mathbf{M}$'s memory is $\emptyset$. Whenever $\mathbf{M}$ receives $y \in F_x$ as input, it adds x to the memory if x is not already in the memory and $[x \leq h(i)$ or $g(y) > 2^x]$; otherwise the memory is unchanged. If the memory does not contain 0, then $\mathbf{M}$ outputs a canonical grammar for L_0. Otherwise it outputs a canonical grammar for L_{e+1}, where $D_e = S - \{0\}$, where S is the set in $\mathbf{M}$'s memory. It is easy to verify that for any text T, $\mathbf{cost}_{\mathbf{M},g}(T)$ is finite and $\mathbf{M}$ $\mathbf{Priced}_g\mathbf{Ex}$-learns $\mathcal{L}$.

Claim. $\mathcal{L} \notin \mathbf{Priced_{FF}Ex}$.

To see this, let g be such that $g(y) = x$, for $y \in F_x$. Suppose by way of contradiction that **M** $\mathbf{Priced}_g\mathbf{Ex}$-learns $\mathcal{L}$. Now, consider **M**'s behaviour on text for L_0, where the elements are presented to it in order of elements from $F_1, F_2, \ldots$. If for every x, **M** makes mind change on some input from F_x, then clearly the cost is infinite. Otherwise, for some x, **M** did not make any mind change on the above text when presented with elements from F_x. Now, if after elements of F_x, **M** only gets elements from F_0, then **M** cannot distinguish between the input being L_e or $L_{e'}$, where $D_e = \{1, 2, \ldots, x\}$ and $D_{e'} = \{1, 2, \ldots, x - 1\}$.

The above two claims prove the theorem. □

The following theorem shows that some class of languages $\mathcal{L}$ can be learnt with respect to every price function in **FF**, but one cannot effectively obtain a learner for even monotonic price functions.

Theorem 19. $\mathbf{Priced_{FF}Ex} \not\subseteq \mathbf{EffPriced_{MF}Ex}$.

Proof. We will define the class $\mathcal{L}$ in dependence of a set $A \subseteq \{\langle x, r\rangle : x \geq 1, r \in \mathbb{N}\}$ to be constructed below by letting

$$\begin{aligned} L_m &= \{0\} \cup (A \cap \{\langle x, r\rangle : x \leq m, r \in \mathbb{N}\}); \\ \mathcal{L} &= \{A\} \cup \{L_m : m \in \mathbb{N}\}. \end{aligned}$$

We will have the property that for all $g \in \mathbf{FF}$, for all but finitely many x, there exists an r such that $\langle x, r\rangle \in A$ and $g(\langle x, r\rangle) > 2^x$.

The above property allows for easy learning of $\mathcal{L}$ in the model $\mathbf{Priced_{FF}Ex}$. To see this, for $g \in \mathbf{FF}$, let R_g be the set of finitely many x such that for all r, $g(\langle x, r\rangle) \leq 2^x$. Then, the $\mathbf{Priced}_g\mathbf{Ex}$-learner **M** is defined as follows. The memory of **M** is a finite set. The memory mem_σ of **M** after having seen input σ consists of the following elements: (i) 0 if $0 \in \mathrm{content}(\sigma)$, (ii) all $x \in R_g$ such that for some $r \in \mathbb{N}$, $\langle x, r\rangle \in \mathrm{content}(\sigma)$, and (iii) all $x \geq 1$ such that for some $r \in \mathbb{N}$, $\langle x, r\rangle \in \mathrm{content}(\sigma)$ and $g(\langle x, r\rangle) > 2^x$. Furthermore, the hypothesis of the learner **M** after having seen input σ is: (a) a canonical index for A if $0 \notin mem_\sigma$, and (b) a canonical index for L_m, if $0 \in mem_\sigma$ and $m = \max(mem_\sigma)$. It is easy to verify that **M** is a $\mathbf{Priced}_g\mathbf{Ex}$-learner for $\mathcal{L}$.

We will now construct A. For this we will define functions $h_{i,j}$ along with sets $B_{i,j}$, for $j < 2^{i+1}$. The functions $h_{i,j}$, if total, will be in **MF**. $B_{i,j}$ will be finite, but may change over time. Let $B^t_{i,j}$ denote its value at time t, where we will only start with $B_{i,j}$ at time $t \geq i$ (so it is assumed to be empty before that).

Intuitively, aim of $B_{i,j}$ and $h_{i,j}$, $j < 2^{i+1}$ is to make sure that φ_i does not witness that $\mathcal{L} \in \mathbf{EffPriced_{MF}Ex}$. For $j < 2^{i+1}$, $X_{i,j}$ denotes the j-th subset of $\{0, 1, 2, \ldots, i\}$.

In the definitions below, we assume some steps to be atomic so that there is no interference between different parts. For ease of notation we will write $\varphi_i(h_{i,j})$ below rather than φ_i(index for $h_{i,j}$).

Definition of A.

Stage s

Each stage s is supposed to be atomic.

1. Enumerate $\langle s, 0\rangle$ into A.
2. For each $i < s$,

   2.1. Let $x < s$ be least, if any, such that
   $x \notin \bigcup_{j<i,k<2^{j+1}} B^s_{j,k}$ and
   for all r such that $\langle x, r\rangle \in A$ enumerated up to now
   $[\varphi_{i,s}(\langle x, r\rangle)\downarrow \leq 2^x$ or $\varphi_{i,s}(\langle x, r\rangle)\uparrow]$.

   2.2. If there is such an x as above and $\exists r \leq s\, [\varphi_{i,s}(\langle x, r\rangle)\downarrow > 2^x]$
   Then enumerate one such $\langle x, r\rangle$ into A.

   End For

   Go to stage $s + 1$

End Stage s

Note that the construction of different $h_{i,j}$ is run in parallel (in dovetailing way, along with procedure for A above) for all $i \in \mathbb{N}$, $j < 2^{i+1}$.

Definition of $h_{i,j}$.

Let $h_{i,j}(0) = 1$.

Initially $\sigma_{i,j} = \Lambda$.

Loop

0. Let $\sigma_{i,j}$ be extension of previous $\sigma_{i,j}$ so that it contains all of A enumerated up to now.
   (Above step is assumed to be atomic.)
1. Suppose $h_{i,j}$ has already been defined on inputs $\leq w$ up to now, and the maximum value in the range of $h_{i,j}$ up to now is z.
2. Pick $B_{i,j}$ of size $z + 2$ containing elements $> i + w$, such that none of the elements of $B_{i,j}$ have been used in any $B_{i',j'}$ at any earlier point in the construction up to now and all the elements of $B_{i,j}$ are larger than any element enumerated by A up to now.
   (We assume the above step to be atomic.)
3. Wait until for each $y \in B_{i,j}$,
   (i) $\langle y, 0\rangle$ is enumerated into A and
   (ii) for each $i' \in X_{i,j}$, a pair $\langle y, r\rangle$ is enumerated into A with $\varphi_{i'}(\langle y, r\rangle) > 2^y$.
4. Extend $h_{i,j}$ monotonically (non-decreasing) so that $h_{i,j}(\langle y, r\rangle) = z + 1$, for each $\langle y, r\rangle \in A$ enumerated up to now such that $y \in B_{i,j}$.
   (We assume above step to be atomic.)
5. Keep extending $h_{i,j}$ monotonically (non-decreasing), while searching for an extension τ of $\sigma_{i,j}$ such that
   (i) content(τ) is contained in A, and
   (ii) for each $y \in B_{i,j}$, there exists a $\langle y, r\rangle$ in content(τ) such that $\mathbf{M}_{\varphi_i(h_{i,j})}$ made a mind change / memory change when it received $\langle y, r\rangle$ as input.

6. If such an extension τ is found, update $\sigma_{i,j}$ to τ and go to the next iteration of the loop. If no extension as above is found, then this iteration of the loop never ends.

End Loop

For each j and $k < 2^{j+1}$, let $B_{i,j}$ denote the eventual value of $B_{i,j}$ (if this value does not exist, then let $B_{i,j}$ be $\emptyset$).

Now consider any i such that $\varphi_i \in \mathbf{FF}$. Let $C_i = \bigcup_{i'<i, j'<2^{i'+1}} B_{i',j'}$. Note that for $x \notin C_i$, eventually some element of the form $\langle x, r\rangle$, with $\varphi_i(\langle x, r\rangle) > 2^x$ will be enumerated into A. Thus, we have that $\mathcal{L} \in \mathbf{Priced_{FF}Ex}$.

Now we show that $\mathcal{L} \notin \mathbf{EffPriced_{MF}Ex}$. Suppose by way of contradiction that φ_i witnesses $\mathcal{L} \in \mathbf{EffPriced_{MF}Ex}$, that is, for all k, if $\varphi_k \in \mathbf{MF}$, then $\mathbf{M}_{\varphi_i(k)}$ is a $\mathbf{Priced}_{\varphi_k}\mathbf{Ex}$-learner for $\mathcal{L}$. Let j be such that $X_{i,j}$ is the set of $i' \le i$ satisfying for all $y \notin \bigcup_{i''<i', j''<2^{i''+1}} B_{i'',j''}$, some $\langle y, r\rangle$ with $\varphi_{i'}(\langle y, r\rangle) > 2^y$ is enumerated into A. Now consider the construction of $h_{i,j}$. Now consider the following cases.

Case 1: Some iteration of the Loop for $h_{i,j}$ does not terminate.

In this case we have that $\mathbf{M}_{\varphi_i(h_{i,j})}$ fails to $\mathbf{Priced}_{h_{i,j}}\mathbf{Ex}$ learn $\mathcal{L}$. To see this let $\sigma_{i,j}$ be as in the iteration of the loop (for $h_{i,j}$) which starts but does not terminate. Consider τ which extends $\sigma_{i,j}$, and contains elements exactly from $\mathrm{content}(\sigma_{i,j}) \cup (A \cap \{\langle y, r\rangle : y \in B_{i,j}\})$, where these $\langle y, r\rangle$ appear in τ in increasing order of y. Then by step 5 not succeeding in the loop, there exists a $y \in B_{i,j}$, for which $M_{\varphi_i(h_{i,j})}$ did not make a mind change / memory change when receiving $\langle y, r\rangle$ as input, for all $\langle y, r\rangle \in A$. Thus, we can fool the learner $\mathbf{M}_{\varphi_i(h_{i,j})}$ by giving 0 just after giving the elements $\langle y, r\rangle$ corresponding to the y above, and then extending it using all elements from L_{y-1}. Then, the learner $\mathbf{M}_{\varphi_i(h_{i,j})}$ is not able to distinguish between input being L_{y-1} or L_y.

Case 2: All iterations of the loop for $h_{i,j}$ terminate.

In this case consider the loop executions in which for each $y \in B_{i,j}$ none of the $i' \le i$, $i' \notin X_{i,j}$ have a $\langle y, r\rangle \in A$ with $\varphi_{i'}(\langle y, r\rangle) > 2^y$. Note that all but finitely many loop executions have this property. In each of such executions, the construction would have found a mind change / memory change which costs in total at least 1 to the learner (as in each iteration, size of $B_{i,j}$ is $z + 2$, and the cost of memory / mind change on $\langle y, r\rangle$, $y \in B_{i,j}$ is $\frac{1}{z+2}$).

Thus, $\mathbf{M}_{\varphi_i(h_{i,j})}$ does not $\mathbf{Priced_{MF}Ex}$-learn $\mathcal{L}$. □

7 Other Uniformity Criteria

In this section we consider some uniformity criteria for priced learning.

Definition 20. Suppose $\mathcal{C} \subseteq \mathcal{R}$.

(a) $\mathbf{UniPriced}_{\mathcal{C}}\mathbf{Ex} = \{\mathcal{L} : \exists \mathbf{M}\, \forall f \in \mathcal{C}\, [\mathbf{M}\ \mathbf{Priced}_f\mathbf{Ex}\text{-learns } \mathcal{L}]\}$.
(b) $\mathbf{EffBPriced}_{\mathcal{C}}\mathbf{Ex} = \{\mathcal{L} : \exists g \in \mathcal{R}\, \forall i\ [\text{if } \varphi_i \in \mathcal{C} \text{ then } \forall b \ge i\, [\mathbf{M}_{g(b)}\ \mathbf{Priced}_{\varphi_i}\mathbf{Ex}\text{-learns } \mathcal{L}]]\}$.
(c) $\mathbf{UniFPriced}_{\mathcal{C}}\mathbf{Ex} = \{\mathcal{L} : \exists \mathbf{M}\, [\mathbf{M}\ \mathbf{Ex}\text{-learns } \mathcal{L} \text{ and } \forall f \in \mathcal{C}\, [\mathrm{card}(\mathcal{L} - \mathbf{Priced}_f\mathbf{Ex}(\mathbf{M})) \text{ is finite }]]\}$.

Intuitively, for $\mathbf{UniPriced}_{\mathcal{C}}\mathbf{Ex}$, the same learner $\mathbf{M}$ succeeds for all price functions in $\mathcal{C}$. For $\mathbf{EffBPriced}_{\mathcal{C}}\mathbf{Ex}$-learning the learner can be obtained effectively from a bound on the index for the price function.

For $\mathbf{UniFPriced}_{\mathcal{C}}\mathbf{Ex}$-learning, the same learner $\mathbf{M}$ $\mathbf{Ex}$-learns the class $\mathcal{L}$ and for all price functions f in $\mathcal{C}$, $\mathbf{M}$ $\mathbf{Priced}_f\mathbf{Ex}$-learns almost all the members of the class $\mathcal{L}$. Thus, the cost of learning can go infinite only for finitely many languages in the class for each price function $f \in \mathcal{C}$.

In addition we consider the case where the learner gets the cost function f as input rather than via index for it. In $\mathbf{ValPriced}_{\mathcal{C}}\mathbf{Ex}$ learning of $\mathcal{L}$, for all $f \in \mathcal{C}$, the same learner $\mathbf{M}$ $\mathbf{Priced}_f\mathbf{Ex}$-learn each member of $\mathcal{L}$ when, in each iteration, instead of a single datum x the pair $(x, f(x))$ is presented to the learner in the iteration, so that the learner knows how expensive it is to process a datum.

Theorem 21. (a) *Let $\mathcal{C}$ be a class of price functions.*
$\mathbf{ValPriced}_{\mathcal{C}}\mathbf{Ex} \subseteq \mathbf{EffPriced}_{\mathcal{C}}\mathbf{Ex} \subseteq \mathbf{Priced}_{\mathcal{C}}\mathbf{Ex}$.
(b) $\mathbf{EffPriced}_{\mathbf{FF}}\mathbf{Ex} \not\subseteq \mathbf{ValPriced}_{\mathbf{MF}}\mathbf{Ex}$.

Part (a) follows by definition. Part (b) can be proved using a technique similar to the one used in the proof of Theorem 19. Consider the class $\mathcal{L}$ and set A as defined in that proof. We modify that proof to make sure that for each e, if $\varphi_e \in \mathbf{FF}$, then for all $x > e$, A contains some element of the form $\langle x, y\rangle$ with $\varphi_e(\langle x, y\rangle) > 2^x$. This allows one to show that $\mathcal{L} \in \mathbf{EffPriced}_{\mathbf{FF}}\mathbf{Ex}$. Furthermore, one can fool any learner using only cost function from input, as it has to make a memory / mind change on some input $\langle x, y\rangle$ in A for each x. This allows us to construct a monotonic cost function h such that the memory / mind change of the learner is forced for a cost of 1 at every stage. The class $\mathcal{L}$ can be $\mathbf{EffPriced}_{\mathbf{FF}}\mathbf{Ex}$-learnt by memorising all x where $\langle x, y\rangle$, with $x \leq e$ appears in the input or $x > e$ and $\varphi_e(\langle x, y\rangle) > 2^x$.

Theorem 22. $\mathbf{ValPriced}_{\mathbf{FF}}\mathbf{Ex} \not\subseteq \mathbf{EffBPriced}_{\mathbf{MF}}\mathbf{Ex}$.

A construction similar to one used in the proof of Theorem 17 can be used to show this result. The class $\mathcal{L}$ used in Theorem 17 proof is in $\mathbf{ValPriced}_{\mathbf{FF}}\mathbf{Ex}$. However, it is not in $\mathbf{EffBPriced}_{\mathbf{MF}}\mathbf{Ex}$ as using recursion theorem we can construct price function φ_r (with bound r being given to the learner) such that φ_r searches for a mind change in each stage, and then puts the costs for these to sum up to one in each stage.

Theorem 23. $\mathbf{UniFPriced}_{\mathbf{FF}}\mathbf{Ex} \not\subseteq \mathbf{Priced}_{\mathbf{MF}}\mathbf{Ex}$.

Proof. Let $\mathcal{L} = \{\mathbb{N} - \{0\}\} \cup \{\{0\} \cup S : S \text{ is finite}\}$. Then, clearly $\mathcal{L}$ is in $\mathbf{UniFPriced}_{\mathbf{FF}}\mathbf{Ex}$ as witnessed by a learner which memorises every input element and conjectures an index for $\mathbb{N} - \{0\}$, if the input does not contain 0; otherwise the learner conjectures the canonical index for the finite input seen so far.

However, $\mathcal{L}$ is not $\mathbf{Priced}_f\mathbf{Ex}$-learnable for $f(x) = x$. Thus, $\mathcal{L} \notin \mathbf{Priced}_{\mathbf{MF}}\mathbf{Ex}$. □

Corollary 24. $\mathbf{UniFPriced}_{\mathbf{FF}}\mathbf{Ex} \not\subseteq \mathbf{EffPriced}_{\mathbf{MF}}\mathbf{Ex} \cup \mathbf{ValPriced}_{\mathbf{MF}}\mathbf{Ex}$.

Theorem 25. $\mathbf{UniFPriced_{MF}Ex} \not\subseteq \mathbf{Priced_{FF}Ex} \cup \mathbf{UniFPriced_{FF}Ex}$.

One can prove the above theorem by using a modification of the proof of Theorem 18, where we fix the cost function $\varphi_{e'} \in \mathbf{FF}$ for which the class $\mathcal{L}$ is not learnable. Then, for each e pad each language in $\mathcal{L}$ by $\langle e, i_e \rangle$, where $D_{i_e} \subseteq \{0, 1, 2, \ldots, e\}$ is the set of indices $j \leq e$ such that $\varphi_j \in \mathbf{MF}$. Then, for each e we only take one language padded by $\langle e, i_e \rangle$ which is not $\mathbf{Priced_{FF}Ex}$-learnt by the e-th learner (and similarly one language each for the i-th $\mathbf{UniFPriced_{FF}Ex}$-learner, for $i < e$), using the price function $f(\langle\langle s, i_s \rangle, x\rangle) = \varphi_{e'}(\langle\langle s, i_s \rangle, x\rangle) * 2^{\langle s, i_s \rangle}$. This makes the class not in $\mathbf{Priced_{FF}Ex} \cup \mathbf{UniFPriced_{FF}Ex}$. However, the class is in $\mathbf{UniFPriced_{MF}Ex}$, as the learner can use the padding part $\langle e, i_e \rangle$ of the input language and then learn it as the learner can learn it with respect to the monotonic functions given by the set D_{i_e}. Thus, for any monotonic function $\varphi_{e'}$, the learner will only have large cost for the padded functions $\langle e, i_e \rangle$, where $e < e'$.

Theorem 26. $\mathbf{ValPriced_{FF}Ex} \not\subseteq \mathbf{UniFPriced_{MF}Ex}$.

This result can be proven using a modification of the proof of Theorem 17. The class is not $\mathbf{UniFPriced_{MF}Ex}$-learnable, by modification of the diagonalisation in the proof of Theorem 17. Where, if $\mathbf{M}$ $\mathbf{UniFPriced_{MF}Ex}$-learns the class, then we construct in the diagonalisation additionally a monotonic function h which searches for mind changes (on enough inputs) for different languages corresponding to infinitely many W_e (in dovetail manner) in each stage: each time memory / mind changes are found, the cost is increased by 1 for learning the corresponding language. If it gets stuck at any point, then learner is not learning a corresponding cylindrification of a finite language from $\mathcal{L}_2$.

It is open whether $\mathbf{UniPriced_{MF}Ex}$ or $\mathbf{UniPriced_{FF}Ex}$ contain a class which cannot be $\mathbf{ItEx}$-learnt. Similarly, it is open whether $\mathbf{EffBPriced_{MF}Ex}$ or $\mathbf{EffBPriced_{FF}Ex}$ contain a class which cannot be $\mathbf{ItEx}$-learnt.

8 Uniformly Recursive Classes of Price Functions

Theorem 12 provided a class for which there is one single learner $\mathbf{M}$ which succeeds for all price functions for which the class is priced learnable. The next result shows that there are classes $\mathcal{L}$ for which this is impossible, although one can still find a single learner which succeeds for every price function from any enumerable class $\mathcal{C}$ of price functions which only consists of price functions with respect to which $\mathcal{S}$ is priced learnable. This is formulated more precisely in the statement of the following theorem.

Theorem 27. *There is an r.e. set A such that the class $\mathcal{L} = \{\mathbb{N} - \{x\} : x \notin A\}$ satisfies the following conditions:*

(1) *For every uniformly recursive class $\mathcal{C}$ of price functions, either there is a price function $f \in \mathcal{C}$ for which $\{\mathbb{N} - \{x\} : x \notin A\}$ is not* $\mathbf{Priced}_f\mathbf{Ex}$*-learnable or there is a single learner* $\mathbf{M}$ *which* $\mathbf{Priced}_f\mathbf{Ex}$*-learns $\{\mathbb{N} - \{x\} : x \notin A\}$ for all $f \in \mathcal{C}$;*

(2) *There is no single learner* $\mathbf{M}$ *which* $\mathbf{Priced}_f\mathbf{Ex}$*-learns* $\{\mathbb{N} - \{x\} : x \notin A\}$ *for all price functions* f *for which* $\{\mathbb{N} - \{x\} : x \notin A\} \in \mathbf{Priced}_f\mathbf{Ex}$.

(3) $\{\mathbb{N} - \{x\} : x \notin A\}$ *is not* $\mathbf{ItEx}$*-learnable.*

Proof. For the construction let $\psi_{d,k}$ be an acceptable enumeration of all numberings of price functions. Here, $\psi_{d,k}$ denotes the k-th function in the d-th enumeration of partial recursive functions. The construction uses markers with e-states. The e-state of a i-th marker sitting at time s on $x_{i,s}$ is $b_0 b_1 \ldots b_e$ where for $d \leq e$ the value $b_d = 1$ iff the first i values $\psi_{d,0}(x_{i,s}) \ldots$, $\psi_{d,i}(x_{i,s})$ are all defined and take a value above $2^{(i+1)^2+i}$. At the beginning the position $x_{i,0}$ of the marker x_i is i.

At stage s one looks for the least $i \leq s$ such that there is $j \in \{i+1, i+2, \ldots, i+s\}$ and $e \leq i$ such that the i-th marker x_i can increase its e-state lexicographically by moving from its current position $x_{i,s}$ to the position $x_{j,s}$. If such an i is found then one updates $x_{k,s+1} = x_{k+(j-i),s}$ for all $k \geq i$ and enumerates the places vacated by the markers into A, that is,

$$A_{s+1} = A_s \cup \{x_{i,s}, x_{i+1,s}, \ldots, x_{j-1,s}\}.$$

If no such i is found then all markers remain at the same position and $A_{s+1} = A_s$.

So A is the set of all places eventually vacated by markers. Note that if the marker increases its e-state for some $e < i$ then it also increases its e-state for $e = i$. Now one can see by induction that each marker x_e moves at most $2^{(e+1)^2} - 1$ times. The marker x_0 moves at most once as its e-state is either 0 or 1 and no higher-order marker makes it move. The marker x_1 makes up to three moves before and up to three moves after this move, resulting in 7 moves (note that $7 \leq 2^4 - 1$). Assuming now that x_e moves at most $2^{(e+1)^2} - 1$ times. Then x_{e+1} moves at most $2^{e+2} - 1$ before, between each two moves and after the last move of x_e plus the at most $2^{(e+1)^2} - 1$ moves of x_e; this gives in total $2^{(e+1)^2+e+2} - 1 \leq 2^{(e+2)^2} - 1$ moves of the marker x_{e+1}. Thus, by induction, every marker x_e moves at most $2^{(e+1)^2} - 1$ many times.

The next step is to show that for each e almost all markers have the same e-state, a fact known from the construction of maximal sets (see [11]). The reason is that if $e \leq i < j$ and x_i and x_j are both settled by stage s and reached their final e-states, then the e-state of x_i at s must be at least as large as the e-state of x_j: Otherwise x_i would move from $x_{i,s}$ to $x_{j,s}$ as the e-state of x_i at position $x_{j,s}$ is at least as large as the e-state of x_j at position $x_{j,s}$, since the e-state of x_i involves only the functions $\psi_{d,k}$ with $d \leq e$ and $k \leq i$ while the e-state of x_j involves the functions $\psi_{d,k}$ with $d \leq e$ and $k \leq j$ and $i < j$. Thus, as x_i does not move anymore, the e-state of x_j is at most the e-state of x_i and so the e-states form a descending sequence which eventually stabilises. Let B be the limit set of these bits, so $B(e)$ satisfies for almost all e that the final e-state of x_e has at position e the bit $B(e)$.

(1): Assume now that $(\psi_{e,k})_{k \in \mathbb{N}}$ is a uniformly recursive family of price functions. One makes a case distinction over the value of $B(e)$.

Case (a) $B(e) = 0$: some price function $f = \psi_{e,k}$ does not permit to $\mathbf{Priced}_f\mathbf{Ex}$-learn the class $\mathcal{L}$. The reason is that there is an i such that x_i does not move and the final e-state $b_0 b_1 \ldots b_e$ of x_i equals $B(0)B(1)\ldots B(e)$. Then for every $j > i$ and the stage s where x_j has reached the final state and position there is a $k \leq i$ such that $\psi_{e,k}(x_{j,s})$ is not above 2^{i^2+i}, as otherwise x_i would move at stage s or later to the position $x_{j,s}$. As this is true for all $j > i$ and the corresponding s (depending on j), it follows that there is a k such that, for infinitely many $x \notin A$, $\psi_{e,k}(x) \leq 2^{(i+1)^2+i}$. Let $f = \psi_{e,k}$. Assume by way of contradiction that $\mathbf{M}$ is a $\mathbf{Priced}_f\mathbf{Ex}$-learner for the class $\mathcal{L}$.

Subcase (aa): There is a finite sequence σ containing two non-elements x, y of A such that $\mathbf{M}$, while processing σ, does not update the memory on these two elements. Thus $\mathbf{M}$ cannot know whether these two elements were in the input or not. Then let T be a sequence consisting of all numbers different from x, y and let σ_x, σ_y be the sequences obtained by deleting either x or y, respectively, from σ. Now the learner $\mathbf{M}$ has, after σ_x and σ_y, the same memory and will therefore on the texts $\sigma_x T$ and $\sigma_y T$ converge to the same hypothesis; thus $\mathbf{M}$ fails to learn at least one of the sets $\mathbb{N} - \{x\}$ and $\mathbb{N} - \{y\}$.

Subcase (ab): There is a text for some language L in the class such that $\mathbf{M}$ updates its memory on all but at most one element of $L - A$. Then $\mathbf{M}$ incurs on $\psi_{e,k}$ infinitely often a cost of $1/(2^{(i+1)^2+i} + 1)$ or more which causes the overall costs to go to ∞; hence $\mathbf{M}$ is not $\mathbf{Priced}_f\mathbf{Ex}$-learning the set L in this case.

As one of the two subcases must apply, $\mathbf{M}$ cannot be a $\mathbf{Priced}_f\mathbf{Ex}$-learner for the given class $\mathcal{L}$.

Case (b) $B(e) = 1$: There is a learner $\mathbf{M}$ which $\mathbf{Priced}_f\mathbf{Ex}$-learns $\{\mathbb{N} - \{x\} : x \notin A\}$ for every $f \in \{\psi_{e,k} : k \in \mathbb{N}\}$, without knowing which of these f the price function is. Let d be so large that all the markers $x_d, x_{d+1}, \ldots$ eventually stabilise on an e-state equal to $B(0)B(1)\ldots B(e)$. The learner $\mathbf{M}$ now can memorise a datum x iff there is a stage $s > x$ such that either $x \in \{x_{0,s}, \ldots, x_{d,s}\}$ or there exists a $k > d$ with $x_{k,s} = x$ and the e-state of the marker x_k at stage s is $B(0)B(1)\ldots B(e)$. The learner $\mathbf{M}$ can ignore a datum x iff there is a stage s such that $x \in A_s$. Note that at least one of these conditions is eventually satisfied, as every x is either enumerated into A or has eventually a marker sitting on it from $x_0, \ldots, x_d$ or has eventually a marker sitting on it whose $(e+1)$-stage converges to $B(0)B(1)\ldots B(e)$; so it depends on what happens first in order for the learner to decide whether to memorise the datum or not; as it only ignores data enumerated into A, it never makes an error by ignoring a datum. Furthermore, as it reads all data except one in the complement of A, it can check infinitely often on whether to update the hypothesis. When reading a datum z, the learner updates the hypothesis to $\{N - \{y\}\}$ whenever y changes, where y is the minimal non-element of $E \cup \{x : \exists s \in E : x \in A_s\}$ and E is the set of data memorised. These updates happen only finitely often, as the algorithm converges to $\mathbb{N} - \{y\}$ for the least y not in the text and it is therefore an explanatory learner.

So it remains to show that the costs are bounded. So assume that $f = \psi_{e,k}$ is the price function used. If the learner memorises a datum $x_{h,s}$ with $h \leq \max\{k, d\}$ then an unknown amount of cost is incurred; this happens, however, only finitely often. If the learner memorises a datum $x_{h,s}$ such that $h > \max\{k, d\}$ and the current e-state being $B(0)B(1)\ldots B(e)$, then $f(x_{h,s}) > 2^{-(h+1)^2-h}$; this happens for the marker x_h at most $2^{(h+1)^2}$ times and therefore at most cost 2^{-h} are incurred. So the overall cost is some constant plus the sum of all 2^{-h} which results in a finite real number. Therefore the price condition for f is kept; it follows that in

Case (b) the class has indeed a learner **M** which $\mathbf{Priced}_f\mathbf{Ex}$-learns $\mathcal{L}$ for any price function $f \in \{\psi_{e,k} : k \in \mathbb{N}\}$.

(2): Assume now by way of contradiction that there is a learner **M** which $\mathbf{Priced}_f\mathbf{Ex}$-learns $\{\mathbb{N} - x : x \notin A\}$ for all price functions f where this is possible (without knowing what f is). Then one can argue that there is a recursive text T for some set $\mathbb{N} - \{x\}$ with $x \notin A$ on which **M** updates the memory on infinitely many $y \in A$; otherwise A would be recursive, as the set of $\{x\}$ on which the learner makes an update is r.e. and contains almost all elements of $\overline{A}$. This is a contradiction as by construction A is not recursive. Hence there is an infinite recursive set $R \subseteq A$ such that the learner **M** on the text T updates the memory when processing an $y \in R$. Now let $f(y) = 0$ for $y \in R$ and $f(y) = 2^y$ for $y \notin R$; the learner **M** pays infinitely often the price of 1 for archiving the elements of R when processing T and therefore incurs infinite cost. Thus, the learner **M** does not $\mathbf{Priced}_f\mathbf{Ex}$-learns $\{\mathbb{N} - \{x\} : x \notin A\}$. However, there is a learner **N** for the class which just memorises the elements outside R – those in R are not relevant – and therefore $\mathbf{Priced}_f\mathbf{Ex}$-learns $\{\mathbb{N} - x : x \notin A\}$ by always conjecturing $\mathbb{N} - \{y\}$ for the least $y \notin R$ which has not yet been memorised.

(3): This is a corollary of **(2)**, as an iterative learner would satisfy the assumptions made on **M** in this case. Thus the class is not iteratively learnable. ☐

9 Conclusions

In this paper we considered a generalization of memory limited learning, called priced learning, where the learner can update its memory based on any datum it receives, but it has a cost $\frac{1}{f(x)+1}$ associated with it, where f is the price function. The learning is said to be successful if the overall cost in the learning process is finite. We gave a characterization that the classes learnable as above with respect to some cost function are exactly the classes learnable by a set-driven learner. We also gave complete picture of when different price functions lead to different learnable classes. In addition we considered when priced learning is possible for all price functions from a class of functions, and when a learner can be effectively found from an index for a price function. It is open at present whether there exists a single learner which learns a non iteratively learnable class from all price functions which are monotonically non-decreasing and unbounded. Similar question is also open for the price functions having each number in the range only finitely often. These questions are open even for the case when the learner is given a bound on an index for the price function.

Acknowledgements: We thank the referees for several helpful comments.

References

1. Lenore Blum and Manuel Blum. Toward a mathematical theory of inductive inference. *Information and Control*, 28:125–155, 1975.

2. Manuel Blum. A machine independent theory of the complexity of recursive functions. *Journal of the Association of Computing Machinery*, 14:322–336, 1967.
3. John Case, Sanjay Jain, Steffen Lange and Thomas Zeugmann. Incremental concept learning for bounded data mining. *Information and Computation*, 152:74–110, 1999.
4. Rusins Freivalds, Efim Kinber and Carl H. Smith. On the impact of forgetting on learning machines. *Journal of the ACM*, 42:1146–1168, 1995.
5. Mark Fulk. Prudence and other conditions on formal language learning. *Information and Computation*, 85:1–11, 1990.
6. E. Mark Gold. Language identification in the limit. *Information and Control*, 10:447–474, 1967.
7. Efim Kinber and Frank Stephan. Language learning from texts: mind changes, limited memory and monotonicity. *Information and Computation*, 123:224–241, 1995.
8. Steffen Lange and Thomas Zeugmann. Incremental learning from positive data. *Journal of Computer and System Sciences*, 53:88–103, 1996.
9. Daniel Osherson, Michael Stob and Scott Weinstein. Learning strategies. *Information and Control*, 53:32–51, 1982.
10. Daniel Osherson, Michael Stob and Scott Weinstein. *Systems That Learn, An Introduction to Learning Theory for Cognitive and Computer Scientists.* Bradford — The MIT Press, Cambridge, Massachusetts, 1986.
11. Hartley Rogers. *Theory of Recursive Functions and Effective Computability.* McGraw-Hill, 1967. Reprinted by MIT Press in 1987.
12. Gisela Schäfer-Richter. Uber Eingabeabhängigkeit und Komplexität von Inferenzstrategien. Ph.D. Thesis, RWTH Aachen, 1984.
13. Kenneth Wexler and Peter W. Culicover. *Formal Principles of Language Acquisition.* Cambridge, Massachusetts, The MIT Press, 1980.
14. Rolf Wiehagen. Limes-Erkennung rekursiver Funktionen durch spezielle Strategien. *Elektronische Informationsverarbeitung und Kybernetik* (EIK), 12:93–99, 1976.