Brouwer's Fan Theorem: An Overview

Crystal Cheung, ctc89@cornell.edu
Fall 2015

Introduction

Beginning in the early 1920's, L.E.J. Brouwer proposed an intuitionsitic reformation of the foundations of mathematics.[1] Among the key elements underlying this reconstruction were: (1) a complete rejection of the principle of the excluded middle, and (2) a collection of axioms culminating in what is commonly referred to as Brouwer's Thesis on bars. The latter arose from Brouwer's effort to redefine the continuum from an intuitionistic standpoint, that is, to characterize and analyze the continuum as a constructively defined object. Brouwer's Thesis, which he proposed as an axiom on the nature of bars, resulted in the proof of the Bar Theorem and its corollary, the Fan Theorem. Interestingly, most mathematically significant consequences of the Bar Theorem can be proved using the Fan Theorem alone [14].

The fan, or finitely branching tree of potentially unbounded depth, can be viewed as the well-studied data structure from computer science, but it is important to understand that in the context of the Fan Theorem, the relevant "objects" within the tree are not the individual nodes and their hierarchical relations to one another, but rather the *paths* originating from the root of the tree and in particular, the paths of unbounded length. Such paths can be interpreted as unbounded sequences of successive choices, which can in turn be used to define the real numbers. Brouwer's goal was to analyze the nature of these unbounded paths, and, in particular, to effectively characterize them using only collections of finite objects. Roughly speaking, the Fan Theorem states that the space of all unbounded binary sequences, i.e. the Cantor space C, can be "covered" by a finite set of finite sequences. In essence, this guarantees a form of compactness for C. As van Atten writes in [15, p. 41]:

> The philosophical value of the bar theorem lies both in its content – it makes the full continuum, which had always been intractable for constructivists, constructively manageable – and in the way its proof fully exploits the intuitionistic conceptions of truth as experienced truth and of proofs as mental constructions.

[1] Published in 1913, [6] is among Brouwer's earlier discussions on intuitionism. Some key ideas of interest for this paper, e.g. bar induction, were not fully developed by Brouwer until much later; some of his later works can be found in [7]. More comprehensive information on intuitionism and constructivism in general can be found in [16] and [12], respectively.

This paper's primary focus will be to discuss the Fan Theorem in a self-contained manner, including the mathematical context of its origins, the nuances in its various expressions, and some of its intriguingly far-reaching philosophical implications.

1 Background

Choice Sequences

Central to Brouwer's redefinition of the continuum is the concept of choice sequences, which he used to construct the real numbers. In particular, he identified each real number as an unbounded sequence of rational intervals, with the property that each successive interval is a strict subinterval lying within the last. While the Fan Theorem is a statement about the nature of unbounded trees, it can equally be seen as a statement about choice sequences – this is, after all, the original motivation behind its conception.

Define a sequence as an ordered list of natural numbers [2], some being finite in length and some being infinitely long. A binary sequence is one whose elements are restricted to {0, 1}. The term *choice sequence* is meant to denote any sequence from the realm of all possible unbounded sequences. These include: (1) the *lawlike* sequences, or those whose elements can be given by an algorithm (these can of course be described by computable functions in the intuitive sense), (2) the *lawless* or *free choice* sequences, each of whose generation may involve some non-algorithmic process, and (3) "hybrid" sequences which are neither strictly lawless nor lawlike, but can only be described as some combination of both.

While it is evident that a conceptual distinction can be made between the lawless and lawlike sequences in terms of the process by which they are generated, it is theoretically possible that every lawless (or hybrid) sequence coincides at every position with some lawlike sequence. Here, an obvious question arises: accepting the choice sequences as a mathematical construct, can all unbounded sequences still be considered as arising from a "computable" function (in the sense of Church and Turing, i.e. the recursive functions), or from a constructive function as defined by Bishop? We return to this issue later.

As mathematical objects, the choice sequences in general must be thought of as perpetually growing and forever incomplete: at any given point in time, we can only grasp a choice sequence by its hitherto known initial segment. On the other hand, each finite sequence can be thought of as the initial segment of infinitely many choice sequences, each yet to be further determined in the future. It is in

[2]More generally speaking, the objects permitted in the sequence can come from any collection for which there exists a bijection from the natural numbers to that collection. Thus there is no loss of generality in dealing only with unbounded sequences of natural numbers when the intended application is on sequences of rational intervals.

this sense that every finite sequence can be thought of as "covering" an infinite collection of associated choice sequences.

Here it is necessary to introduce some basic notation. First, let the Baire space $\mathcal{B}$ denote the space of all unbounded sequences of natural numbers, and let the Cantor space $\mathcal{C}$ denote the space of all unbounded binary sequences. Denote by variables such as α, β an unbounded sequence, and by variables such as σ, τ a finite sequence. Unless otherwise noted, let variables such as a, m, n (i.e. the Latin lowercase lettters) stand for natural numbers. Let $< a_0, a_1, ..., a_{n-1} >$ denote a finite sequence of length n whose ordered elements are $a_0, ..., a_{n-1}$; $<>$ denotes the empty sequence of length 0.

Let $\cdot$ denote concatenation, so that if $\sigma =< a_0, ..., a_{n-1} >$, then $\sigma \cdot a_n = < a_0, ..., a_{n-1}, a_n >$. If there exists some number a_n such that $\tau = \sigma \cdot a_n$, then say that τ is an *immediate descendant* or *extension* of σ, or equivalently, σ is an *immediate ascendant* of τ, both denoted by $\sigma \sqsubseteq \tau$. More generally, if there exists a finite sequence σ' such that $\tau = \sigma \cdot \sigma'$, then say that τ is a *descendant* of σ, or equivalently, σ is an *ascendant* of τ, both denoted by $\sigma \prec \tau$. If σ, τ are such that neither sequence is a descendant of the other, then say that σ and τ are *incompatible*.

Let $\alpha(i)$ be the ith element of α. Let $\overline{\alpha}n$ denote the n-length initial segment of α, i.e. $\overline{\alpha}n :=< \alpha(0), \alpha(1), ..., \alpha(n-1) >$. Given any finite sequence σ and any unbounded sequence α, if there exists some number n such that $\overline{\alpha}n = \sigma$, then say that α *passes through* σ, denoted $\alpha \sqsupseteq \sigma$. Finally, for a fixed σ and collection of unbounded sequences $\mathcal{U}$, let $\mathcal{U} \cap \sigma$ denote the the collection of all $\alpha \in \mathcal{U}$ such that α passes through σ.

Brouwer defined collections of choice sequences using the concept of a *spread*. Crucially, the continuum itself (as well as $\mathcal{B}$ and $\mathcal{C}$) can be modeled as a spread. Broadly speaking, the elements of a given spread are determined by its corresponding *spread law*, a decidable criterion which either admits or rejects finite sequences, under the following conditions: (1) for each admitted sequence, all of its initial segments must also be admitted, and (2) every admitted sequence must have at least one immediate successor that is also admitted.

Notice that every spread can naturally be represented as an unbounded tree, with each node representing an admitted finite sequence, and the root of the tree in particular corresponding to the empty sequence.[3] The unique tree corresponding to a given spread is referred to as its *underlying tree*. Given a spread, the unbounded paths[4] of its underlying tree correspond exactly to the choice sequences captured by the spread. For instance, the spread consisting of all possible binary sequences is represented by the binary tree, and the *universal spread*, i.e. the spread consisting of all possible sequences of natural numbers, is represented by the *universal tree*. The former can be viewed as a constructive

[3]Equivalently, the admitted finite sequences can be thought of as corresponding to the finite paths in the tree – each node can be addressed by the path leading up to it.

[4]As is the case for unbounded sequences, it is crucial to consider the infinitely long paths of a tree as fundamentally incomplete objects.

definition of the space $\mathcal{C}$, and the latter of the space $\mathcal{B}$.

Though strictly speaking a distinction can be made between the concepts of a spread and its underlying tree, we will frequently refer directly to the tree itself as a collection of sequences (both finite and unbounded), sometimes conflating the spread with its underlying tree, as is commonly done in the literature (e.g. in [14]). It is worth emphasizing that the *nodes* (or finite paths) in the tree, by which it is explicitly defined, correspond simply to finite sequences, whereas the primary objects of interest in this case are the *unbounded paths* in the tree – their admittance to the spread is implied by the collection of admitted nodes.

Now it is possible to elaborate on the aforementioned notion of "covering" a spread of choice sequences. Say for example that T is the underlying tree of $\mathcal{C}$, i.e. T is the binary tree. We can think of its "cover" as a set B such that every unbounded sequence represented by a path in T belongs to some collection $\mathcal{C} \cap \sigma$, where σ is in B. In other words, if we could constructively define B, then we would have a collection of finite objects which, in a sense, account for all the unbounded sequences in $\mathcal{C}$.[5] Such a set B can be referred to as a *bar in* $\mathcal{C}$. We will return to a more precise definition of bars in our discussion of Brouwer's Thesis.

What is the significance of being able to describe the universe of all possible choice sequences in terms of finite objects? Since such a space consists of fundamentally incomplete objects, it is not immediately clear how the space can be constructively defined in its totality, much less reasoned about. A primary concern of the Fan Theorem will be to show that there always exists, in the constructive sense, a set of finite sequences that sufficiently captures certain key properties of *all* the choice sequences.

The Weak Continuity Principle for Numbers

The issue of selecting collections of choice sequences based on their inherent properties is theoretically problematic – since any given choice sequence is always incomplete, it is impossible to reason about the object in its entirety. For example, say that we have the arbitrary choice sequences α and β, and we wish to determine whether the number of 1's in α is greater than the number of 1's in β. Despite the premise of the problem being simple, the answer must be undeterminable – at any given point in time, only the initial segments of α and β are known.

In general, if any claim is to be made about a choice sequence as a whole, it must be provable using only an initial segment of that sequence. In other words, any property that can be proved about a collection of choice sequences must be demonstrable using finite sequences alone. This raises the question: is it possible for a choice sequence to have inherent properties that do not hold for *any* of its initial segments?

[5]One preliminary "solution" might be to construct a trivial cover by taking, for example, a binary tree of depth 1, but the true purpose here is to find a cover that sufficiently captures certain properties of $\mathcal{C}$.

Brouwer addressed this concern by adopting WC-N, the Weak Continuity Principle for Numbers, which in essence states that any decidable property of a given choice sequence can be determined entirely from one of its initial segments:

$$\forall \alpha . \exists p . A(\alpha, p) \rightarrow \forall \alpha . \exists n . \exists p . \forall \beta . (\overline{\beta} n = \overline{\alpha} n \rightarrow A(\beta, p)) \qquad \text{(WC-N)}$$

where α, β range over all choice sequences. That is, if A is a relation that associates every choice sequence with some number p, then by WC-N there must exist some number n such that for a fixed α, every β with the same n-length initial segment as α is mapped to the same p by A. In other words, WC-N guarantees that every total function from choice sequences to natural numbers can be computed using only initial segments as input.

Interestingly, Brouwer himself never gave a full justification for his acceptance of the Weak Continuity Principle [15]. It is straightforward to see that WC-N holds for all the lawless sequences: Since α is incomplete, the proof of $A(\alpha, p)$ must rely solely on some intial segment of $A(\alpha, p)$, in which case $A(\beta, p)$ must also hold for every lawless sequence β with the same initial segment. This line of reasoning, however, is not valid for all choice sequences in general. Say for example that we have a lawlike sequence α for which $A(\alpha, p)$ holds. We cannot say in general that the proof of $A(\alpha, p)$ relies only on an initial segment of α – in this case, every fixed position in α is predetermined by some law, so we have access to intensional[6] information about α that may not be reflected by any of its intial segments. Thus, the justification for accepting WC-N over all choice sequences is not immediately obvious.

One approach (as in [15]) is to consider a *hesitant* sequence, defined as a choice sequence that starts out as a lawless sequence, with the provision that at any point in time, it may permanently transition into a lawlike sequence by imposing some restriction on all of its future positions. Of particular note is that fact that for hesitant sequences, the intensional information can change over time. Furthermore, for as long as the hesitant sequence remains lawless, the intensional information has no impact on the validity of WC-N.

Now consider the relation A in WC-N. Intuitively, it is reasonable to expect that if $A(\alpha, p)$ holds, then it should hold throughout the entire ongoing construction of α – that is, we limit the scope of A to relations that are guaranteed to remain constant throughout time. Bearing the existence of hesitant sequences in mind, we can then argue that since intensional information is potentially unstable, the proof of $A(\alpha, p)$ can never be dependent on it, and hence WC-N should hold in general for all choice sequences.

Notice that constructively speaking, if the antecedent of WC-N is fulfilled, then given any α we must have a way to compute the exact p such that $A(\alpha, p)$

[6] Viewing choice sequences as functions in the most general sense, *intensional* information refers to the procedure by which the sequence is generated. This is in contrast to *extensional* information, which refers strictly to the function's output at each position, irrespective of the procedure used to generate it.

holds; thus there must exist some *constructive* function f such that on any input α, $f(\alpha) = p$. Furthermore, we can assume the following about f: (1) f requires only an initial segment of α to compute $f(\alpha)$, (2) $f(\alpha)$ is constant for every initial segment of sufficient length to compute $f(\alpha)$, and (3) the question of whether or not a given intial segment of α is of sufficient length to compute $f(\alpha)$ is decidable. Denote by K_0 the class of all such functions f.

This intuition gives us the following strengthening of WC-N, referred to as the (Strong) Continuity Principle for Numbers:

$$\forall\alpha.\exists p.A(\alpha, p) \rightarrow \exists f \in K_0.\forall\alpha.A(\alpha, f(\alpha)) \tag{C-N}$$

Note that neither WC-N nor C-N are classically valid, though it can be shown that WC-N and CN can be generalized to arbitrary spreads, as shown in [14]. The axioms WC-N and C-N are a relatively simple embodiment of some key issues regarding choice sequences – these concepts reappear in Brouwer's most studied approach to the Fan Theorem. For a more complete discussion of the continuity axioms, see [15] and [19].

2 Brouwer's Thesis on Bars

First we give a more precise definition for the concept of a bar. Let $\mathcal{U}$ be a collection of choice sequences as defined by some spread, e.g. the spaces $\mathcal{C}$ or $\mathcal{B}$. Let B be a set of finite sequences. Say that *B is a bar in $\mathcal{U}$* if and only if, for *every* sequence α in $\mathcal{U}$, we have $\overline{\alpha}n \in B$ for some n. Now let σ be an arbitrary finite sequence. Say that B is a *bar in $\mathcal{U}$ above σ* (or equivalently, *B bars σ* in $\mathcal{B}$[7]) if and only if, for all α passing through σ, we have $\overline{\alpha}n \in B$ for some n. That is, B bars σ if and only if B is a bar in $\mathcal{U} \cap \sigma$. Note that the statement "B is a bar in $\mathcal{U}$" is equivalent to the statement "B bars <> (in $\mathcal{U}$)". With respect to particular unbounded sequences α, say that *B bars α* if and only if B contains some initial segment of α.

Suppose B is a bar in $\mathcal{U}$. Say that B is a *monotone* bar if and only if for every finite sequence σ and number n, if $\sigma \in B$, then $\sigma \cdot n \in B$. Say that B is an *inductive* bar if and only if for every finite sequence σ, if $\sigma \cdot n \in B$ for all possible n, then $\sigma \in B$.

Now suppose we have a set B and we are interested in whether or not B constitutes a legitimate bar for the universal spread corresponding to $\mathcal{B}$. If B is indeed a bar in $\mathcal{B}$, how can this fact be verified in general? Brouwer's Thesis asserts the following: If B is a bar, there must exist a *canonical proof* of the fact.

The canonical proof of "B is a bar in $\mathcal{B}$" is outlined as follows. If B is a bar, then there exists a proof-tree from which the statement "B bars <>" can be derived. The proof-tree is formed from inferences of the following types:

(η-inference) $\sigma \in B$, therefore B bars σ.

[7] This will sometimes be abbreviated to "B bars σ" if the context is clear.

(F-inference) B bars $\sigma \cdot n$ for all n, therefore B bars σ.

(ρ-inference) B bars σ, therefore B bars $\sigma \cdot n$.

Note that if B is assumed to be monotone or decidable, then the canonical proof does not require any ρ-inferences [14].

Brouwer took care to distinguish between any finite description of the canonical proof-tree and the proof-tree proper, which should be regarded as a mathematical object of potentially unbounded size [18]. For example, by ranging over arbitrary numbers n, infinitely many ρ-inferences can be instantiated for any fixed σ – each of these statements is integral to the proof-tree. Both the Bar Theorem and Fan Theorem follow from Brouwer's Thesis in the sense that their proofs are relativized on the guaranteed existence of the canonical proof for bars.

3 A Brief Look at the Bar Theorem

We first give a brief summary of the Bar Theorem, of which the Fan Theorem is a main corollary.

First say that a bar is *well-ordered* if its elements can be constructively enumerated.[8] Define a *thin bar* as a minimal bar in the sense that if B^* is a thin bar, then there does *not* exist any set $B \subset B^*$ such that B is also a bar. Note that if B^* is a thin bar, then for all $\sigma, \tau \in B^*$, σ and τ must be incompatible with one another. With these added definitions, the Bar Theorem can be stated as follows [18]:

Theorem 1 (Bar Theorem 1)
If B is a decidable bar, then there exists a set $B^* \subseteq B$ such that B^* is a thin, well-ordered bar.

While in his original proof of the Bar Theorem Brouwer did not explicitly state that B must be decidable, Kleene later showed that this assumption is necessary [15]. The condition that B must be decidable holds if we restrict the bars in question to those implicitly defined by the continuity principle C-N.

To see this, suppose we have a collection $\mathcal{U}$ of choice sequences and an arbitrary function f mapping every α in $\mathcal{U}$ to some number p. By C-N, we have the following: (1) For every α there must exist some n such that $f(\alpha) = p$ can be computed using only the initial segment $\overline{\alpha}n$, (2) $f(\beta) = f(\alpha)$ for all choice sequences β such that $\overline{\beta}n = \overline{\alpha}n$, and (3) For every initial segment $\overline{\alpha}m$ such that $m < n$, we can decide whether or not $\overline{\alpha}m$ is of sufficient length to compute $f(\alpha)$.[9] Therefore, to every such function f we can associate a decidable bar B defined as follows: Given any finite sequence σ, decide whether or not σ is of sufficient length to compute $f(\alpha)$, where α is any choice sequence of which σ is an initial segment. If σ is of sufficient length, then let $\sigma \in B$, otherwise let $\sigma \notin B$. All the bars that can be defined in this manner form a class of decidable bars implicated by the assumption of C-N.

[8] Here, "well-ordered" is meant in the intuitionistic sense, as opposed to its classical definition.
[9] For all $\overline{\alpha}m$ such that $m > n$, we obviously have sufficient information to compute $f(\alpha)$.

Kleene's Interpretation of the Bar Theorem

Based on a footnote mentioned by Brouwer in 1927, Kleene developed an alternate view of the Bar Theorem that is more classically convincing in that it does not rely on Brouwer's Thesis [10].

Where α is an unbounded sequence, say that $\overline{\alpha}n$ is *secured* with respect to a decidable predicate A if and only if $A(\overline{\alpha}m)$ holds for some $m \leqslant n$. If $\overline{\alpha}n$ is the shortest initial segment of α for which A holds (i.e. if $A(\overline{\alpha}n)$ and yet $\neg A(\overline{\alpha}m)$ for all $m < n$), then say that $\overline{\alpha}n$ is *immediately secured*. If $A(\overline{\alpha}n)$ and we have some $m < n$ such that $A(\overline{\alpha}m)$ also holds, then say that $\overline{\alpha}n$ is *past secured*. Say that a finite sequence σ is *securable* if for all α passing through σ, $A(\overline{\alpha}n)$ holds for some n.

The following figure (adapted from [10, p. 49]) makes the above definitions clear, using the unbounded binary tree as an example.

Figure 1

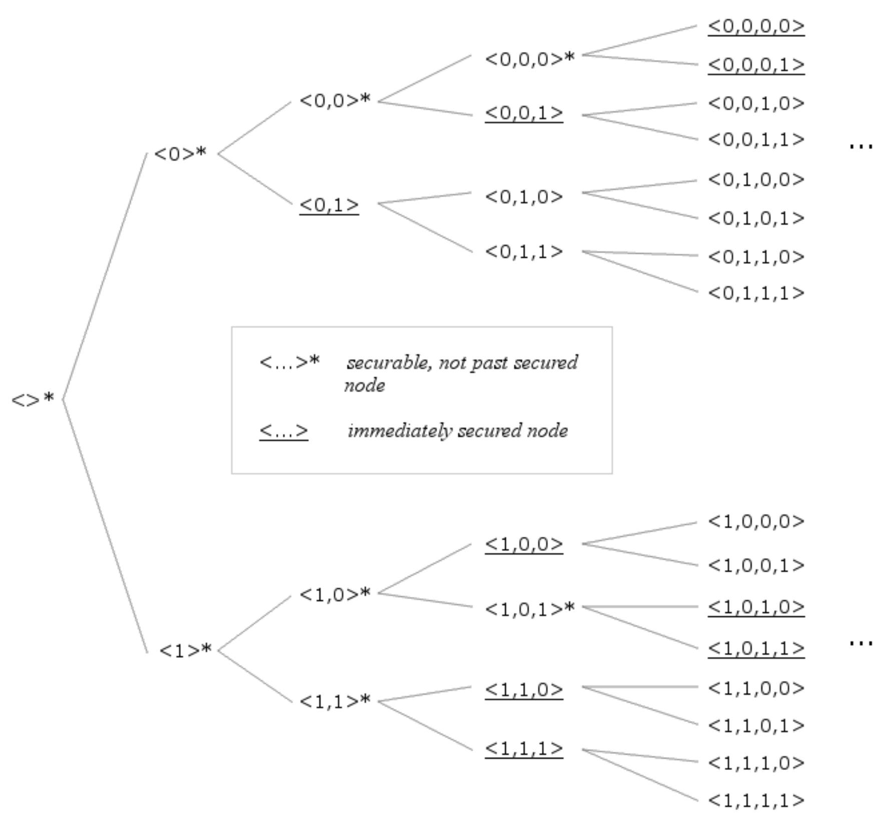

By inspection, it is easy to see that (1) the immediately secured nodes in Figure 1 can be identified as a thin bar in C, and (2) by taking the set of immediately secured and securable (not past secured) nodes in Figure 1, we have a bar in C whose elements form a tree rooted at the empty sequence. If we generalize this example to arbitrary selections of underlined nodes, then it will not necessarily be the case that the set of all immediately secured and securable (not past secured) nodes includes the empty sequence. It is intuitively clear, however, that whenever $\sigma \cdot 0$ and $\sigma \cdot 1$ are both either immediately secured or securable, then it must be the case that σ is also securable. For now, let us inductively define the set I by taking $\sigma \in I$ whenever (1) σ is immediately secured, or (2) $\sigma \cdot 0 \in I$ and $\sigma \cdot 1 \in I$.

We now distinguish the two alternate notions of "securability" which form the crux of Kleene's observations regarding the bar theorem. First, say that a node σ is securable in the *explicit* sense if and only if it fulfills our aforementioned notion of securability, i.e. for all α such that $\alpha \sqsupseteq \sigma$, we are guaranteed $A(\overline{\alpha}n)$ for some n. Intuitively, we can think of σ as explicitly securable whenever every possible unbounded sequence of choices beginning with σ must eventually fulfill the property A – in Figure 1, this amounts to saying that every path passing through σ eventually hits an underlined node.

Second, say that a node σ is securable in the *inductive* sense if and only if σ belongs to the set I. As described earlier, the set I is generated by beginning with the immediately secured nodes and inductively adding nodes in the direction towards the root.

Given a spread such as C and an arbitrary decidable predicate A, let Σ_E be the set of nodes that are explicitly securable with respect to A, and let Σ_I be the set of nodes that are inductively securable with respect to A. Kleene's pivotal observation based on Brouwer's Footnote can then be stated as the following: for all finite sequences σ, $\sigma \in \Sigma_E$ if and only if $\sigma \in \Sigma_I$. In other words, the two definitions of securability are equivalent. The truth of this statement is strongly intuitive when one considers an example such as Figure 1. Here, Σ_E is precisely the set of immediately secured and securable (not past secured) nodes, whereas Σ_I is the set generated by starting with the immediately secured nodes and adding a node towards the root to Σ_I whenever both of its immediate descendants are in Σ_I – in the figure, one can easily see that these two sets are the same. Both classical and constructive proofs of the general equivalence can be found in [10].

Kleene notes that the Bar Theorem can simply be viewed as the implication $\sigma \in \Sigma_E \rightarrow \sigma \in \Sigma_I$. More formally, this can expressed as:

Theorem 2 (Bar Theorem 2)
If A is a decidable predicate on finite sequences, then
$\forall\alpha.\exists n.A(\overline{\alpha}n) \;\rightarrow\; (\,(\forall\sigma.(A(\sigma) \rightarrow Q(\sigma)) \,\wedge\, \forall\sigma.(\forall a.Q(\sigma \cdot a) \rightarrow Q(\sigma))\,) \;\rightarrow\; Q(<>)\,)$

Note that the predicate Q in the consequent essentially captures the property of belonging to the set of inductively securable nodes. Of course, this formulation of the Bar Theorem follows trivially from the equivalence of Σ_E and Σ_I. Furthermore, as Kleene and Vesley state in [10, p. 48]:

> ...Brouwer's Footnote 7 says that securability is that property (of sequence numbers not past secured) which originates at the immediately secured sequence numbers, and propagates back to the unsecured but

> securable numbers across the junctions between a sequence number... and its immediate extensions...

This notion of "backwards propagation" is key to the principle of bar induction, which we shall revisit after discussing the Fan Theorem.

4 The Fan Theorem

As mentioned earlier, the term *fan* is used to denote any finitely branching tree of unbounded depth. That is, T is a fan if every path in T is unbounded in length and every node in T has only a finite number of immediate descendants; the binary fan T_{01} in particular is the fan where every node has exactly two immediate descendants. In the context of choice sequences, we can define a fan as the underlying tree T of a spread in which all the positions in a sequence are restricted to a finite number of choices.

More formally, define a tree T as a decidable set of finite sequences of natural numbers such that: (1) The empty sequence <> is the root of T, and (2) If $\tau \in T$ and $\sigma \sqsubseteq \tau$, then $\sigma \in T$, i.e. T is closed under ascendants. By identifying the finite elements of T as its nodes and taking (σ, τ) as an edge whenever $\sigma \sqsubseteq \tau$, this definition of T can easily be identified with the usual definition of trees under graph theory.

If T is the underlying tree of some spread, then we are guaranteed that every path in T is unbounded in length, i.e. if $\sigma \in T$, then we must have $\tau \in T$ for some $\tau \sqsupseteq \sigma$. Thus if T represents a spread, then the unbounded paths in T correspond exactly to the choice sequences captured by that spread – we can view these choice sequences as the *unbounded* elements of T. Say that the choice sequence α is an element of T if and only if *every* initial segment of α is a node in T, i.e. $\overline{\alpha}n \in T$ for all n.

Let T_U denote the universal tree whose nodes consist of all possible finite sequences of natural numbers. It is obvious that the unbounded elements of T_U are precisely the elements of the Baire space $\mathcal{B}$ – it is in this sense that T_U captures the space of all possible unbounded sequences of natural numbers. Similarly, T_{01} captures the Cantor space $\mathcal{C}$; T_{01} can be viewed as the subtree of T_U consisting exactly of the paths corresponding to the binary sequences. In general, every fan can be viewed as a subtree of T_U. For reasons that will become clear later, we limit our initial discussion of the Fan Theorem to its application on the binary fan.

Formulating the Fan Theorem

A preliminary, intuitive statement of the Fan Theorem can be given as follows [18]:

Theorem 3 (General Fan Theorem 1)
If B is a bar in $\mathcal{C}$, then there exists a finite *set $B^* \subseteq B$ such that B^* is also a bar in $\mathcal{C}$.*

Since B^* is finite, it is clearly decidable. Furthermore, there must be some upper bound N on the length of its elements. This observation gives us the following equivalent statement:

Theorem 4 (General Fan Theorem 2)
If B is a bar in $\mathcal{C}$, then $\exists N.\forall \alpha \in T_{01}.\exists n < N.\overline{\alpha}n \in B$.

Note that if this is to hold constructively, then there must be a method to compute the exact N described. To see that these two statements are equivalent: given N, define $B^* := \{\sigma : \sigma \in B, length(\sigma) \leqslant N\}$; given B^*, let $N = \max_{\sigma \in B^*}(length(\sigma))$.

While Theorem 3 can be viewed as a consequence of the Bar Theorem (by which there exists a thin, well-ordered bar $B^* \subseteq B$ that is finite since T_{01} is a fan), the following proof shows how it can be seen directly from Brouwer's Thesis.[10]

Proof. Assuming Brouwer's Thesis, there exists a canonical proof for every bar B' that is a bar in $\mathcal{B}$. Suppose $B \subset B'$ is a bar in $\mathcal{C}$. Viewing $\mathcal{C}$ as a subspace of $\mathcal{B}$, from the existence of a canonical proof that B' is a bar in $\mathcal{B}$, we can deduce that a proof-tree of the following form exists[11]:

– The root of the proof-tree, which is the conclusion of the proof, is the statement: "B bars <> in $\mathcal{C}$".

– For every $\sigma \in B$, we have available the η-inference: "$\sigma \in B$, therefore B bars σ in $\mathcal{C}$". Since T_{01} is a fan, only a finite number of such η-inferences are needed to complete the proof-tree.

– Whenever B bars both $\sigma \cdot 0$ and $\sigma \cdot 1$, we have a F-inference with the conclusion: "B bars σ in $\mathcal{C}$".

– Whenever we have the statement "B bars σ", we have two ρ-inferences with the conclusions "B bars $\sigma \cdot 0$" and "B bars $\sigma \cdot 1$", respectively.

Note that this proof-tree proves the statement "B is a bar in $\mathcal{C}$" and is a sub-tree of the canonical proof that B' is a bar in $\mathcal{B}$. Furthermore, it requires only finitely many η-inferences. Construct the finite set $B^* \subseteq B$ by defining $\sigma \in B^*$ if and only if "$\sigma \in B^*$" appears in an η-inference of the proof-tree. Since the same proof-tree is also sufficient for proving that B^* bars <>, we have a proof of the statement "B^* is a bar in $\mathcal{C}$", as desired.

□

The following variations in the expression of the Fan Theorem are useful in exploring the impact of possible assumptions that can be made about the premises of the theorem. We express these for now as propositions and discuss the logical relations between them.

[10] This proof is adapted from [18].

[11] Note that for every conclusion "B' bars σ in $\mathcal{B}$" in the canonical proof, where σ is a finite binary sequence, we have a fortiori "B bars σ in $\mathcal{C}$".

Proposition 5 (FAN_D)
$\forall \sigma \in T_{01}.(A(\sigma) \vee \neg A(\sigma)) \wedge \forall \alpha \in T_{01}.\exists n.A(\overline{\alpha}n) \rightarrow \exists N.\forall \alpha \in T_{01}.\exists n \leqslant N.A(\overline{\alpha}n)$

Fan_D says that if: (1) A is a decidable property with respect to all finite binary sequences, and (2) every unbounded binary sequence has some initial segment satisfying A, then we can conclude that there is a *uniform upper bound* N on the length of the initial segments posited by (2). That is, for every α in T_{01}, we have $\overline{\alpha}n$ for some number $n \leqslant N$.

Note the relation of FAN_D to Theorem 4. Assuming that FAN_D holds, we can define a bar B in C by taking a subset of the finite sequences posited by the assumption $\forall \alpha \in T_{01}.\exists n.A(\overline{\alpha}n)$, specifically the set $\{\sigma : A(\sigma) \wedge (\forall \sigma' \prec \sigma.\neg A(\sigma))\}$. In fact, such a set can be identifed with the thin bar B^* guaranteed by Theorem 3.

It is also straightforward to see that FAN_D follows from Kleene's version of the Bar Theorem. Intuitively, any bar B in C will be formed by some collection of immediately secured and securable (past secured) nodes with respect to some decidable property. In terms of the bar, all past secured nodes are redundant – by removing all such nodes we obtain a thin bar $B' \subseteq B$. Furthermore, since T_{01} is a fan, B' can not be infinite in size without eventually having redundant nodes. We must therefore conclude that B' is finite, i.e. there exists an upper bound on the length of sequences in B'. Kleene's proof that the explicit and inductive definitions of securability are equivalent thus provides us with a method of proving FAN_D that does not rely on Brouwer's Thesis – moreover, Kleene's version of the proof is both classically and constructively true.

The following strengthening of Fan_D omits the assumption that A is decidable for *all* nodes in T_{01}, requiring only that A holds for some initial segment of every unbounded sequence:

Proposition 6 (FAN)
$\forall \alpha \in T_{01}.\exists n.A(\overline{\alpha}n) \rightarrow \exists N.\forall \alpha \in T_{01}.\exists n \leqslant N.A(\overline{\alpha}n)$

This version is close in spirit to Brouwer's original description of the Fan Theorem, in which the property A is instead viewed as a relation associating each initial segment with some natural number [15]:

$$\forall \alpha \in T_{01}.\exists n.A(\overline{\alpha}n, p) \rightarrow \exists N.\forall \alpha \in T_{01}.\exists n \leqslant N.A(\overline{\alpha}n, p)$$

Now suppose that FAN holds. If we assume the continuity principle C-N, then whenever we can determine $A(\alpha, p)$ based only on $\overline{\alpha}n$, we can conclude that $A(\beta, p)$ holds for all β that pass through $\overline{\alpha}$n. Furthermore, by FAN we have that if $A(\alpha, p)$ can be determined for every α in T_{01} based on some initial segment, then there must be a universal upper bound N on the minimum lengths of the initial segments needed to compute $A(\alpha, p)$ for any α. Thus by accepting C-N, we can generalize FAN to the following statement:

Proposition 7 (FAN^*)
$\forall \alpha \in T_{01}.\exists p.A(\alpha, p) \rightarrow \exists N.\forall \alpha \in T_{01}.\exists p.\forall \beta \sqsupseteq \overline{\alpha}N.A(\beta, p)$

It can be shown that FAN* is not classically valid – intuitively, this result is unsurprising since C-N itself is not classically valid. Note also that in the presence of C-N, FAN_D, FAN, and FAN* are all equivalent [14].

Up until now we have focused on binary fans in particular. As it turns out, there is no loss of generality in limiting the discussion of the Fan Theorem to binary fans only. The following result shows that the Fan Theorem, as stated for arbitrary fans, can always be reduced to a statement about binary fans.

Proposition 8
Every fan T can be homeomorphically mapped to a subfan T'_{01} of the binary fan T_{01}.

That is, there exists a (topologically) continuous function Φ such that: (1) Φ forms a bijection between the paths in T and the paths in T'_{01}, and (2) Φ is invertible [4]. The idea is to uniquely encode every unbounded path α in T as the binary sequence

$$1^{\alpha(0)+1}, 0, 1^{\alpha(1)+1}, 0, 1^{\alpha(2)+1}, 0, 1^{\alpha(3)+1}, 0, ...$$

where 1^n denotes a sequence of n consecutive 1's. This reversible encoding process is then taken to be the mapping Φ.

Given any arbitrary fan T, we can effectively compute its unique corresponding binary subfan T'_{01}. It can be shown that if FAN_D (or FAN, FAN*, respectively) is valid for T_{01}, then it must also hold for T'_{01}, and thus we can conclude that FAN_D, FAN, and FAN* hold for all fans in general [14].

Relating the Fan Theorem to the Classical König's Lemma

In its application on trees, König's (Infinity) Lemma, a classical result from graph theory, states that if T is a finitely branching tree with infinitely many nodes, then there exists at least one unbounded path in T. Classically, FAN is equivalent to König's Lemma by way of contraposition – it is possible to equate FAN with the statement, "If T is a finitely branching tree and every branch of T is finite, then T must have a finite number of nodes".

To see this interpretation of FAN, consider the following. Let T_B be a finitely branching binary tree in which every branch is finite. Let T'_{01} be a subfan of the binary fan such that for every unbounded sequence α, $\alpha \in T'_{01}$ if and only if α passes through some finite sequence $\sigma \in T_B$. Define the property A so that for any σ, $A(\sigma)$ holds if and only if $\sigma \notin T_B$. By FAN, there is an upper bound on the length of the paths in T_B. In this sense, FAN can be viewed as asserting that the number of nodes in T_B is finite.

That FAN and König's Lemma are contrapositives of one another can also be seen from the following expression of the latter:

Theorem 9 (KL)
If T is a fan, then $\forall N.\exists \alpha \in T.\forall n \leqslant N.\neg A(\overline{\alpha}n) \rightarrow \exists \alpha \in T.\forall n.\neg A(\overline{\alpha}n)$

For every bound N, define $B_N := \{\sigma : (length(\sigma) \leqslant N) \wedge A(\sigma) \wedge (\forall \sigma' \prec \sigma. \neg A(\sigma))\}$. With respect to fans in particular, KL says that if for *every* N there is some $\alpha \in T$ such that B_N does not bar α, then we can conclude that there exists at least one particular path α_0 in T such that A does not hold for *any* initial segment of α_0, i.e. α_0 is not barred by any B_N.

The following is a classical proof of KL[12]:

Proof. Let T be an arbitrary fan and fix the property A. For every number N, define B_N as above. Let $B := \bigcup_{N \in \mathbb{N}} B_N$, and note that for any finite sequence σ of length N or less such that $\sigma \notin B_N$, we can conclude that $\sigma \notin B$.

Assume the antecedent of KL, so that for every N, there is some $\alpha \in T$ such that B_N does not bar α. In particular, for every N, there is at least one finite sequence σ of length N such that we have: (1) $\neg A(\sigma)$, and (2) $\neg A(\tau)$ for all $\tau \prec \sigma$. Since there may be more than one such σ for any fixed N, denote by Σ_N the finite set of all N-length σ satisfying criteria (1) and (2).[13] By definition, if $\sigma \in \Sigma_N$, then $\sigma \notin B_N$.

Fix N and Σ_N. Now consider the set Σ_{N+1}: for all $\sigma' \in \Sigma_{N+1}$, it must be the case that σ' extends some $\sigma \in \Sigma_N$, otherwise criteria (2) would not hold for σ'. In general, for any arbitrary Σ_N and $\Sigma_{N'}$ such that $N' > N$, there must exist $\sigma \in \Sigma_N$ and $\sigma' \in \Sigma_{N'}$ such that $\sigma \prec \sigma'$.

Thus we can classically define a set K as the set of all sequences σ such that for any arbitrary length $\ell > length(\sigma)$, there exists some ℓ-length sequence $\sigma' \in \Sigma_\ell$ that is a descendant σ. Clearly, K is a subset of $\overline{B}$. From the assumed antecedent of KL, we must have $<> \in K$, since $<>$ is the only node at depth 0. For every $\sigma \in K$, let S_σ denote its finite set of immediate descendants. By definition, if $\sigma \in K$, then $\sigma' \in K$ for some $\sigma' \in S_\sigma$.

Therefore, beginning with $\alpha_0(0) = <>$, there must exist a particular unbounded sequence α_0 with the property that for any choice of n, $\overline{\alpha}n \in K$ and hence $\overline{\alpha}n \notin B$. Such an α_0 fulfills the consequent, $\exists \alpha \in T. \forall n. \neg A(\overline{\alpha}n)$, as desired. □

From the classical proof of KL, it follows that FAN is also classically true. Since every decidable bar can be viewed as arising from some decidable property of all finite sequences, KL essentially describes the conditions under which a given decidable property can not be successfully used to define a bar. Note however that the classical proof of KL does not hold constructively – despite showing that a particular α_0 must exist for which A is never satisfied, the proof does not provide a constructive method for exhibiting this exact sequence.[14]

[12] This proof is loosely adapted from [9].

[13] Since T is a fan, the number of nodes at any fixed depth N is finite, thus for any N, Σ_N must be finite.

[14] Based on the classical proof, one way to encounter the posited α_0 is to perform a depth-first search of T, backtracking whenever we encounter a node in B – this method will of course fail in practice to actually identify α_0, since at any given point in the search there is no way to distinguish α_0 from a path that will eventually "hit B" at some depth in the future.

If we accept the principle of the excluded middle, $(p \vee \neg p)$, as in classical logic, then for any property A of finite sequences σ, $(A(\sigma) \vee \neg A(\sigma))$ always holds. Thus classically speaking, there is no meaningful distinction between the forms of the Fan Theorem FAN_D and FAN, so FAN_D must also be classically true. It can be shown directly that FAN_D and König's Lemma (in its general form as applied to trees) are classically equivalent[15]:

Proof. (König's Lemma $\Rightarrow$ FAN_D)

Let T be an arbitrary fan. Assume that we have $\forall \alpha \in T.\exists n.A(\overline{\alpha}n)$. Classically, we can take for granted that $\forall \sigma \in T.(A(\sigma) \vee \neg A(\sigma))$. It follows that the tree $T_A := \{\sigma \in T : \forall \tau \preceq \sigma.\neg A(\tau)\}$ has only finite paths. Intuitively, T_A is the maximal subtree of T such that: (1) T_A and T share the same root, and (2) every node of $\sigma \in T_A$ satisfies $\neg A(\sigma)$.

Suppose by contradiction that the consequent of FAN_D does not hold, i.e. we have $\neg(\exists N.\forall \alpha \in T.\exists n \leqslant N.A(\overline{\alpha}n))$, which is classically equivalent to the statement: $\forall N.\exists \alpha.\forall n \leqslant N.\neg A(\overline{\alpha}n)$. This means that there is no upper bound on the length of paths in T_A – for every number n we can find some (finite) path in T_B of length n, and thus there is no upper bound on the number of nodes in T_A. By König's Lemma, there exists an unbounded path in T_A, i.e. $\exists \alpha.\forall n.\neg A(\overline{\alpha}n)$, which contradicts our original assumption.

Therefore it must be the case that $\exists N.\forall \alpha \in T.\exists n \leqslant N.A(\overline{\alpha}n)$, so FAN_D holds as desired. □

Proof. (FAN_D $\Rightarrow$ König's Lemma)

Classically, FAN_D, FAN, and the contrapositive of FAN (i.e. KL) are equivalent. We assume KL. Let T_A be a finitely branching tree with an unbounded number of nodes. Let T be a fan such that every (unbounded) path in T passes through some finite path in T_A. Define the property A so that $A(\sigma)$ holds if and only if $\sigma \in T$ but $\sigma \notin T_A$. In particular, for every $\sigma \in T$, $\neg A(\sigma)$ holds if and only if $\sigma \in T_A$.

Notice that there must exist paths of arbitrary length in T_A. Thus for any choice of N, we have $\exists \alpha \in T.\forall n \leqslant N.\neg A(\overline{\alpha}n)$. By KL, there must exist an unbounded path $\alpha \in T$ such that $\neg A$ holds for every initial segment of α, so by definition there must exist some unbounded branch in T_A. □

It is apparent that classically speaking, FAN_D, FAN, KL, and the general König's Lemma as applied to trees are all equivalent. On the other hand, from the intuitionistic perspective, that A is a decidable property for every finite sequence cannot be taken for granted. Therefore FAN_D (rather than FAN) should be viewed as the intuitionistic counterpart to König's Lemma. Interestingly, Brouwer's original proof of the Fan Theorem in 1924 predates the first appearance of König's Lemma [15].

[15] The proof of the forward implication (König's Lemma $\Rightarrow$ FAN_D) is adapted from [14]

5 Kleene's Counterexample to the Fan Theorem

From the preceding formulations of the Fan Theorem, it was implied that the statements "B is a bar in C" and "$\forall\alpha \in T_{01}.\exists n.A(\overline{\alpha}n)$" are equivalent. It turns out that the precise interpretation of the intuitive statement "B is a bar" has a significant impact on whether or not the Fan Theorem can be proved constructively. In particular, Kleene showed via counterexample that if all the unbounded sequences of natural numbers are viewed as arising from recursive functions (or constructive functions as defined by Bishop), then the Fan Theorem is *not* constructively true.

To see this, we first define *Kleene's T-predicate*. Say that a function is "computable (by algorithm)" or "algorithmic" if it belongs to the class of recursive functions. The term "algorithmic" is by necessity informal; intuitively, such a function can be viewed as fully describable by some *finite* specification. Formally, the class of (partial) recursive functions is equivalent to: (1) the Turing-computable functions, (2) the class of functions described by Church's lambda calculus, and (3) the formal characterization of computable functions according to Kleene's T-predicate [11]. Furthermore, Bishop's constructive definition of functions, which hinges on the existence of a "finite routine" mapping every input to its output, can intuitively be viewed as describing the same class of recursive functions [5].

Note that there exists an algorithmic mapping that assigns a unique code number x to each intensional instruction set or "program" used to compute some recursive function.[16] These code numbers are referred to as the *Gödel numbers* of the programs. Given any number x, denote by P_x the program whose assigned Gödel number is x, and denote by φ_x the partial function computed by P_x.

Where A is any predicate or relation, let χ_A denote the characteristic function of A.[17] Say that A is algorithmic if and only if χ_A is algorithmic. Kleene's T-predicate is an algorithmic predicate defined as follows: $T(x, y, z)$ holds if and only if (1) x is the Gödel number corresponding to a program P_x that computes some recursive function φ_x, (2) y is a number accepted as input by P_x, and (3) z is the Gödel number corresponding to a program that simulates P_x on input y and successfully terminates. In other words, $T(x, y, z)$ holds exactly when $\varphi_x(y)$ is computable. Define the algorithmic function U such that if $T(x, y, z)$ holds, then $U(z) = \varphi_x(y)$. In the literature, U is called the *result-extracting function* corresponding to T [14].

Say that an unbounded sequence α is recursive if and only if there exists a total recursive function φ_α such that for all n, $\alpha(n) = \varphi(n)$. Kleene's counterexample to the Fan Theorem can now be expressed as follows [14]:

Theorem 10 (Kleene's Counterexample to FAN_D)
Suppose that every unbounded sequence α is recursive. There exists a decidable pred-

[16] In general, each recursive function may be given by more than one program. The existence of such a mapping implies (classically) that the recursive functions are denumerable [11].

[17] I.e. If A is a predicate, for any input x: $\chi_A(x) = 1$ if and only if $A(x)$ holds, otherwise $\chi_A(x) = 0$. Similarly, if A is a relation, for any input N: $\chi_A(x, y) = 1$ if and only if $(x, y) \in A$, otherwise $\chi_A(x, y) = 0$.

icate R such that $\forall \sigma, \tau \in T_{01}.(R(\sigma \cdot \tau) \to R(\sigma))$, *and: (1)* $\forall \alpha.\exists n.\neg R(\overline{\alpha}n)$, *and (2)* $\forall n.\exists \sigma.((length(\sigma) = n) \wedge R(\sigma))$.

Proof. [18] Let T be Kleene's T-predicate and define the predicate R such that for any n:

$$R(\overline{\alpha}n) := \forall m < n.((\exists u < n.p_{01}) \wedge (\exists u < n.p_{10})),$$
$$\text{where } p_{01} := T(m, m, u) \wedge (U(u) = 0) \to \alpha(m) = 1$$
$$\text{and } p_{10} := T(m, m, u) \wedge (U(u) = 1) \to \alpha(m) = 0.$$

Intuitively, the proposition p_{01} says that if $\varphi_m(m) = 0$, then $\alpha(m) = 1$; p_{10} says that if $\varphi_m(m) = 1$, then $\alpha(m) = 0$. Thus $R(\overline{\alpha}n)$ holds exactly when every position $m < n$ in $\overline{\alpha}n$ fulfills the properties: (i) $\alpha(m) \neq \varphi_m(m)$, and (ii) $\varphi_m(m)$ can be computed by a program whose Gödel number is less than n.

(2) Given any n, choose the (recursive) unbounded binary sequence α such that $\alpha(m) = 1$ if and only if: (i) $m < n$, and (ii) for some $u < n$, we have $T(m, m, u) \wedge U(u) = 0$. Thus $\forall n.\exists \sigma.((length(\sigma)) = n \wedge R(\sigma))$ holds.

(1) Let φ_m be a function corresponding to α and suppose we have $T(m, m, u)$. Suppose that we have $R(\overline{\alpha}(u + m + 1))$. Since α is binary, either $U(u) = 0$ or $U(u) = 1$. But if $U(u) = 0$, then $\exists v < u + m + 1.(T(m, m, v) \wedge U_v = 0)$. Thus according to the program P_v, we have $\alpha(m) = 1$, i.e. $\varphi_m(m) = 1$, which contradicts our assumption that $T(m, m, u)$. If $U(u) = 1$, we reach an analogous contradiction. Therefore it must be the case that $\neg R(\overline{\alpha}(u + m + 1))$, and in general we have $\forall \alpha.\exists n.\neg R(\overline{\alpha}n)$, as desired. □

Notice that the proof of Kleene's counterexample utilizes a contradiction arising from $T(m, m, u)$, i.e. the situation when the program P_m receives as input its own Gödel number. In a sense, we can view this proof as a topological form of diagonalization.

An immediate consequence of Theorem 10 is that FAN_D does not hold if every α in T_{01} (and a fortiori, if every unbounded sequence of natural numbers) can be given by a total recursive function. Define the decidable predicate A such that for all finite binary sequences σ, $A(\sigma) \leftrightarrow \neg R(\sigma)$. From property (1) of R, we have $\forall \alpha.\exists n.A(\overline{\alpha}n)$, and yet by (2), there are finite sequences of arbitrary length for which A is not satisfied. Thus if we make the assumption that every sequence α is recursive, then there is no universal upper bound N such that $\forall \alpha.\exists n < N.A(\overline{\alpha}n)$.

An Alternate Interpretation of "B is a Bar"

One possible solution to the issues that arise from equating the intuitive statement "B is a bar" with the formulation "$\forall \alpha.\exists n.A(\overline{\alpha}n)$" is to define the former without referencing any unbounded sequences at all, thereby sidestepping the issue of

[18] This version of the proof is from [14].

how the quantifier ”$\forall\alpha$” should be interpreted in the antecedent of FAN_D, FAN, and FAN*. This approach is discussed explicitly in [9] and implicitly in [14].

Let B be a set of finite binary sequences. Define the predicate $B \mid \sigma$ inductively as follows: (1) if $\sigma \in B$, then $B \mid \sigma$, and (2) if $B \mid \sigma \cdot 0$ and $B \mid \sigma \cdot 1$, then $B \mid \sigma$. The idea then is to say that "B is a bar" if and only if a proof-tree for $B |{<>}$ can be constructed, using statements of type (1) as starting points. As discussed in the proof of Theorem 3, if B is a bar (in the intuitive sense), then there exists a canonical proof of the fact – by inspection it is evident that the proof-tree for $B |{<>}$ can be identified with a finite rooted sub-tree of the canonical proof-tree. In particular, the terminal nodes of the proof-tree for $B |{<>}$ can be identified with the statements in the canonical proof of the form "$\sigma \in B$". Notice also that the proof-tree of $B |{<>}$ can be viewed as a *well-ordered* or *well-founded inductive binary tree* as defined in [14]. Moreover, as we observed earlier, this proof-tree is finite since B bars a fan. Hence, if we interpret the Fan Theorem as:

Theorem 11 (Inductive Fan Theorem)
$B |{<>} \;\rightarrow\; \exists N.\forall\alpha \in T_{01}.\exists n < N.\overline{\alpha}n \in B,$

then N can easily be found by determining the depth of the proof-tree for $B |{<>}$. This expression of the Fan Theorem is thus constructively true. Crucially, the definition of "B is a bar" in this case is based on a predicate defined solely from finite sequences. There is again no loss of generality in considering the binary fan in particular rather than arbitrary fans.

This interpretation of what it means for a set to constitute a bar informs the concept of *bar induction*, an induction principle of practical use for proving certain properties of unbounded sequences [14]:

Proposition 12 (Decidable Bar Induction)

$$\forall\alpha.\exists n.P(\overline{\alpha}n) \wedge \forall\sigma.(P(\sigma) \vee \neg P(\sigma)) \wedge \forall\sigma.(P(\sigma) \rightarrow Q(\sigma)) \\ \wedge \forall\sigma.(\forall a.Q(\sigma \cdot a) \rightarrow Q(\sigma)) \;\rightarrow\; Q(<>)$$

That is, a property Q can be shown to hold for the empty sequence $<>$ if the following four conditions are fulfilled: (1) P is a decidable property for all finite sequences σ (2) Every unbounded sequence has some initial segment that satisfies P, (3) For any σ, Q holds if P holds, and (4) Q holds for σ if Q holds for every immediate descendant of σ. Practically speaking, whenever $Q(<>)$ is a predicate that holds if and only if Q holds for every unbounded sequence, bar induction can be used as a method for constructively proving that certain properties hold for all the choice sequences captured by the underlying tree of a spread.

Notice the close resemblance between decidable bar induction and Kleene's version of the Bar Theorem: the latter is a statement of the form $p \rightarrow q$, while the former takes the form $p \wedge q$. Kleene's proof of the equivalence between the explicit and inductive definitions of node securability thus serves additionally as a proof of the bar induction schema.

The following is a generalized version of bar induction in which we assume that P is monotone rather than decidable and conclude that P itself holds at the root:

Proposition 13 (Monotone Bar Induction)

$$\forall\alpha.\exists n.P(\overline{\alpha}n) \wedge \forall\sigma\tau.(P(\sigma) \rightarrow P(\sigma \cdot \tau)) \wedge \forall\sigma.(\forall a.P(\sigma \cdot a) \rightarrow P(\sigma)) \;\rightarrow\; P(<>)$$

It can be shown that in the presence of C-N, these two forms of bar induction are equivalent [14].

Far from being an axiom schema of purely theoretical interest, the practical relevance of bar induction has been shown by the role it plays in constructive type theories such as Nuprl[19] – a recent description of the realizer for strong bar induction in Nuprl, inspired by Kleene's proof of bar induction, can be found in [4]. This version of bar induction is strong in the sense that while the bar must be decidable, the spread of interest itself need not be. From this, we can see that bar induction is implementable in the most practical sense of the word – it is possible to (constructively) define a program for drawing conclusions of the form given by the bar induction principle.

6 On the Incompatibility of the Fan Theorem and Church's Thesis

Informally, *Church's Thesis*[20] refers to the notion that the class of recursive functions sufficiently captures the the informal notion of what it means for a function to be "algorithmic". Since "algorithm" is an intuitive concept rather than a formal one, it is not possible to prove or disprove Church's Thesis – it is largely accepted in computer science, and in particular in computability theory, due to a combination of empirical evidence and the fact that the acceptance of Church's Thesis as an axiom has yielded many fruitful outcomes [11]. The constructive version of Church's Thesis, which is a strengthening of the aforementioned concept, is as follows[21]:

$$\forall\alpha.\exists x.\forall y.\exists z.(T(x, y, z) \wedge U(z) = \alpha(y)) \tag{CT}$$

That is, CT makes explicit the idea that the lawlike sequences can be identified with the recursive functions. Moreover, the quantifier $\forall\alpha$ implies that *every* infinite sequence of natural numbers can be given by some total recursive function. An

[19] Nuprl is currently the only fully implemented intuitionistic foundational theory, both for mathematics and computer science. Further information on Nuprl can be found in [1] and [8].

[20] i.e. the Church-Turing Thesis from computability theory

[21] In the rest of this paper, the term "Church's Thesis" will generally be used to denote the constructive version.

intriguing consequence of Kleene's Theorem 10 is the revelation of the following fact: the Fan Theorem (FAN_D) is incompatible with Church's Thesis (CT).

If FAN_D holds for the space of all unbounded sequences in a fan, yet fails to hold for recursive sequences in particular, then there must exist paths in the fan that do not correspond to any total recursive function. In retrospect, this concept can be seen as a classical justification of Brouwer's use of choice sequences to define the real numbers: since the real numbers are non-denumerable, the number of choice sequences available must be greater than the number of total recursive functions, of which there are (classically) countably many.

Nevertheless, the conflict between the Fan Therorem and CT brings to light a caveat regarding the intuitionistic conception of unbounded sequences – certain significant properties of choice sequences cannot be assumed in the presence of CT. This is perhaps surprising, considering that CT (viewed as applying to the lawlike sequences only) has been shown to be compatible with most other aspects of intuitionism [14].

Essentially, proofs that are relativized on CT are as a whole incompatible with the Fan Theorem, which is itself relavitized on Brouwer's Thesis – this alludes to a deep incommensurability between the two systems of thought. Not only does this conflict apply to the way in which unbounded sequences are visualized, but it also raises the question of whether or not the real numbers can be appropriately represented by total recursive functions. In fact, in Bishop's seminal contribution to constructive mathematics (published in 1967), he rejected the concept of choice sequences in general [2]. It was later shown that Bishop's development of constuctivism is compatible with CT [2], though without the lawless sequences, his focus had to shift from describing the continuum in its totality towards defining individual real numbers [15].

7 The Continuum as a Fan

Brouwer's original motivation for developing the Fan Theorem was to achieve, intuitionistically, the result that the continuum is uniformly continuous. By modeling the real numbers as unbounded sequences of nested rational intervals, Brouwer was able to avoid describing the continuum as a collection of discrete objects, and as the Fan Theorem would show, the property of uniform continuity could be intuitionistically derived. As van Atten describes in [15, p. 33-34]:

> Brouwer's achievement is to have found a way to analyze the continuum that does not let it fall apart into discrete elements (as extensional equality of choice sequences is not decidable), and is constructive to boot.

The issue of showing that the continuum is uniformly continuous reduces to the issue of showing this property for reals in the interval [0, 1], since the same argument can be applied to the real interval between any two consecutive numbers. In particular, the goal is to show that every total function on [0, 1] is (uniformly) continuous.[22]

First define a real number as an unbounded sequence of rational intervals $I_0, I_1, I_2, ...$, such that every interval has the form

$$\left[\frac{a}{2^{k+1}}, \frac{a+2}{2^{k+1}}\right]$$

where $2 \leqslant a + 2 \leqslant 2^{k+1}$. Note the following observations: (1) Every such interval lies within the interval [0 ,1], (2) Every interval I_j can be uniquely associated with a pair of numbers <a, k>, and thus intervals of this form are denumerable, and (3) For a fixed interval I, there are only finitely many intervals I_j such that I_j is a proper subinterval of I.

Now define the spread J by taking as its spread law the condition that if $< a_0, a_1, ..., a_n >$ is an admitted sequence, then its extension $< a_0, a_1, ..., a_n, a_{n+1} >$ is admitted if and only if I_{a_n+1} is a proper subinterval of I_{a_n}. Let T_J denote the underlying tree of J; from observation (3), T_J is a fan, and thus there is a corresponding subfan of the binary fan that also represents J. It can be shown that every unbounded path in T_J represents a real number in $[0, 1]$, and conversely, every real number in [0,1] corresponds to some unbounded path in T_J [15].

Therefore, the spread J can be viewed as a sufficient representation of the continuum $[0, 1]$. From this, it is clear that the Fan Theorem can be invoked to show that if uniform continuity holds over certain finite sequences of J (i.e. those that are barred), then the same property must hold over *all* the reals in $[0, 1]$.

Thus by applying the Fan Theorem to the fan T_J, Brouwer achieved the following results [15]:

Theorem 14 (Uniform Continuity Theorem)
$\forall \varepsilon . \exists \delta . \forall x_1 x_2 . (\mid x_1 - x_2 \mid < \delta \ \rightarrow \mid f(x_1) - f(x_2) \mid < \varepsilon)$, *where* δ, ε *are positive reals and* $x_1, x_2 \in [0, 1]$

which implies the corollary:

Theorem 15 (Continuity Theorem)
$\forall \varepsilon . \forall x_1 . \exists \delta . \forall x_2 . (\mid x_1 - x_2 \mid < \delta \ \rightarrow \mid f(x_1) - f(x_2) \mid < \varepsilon)$, *where* δ, ε *are positive reals and* $x_1, x_2 \in [0, 1]$

The first theorem says that every total function from the interval $[0, 1]$ to $\mathbb{R}$ is uniformly [23] continuous, while the second, weaker statement asserts only that the functions are continuous.

[22]The following description is adapted from [15].
[23]Here, the term "uniform" can be seen as applying to the condition that there exists a *single* δ such that $(\mid x_1 - x_2 \mid < \delta \ \rightarrow \mid f(x_1) - f(x_2) \mid < \varepsilon)$ holds for any arbitrary x_1, x_2.

8 Discussion

In its historical context, one of the most important aspects of the Fan Theorem is the role that it played in allowing Brouwer to show that the intuitionistically conceived continuum is not only continuous, but uniformly so. There are other notions of computational and philosophical interest that can be gleaned from the Fan Theorem. In particular, the Fan Theorem highlights the dichotomy between a deterministic and non-deterministic conception of "choice". To accept the lawless sequence as a mathematical construct extensionally distinct from the lawless sequence is to accept that in the universe of unbounded sequential choices, there exist successions of choices that cannot be described by any pre-determined law or "algorithm". This is a key distinction and the underlying souce of conflict between the Fan Theorem and Church's Thesis. In a sense, the Fan Theorem embodies a crossroads with respect to constructive analysis. Kleene showed that whether or not we accept the Fan Theorem, its existence demonstrates that it is not possible to simultaneously believe that every sequence can be algorithmically generated, and that there exists a computable covering for the space of all unbounded sequences when it comes to certain "universal" properties – the two situations must be mutually exclusive. Despite the troubling implications of Kleene's counterexample, there are nevertheless reasons both intuitive and practical to accept some form of the Fan Theorem: at once it can be viewed as an intuitionistic version of König's Lemma, an expression of compactness for $\mathcal{C}$ (and even $\mathcal{B}$), and a constructive method of effectively covering a space otherwise difficult, perhaps impossible, to fully envision.

References

[1] Allen, Stuart, Mark Bickford, Robert Constable, Richard Eaton, Christoph Kreitz, Lori Lorigo, and Evan Moran (2006). "Innovations in Computational Type Theory using Nuprl." *Journal of Applied Logic* 4.4: 428-469. 19

[2] Beeson, M. (2005). "Constructivity, Computability, and the Continuum." *Essays on the Foundation of Mathematics and Logic*, Polimetrica S.a.s.: 9-37. 20

[3] Berger, J. (2012). "Aligning the weak König lemma, the uniform continuity theorem, and Brouwer's fan theorem." *Annals of Pure and Applied Logic* 163: 981-985.

[4] Bickford, Mark S. and Robert Constable. "Inductive Construction in Nuprl Type Theory Using Bar Induction." *Types 2014: Types for Proofs and Programs*. Paris, France. 19

[5] Bridges, D. S. (1999). "Constructive mathematics: a foundation for computable analysis." *Theoretical Computer Science* 219: 95-109. 16

[6] Brouwer, L. E. J. (1913). "Intuitionism and formalism." *Bull. Amer. Math. Soc.* 20.4: 81-96. 1

[7] Brouwer, L. E. J. (1975). *Collected Works, Volume 1*. Ed. A. Heyting. North-Holland Publishing Co., Amsterdam. 1

[8] Constable, Robert L., Stuart F. Allen, H. M. Bromley, W. R. Cleaveland, J. F. Cremer, R. W. Harper, Douglas J. Howe, T. B. Knoblock, N. P. Mendler, P. Panangaden, James T. Sasaki, and Scott F. Smith. (1986). *Implementing Mathematics with the Nuprl Proof Development System*. Prentice-Hall, New Jersey. 19

[9] Coquand, T. About Brouwer's Fan Theorem. (2004). *Revue internationale de philosophie* 4 no. 230: 483-489. 14, 18

[10] Kleene, S. C. and R. E. Vesley. (1965). *Foundations of Intuitionistic Mathematics*. North-Holland Publishing Co.: 42-59. 8, 9

[11] Rogers, Jr., H. (1987). *Theory of Recursive Functions and Effective Computability*. McGraw-Hill Book Company: 18-31. 16, 19

[12] Troelstra, A.S. (1991). "History of constructivism in the 20th century." *ITLI Publication Series* ML-91-05: 1-32. 1

[13] Troelstra, A. S. (1996). "Choice sequences: a retrospect." *CWI Quarterly* 9.1-2: 143-149.

[14] Troelstra, A. S. and van Dalen, D. (1988) *Constructivism in Mathematics: An Introduction, Volume 1 (Studies in Logic and the Foundations of Mathematics, Volume 121)*. North-Holland Publishing Co., Amsterdam: 139-243. 1, 4, 6, 7, 13, 15, 16, 17, 18, 19, 20

[15] van Atten, M. (2004). *On Brouwer*. Wadsworth Philosophers Series. Wadsworth/Thomson Learning, Belmont: 30-63. 1, 5, 6, 7, 12, 15, 20, 21

[16] van Stigt, Walter P. (1990). *Brouwer's Intuitionism*. North-Holland Publishing Co., Amsterdam. 1

[17] Veldman, W. (2006). "Brouwer's Real Thesis on Bars." *Philosophia Scientiae* Cahier spécial 6: 21-42.

[18] Veldman, W. (2008). "Some applications of Brouwer's Thesis on Bars." *One Hundred Years of Intuitionism (1907-2007)*. Birkhäuser Basel: 326-340. 7, 10, 11

[19] Veldman, W. (2014). "Brouwer's Fan Theorem as an axiom and as a contrast to Kleene's alternative." *Archive for Mathematical Logic* 53.5-6: 621-693. 6