Convergence of preference functions

Achim Jung, Jonathan E. Rowe[1]

School of Computer Science, University of Birmingham, Birmingham, B15 2TT, United Kingdom
A.Jung@cs.bham.ac.uk, J.E.Rowe@cs.bham.ac.uk

Abstract

A *preference function* is a function which selects a subset of objects based on (partial) information. As information increases, different objects may be selected. We examine conditions under which the selection of objects converges to the choice that would be made if full information were available, making use of tools from domain theory. The work is motivated by previous research on *co-evolutionary algorithms* in which an evolving population of agents interact with each other and, it is hoped, produce better and better quality behaviour. The formalisation of how quality can be measured in this context has introduced the concept of a *convex* preference function (or "solution concept"). We simplify and extend the scope of this previous work, examining the relationship between convexity and convergence properties.

Keywords: preference functions, convergence, continuity, convex function, co-evolutionary algorithms

1. Introduction

We imagine that there is some process by which we produce, and accumulate, information of some kind. Based on the current information acquired, we make decisions about which members of a possible solution set we prefer. Our preference changes as our information increases and perhaps not in a particularly coherent way: we may change our minds many times about whether or not some solutions are to be preferred. What we would like, though, is some assurance that as our information increases, so our preferences converge to the "right" answer, namely, that set we would select if we had access to all relevant information. In this paper we formalise, and study, such convergence properties.

We will model our states of information as a partially ordered set A, so that $a \leq b$ means that state b contains at least as much information as state a. The aggregation of information will be given by a directed subset $S \subseteq A$. By "directed" is meant that if a and b are in S, then there is some $c \in S$ greater than

[1]Corresponding author

both a and b. So if information is obtained independently, there is some point at which it is brought together. Our preferences will concern objects (solutions) in a set X. Thus we are interested in functions $f : A \to 2^X$, which we will call *preference functions*. Given a directed set $S \subseteq A$, we will be asking under what conditions $f(S)$ converges to $f(\bigvee^{\uparrow} S)$, where $\bigvee^{\uparrow} S$ is the limit (directed join) of S. To simplify matters, we will assume that all directed sets have joins, and further that A is in fact *algebraic*. This means, roughly, that any state of information in A can be approximated by a directed set of "finite" pieces of information. We give formal definitions in Section 2.

A typical class of examples, which will occur throughout, concerns having a set of objects X and a number of tests T which can be performed on the objects. For example X may be a set of algorithms and T may be a set of benchmark problems. Information is generated when we apply a test to an object and get some result. Given what we learn from a set of such experiments, we would like to choose a subset of objects which we prefer (e.g. the algorithms which we think are performing best). In this case $A = 2^{X \times T}$.

A related example comes from non-monotonic reasoning, in which the beliefs of an agent change according to the knowledge so far accumulated. In this case X may be thought of as a set of possible propositions, and $f(a)$ the particular set of beliefs corresponding to the state of knowledge a. This is related to Putnam's and Gold's concept of "limiting recursion", in which a set is *limit recursive* if there is a decision procedure for the set which is allowed to change its mind a finite number of times [1, 2].

Another class of examples arises from observing a dynamical system as it evolves, and making decisions based on the trace of states seen. A situation which combines this with the previous ideas arises in the study of *co-evolutionary* algorithms [3]. Here we have a population of agents (or strategies) which interact with each other (usually via some game matrix). The population evolves, over time, as a function of the relative success agents have in their interactions. The idea in using such algorithms is the hope that more advanced solutions to certain problems will be evolved through this process. However, just because an agent performs well in some cases, does not mean it will in others, and so the decision as to which agents are doing well at any particular time is a difficult one. In this context, researchers such as Ficici [4, 5], Bucci [6, 7] and Popovici [8] have developed a theory of *solution concepts* to study this question. Our current work arises as an attempt to simplify and generalise this work.

Within the co-evolutionary algorithm literature, the formalism used is rather more specific and complex than ours. Typically, there will be some set of potential solutions (often called *complexes*), which is denoted C. A certain (fixed) subset $D \subseteq C$ is given which specifies which complexes are, in general, to be preferred. The set D is called the *solution concept*. It is then supposed that there is a method $c : A \to 2^C$ which, given a state of information $a \in A$, returns a set of possible complexes. One is then interested in those which are in the solution concept, and defines $d(a) = c(a) \cap D$ (which is assumed to be always non-empty). In addition, there is then some heuristic function $h : 2^C \to C$ which chooses which of the preferred constructed complexes will actually be observed. Complexes themselves may have a variety of structures, for example they may be subsets of some set of components, or they may be probability distributions over a set of strategies. For this paper we would like to abstract from such detail, which is possible because d (or $h \circ d$) can be seen as a preference function in our sense.

A further complication seemingly arises from the fact that the co-evolutionary literature allows the function h to be dependent on the sequence by which a state of information a was attained. In [4, 5], Ficici furthermore proposes that the construction d could be similarly history dependent. These approaches, however, are easily captured in our framework because if A is an algebraic domain (in the sense made precise below) then so is the set of monotone sequences over A, ordered componentwise. In other words, history-dependence can be encoded in the way the elements of A are interpreted as "states of knowledge."

Finally, we would like to note that Ficici et al typically assume that all sets are finite and indeed this is essential for some of their definitions. However, since our aim is to study convergence properties of preference functions, we must remove this restriction. Indeed, one of the contributions of this paper is to open up the possibility of co-evolutionary systems over an infinite set of strategies, and to consider some of the new issues this raises. Also, the notion of "compact element," which underlies the definition of an algebraic domain, is a very good substitute for concrete notions of finiteness, and so it is likely that finiteness restrictions are expressible in our context though we do not explore this here.

2. Notation and definitions

Let $(A, \leq)$ be a partially-ordered set. A non-empty subset $S \subseteq A$ is called *directed* if whenever $a, b \in S$, there exists some $c \in S$ greater than both a and b. If a directed set S has a least upper bound (supremum) in A it is denoted by $\bigvee^{\uparrow} S$, where the arrow reminds us that S is directed. We call $\bigvee^{\uparrow}$ a *directed join.* $(A, \leq)$ is called a *directed complete partial order* or *dcpo* if all directed subsets have a supremum in A.

An element $c \in A$ is called *compact* if for any directed set $S \subseteq A$ with $c \leq \bigvee^{\uparrow} S$ there exists some $s \in S$ with $c \leq s$. We denote the set of all compact elements less than or equal to an element $a \in A$ by K_a. We call $(A, \leq)$ an *algebraic dcpo* or an *algebraic domain* if $a = \bigvee^{\uparrow} K_a$ for all $a \in A$. Note that by our convention this definition subsumes the requirement that K_a is always a directed set in an algebraic dcpo.

A prime example of an algebraic dcpo is the power set 2^X of subsets of some set X, ordered by set inclusion. Suprema are given by set union. Compact elements are the finite subsets. This is an algebraic domain because any set is the directed union of its finite subsets. In fact, 2^X is a complete lattice as all suprema of subsets exist (realised by set union). A more typical example of a dcpo is given by the set of partial functions from $\mathbb{N}$ to $\mathbb{N}$, ordered by graph inclusion. While directed joins exist (and are given by the union of the graphs of the functions) general suprema do not since the union of two graphs may no longer be single-valued. This, too, is an algebraic dcpo with the compact elements given by partial functions with finite graphs.

In the following, A is an algebraic dcpo, and X is some set. We consider convergence properties of functions $f : A \to 2^X$. The algebraicity of A will be important but that of 2^X will not; it will be best to think of 2^X very concretely as a powerset. As described in the introduction, the elements of A are meant to represent "states of information" while the elements of X are objects (e.g., algorithms) which are "preferred" or not at a given state, as expressed by a "preference function" $f : A \to 2^X$. A directed set in A, finally, should be thought of as a "process of accumulating information." If only one agent is involved in this, then it would be reasonable to assume that the directed set has the shape of a *chain* $a_1 \leq a_2 \leq a_3 \leq \cdots$, but for our results below there is no need to make this restriction.

We make use of a number of different topologies on algebraic dcpos:

Discrete topology All subsets are considered open.

d-topology A set U is open if whenever $\bigvee^{\uparrow} S \in U$, there is some $s \in S$ with $s \in U$. We say that such sets are inaccessible by directed joins.

Scott topology A set is open if it is inaccessible by directed joins and is upper closed (i.e. if $a \in U$ and $a \leq b$ then $b \in U$). A basis for this topology is the collection of all sets of the form $\uparrow k$, for compact elements k.

Lawson topology A sub-basis for the Lawson topology is given by the collection of all sets of the form $\uparrow k$ and $(\uparrow k)^c$ for compact elements k. It includes all Scott-opens.

For further information, the reader is referred to standard texts such as [9, 10] though this paper will make very little use of the theory of (algebraic) domains.

3. Convergent preference functions

We begin by defining what we mean for a preference function to be convergent. Our definition asks that for any object which is preferred in the limit of a process of information accumulation, there must be some point in the process at which we find it and continue to prefer it from then on. Similarly, if there is an object we do not prefer in the limit, there must be a point in the process beyond which we never prefer it.

Definition 1. *A function $f : A \to 2^X$ is* convergent *if the following holds for all $x \in X$ and directed sets $S \subseteq A$:*

1. *If $x \in f(\bigvee^{\uparrow} S)$ then there exists $s \in S$ such that*

$$s \leq b \leq \bigvee\nolimits^{\uparrow} S \Longrightarrow x \in f(b)$$

for all $b \in S$.

2. *If $x \notin f(\bigvee^{\uparrow} S)$ then there exists $s \in S$ such that*

$$s \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \notin f(b)$$

for all $b \in S$.

Example (Pareto dominance)

Consider a set of objects, X, and a set of tests, T, for example, a set of algorithms and a set of benchmark problems. We suppose there is an evaluation function $e : X \times T \to \mathbb{R}$ which evaluates objects using the tests, for example, $e(x,t)$ could be the amount of resource consumed by algorithm x in the course of solving problem t. At any given time, we assume that we have the results of some of these evaluations, that is, a "state of information" is given by the subset a of $X \times T$ for which the values $e(x,t)$ have been established. Thus, couching the situation in the language of this paper, the algebraic domain A is $2^{X \times T}$. In order to describe the preference function, for $a \in A$, $x, y \in X$, let

$$C_a(x,y) = \{t \in T \mid (x,t) \in a \text{ and } (y,t) \in a\}$$

be the set of tests for which we have the results for both x and y at stage a. Now say that "x is dominated by y at stage a" if

$$\begin{array}{ll} (i) & \forall t \in C_a(x,y).e(x,t) \leq e(y,t), \text{ and} \\ (ii) & \exists t \in C_a(x,y).e(x,t) < e(y,t). \end{array}$$

For $a \in A$, say that "x occurs in a" (and write $x \in o(a)$) if there exists at least one $t \in T$ such that $(x,t) \in a$. Now say that "x is preferred at stage a" if $x \in o(a)$ and is not dominated by any other object occurring in a. We then let $f(a)$ be the set of all objects preferred at stage a. Under the assumption that X is finite we show that $f : A \to 2^X$ is convergent.

To start, let S be a directed set in A and $a = \bigvee^{\uparrow} S$. To establish the first half of convergence, assume $x \in f(a)$ for some $x \in X$. By definition, x is not dominated by any other object occurring in a, which means that for any such $y \in o(a)$, either $e(x,t) > e(y,t)$ for some $t \in C_a(x,y)$, or $e(x,t) = e(y,t)$ for all $t \in C_a(x,y)$. The latter situation is not interesting because it will be true at all stages $a' \subseteq a$ and so in particular at the stages mentioned in S; what we really need to show is that the former situation is established at some stage $s \in S$ already, and it is here where we need that X is finite. Note that any $(y,t) \in a$ must appear at some stage $s \in S$ already because $\bigvee^{\uparrow} S = \bigcup S$ in the powerset $A = 2^{X \times T}$. So for any $y \in o(a)$ such that $e(x,t) > e(y,t)$ for some $t \in C_a(x,y)$ choose a particular test t_y witnessing this situation and then pick s_{x,t_y} and s_{y,t_y} in S such that $(x,t_y) \in s_{x,t_y}$ and $(y,t_y) \in s_{y,t_y}$. Since X is finite, there are only finitely many such elements to be chosen from S, and we find an upper bound s for them as S is directed. By construction, x is non-dominated by such y at any stage above s, and in particular at all $b \in S$ with $s \subseteq b \subseteq a$.

The second part of Definition 1 is easier to establish and does not rely on the finiteness of X. If x occurs in a but $x \notin f(a)$ then x is dominated by some

$y \in o(a)$. By definition this means that $e(x,t) < e(y,t)$ for some particular $t \in C_a(x,y)$. As before, there must then exist a stage $s \in S$ at which both (x,t) and (y,t) are present. Since $s \subseteq a$ we know that for any other tests $t' \in C_a(x,y)$ and for $s \subseteq b \subseteq a$ we must have $e(x,t') \leq e(y,t')$, which means that x is dominated by y at b and hence not preferred.

In the case where X is infinite we may not get convergence as the following example shows. Suppose $X = \mathbb{N}$ and $T = \{0,1\}$. Our evaluation function is $e(1,0) = 1$, $e(1,1) = 0$ and $e(i,j) = j$ for $i > 1$. We test each object in turn in ascending numerical order, performing first test 1 and then test 0. Initially, having only tested object 1, it is selected by f. However, each time we evaluate a new object k with test 0, we exclude object 1 as it appears dominated by the new object. When we then evaluate k on test 1, however, x is now non-dominated once more. In the limit, x is non-dominated, but this cannot be established by a finite set of tests.

One can wonder whether the failure of convergence in the infinite case is the result of a poor choice of the notion of dominance. As an alternative, one could add to conditions (i) and (ii) above the stipulation

$$\forall t \in T.(x,t) \in a \Longrightarrow (y,t) \in a$$

meaning that y is accepted as better than x only if it has been tested at least as widely as x. However, as closer analysis shows, the resulting preference function is only convergent if T is finite, though X can now be infinite. It is a challenge to come up with meaningful notions of dominance and preference which yield a convergent preference function in the case where both X and T are allowed to be infinite. One proposal is to fix some threshold of acceptable performance θ and then define

$$f(a) = \{x \text{ occurs in } a \,|\, e(x,t) \geq \theta \text{ for all } (x,t) \in a\}$$

We will show in Section 6 that this function is indeed convergent.

Example (worst case behaviour)

Let X be a finite set of objects and T be a (possibly infinite) set of tests, and suppose we have an evaluation function $e : X \times T \to \mathbb{R}$. Define a preference function $f : 2^{X \times T} \to 2^X$ as follows:

$$f(a) = \arg\min_{x \in X} \sup\{e(x,t) \,|\, (x,t) \in a\}$$

That is, for each object we find the maximum (or supremum) of test results across all tests for which it has been evaluated, and then select the object (or objects) minimising this supremum. This kind of minimax preference function would seem well suited to comparing algorithms on how well they perform in the worst (known) case. For example, we could measure the time complexity of algorithms on a set of benchmark problems, or the quality of solutions they find within a fixed amount of time (see chapter 4 of [11] for further investigation of such empirical testing). However, in general it is not convergent. Consider the case of two objects $X = \{x, y\}$ and an infinite set of tests indexed by the natural numbers, $T = \{t_1, t_2, t_3, \ldots\}$. Suppose that both x and y evaluate to n on test t_n. It is evident that $f(X \times T) = \{x, y\}$ since the suprema of their evaluations across all tests is equal to ∞ for both x and y, but whenever x or y is preferred in some *finite* set of tests, there is always a larger set of tests in which they are rejected. Thus convergence fails for the directed set of finite subsets of $X \times T$.

In both examples, compact elements of A (appearing as finite subsets of $X \times T$) play a prominent role in the argument. The following lemma establishes that this is no coincidence.

Lemma 1. *The following are equivalent for preference functions $f : A \to 2^X$.*

1. *f is convergent.*
2. *f satisfies conditions (1) and (2) of Definition 1 for directed sets consisting of compact elements only.*
3. *f satisfies the two conditions:*
   (a) *If $x \in f(a)$, then there exists $c \in K_a$ such that $c \leq b \leq a \Rightarrow x \in f(b)$, for all $b \in A$.*
   (b) *If $x \notin f(a)$, then there exists $c \in K_a$ such that $c \leq b \leq a \Rightarrow x \notin f(b)$, for all $b \in A$.*
4. *For all directed sets $S \subseteq A$ and all $x \in X$, f satisfies the two conditions:*
   (a) *If $x \in f(\bigvee^{\uparrow} S)$ then there exists some $a \in S$ such that $x \in f(a)$.*
   (b) *If $x \notin f(\bigvee^{\uparrow} S)$ then there exists some $a \in S$ such that $x \notin f(a)$.*

Proof

It is clear that $1 \Rightarrow 2$. For $2 \Rightarrow 3$ assume $x \in f(a)$ for some $a \in A$, $x \in X$. Since $a = \bigvee^{\uparrow} K_a$ there exists $c \in K_a$ such that $c \leq b \leq a \Rightarrow x \in f(b)$, for all $b \in K_a$. Now consider any $w \in A$ with $c \leq w \leq a$. If $x \notin f(w)$ then there exists some $v \in K_w$ such that $v \leq b \leq w \Rightarrow x \notin f(b)$ for all $b \in K_w$. Now $c \leq w$ implies $c \in K_w$ so we can find some $z \in K_w$ greater than both c and v. But $z \in K_w \subseteq K_a$ with $c \leq z$ implies $x \in f(z)$ whereas $v \leq z$ implies $x \notin f(z)$. Consequently, it must be the case that $x \in f(w)$ for any $w \in A$ with $c \leq w \leq a$.

The case $x \notin f(a)$ is similar, so next consider $3 \Rightarrow 1$. For this, let $S \subseteq A$ be directed and let $x \in f(\bigvee^{\uparrow} S)$. Then there exists $c \in K_{\bigvee^{\uparrow} S}$ such that $c \leq b \leq \bigvee^{\uparrow} S$ implies $x \in f(b)$ for all $b \in A$. Since S is directed, there is some $s \in S$ with $s \geq c$, and then it follows that $s \leq b \leq \bigvee^{\uparrow} S$ implies $x \in f(b)$ for all $b \in S$.

The case $x \notin f(\bigvee^{\uparrow} S)$ is again similar.

We finally show $1 \Leftrightarrow 4$. It is clear from the definitions that $1 \Rightarrow 4$. For the converse, consider the case $x \in f(\bigvee^{\uparrow} S)$. Suppose that for all $a \in S$ there exist some $b \in S$ with $a \leq b$ and $x \notin f(b)$. That is, if we let $Z_a = \{b \in S \,|\, a \leq b \text{ and } x \notin f(b)\}$, then we are assuming that Z_a is non-empty for all $a \in S$. Now consider the set $B = \bigcup_{a \in S} Z_a$. We claim first that this is a directed set. Taking any two $b_1, b_2 \in B$ we can certainly find $c \in S$ bigger than both. Then by assumption, Z_c is non-empty, and so there is an element of B bigger than

c. Secondly, we claim that $\bigvee^{\uparrow} B = \bigvee^{\uparrow} S$. For any element of S our assumption is that there is an element of B greater than it. Therefore $\bigvee^{\uparrow} B$ is an upper bound for S and since $B \subseteq S$, it must be the least upper bound. We now have a directed set $B \subseteq A$, with $x \in f(\bigvee^{\uparrow} B)$. Our condition states that there therefore exists some element $s \in B$, such that $x \in f(s)$, contradicting the definition of B. Our assumption is therefore false, and it follows that there must exist some $a \in S$ for which Z_a is empty.

The case $x \notin f(\bigvee^{\uparrow} S)$ is, yet again, similar. □

4. Uniform convergence

For f to be convergent, then for each individual object $x \in X$ there must be a point in the process by which we have made up our minds as to whether or not we prefer x. But this point may be different for each object. We provide a stronger condition, *uniform* convergence, for which the same point will do for all objects in X.

Definition 2. *A function $f : A \to 2^X$ is* uniformly convergent *if for any directed set $S \subseteq A$, there exists an $a \in S$ such that*

$$a \leq b \leq \bigvee^{\uparrow} S \Longrightarrow f(b) = f(\bigvee^{\uparrow} S)$$

for all $b \in S$.

We have the following alternative characterisations, in analogy to Lemma 1 and with analogous proofs.

Lemma 2. *The following are equivalent for preference functions $f : A \to 2^X$.*

1. *f is uniformly convergent.*
2. *f satisfies the condition of Definition 2 for directed sets consisting of compact elements only.*
3. *For all $a \in A$, there exists $c \in K_a$ such that $c \leq b \leq a \Longrightarrow f(b) = f(a)$ for all $b \in A$.*
4. *For all directed sets $S \subseteq A$, there exists some $a \in S$ such that $f(a) = f(\bigvee^{\uparrow} S)$.*

Example

It is clear that if $f : A \to 2^X$ is convergent, and X is finite, then f is uniformly convergent. The Pareto dominance example on a finite set of objects from above is therefore an example of uniform convergence.

5. Convergence and continuity

We now ask how we can characterise convergence and uniform convergence in terms of the continuity of f with respect to various topologies.

Theorem 3. *Let A have the d-topology and 2^X have the Lawson topology. Then $f : A \to 2^X$ is convergent if and only if it is continuous w.r.t. these topologies.*

Proof
Assume that f is convergent. The Lawson topology on 2^X is generated by sets of the form $\uparrow\{x\}$ and $(\uparrow\{x\})^c$ for elements $x \in X$. We show that the inverse image of such sets are d-open in A. So let $S \in A$ be directed. Firstly, suppose $\bigvee^{\uparrow} S \in f^{-1}(\uparrow\{x\})$. That is, $x \in f(\bigvee^{\uparrow} S)$. Then there exists $c \in S$ such that $x \in f(c)$. That is $c \in f^{-1}(\uparrow\{x\})$. Therefore $f^{-1}(\uparrow\{x\})$ is d-open. Secondly, suppose $\bigvee^{\uparrow} S \in f^{-1}((\uparrow\{x\})^c)$. That is, $x \notin f(\bigvee^{\uparrow} S)$. There there exists $c \in S$ such that $x \notin f(c)$. That is, $c \in f^{-1}((\uparrow\{x\})^c)$. So $f^{-1}((\uparrow\{x\})^c)$ is d-open. This shows that f is continuous.

Conversely, assume that f is continuous. Let $S \subseteq A$ be directed. Suppose $x \in f(\bigvee^{\uparrow} S)$. That is, $\bigvee^{\uparrow} S \in f^{-1}(\uparrow\{x\})$, which is d-open, by continuity of f. Therefore, there exists some $s \in S$ such that $s \in f^{-1}(\uparrow\{x\})$. That is, $x \in f(s)$. Similarly, suppose $x \notin f(\bigvee^{\uparrow} S)$. Then $\bigvee^{\uparrow} S \in f^{-1}(\uparrow\{x\}^c)$, which is d-open. Therefore, there must be some $s \in S$ for which $s \in f^{-1}(\uparrow\{x\}^c)$, which means $x \notin f(s)$. Appealing to Lemma 1 completes the proof. $\square$

Recall that a *convergent net* in a topological space T is a function $n: S \to T$ such that S is directed and for every open set O of T, $n(s)$ is *eventually in* O, which in turn means that there exists $s_O \in S$ such that $n(s') \in O$ for all $s' \geq s$. The following is now an immediate consequence of the preceding theorem:

Corollary 4. *A preference function $f : A \to 2^X$ is convergent if and only if for all directed $S \in A$, $f(S)$ is a convergent net w.r.t. the Lawson topology.*

Theorem 5. *Let A have the d-topology and 2^X have the discrete topology. Then $f : A \to 2^X$ is uniformly convergent if and only if it is continuous w.r.t. these topologies.*

Proof
Assume that f is uniformly convergent. Let $U \subseteq 2^X$ be any set. Let $S \subseteq A$ be directed with $\bigvee^{\uparrow} S \in f^{-1}(U)$. That is, $f(\bigvee^{\uparrow} S) \in U$. There exists $s \in S$ such that $f(s) = f(\bigvee^{\uparrow} S)$ and so $s \in f^{-1}(U)$. Thus $f^{-1}(U)$ is d-open, and f is continuous.

Conversely, suppose f is continuous. Let $S \subseteq A$ be directed. The singleton set $U = \{f(\bigvee^{\uparrow} S)\}$ is open in 2^X. Therefore, by continuity, $f^{-1}(U)$ is d-open, with $\bigvee^{\uparrow} S \in f^{-1}(U)$. Therefore, there is some $s \in S$ for which $s \in f^{-1}(U)$ and so $f(s) = f(\bigvee^{\uparrow} S)$. Appealing to Lemma 2 completes the proof.

□

Corollary 6. *A function $f : A \to 2^X$ is uniformly convergent if and only if for all directed $S \in A$, $f(S)$ is a convergent net w.r.t. the discrete top-logy.*

Proof
Follows from previous theorem and definition of a convergent net. □

6. Convex functions

One of the main contributions made by Ficici was the definition of *convex* preference functions [4, 5]. Ficici actually called them *monotonic* which is misleading, since they are not monotone according to the standard definition for maps between partially ordered sets (in fact they include monotone functions as a proper subset). Bucci [7] noticed that these functions have the property that for any $x \in X$, the set $f^{-1}(\{x\})$ is convex as a subset of A, and this motivates our terminology.

Definition 3. *A function $f : A \to 2^X$ is* convex *if for all $a, b, c \in A$ with $a \leq b \leq c$*

$$f(a) \cap f(c) \subseteq f(b)$$

The idea here is that although, in general, we may change our minds as to whether or not we prefer objects many times, for convex functions we change our minds at most once. When an object has been preferred but is now discarded, it remains discarded forever. Ficici used this idea to try to maintain some idea of progress as information increases (see the following section). Unfortunately, in the case when A is infinite, convexity is not sufficient to ensure convergence. We require some extra "finiteness" conditions as follows.

Definition 4. *A function $f : A \to 2^X$ is* finite *if the following conditions are satisfied for all $a \in A$ and $x \in X$:*

1. *If $x \in f(a)$ then there exists $c \in K_a$ such that $x \in f(c)$.*
2. *If $x \notin f(a)$, then for all $c \in K_a$ with $x \in f(c)$ there exists some $k \in K_a$ with $k > c$ and $x \notin f(k)$.*

Lemma 7. *If $f : A \to 2^X$ is convergent, then it is finite.*

Proof
Part 1 follows directly from the definition, observing that $a = \bigvee^{\uparrow} K_a$. For part 2, suppose $x \notin f(a)$ and there is some $c \in K_a$ with $x \in f(c)$. Because f is convergent, we can find some $c' \in K_a$ such that $c' \leq b \leq a \Longrightarrow x \notin f(b)$ for

all $b \in K_a$. Let $k \in K_a$ be greater than c and c'. Then $k > c$ and $x \notin f(k)$ as required. $\square$

Theorem 8. *Let $f : A \to 2^X$ be a convex function. Then f is convergent if and only if it is finite.*

Proof

From the previous lemma it is clear that convergent implies finite. So now assume that f is finite and convex.

Firstly, suppose that $x \in f(\bigvee^{\uparrow} S)$ for some $x \in X$ and some directed $S \subseteq A$. Then there exists $c \in K_{\bigvee^{\uparrow} S}$ with $x \in f(c)$. Since S is directed, there is some $s \in S$ with $c \leq s$. By convexity, $c \leq s \leq b \leq \bigvee^{\uparrow} S$ implies $x \in f(b)$ for any $b \in S$.

Secondly, suppose that $x \notin f(\bigvee^{\uparrow} S)$. If there are no points $a \leq \bigvee^{\uparrow} S$ with $x \in f(a)$ then we are done. Otherwise, there must be some compact $c \leq \bigvee^{\uparrow} S$ with $x \in f(c)$. By the second finiteness condition, there exists another compact $k \in K_{\bigvee^{\uparrow} S}$ with $k > c$ and $x \notin f(k)$. By convexity, $c \leq k \leq b \leq \bigvee^{\uparrow} S$ implies $x \notin f(b)$ for all $b \in S$. $\square$

Example

Let X be a (possibly infinite) set of objects and T an infinite number of tests, and assume there is an evaluation function $e : X \times T \to [0, 1]$. Given $0 < \theta < 1$ define a preference function $f : 2^{X \times T} \to 2^X$ by

$$f(a) = \{x \text{ appearing in } a \,|\, e(x, t) \leq \theta \text{ for all } (x, t) \in a\}$$

Then f is convergent and convex. To see that it is convex, notice that as soon as a test is found such that $e(x, t) > \theta$ then x can never be preferred by f again. It also only requires a single test with $e(x, t) \leq \theta$ to select x. Hence f is finite and so by the above theorem, also convergent.

Example

The Pareto dominance example on a finite set of objects is convergent, but not convex, since it is possible that a preferred object might become dominated by another, but then later achieve non-dominance through further tests. The same consideration applies in Ficici's scenario, in which a whole set of possible non-dominated subsets is returned. Ficici therefore proposes an adaptation of Pareto dominance to enforce convexity [4, 5]. His method, however, does not readily fit our framework as the choice of preferred object in the case of equal evaluation depends critically on the order in which tests have been performed, and not just on what the tests were. Consequently, it would not be possible to apply his method, say, to a situation where there are two or more co-evolving populations which become merged or share information over time.

Example

A simple but paradoxical example of a convex convergent preference function comes from the "greater than" game [12]. Given a set of natural numbers, we think of larger numbers beating smaller numbers. We want to select players who are unbeaten. Thus we define a preference function $f : 2^{\mathbb{N}} \to 2^{\mathbb{N}}$ by

$$f(a) = \{x \in a \,|\, \text{there are no } y \in a \text{ with } y > x\}$$

If a is finite then $f(a)$ is simply the maximum. If a is infinite then $f(a) = \emptyset$. This preference function is convex, since we only reject a number if we found another bigger than it. From then on, it remains rejected. Any number x is accepted in the singleton subset $\{x\}$ and rejected (permanently) by any subset containing a greater number. Hence it is finite and therefore convergent. The paradox is that at each stage of finite testing, we always prefer exactly one number, but in the limit, we prefer none. The paradox is resolved by the fact that the function is, nevertheless, continuous with respect to the Lawson topology, i.e., topologically the sequence $(\{x\})_{x\in\mathbb{N}}$ converges to $\emptyset$.

7. Weak preference pre-order

Given a preference function $f : A \to 2^X$, Ficici [4, 5] defines a partial order (called *weak preference*) on the elements of X. Unfortunately, his definition assumes that A is finite, and does not work for the infinite case. However, a minor variant provides a corresponding pre-order as follows.

Definition 5. *Given $f : A \to 2^X$, then for $x, y \in X$ we say $x \lesssim y$ if for all $a \in A$ such that $x \in f(a)$, there exists $b \geq a$ such that $y \in f(b)$.*

Lemma 9. *The pair $(X, \lesssim)$ is a pre-order.*

Proof
It is clear that $x \lesssim x$ (by taking $b = a$ in the definition). If $x \lesssim y$ and $y \lesssim z$, then for all $a \in A$ with $x \in f(a)$ there is a $b \geq a$ with $y \in f(b)$ and therefore $c \geq b \geq a$ such that $z \in f(c)$. Hence $x \lesssim z$. $\square$

Theorem 10. *Suppose the $f : A \to 2^X$ is convex. Let $a \leq b$ in A with $x \in f(a)$ and $x \notin f(b)$ for some $x \in X$. Then for all $y \in f(b)$ it follows that $y \not\lesssim x$.*

Proof
Suppose $y \lesssim x$. Then there exists some $c \geq b$ such that $x \in f(c)$. But then $x \in f(a) \cap f(c)$ but $x \notin f(b)$ which contradicts convexity. $\square$

Ficici argues that as exploration proceeds, the search process never re-visits points worse than those currently preferred. If X is finite, this forces the search to progress to increasingly better elements of X. Such a conclusion would not necessarily extend to the case when X is infinite.

8. The Boolean algebra of convergent functions

We now seek to understand the relationship between convex convergent functions and convergent functions in general. To do this, we begin by observing that the set of convergent functions may be turned into a Boolean algebra in a natural way.

Definition 6. *Let $f, g : A \to 2^X$. Define $f \cap g : A \to 2^X$ by*

$$(f \cap g)(a) = f(a) \cap g(a)$$

Similarly, define $f \cup g : A \to 2^X$ by

$$(f \cup g)(a) = f(a) \cup g(a)$$

and $\neg f : A \to 2^X$ by

$$(\neg f)(a) = f(a)^c$$

Theorem 11. *If $f, g : A \to 2^X$ are convergent functions, then so are $f \cap g$, $f \cup g$ and $\neg f$. Hence the set of convergent functions from A to 2^X forms a Boolean algebra, with zero and one being the functions $0(a) = \emptyset$ and $1(a) = X$ for all $a \in A$. We have $f \leq g$ if and only if for all $a \in A$, $f(a) \subseteq g(a)$.*

Proof
Let $S \subseteq A$ be directed.

1. Suppose $x \in (f \cap g)(\bigvee^{\uparrow} S)$. Then $x \in f(\bigvee^{\uparrow} S)$ and there exists some $a \in S$ such that

$$a \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \in f(b)$$

for all $b \in S$. Similarly, $x \in g(\bigvee^{\uparrow} S)$ and there exists $a' \in S$ such that

$$a' \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \in g(b)$$

for all $b \in S$. Since S is directed we can find $s \in S$ greater than a and a' and then

$$s \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \in (f \cap g)(b)$$

for all $b \in S$.
2. Suppose $x \notin (f \cap g)(\bigvee^{\uparrow} S)$. Assume, then, w.l.o.g. that $x \notin f(\bigvee^{\uparrow} S)$. Then there exists $a \in S$ such that

$$a \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \notin f(b) \Longrightarrow x \notin (f \cap g)(b)$$

for all $b \in S$.
3. Suppose $x \in (f \cup g)(\bigvee^{\uparrow} S)$. Assume w.l.o.g. that $x \in f(\bigvee^{\uparrow} S)$. Then there exists $a \in S$ such that

$$a \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \in f(b) \Longrightarrow x \in (f \cup g)(b)$$

for all $b \in S$.

4. Suppose $x \notin (f \cup g)(\bigvee^{\uparrow} S)$. Then $x \notin f(\bigvee^{\uparrow} S)$ and there exists some $a \in S$ such that
$$a \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \notin f(b)$$
for all $b \in S$. Similarly, $x \notin g(\bigvee^{\uparrow} S)$ and there exists $a' \in S$ such that
$$a' \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \notin f(b)$$
for all $b \in S$. Taking $s \in S$ with s greater than a and a' we have
$$s \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \notin (f \cup g)(b)$$
for all $b \in S$.
5. Suppose $x \in (\neg f)(\bigvee^{\uparrow} S)$. Then $x \notin f(\bigvee^{\uparrow} S)$. Then there exists some $a \in S$ such that
$$a \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \notin f(b) \Longrightarrow x \in (\neg f)(b)$$
for all $b \in S$.
6. Suppose $x \notin (\neg f)(\bigvee^{\uparrow} S)$. Then $x \in f(\bigvee^{\uparrow} S)$. Then there exists $a \in S$ such that
$$a \leq b \leq \bigvee^{\uparrow} S \Longrightarrow x \in f(b) \Longrightarrow x \notin (\neg f)(b)$$
for all $b \in S$.

The functions 0 and 1 are clearly convergent, and hence the set of convergent functions forms a subalgebra of the Boolean algebra of all functions from A to 2^X.

Finally, we have $f \leq g$ if and only if $f = f \cap g$. That is, for all $a \in A$
$$f(a) = (f \cap g)(a) = f(a) \cap g(a) \Longleftrightarrow f(a) \subseteq g(a)$$
□

Given any convergent function $f : A \to 2^X$ we can define a corresponding function $m_f : X \to 2^A$ as follows:
$$m_f(x) = \{a \in A \,|\, x \in f(a)\}$$

Lemma 12. *For any $x \in X$ the set $m_f(x)$ is clopen in the d-topology on A.*

Proof

First we show $m_f(x)$ is open. Let $\bigvee^{\uparrow} S \in m_f(x)$. Then $x \in f(\bigvee^{\uparrow} S)$. Since f is convergent, there exists some $s \in S$ such that $x \in f(s)$ and therefore $s \in m_f(x)$.

Second, we show $m_f(x)$ is closed, by showing that $m_f(x)^c$ is open. Let $\bigvee^{\uparrow} S \in m_f(x)^c$ so that $\bigvee^{\uparrow} S \notin m_f(x)$. Then $x \notin f(\bigvee^{\uparrow} S)$. Since f is convergent, these exists some $s \in S$ such that $x \notin f(s)$ and therefore $s \notin m_f(x)$. □

Lemma 13. *The map $f \mapsto m_f$ is an injection.*

Proof
Suppose $m_f = m_g$ for two convergent functions f and g. Then for all $a \in A$ and all $x \in X$ we have $a \in m_f(x)$ if and only if $a \in m_g(x)$. That is, $x \in f(a)$ if and only if $x \in g(a)$. Therefore $f(a) = g(a)$ for all $a \in A$, and so $f = g$. $\square$

Lemma 14. *Let $m : X \to 2^A$ have $m(x)$ is d-clopen for all x. Then there exists a convergent function $f : A \to 2^X$ such that $m = m_f$.*

Proof
Define $f : A \to 2^X$ to be

$$f(a) = \{x \in X \mid a \in m(x)\}$$

Then $a \in m_f(x)$ if and only if $x \in f(a)$ if and only if $a \in m(x)$, so it follows that $m = m_f$. We now show f is convergent by showing it is continuous as a map from A with the d-topology to 2^X with the Lawson topology.

Let $x \in X$ and let $\bigvee^{\uparrow} S \in f^{-1}(\uparrow\{x\})$. Then $\{x\} \subseteq f(\bigvee^{\uparrow} S)$ and so $\bigvee^{\uparrow} S \in m(x)$. Since $m(x)$ is open, there exists some $s \leq \bigvee^{\uparrow} S$ such that $s \in m(x)$. Therefore $f(s) \in \uparrow\{x\}$ and so $s \in f^{-1}(\uparrow\{x\})$. It follows that $f^{-1}(\uparrow\{x\})$ is open in A.

Similarly, let $\bigvee^{\uparrow} S \in f^{-1}((\uparrow\{x\})^c)$, so that $\{x\} \nsubseteq f(\bigvee^{\uparrow} S)$ and so $\bigvee^{\uparrow} S \notin m(x)$. Since $m(x)^c$ is open, there exists some $s \in S$ such that $s \notin m(x)$. Therefore $f(s) \notin \uparrow\{x\}$ and so $s \in f^{-1}((\uparrow\{x\})^c)$. It follows that $f^{-1}((\uparrow\{x\})^c)$ is open in A. $\square$

Lemma 15. *The set of all functions from X to the d-clopens of A forms a Boolean algebra, where operations are defined pointwise.*

Proof
The set of d-clopens of A form a Boolean algebra under union, intersection and complement. The set of functions from X to this Boolean algebra is a product of Boolean algebras, and hence is itself a Boolean algebra under pointwise operations. $\square$

Theorem 16. *The Boolean algebra of convergent functions from A to 2^X is isomorphic to the Boolean algebra of all functions from X to the d-clopens of A.*

Proof
We already have that the map $f \mapsto m_f$ is a bijection between these Boolean

algebras. We now show it is a Boolean homomorphism. So, for all $x \in X$ and $a \in A$

$$\begin{aligned} a \in m_{f\wedge g}(x) &\iff x \in (f \wedge g)(a) \\ &\iff x \in f(a) \cap g(a) \\ &\iff a \in m_f(x) \cap m_g(x) \\ &\iff a \in (m_f \wedge m_g)(x) \end{aligned}$$

Therefore, for all $x \in X$ it follows that $m_{f\wedge g}(x) = (m_f \wedge m_g)(x)$ and so $m_{f\wedge g} = m_f \wedge m_g$. The proof that $m_{f\vee g} = m_f \vee m_g$ is dual. Finally,

$$\begin{aligned} a \in m_{\neg f}(x) &\iff x \in (\neg f)(a) \\ &\iff x \notin f(a) \\ &\iff a \in m_f(x)^c \\ &\iff a \in (\neg m_f)(x) \end{aligned}$$

from which it follows that $m_{\neg f} = \neg m_f$. □

9. Constructing convergent functions from convex functions

We now show that any convergent function can be built up by the set of all (uniformly) convergent convex functions which are beneath it in the Boolean algebra.

Theorem 17. *Let $f : A \to 2^X$ be convergent, and define*

$$D_f = \{d : A \to 2^X \mid d \textit{ is convex and uniformly convergent and } d \leq f\}$$

then

$$f = \bigcup D_f$$

Proof
Clearly $f \geq \bigcup D_f$. Suppose they are not equal. Then there exists an $a \in A$ and $x \in X$ such that $x \in f(a)$ but $d \in D_f \Rightarrow x \notin d(a)$. Since f is convergent then, by Lemma 1, there exists a compact $c \in K_a$ such that $c \leq b \leq a$ implies $x \in f(b)$ for all $b \in A$. Now define a function $g : A \to 2^X$ as follows:

$$g(z) = \begin{cases} \{x\} & \text{if } c \leq z \leq a \\ \emptyset & \text{otherwise} \end{cases}$$

We now show that g is convex, uniformly convergent and that $g \leq f$. This means that $g \in D_f$ with $x \in g(a)$ which gives a contradiction.

1. Let $u \leq v \leq w$ in A. and suppose $y \in g(u) \cap g(w)$. Then $c \leq u \leq v \leq w$ which means $y \in g(v)$. Therefore g is convex.

2. Let $S \subseteq A$ be directed. First suppose $g(\bigvee^{\uparrow} S) = \{x\}$. Then $c \leq \bigvee^{\uparrow} S \leq a$. Since c is compact, there exists some $s \in S$ with $c \leq s$ and then $s \leq b \leq \bigvee^{\uparrow} S$ implies $g(b) = g(\bigvee^{\uparrow} S)$ for all $b \in S$. Secondly, suppose $g(\bigvee^{\uparrow} S) = \emptyset$. Suppose that, for all $s \in S$ there is some $b \in S$ with $b \geq s$ and $g(b) \neq \emptyset$. Then $\bigvee^{\uparrow} S \geq c$ and a is an upper bound for S. But this implies $a \geq \bigvee^{\uparrow} S$ and so $g(\bigvee^{\uparrow} S) = \{x\}$, which is a contradiction. Therefore, there must exist some $s \in S$ such that $s \leq b \Rightarrow g(b) = g(\bigvee^{\uparrow} S)$ for all $b \in S$.
3. We already have seen that $c \leq b \leq a$ implies $x \in f(b)$ and so it follows that $f \geq g$ as required.

$\square$

In the Boolean algebra of convergent functions, it is only guaranteed that finite sets have joins. That is, the algebra is not complete. However, the above theorem tells us that every convergent function does arise as a potentially infinite join. We therefore ask, given a set of functions, if there is a condition describing whether or not the join exists.

Definition 7. *For any $f : A \to 2^X$ and any $x \in X$, define $f_x : A \to 2^X$ to be*

$$f_x(a) = \begin{cases} \{x\} & \text{if } x \in f(a) \\ \emptyset & \text{if } x \notin f(a) \end{cases}$$

Theorem 18. *Let F be a set of convergent functions from A to 2^X. For any $x \in X$, let*

$$F_x = \{f_x : A \to 2^X \mid f \in F\}$$

Then $\bigcup F$ is convergent if and only if for all $a \in A$ and all $x \in X$ we can find a d-open neighbourhood U of a such that

$$(\bigcup F_x)|_U = f_x|_U$$

for some $f \in F$.

Proof
Assume that for all $a \in A$ and all $x \in X$ we can find a d-open neighbourhood U of a such that

$$(\bigcup F_x)|_U = f_x|_U$$

for some $f \in F$. To show $\bigcup F$ is convergent, let $S \subseteq A$ be directed and consider two cases, for any $x \in X$.

1. $x \in (\bigcup F)(\bigvee^{\uparrow} S)$. Then there exists some $f \in F$ such that $x \in f(\bigvee^{\uparrow} S)$. Then there exists $c \in S$ such that for all $b \in S$

$$c \leq b \leq \bigvee\nolimits^{\uparrow} S \Longrightarrow x \in f(b) \Longrightarrow x \in (\bigcup F)(b)$$

2. $x \notin (\bigcup F)(\bigvee^{\uparrow} S)$. There exists a d-open neighbourhood, U of $\bigvee^{\uparrow} S$ with $(\bigcup F_x)|_U = f_x|_U$ for some $f \in F$, and we know $x \notin f_x(\bigvee^{\uparrow} S)$. Therefore there exists some $s \in S$ such that for all $b \in S$, $s \leq b \Rightarrow x \notin f_x(b)$. Since U is d-open, we can find some $c \in S$ such that

$$\{b \in S \,|\, c \leq b\} \subseteq U$$

(using an argument similar to that of Lemma 1, part 4). Choose $c' \in S$ greater than c and s. Then, for all $b \in S$,

$$c' \leq b \Longrightarrow x \notin f_x(b) \Longrightarrow x \notin (\bigcup F_x)(b) \Longrightarrow x \notin (\bigcup F)(b)$$

as required.

Conversely, assume $\bigcup F$ is convergent. Fix $a \in A$ and $x \in X$. Consider two cases:

1. $x \in (\bigcup F)(a)$. Then there exists some $f \in F$ such that $x \in f(a)$. So there exists some $c \in K_a$ such that, for all $b \in A$,

$$c \leq b \leq a \Longrightarrow x \in f(b)$$

Let $U = \{b \in A \,|\, c \leq b \leq a\}$ which is a d-open neighbourhood of a. For any $u \in U$ we have

$$f_x(u) = \{x\}$$

and

$$(\bigcup F_x)(u) = \bigcup_{g \in F} g_x(u) = \{x\}$$

2. $x \notin (\bigcup F)(a)$. Then there exists $c \in K_a$ such that for all $b \in A$

$$c \leq b \leq a \Longrightarrow x \notin (\bigcup F)(b)$$

Let $U = \{b \in A \,|\, c \leq b \leq a\}$ which is a d-open neighbourhood of a. Pick any $f \in F$. Then, for any $u \in U$

$$f_x(u) = \emptyset$$

and

$$(\bigcup F_x)(u) = \bigcup_{g \in F} g_x(u) = \emptyset$$

□

10. Characterising convex convergent functions

Given the importance of convex functions in describing convergence, we seek to further characterise them. We start with some simple definitions.

Definition 8. *Given any two functions* $g, h : A \to 2^X$*, the map* $(g - h) : A \to 2^X$ *is given by*

$$(g - h)(a) = g(a) \cap h(a)^c$$

Definition 9. *Given* $f : A \to 2^X$ *we define corresponding functions* $p, q : A \to 2^X$ *as follows:*

$$p(a) = \{x \in X \mid \exists b \leq a \text{ such that } x \in f(b)\}$$

$$q = p - f$$

One may think of the function p as describing all the objects which could possibly have been preferred up to the current state of information. The function q describes all objects that could possibly have been preferred prior to the current state, but are no longer preferred. The following lemma shows that p always grows larger and that, if f is convex, so does q.

Lemma 19. *The function* p *is monotonic;* q *is monotonic if and only if* f *is convex.*

Proof
Suppose $a \leq b$, and $x \in p(a)$. Then there exists $c \leq a$ with $x \in d(c)$. Therefore $x \in p(b)$. So $p(a) \subseteq p(b)$.

Now suppose f is convex and suppose $x \in q(a)$. Then $x \notin f(a)$ but there exists $c \leq a$ with $x \in f(c)$. From $c \leq a \leq b$ it follows, by convexity of f, that $x \notin f(b)$. Therefore $x \in q(b)$ and so $q(a) \subseteq q(b)$. For the converse, assume f is not convex, so there exist $a \leq b \leq c$ with $x \in f(a) \cap f(c)$ but $x \notin f(b)$. This means that $x \in q(b)$ but $x \notin q(c)$, contradicting monotonicity. $\square$

We now characterise convex convergent functions f in terms of the continuity properties of p and q. Recall that a function between dcpos is called *Scott-continuous* if it is monotone and preserves suprema of directed sets. This is in line with the definition of topological continuity, assuming the Scott topology on the two dcpos. In our concrete setting of functions $p : A \to 2^X$ it means precisely that whenever $x \in p(a)$ then there exists a compact element $c \leq a$ such that $x \in p(c)$.

In order to deal with uniform convergence we will also need a stronger continuity condition, namely where we equip 2^X with the *Alexandrov topology* of *all* upper sets, not just the upper sets inaccessible by directed joins. We call a function *Scott-Alexandrov-continuous* if the inverse image of any upper set is Scott-open. In our setting of functions $p : A \to 2^X$ it means precisely that for any $a \in A$ there exists a compact element $c \leq a$ such that $p(c) = p(a)$.

Theorem 20. *If $f : A \to 2^X$ is convergent then p is Scott-continuous. If f is uniformly convergent, then p is Scott-Alexandrov-continuous.*

Proof
For the first part, assume that $x \in p(a)$ for some $x \in X$, $a \in A$. By definition, there is some $b \leq a$ such that $x \in f(b)$. The claim now follows from Lemma 1(3a).

For the second part, let $a \in A$ and consider the Alexandrov open set $\uparrow p(a)$. We need to show that there is a compact element c below a such that $p(c) \in \uparrow p(a)$ which by monotonicity means that $p(c) = p(a)$. Now by Lemma 2(3), there exists a compact element $c \leq a$ such that $f(b) = f(a)$ for all $c \leq b \leq a$. This means that, likewise, the function p does not change between c and a, in other words, $p(c) = p(a)$, as desired. $\square$

Theorem 21. *If f is convergent and convex, then q is Scott-continuous. If f is uniformly convergent and convex then q is Scott-Alexandrov-continuous.*

Proof
Since f is convex, q is monotone. Now suppose $x \in q(a)$. Therefore $x \notin f(a)$ but there exists $b \leq a$ with $x \in f(b)$. By Lemma 1(3) we find compact elements $c \leq a$ with $x \notin f(c)$ and $c' \leq b$ with $x \in f(c')$. Let c'' be a compact upper bound of $\{c, c'\}$ in K_a. By convexity we have $x \notin f(c'')$, and since $c' \leq c''$ it follows that $x \in q(c'')$.

The second part follows as in the preceding theorem: Since f does not change its value on an interval $[c, a]$ with c compact, neither does q. $\square$

Conversely, we now ask if we have two Scott-continuous (Scott-Alexandrov-continuous) functions, whether it is possible to construct a convex (uniformly) convergent function out of them. The answer is yes.

Lemma 22. *If $g, h : A \to 2^X$ are both monotonic, then $g - h$ is convex.*

Proof
Let $a \leq b \leq c$. Then

$$\begin{aligned}
(g-h)(a) \cap (g-h)(c) &= (g(a) \cap h(a)^c) \cap (g(c) \cap h(c)^c) \\
&= g(a) \cap h(c)^c \\
&\subseteq g(b) \cap h(b)^c \\
&= (g-h)(b)
\end{aligned}$$

$\square$

Theorem 23. *If $g, h : A \to 2^X$ are both Scott-continuous (Scott-Alexandrov-continuous), then $g - h$ is convex and (uniformly) convergent.*

Proof
Scott-continuity implies monotonicity, so the lemma demonstrates convexity. For convergence, assume $x \in (g-h)(a)$, that is, $x \in g(a)$ and $x \notin h(a)$. By continuity, there exists a compact element $c \leq a$ such that $x \in g(c)$. By the monotonicity of g, the same is true for all $c \leq b \leq a$. By the monotonicity of h, we have $x \notin h(b)$ for all such b, hence $x \in (g-h)(b)$.

Next assume $x \notin (g-h)(a)$. We distinguish two cases: If $x \notin g(a)$ then we are done because this implies $x \notin (g-h)(b)$ for all $b \leq a$. Alternatively, it could be the case that $x \in g(a)$ and $x \in h(a)$. By continuity, there exist compact elements $c, c' \leq a$ such that $x \in g(c)$ and $x \in g(c')$. Let c'' be an upper bound for $\{c, c'\}$ in K_a. By monotonicity we have $x \in g(b) \cap h(b)$ for all $c'' \leq b \leq a$, hence $x \notin (g-h)(b)$ for all such b. Appealing to the characterisation of convergence in Lemma 1(3) concludes the proof.

The uniform case is similar: Given $a \in A$ we know that there is a compact element $c \leq a$ such that $g(c) = g(a)$, and in fact $g(b) = g(a)$ for all $c \leq b \leq a$ by monotonicity. Likewise, there is a compact element $c' \leq a$ such that $h(b) = h(a)$ for all $c' \leq b \leq a$. We pick a compact upper bound c'' of $\{c, c'\}$ and get $(g-h)(b) = (g-h)(a)$ for all $c'' \leq b \leq a$. By Lemma 2(3) it follows that $g - h$ is uniformly convergent. □

Corollary 24. *A preference function $f : A \to 2^X$ is convex and (uniformly) convergent if and only if there exist Scott-continuous (resp. Scott-Alexandrov-continuous) functions $g, h : A \to 2^X$ such that $f = g - h$.*

11. Conclusions

We have presented a framework for studying the convergence of preference functions based on aggregating information. The framework simplifies and generalises the approach taken in the theory of co-evolutionary algorithms, and a number of results map over from that area. In particular, the idea of a convex function is found to be very useful. Our main results show that the set of convergent functions forms a Boolean algebra, in which each function is the join of the convex convergent functions below it. Further, we characterise all convergent convex functions as the difference between two Scott-continuous functions.

Acknowledgements. We would like to record our gratitude to Prof M. Andrew Moshier who patiently listened to our attempts to capture the general concepts correctly and who made many perceptive and constructive comments. Financial support from the EPSRC is also acknowledged as it made the extended stay of Prof Moshier in Birmingham possible. Thanks are also due to the referee who pushed us to extend the results in the last section to the uniformly convergent case.

References

[1] H. Putnam, Trial and error predicates and the solution to a problem of Mostowski, Journal of Symbolic Logic 30 (1) (1965) 49–57.

[2] E. M. Gold, Limiting recursion, Journal of Symbolic Logic 30 (1) (1965) 28–48.

[3] E. Popovici, A. Bucci, R. P. Wiegand, E. D. de Jong, Coevolutionary principles, in: G. Rozenberg, T. Bäck, J. N. Kok (Eds.), Handbook of Natural Computing, Springer, 2012, pp. 987–1033.

[4] S. Ficici, Solution concepts in coevolutionary algorithms, Ph.D. thesis, Brandeis University (2004).

[5] S. Ficici, Monotonic solution concepts in coevolution, in: Proc. Genetic and Evolutionary Computation Conference, ACM, 2005, pp. 499–506.

[6] A. Bucci, J. B. Pollack, A mathematical framework for the study of coevolution, in: K. De Jong, R. Poli, J. E. Rowe (Eds.), Proceedings of the 7th Workshop on Foundations of Genetic Algorithms, Morgan Kaufmann, 2003, pp. 221–236.

[7] A. Bucci, J. B. Pollack, Thoughts on solution concepts, in: Proc. Genetic and Evolutionary Computation Conference, ACM, 2007, pp. 434–439.

[8] E. Popovici, K. De Jong, Monotonicity versus performance in co-optimization, in: I. I. Garibay, T. Jansen, R. P. Wiegand, A. S. Wu (Eds.), 10th ACM SIGEVO International Workshop on Foundations of Genetic Algorithms, ACM, 2009, pp. 151–170.

[9] S. Abramsky, A. Jung, Domain theory, in: S. Abramsky, D. M. Gabbay, T. S. E. Maibaum (Eds.), Handbook of Logic in Computer Science, Vol. 3, Clarendon Press, 1994, pp. 1–168.

[10] B. A. Davey, H. A. Priestley, Introduction to Lattices and Order, Cambridge University Press, 2002.

[11] H. H. Hoos, T. Stützle, Stochastic Local Search: Foundations and Applications, Morgan Kaufmann Elsevier, 2004.

[12] R. A. Watson, J. B. Pollack, Coevolutionary dynamics in a minimal substrate, in: Proc. Genetic and Evolutionary Computation Conference, Morgan Kaufmann, 2001, pp. 702–709.