Chapter 4

Brouwer's proof of the bar theorem

4.1 The theorem and its philosophical interest

A formulation of the bar theorem must be postponed until some basic concepts that figure in it have been introduced, but its place in intuitionism can already be indicated: it is the apotheosis of Brouwer's reconstruction of mathematics, both mathematically and philosophically. The bar theorem allowed him to go beyond other varieties of constructivism without betraying the principles of intuitionism. Responsible for this is mainly a corollary of the bar theorem, the fan theorem. Using that, Brouwer proved that all total functions on the continuum are continuous, and even uniformly so. The first result has the important consequence that in dealing with such functions, approximations will always work fine. The second is of vital importance for a satisfactory intuitionistic theory of integration (and, hence, of probability). It also helps considerably in finding intuitionistic counterparts to theorems in classical analysis. On the intuitionistic construals of the continuum and of logic the two continuity theorems are valid; on the respective classical construals, they are not. Yet classically they are valid for certain subclasses of functions, and the intuitionist may try to transform classical theorems that hold for such subclasses into intuitionistic theorems on all total functions. Even though intuitionistic mathematics is autonomous, the aim to find such counterparts is sound from an economical as well as from a missionary point of view.

Classically, the bar and fan theorems themselves are easily established; the fact that, having gone through the more complicated process of establishing them intuitionistically, Brouwer could prove theorems that are not classically valid, is accounted for by the presence of a third ingredient in his proofs, the continuity principle discussed in the previous chapter. It is this classically false principle that connects the bar theorem to the intuitionistic modelling of the continuum using choice sequences.

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.

In this chapter we will go through one of Brouwer's proofs of the bar theorem. Although it is a bit technical, it is not particularly difficult, and moreover, we will find the notion of intentionality at the very heart of it. It comes into play in the deep reflection on the activity of the creating subject that the proof depends on. It was from this reflection that Brouwer's concept of a canonical proof emerged, which we discussed in chapter 2.

In the literature on intuitionism, the reader will find several versions of the proof, e.g. Kleene [95, 6.12], Heyting [75, 3.4], Dummett [49, 3.4]. They all provide different perspectives, and it pays to compare them. To facilitate reading Brouwer in the original, I will discuss one of his own presentations (and often use his own notation). The ones included in the anthologies edited by Mancosu (the proof from 1924 [19]) and van Heijenoort (the proof from 1927 [21]) are probably the most widely available.[28] They are different, although this is mostly a matter of cosmetics; I will discuss the 1927 version, and comment on the differences afterwards.[29]

Now the the key notions can be defined, and the theorem stated.

In general, mathematicians are not as interested in proving theorems about particular real numbers as they are in proving theorems about classes of real numbers and functions of real numbers. It will not do to collect choice sequences in a set in the Cantorian sense, because, intuitionistically, such sets can be no larger than denumerably infinite, which is too small for the continuum and many of its subsets. Rather, choice sequences are held together in a *spread* ('Menge', in Brouwer's original, somewhat confusing terminology, as it is the German word commonly used for Cantor's sets). Brouwer regretted that his definition suffers from a certain prolixity ([20, footnote 2]):

> A spread is a law on the basis of which, if repeated choices of arbitrary natural numbers are made, each of these choices either generates a definite sign series, with or without termination of the process, or brings about the inhibition of the process together with the definitive annihilation of its result; for every $n > 1$, after every unterminated and uninhibited sequence of $n - 1$ choices, at least one natural number can be specified that, if selected as the n-th number, does not bring about the inhibition of the process. Every sequence of sign series generated in this manner by an unlimited choice sequence (and hence generally not representable in a finished form) is called an element of the spread. We shall also speak of the common mode of formation of the elements of a spread M as, for short, the spread M. [20, pp.244–245], translation adapted from [67, p.453]

The last line of the definition indicates that a spread is a special kind of species.

Because, as we saw in chapter 1, Brouwer holds that mathematics is essentially languageless, in this definition 'sign' and 'sign series' are to be understood 'in the sense of *mental* symbols, consisting in previously obtained mathematical concepts' [28].

In a footnote to the above definition, Brouwer adds that the possibility of terminating the process at a certain point can be replaced by the possibility of having all choices from that point on generate 'nothing'. This 'nothing' is an empty sign series, but nevertheless an object, on a par with, for example, the empty set.

By way of explanation, this definition can be rephrased if we introduce the notions of a spread law and a correlation law.

The objects chosen in the choice sequences in the spread may be natural numbers but also any other kind of mathematical object (the condition being that the objects have been constructed prior to the spread). In the latter case, we conceptualize making a choice in a choice sequence in the spread as first choosing a natural number and then obtaining, via a mapping from the natural numbers to a fixed, countable species of objects of the other type, the object desired. This mapping is the correlation law.

For generality, Brouwer always assumes a correlation law; one obtains choice sequences of natural numbers simply by mapping each natural number to itself. Brouwer uses the term 'choice sequence' both for the original sequence of natural numbers and for what you get by applying the correlation law to it.

A spread law either admits a given finite segment of natural numbers, or inhibits it. There are three conditions on a spread law.

1. It should be decidable. That is, we should have a means to tell in finite time whether a given sequence $\langle a_0, a_1, \dots a_n \rangle$ should be used in the construction. For example, we would not be sure how to proceed building the spread if our next step depended on an open problem of which we have no idea when it will be solved.

2. Of each admitted sequence, at least one immediate extension should be admitted as well. Each admissible sequence $\langle a_0, a_1, \dots a_n \rangle$, must have an immediate extension $\langle a_0, a_1, \dots a_n, a_{n+1} \rangle$ which is likewise admissible.

3. If a sequence is admitted, then so should all of its initial segments. This is a natural demand, as we use these sequences to bundle choice sequences, there should be no gaps or jumps. If $\langle a_0, a_1, \dots a_n, a_{n+1} \rangle$ is admissible, then so is $\langle a_0, a_1, \dots a_n \rangle$.

The admitted sequences form a growing tree, hence they are also known as nodes; that is, a node in the tree is identified with the path leading up to it. The root or top (I think of trees as growing downwards) of the tree is the empty sequence. Brouwer calls the paths through a tree, whether finite or infinite, its elements. By 'choice sequences' in the pregnant sense of the word, the infinite paths are meant.

Of a pair of nodes $\langle a_0, a_1, \dots a_n \rangle$ and $\langle a_0, a_1, \dots a_n, \dots a_{n+k} \rangle$, we call the second a descendant of the first, and the first an ascendant of the second. If $k = 1$, we speak of an immediate descendant and an immediate ascendant.

The tree of admitted sequences may be called the underlying tree of the spread (Brouwer does not use this term, but speaks of 'the species of choice sequences upon which the spread is based'). The correlation law maps the elements of the underlying tree of natural numbers to the desired objects; like the underlying tree, the spread consists of nodes that form a tree.

To an admissible sequence, the correlation law assigns either an object, or 'nothing'. The latter option lets us simplify the definition of a spread by dropping the possibility of termination: one gets the same effect of a finite spread by always assigning, to each admissible sequence, 'nothing' from a certain point onwards.

An inadmissible sequence is inhibited. It will not play a role in the generation of the elements in the spread anymore, as the correlation law does not apply to such a sequence.

A special spread is the universal spread, the spread of choice sequences of natural numbers which inhibits no sequences, therefore admitting all.

As another example, consider the following spread J. Let $I_0, I_1, I_2, \dots$ be an enumeration of the intervals of the continuum of the form

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

where $2 \leq a + 2 \leq 2^{k+1}$, so that the boundaries of these intervals are ≥ 0 and ≤ 1. (That they can be enumerated follows from the fact that each such interval is determined by a pair of natural numbers $\langle a, k \rangle$, and such pairs are enumerable.) The correlation law of J assigns to n the interval I_n. Let the spread law be: $\langle a_0, a_1, \dots a_n, a_{n+1} \rangle$ is an admissible extension of $\langle a_0, a_1, \dots a_n \rangle$ exactly if interval $I_{a_{n+1}}$ is properly contained in interval I_{a_n}. The elements of J are convergent sequences of intervals of $[0, 1]$. One can prove that J coincides with that interval; we will use J in our discussion of the fan theorem below.

Any node in a tree determines or dominates a particular subtree, namely, the subtree consisting of all paths that pass through that node. Equivalently, the subtree consists of the sequences that share the initial segment defining that node. We will call this subtree the subtree dominated by that node. A species of nodes determines a species of subtrees. As a spread is also a tree (in particular, it is a tree in which every node has at least one immediate descendant) we can also speak of subspreads.

The notion of a bar is defined as follows. If B is a bar for a spread M, this means that each of the infinite choice sequences in the underlying tree of the spread (call it U) has a finite initial segment which is an element of B:

$$\forall \alpha \in U \exists n (\bar{\alpha} n \in B)$$

A species of nodes is a bar, that is, if every infinite path though the tree has at least one node in common with it ('hits the bar'). A bar determines a subtree, namely the tree starting at the root of the underlying tree and ending in the nodes that have just hit the bar.

A first statement of the bar theorem as proved by Brouwer in 1927 would be: if B is a bar, then the species of nodes that have just reached B can be well-ordered. We will see a more precise statement in a moment, but this formulation suffices to give a clue why this theorem is useful.

Generally, the number of immediate successors of a node in a tree may be finite or infinite. A tree in which every node has only finitely many immediate successors is called a finitary tree or a fan. One also says that such a tree is finitely branching. The same goes for the underlying tree. The spread J that we just saw is in fact a fan. As we will see later, a corollary of the bar theorem is the fan theorem: if B is a bar for a fan, we can effectively determine a bound on the length of the paths to the bar. As Brouwer also showed how to represent the continuum by a fan, the fan theorem enables him to prove, given some suitable bar, theorems about the continuum by proving theorems about choice sequences of a certain bounded length. The latter certainly are much more managable than the continuum itself.

Brouwer seems to have proved the bar theorem solely for the purpose of establishing this corollary. From a classical point of view, the fan theorem is the contraposition of König's lemma, which was proved later: 'If a fan contains infinitely many nodes, it contains an infinite path'. But intuitionistically, König's lemma is not valid, so it cannot be used to obtain the fan theorem. The problem with the lemma is that it doesn't enable one actually to construct an infinite path through the fan. I will come back to the fan theorem and its relation to König's lemma later.

In order to give a more precise statement of the bar theorem, we need the notion of a thin bar, and make explicit a condition on B.

A thin bar is a bar that contains no more elements than necessary to be a bar:

$$b \in B \wedge a < b \rightarrow a \notin B$$

where a and b are variables for elements, and $a < b$ means 'b is an extension of a'. That is, a thin bar never contains a pair of nodes one of which is a descendant of the other. With respect to its property of being a bar, such descendants are superfluous. (Brouwer did not use the term 'thin bar' in print before 1954 [34]).

The condition is that the bar B should be decidable, that is, of any node we should be able to tell in finite time whether it belongs to the bar or not. The condition of decidability is not explicit in Brouwer's proofs of 1924 and 1927, but it is essential, as Kleene has shown; we will see a version of his argument in the comments on the proof of the bar theorem below. If one thinks of bars that are implicitly determined by the continuity principle, as Brouwer does in the two proofs mentioned, this condition is surely fulfilled. His point of departure is that of a spread M and a function that assigns to each element of M a natural number β. The continuity principle then says that for every sequence a number n can be found such that you need only the first n choices in the sequence to calculate the number β that the function assigns to it. Given an initial segment

of a choice sequence—which corresponds to a node in the tree—determine n for that sequence; the segment belongs to the bar exactly if its length is equal to or greater than n, and this we can decide.

The bar theorem can now be stated as: if B is a decidable bar, then it contains a well-ordered thin bar.

So far, I have not said how the notion of well-ordering is defined in intuitionistic mathematics. An example will show that it cannot be the classical definition.

In the classical definition, a set is well-ordered if it is ordered and every non-empty subset has a first element. That every set can be well-ordered in this sense is implied by the axiom of choice. But that axiom does not hold in intuitionistic mathematics, for it does not tell you how to construct that first element. Here is an example of a species which can intuitionistically be shown non-empty, but of which we have yet no construction for an element. Let p be an as yet undecided proposition, say, Goldbach's conjecture that every even number is the sum of two odd primes. Now define the species A as follows:

$$x \in A \equiv (x = 0 \wedge p) \vee (x = 1 \wedge \neg p)$$

It is easy to see that A cannot be empty. Assume A contained no elements. On that assumption, both the condition for the inclusion of 0 and the condition for the inclusion of 1 must have failed. Then both p and $\neg p$ must be false, which is to say, $\neg p \wedge \neg\neg p$ is true. But that is a contradiction, so the assumption that A is empty must be false.

It is also easy to see that A cannot contain any element other than 0 or 1, as such an element would simply fail both conditions. So if x is an element of A, then it must be 0 or 1. It follows, therefore, that A is a subspecies of, for example, the species $\{0, 1, 2\}$.

However, we cannot indicate an element of A. To do so requires that we have established p or $\neg p$, which, by hypothesis, we have not.

All this means that there is a non-empty subspecies of $\{0, 1, 2\}$ of which we cannot say (now) that it has an element, let alone a first element, which would be required by a classical well-ordering. If we adopt that definition of a well-ordering in intuitionism, then the species $\{0, 1, 2\}$ would not be well-ordered.

Hence Brouwer had to resort to another definition of the notion. For this, he reached back to Cantor's original suggestion from 1883 (the now usual classical definition was a later suggestion of his) and defined well-ordering by induction. According to this definition, a species is well-ordered if it can be generated inductively, as follows:

Induction basis

Any one-element species A is well-ordered. Brouwer calls such a species a primitive species.

Induction step

(a) Generating operation of the first kind
If $A_0, \ldots, A_n$ are a positive, finite number of disjunct well-ordered species, then their ordered sum is a well-ordered species. The 'ordered sum' of the A_i is their union, where each species remains ordered in its original way, but the clause is added that, if $j < k$, each element of A_j precedes each element of A_k: $a \prec b$ for $a \in A_j, b \in A_k$.

(b) Generating operation of the second kind
If $A_0, A_1, A_2, \ldots$ is a denumerable sequence—hence, on the intuitionistic understanding of infinity, given by a construction method—of well-ordered species, then their infinite ordered sum is also a well-ordered species. Here, $a \prec b$ is defined in the same way as in the generating operation of the first kind.

It is easily shown that every well-ordered species has a first element (in the intuitionistic sense, i.e., we can exhibit it), and that every element in a well-ordered species either has an immediate successor or is the last element. Together these properties ensure that for a well-ordered species we always have an effective method to run through its elements. Hence every well-ordered species is decidable, i.e., for any well-ordered species we have a method to tell whether a given element belongs to it or not.

Moreover, from this definition it is easily proved (by induction) that of every well-ordered species it can be indicated either that it is finite or that it is denumerably infinite.

For reasons that will become clear later, it is convenient to have the option of labelling the elements in a well-ordering either 'full' or 'null'. There is no intrinsic connection between an element and its label; the labelling depends on the particular use one wants to make of such a well-ordering. In our case, the labelling will help to distinguish admissible sequences from certain unadmissible ones.

We can now verify that $\{0, 1, 2\}$ is a well-ordered species according to this new definition. The induction basis says that every one-element species is well-ordered, and accordingly, $\{0\}$, $\{1\}$ and $\{2\}$ each are well-ordered. So by the first induction step, their ordered sum is also a well-ordered species. This construction is not unique: one might also first add $\{0\}$ and $\{1\}$, and then, again by induction step (a), add $\{0, 1\}$ and $\{2\}$. The order relations are the same in both cases. Note that species such as A, which we used above to show that not every non-empty subspecies of $\{0, 1, 2\}$ has a first element, need not bother us anymore. In the new definition, that no longer is the defining characteristic of a well-ordering.

Of course, if we somehow have an immediate view of all the nodes in a thin bar, then we can well-order it by sight. But such an overview is, but in the simplest of cases, out of the question, either because the bar contains a large and perhaps even infinite number of nodes, or because the proof that all infinite paths through a node of the underlying tree have an initial segment in the bar is

complicated enough to leave it opaque what the species of these initial segments will look like; and often because of both. The whole point of the bar theorem is that in spite of this, it can still be shown that the thin bar can be well-ordered.

Here is a simple-minded attempt at a proof of the bar theorem. Certainly the infinite species of finite sequences of natural numbers, call it μ, can be well-ordered. For example, one can verify that the following rule will do. Let $a \prec b$ whenever a is shorter than b, and if a and b are of equal length, order them according to smallest first differing number: e.g., $\langle 1, 2, 3, 4\rangle \prec \langle 1, 2, 4, 3\rangle$. The first element in the whole species will be the empty sequence $\langle\rangle$. By hypothesis, the predicate 'm is in the bar', defined on μ, is decidable. Then so is the predicate 'm is in the thin bar', as it is then decidable for a given m in the bar whether any of its ascendants is too. This second predicate defines a subspecies μ_1 of μ; and we know that μ is well-ordered. Doesn't this prove the bar theorem?

It does not. We do not know now that μ_1 is also well-ordered. For one thing, it is a property of well-ordered species that we can determine either that it is finite or that it is infinite, but of μ_1 we cannot, in any case not yet. The knowledge that a given predicate defined on an infinite well-ordered species is decidable does not by itself give us a construction method for the species of elements for which that predicate holds. One way of looking at it is that the knowledge that the predicate provides when it holds is too local to derive something global from it. What is called for, then, is a deeper consideration of the nature of bars, so as to find a principle that somehow unifies the elements that make up a bar and from which a construction method for the thin bar can be derived. Thinking things through, Brouwer arrived at a method of induction, and moreover, an induction that does not work its way from the root of the tree to the thin bar, but from the thin bar to the root. Let us now see how this works.

4.2 Brouwer's proof

Brouwer's strategy is to divide the proof of the bar theorem,

> if B is a decidable bar, then it contains a well-ordered thin bar

into two parts:

1. Show that any proof of the antecedent, 'B is a decidable bar', can itself be rendered as a certain well-ordered species;

2. Show how, given this well-ordered species, one constructs a well-ordered thin bar, thus proving the consequent.

In intuitionistic proofs of implications, one usually doesn't need much more information about a proof of the antecedent beyond the fact that, in the case of a conjunction for example, one indeed has a proof of each conjunct. In the proof of the bar theorem, however, Brouwer analyses what a proof of the antecedent could be like in great detail.

4.2.1 Part 1

What information is available to us to prove that B is a (decidable) bar? First of all the spread M, and in particular, its underlying tree. The elements of the underlying tree are infinitely proceeding choice sequences of natural numbers. In the proof of the bar theorem we might as well work with initial segments of these elements, precisely because the nodes that we want to well-order are all in the bar. The species of initial segments of choice sequences of natural numbers is the species μ that we saw earlier. The elements of μ can be divided into two species: those that are admissible in M and those that are not admissible (the latter species might be empty, i.e., in the case of the universal spread). The elements of the bar B will all be in the first species.

Brouwer defines μ_1 to be the thin bar contained in B. μ_1 is the species of those elements of μ such that they are in B and their presence in B is not redundant, because they have no proper initial segment that is also in B. Every sequence that hits B also hits μ_1. We saw that B, as it is defined on the basis of the continuity principle, is a decidable bar; therefore, so is μ_1.

An element of μ is called *secured* (relative to the spread M) when we are sure of its status with respect to the thin bar μ_1: that is, when we either know that it has an initial segment in μ_1, or that it never will, because it is inhibited. Thus we can split the species μ in two, namely, into the species τ of secured elements (relative to M), and the species σ of unsecured elements (relative to M), that is, of the elemtents that are admissible but that have not yet hit the bar.

To have a proof that B is a thin bar means that, in particular, we have a proof h that shows that, for any element of σ, any infinite sequence α that is nowhere inhibited and of which this element is an initial segment will at some point n have hit B:

$$\forall \alpha \in M \exists n (\bar{\alpha} n \in B)$$

In accordance with the intuitionistic interpretation of the logical constants, any proof h of this statement supplies us with a method to construct, for a given α in the spread, this number n. The method specified by a particular h need not be the most efficient one, and even if it is, it may be still quite complex. Brouwer imagines possibilities such as the following:

> The algorithm in question may indicate the calculation of a maximal order n_1 at which will appear a finite method of calculation of a further maximal order n_2 at which will appear a finite method of calculation of a further maximal order n_3 at which will appear a finite method of calculation of a further maximal order n_4 at which the postulated node of intersection must have been passed.

'Much higher degrees of complication are thinkable,' he adds. As mentioned earlier, there can be a huge gap between having a proof that there is a bar and knowing exactly what the bar looks like.

We saw that in the setting in which Brouwer proves the bar theorem we have a spread M and a function or algorithm that assigns to every choice sequence

in M a natural number β. Because of the continuity principle, this implicitly defines a thin bar μ_1 in the underlying tree. A sufficient condition for a node in M to have the property that any choice sequence passing through it will be assigned the same number β is that this node, or an ascendant of it, was obtained by the correlation law from a node in μ_1. But this is not a necessary condition: perhaps one can find, once a sufficient number of such assignments of numbers β to choice sequences have been determined, a node above the bar (i.e., in σ) such that its correlated node in M also has that property. The situation, when it occurs, is accounted for by the fact that the algorithm used to calculate β may not be clever enough to extract this number from shorter segments [21, pp.459–460]. (Under certain conditions, it is in fact necessary to put off the assignment until later [95, p.71].)

Of course, whether a given thin bar could have been placed higher up in the tree if we had a different algorithm does not influence the fact that any admissible sequence that passes through a node $\alpha(n)$ in the given thin bar is secured. In the presence of a proof h, an element of σ is therefore called 'securable'. (An element of σ cannot be inhibited altogether, for then the element would already have been in τ.) To say that there is a bar in a tree and to say that the root of that tree is securable are equivalent.

The essential statement Brouwer wants to prove is in paragraph 4 of §2: 'We now assert that every element $F_{sn_1 \dots n_r}$ of σ possesses the *well-ordering property*.' (Brouwer uses the notation $F_{sn_1 \dots n_r}$ for the element $\langle s, n_1, \dots, n_r \rangle$ of σ.)

What Brouwer means by this is, essentially, that the thin bar of the subtree dominated by $F_{sn_1 \dots n_r}$, that is, the thin bar that bars exactly those sequences that share the initial segment $F_{sn_1 \dots n_r}$, admits of a well-ordered construction. Once this claim has been proved, it can then be instantiated for the subtree dominated by the root node, which is just the whole tree, barred by the whole thin bar.

Let us now take a look at Brouwer's own, precise definition of the well-ordering property of $F_{sn_1 \dots n_r}$, in which also the spread M is taken into account. Consider the following two species. (See figure 4.1.)

- The subspread (after all, a species defined on a spread) $M_{sn_1 \dots n_r}$ of M determined by $F_{sn_1 \dots n_r}$. This subspread consists of all the infinite elements of M that share the initial segment obtained by successively applying the correlation law to the initial segments of $F_{sn_1 \dots n_r}$ (i.e., first to F_s, then F_{sn_1}, then $F_{sn_1 n_2}$, and so on).
- The species of subspreads M_α, where α is an extension of $sn_1 \dots n_r$ such that F_α hits the thin bar but does not go beyond it. In other words, an M_α is a subspread of M determined by an F_α, where F_α is a descendant of $F_{sn_1 \dots n_r}$ and is an element of μ_1. A species of such F_α defines a little thin bar, namely, that part of the whole thin bar which is responsible for barring the sequences with initial segment $F_{sn_1 \dots n_r}$.

Clearly, the union of all M_α (the union of all the elements of the second species) is identical to $M_{sn_1 \dots n_r}$ (the first species). In Brouwer's terms, the first species

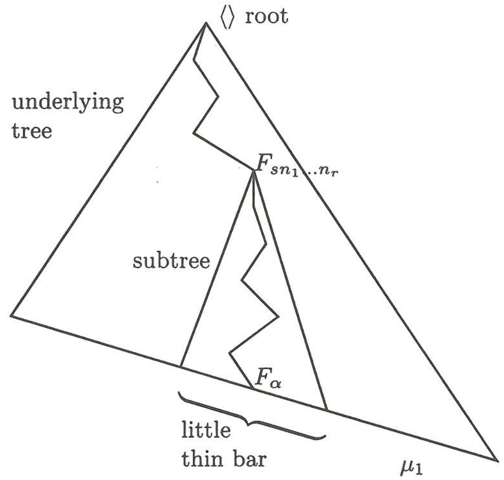

Figure 4.1: Little thin bar determined by $F_{sn_1...n_r}$

splits into the second species.

Now Brouwer defines the claim that $F_{sn_1...n_r}$ has the well-ordering property as the claim that the second species, the species of subspreads M_α, is similar to the species of full elements of a well-ordered species $T_{sn_1...n_r}$ (as we will see, the construction of this species T brings with it that T will contain null elements as well). It is claimed, that is, that we can construct a one-one relation between these two species that leaves the ordering relations invariant. In effect this means that the species of subspreads M_α can be well-ordered: first construct $T_{sn_1...n_r}$ and then use the one-one relation. As to each M_α corresponds F_α, this also yields a well-ordering of the F_α.

As mentioned, the aim is to show that the root of the whole tree has the well-ordering property. In that case, the corresponding F_α together make up the whole thin bar, because now α is any path reaching the thin bar from the root. Moreover, we can then think of T as a well-ordered construction of the whole thin bar.

Incidentally, Brouwer does not have the empty node as the root of the whole tree; instead, he proves the result for trees with arbitrary non-empty top, and the bars of all these trees taken together form the whole thin bar μ_1. But it is convenient to add all these trees together (in a generating operation of the second kind) with the empty node as top. As a matter of notation, F_s then is not, as in Brouwer, the one-number sequence $\langle s\rangle$, but the empty sequence $\langle\rangle$, and we will take $F_{sn_1...n_r}$ to refer to the element $\langle n_1,\ldots,n_r\rangle$ of σ.

If you think things through intuitionistically, Brouwer claims in footnote 7 of the 1927 paper, then you can actually see that the root has the well-ordering property, and no further proof is needed:

> When carefully considered from the intuitionistic point of view, this securability is seen to be nothing but the property defined by the stipulation that it shall hold for every element of μ_1 and for every inhibited element of μ, and that it shall hold for an arbitrary $F_{sn_1...n_r}$ as soon as it is satisfied, for every ν, for $F_{sn_1...n_r\nu}$. This remark immediately implies the well-ordering property for an arbitrary $F_{sn_1...n_r}$.

The well-ordering property follows immediately, as the stipulations defining securability translate immediately to the induction clauses for one-element species and the second generating operation.

It has been questioned whether Brouwer's full proof of the bar theorem really is more evident than the principle formulated in this footnote; I will come back to that issue in the comments on Brouwer's proof below. But what can be said now is that the full proof certainly is more fundamental in the sense that it shows a finer-grained constructivism, turning, as we will see, on the issue what exactly a mental proof is.[30] This must be why Brouwer continues the footnote by saying

> The proof carried out in the text for the latter property, however, seems to me to be of interest nevertheless on account of the propositions contained in its elaboration.

In chapter 2 it was explained that proofs can always be thought of as consisting of just immediate facts and elementary inferences; proofs have a canonical form. We now see that these canonical forms are well-ordered trees. The immediate facts correspond to null elements as they are considered to be 0-step inferences. Elementary inferences correspond to generating operations of the first kind if they have a finite number of premises, and to operations of the second kind if they have infinitely many premises. A canonical form makes fully explicit the intentional structure of a proof.

Brouwer's intentional analysis of proofs of securability led him to the following conclusions. Any proof of a node's securability must turn on the relations between the various elements of the underlying tree, and these relations can be decomposed into basic relations that relate an element $F_{mm_1...m_g}$ to its immediate ascendant and to one or more of its immediate descendants $F_{mm_1...m_gm_{g+1}}$. In a proof in which only these basic relations play a role, these relations are established by elementary inferences. According to Brouwer, these are:

1. Immediate facts or 0-step inferences which show the securability of a secured element. Brouwer didn't give these a name. If an element is secured because it is in the bar, Kleene called the 0-step proof of this an η-inference (pronounced 'eta-inference'); Dummett follows this. The other possibility, that the element is secured because it is inhibited, does not occur in Kleene's proof because he proves the bar theorem for just the universal spread (which means no loss of generality, as any spread can be embedded in it).

2. ζ-inferences (pronounced 'zeta-inferences'). From the securability of $F_{mm_1 \ldots m_{g-1}}$, conclude the securability of $F_{mm_1 \ldots m_g}$. (Rationale: if all paths through a node hit the bar, then surely a path through one of that node's immediate descendants hits the bar.)

3. $\digamma$-inferences.[31] From the securability of $F_{mm_1 \ldots m_g \nu}$ for all ν, conclude the securability of $F_{mm_1 \ldots m_g}$. (Rationale: if all paths through any of a node's immediate descendants hit the bar, then all paths through that node hit the bar.) This inference has infinitely many premises. It therefore depends on the existence of a construction method for the species of its premises.

Canonical proofs are built out of these elementary inferences; and the way they are combined makes such a proof a well-ordered species. The 0-step-inferences are the primitive species; ζ-inferences correspond to the generating operation of the first kind; and $\digamma$-inferences to that of the second kind. That these proofs have a well-ordered structure is at the basis of Brouwer's proof of the bar theorem.

A canonical proof for the securability of the element $F_{sn_1 \ldots n_r}$ is indicated by, for example, $k_{sn_1 \ldots n_r}$.

Canonical proofs need not be as efficient as possible. They may contain detours and redundancies by proving securability of a given node more than once. We will see that this does not influence the proof of the bar theorem in any essential way.

4.2.2 Part 2

The idea behind the second part is that we first construct proofs of the bar theorem for little bars near the bottom of the tree, and then show that if such proofs for small bars are combined, we obtain a proof of the bar theorem for a larger subtree. Joining ever larger subtrees in this way finally brings one back to the top of the whole tree, yielding a proof of the bar theorem. Thus, the proof is by induction on proofs for subtrees; but it differs from more common forms of induction in two ways.

First, it is induction in the reverse direction from the usual: it proceeds from the bottom of the tree to the top ('backward', [95, p.65]).

Second, this induction is transfinite because, as we will see, the number of premises in the induction step (corresponding to the number of subtrees involved in the step upwards in the tree) is infinite.

The data available to us to carry out the proof that each element of σ has the well-ordering property are μ_1, the inhibited sequences, and the canonical proofs.

In order to be able to make use of the canonical proofs, Brouwer defines two properties on them, based on the following circumstances.

- A canonical proof $k_{sn_1 \ldots n_r}$ is a well-ordered species; it is built up inductively from subspecies which are themselves proofs. These are called the constructive underspecies of the species, and they include as special case

the species itself. Each of these constructive underspecies is a proof ascertaining the securability of a particular element.

- Because of the first circumstance, one can construct from $k_{sn_1 \ldots n_r}$ the well-ordered species $f_{sn_1 \ldots n_r}$ of the elements of σ whose securability is established in the course of it. This species in particular contains $F_{sn_1 \ldots n_r}$ itself.

The two properties are:

the well-ordering property From the already defined well-ordering property for elements, one for canonical proofs is derived. A constructional underspecies u of $k_{sn_1 \ldots n_r}$ (that is, one of the subproofs out of which $k_{sn_1 \ldots n_r}$ is built) has the well-ordering property for canonical proofs if every element σ of f_u (that is, every element of which the securability is established in the course of u) has the well-ordering property for elements.

the preservation property[32] A constructional underspecies u of $k_{sn_1 \ldots n_r}$ has the preservation property if the following holds: if every element of $f_{sn_1 \ldots n_r}$ of which the securability functions as a premise in u has the well-ordering property, then every element of $f_{sn_1 \ldots n_r}$ of which the securability is derived in the course of u also has the well-ordering property. It is the preservation property that guarantees that, if we join proofs of the bar theorem for little bars, we get a proof of the bar theorem for the bar that combines these little bars.

In Brouwer's text, the whole proof that every element $F_{sn_1 \ldots n_r}$ of σ has the well-ordering property passes by in one short paragraph. It consists of three steps:

1. The preservation property holds for $k_{sn_1 \ldots n_r}$

2. Therefore, the well-ordering property (for proofs) holds for $k_{sn_1 \ldots n_r}$

3. Therefore, the well-ordering property (for elements) holds for $F_{sn_1 \ldots n_r}$

One sees from this proof that the well-ordering property of a node rides pickaback on the preservation property of canonical proofs of its securability. That a node has a certain property is derived here from a property of certain proofs; this move is typical for intuitionism. Let's go through these three steps now.

1. The preservation property holds for $k_{sn_1\dots n_r}$ By induction along the construction of $k_{sn_1\dots n_r}$ (which, after all, is a well-ordered species):

   induction basis

   The primitive species in a canonical proof are the 0-step-inferences, proving the securability of a secured element. These are 0-step proofs and have 0 premises. So, vacuously, a 0-step-inference preserves the well-ordering property.

   induction step

   (a) $\digamma$-inference
   As this elementary inference has infinitely many premises, this is where we will use transfinite induction. The premises are canonical proofs of the securability of nodes $F_{mm_1\dots m_g\nu}$, for all ν. What we have to prove is that, if each of these nodes has the well-ordering property, so has the the node of which the $\digamma$-inference in question proves the securability, $F_{mm_1\dots m_g}$; this node is the top of the subtree in which the subtrees dominated by the nodes in the premise are combined. Assume that the nodes $F_{mm_1\dots m_g\nu}$ have the well-ordering property for every ν. Then to each corresponds a well-ordered species of nodes $T_{mm_1\dots m_g\nu}$ forming a little thin bar. These species are all disjoint, as the initial segments $\langle m, m_1, \dots m_g\nu\rangle$ of the nodes in different species will differ in their value for ν. But then they can be added in a generating operation of the second kind, giving the well-ordered species $T_{mm_1\dots m_g}$, which shows that $F_{mm_1\dots m_g}$ has the well-ordering property.

   (b) ζ-inference
   Its premise is a canonical proof of the securability of $F_{mm_1\dots m_{g-1}}$. Assume that $F_{mm_1\dots m_{g-1}}$ has the well-ordering property. Then to this node corresponds a well-ordered species of nodes $T_{mm_1\dots m_{g-1}}$ forming a little thin bar. But a decidable subspecies of this one is that of all elements having initial segment $F_{mm_1\dots m_g}$; then this subspecies is $T_{mm_1\dots m_g}$, establishing that $F_{mm_1\dots m_g}$ has the well-ordering property.

   We see that a canonical proof $k_{sn_1\dots n_r}$ always has the preservation property.

2. Therefore, the well-ordering property holds for $k_{sn_1\dots n_r}$

   That $k_{sn_1\dots n_r}$ has the preservation property implies that it has the well-ordering property, as follows. $k_{sn_1\dots n_r}$ as a whole has as given only the securability of the elements in μ_1 and the inhibited sequences. In other words, the well-ordered construction of $k_{sn_1\dots n_r}$ has to start from proofs that are null elements, corresponding to the fact that secured elements

(elements of τ, i.e., in the bar or inhibited) are trivially securable (0-step proof). These are its primitive species. A canonical proof determines a species of elements, namely the species of those elements of which it establishes their securability. In the case of a null proof, this species has exactly one element, and is therefore well-ordered by definition. By preservation, $k_{sn_1\dots n_r}$, built up starting from these primitive species, has the well-ordering property.

3. Therefore, the well-ordering property holds for $F_{sn_1\dots n_r}$

   The fact that the well-ordering property (for proofs) holds for $k_{sn_1\dots n_r}$ means that, in particular, the element of σ of which this proof as a whole establishes the securability, $F_{sn_1\dots n_r}$, has the well-ordering property (for elements).

This enables us to conclude that the well-ordering property holds for the root of the tree, $\langle\rangle$: $F_{sn_1\dots n_r}$ was arbitrarily chosen from the nodes in σ, so we are allowed to infer that *every* element of σ has the well-ordering property. Instantiating that conclusion for the root of the tree, we conclude that the thin bar of the subtree with the root as top, that is, the thin bar of the whole tree, admits of a well-ordered construction. This proves the bar theorem.

We will now have a closer look at how this proof specifies the construction, for an arbitrary element $F_{sn_1\dots n_r}$ of σ, of $T_{sn_1\dots n_r}$.

The species T is constructed using generating operations of the second kind. This operation corresponds, as Brouwer puts it in the fourth paragraph of §2, 'to the inversion of the continuation, by a new free choice, of a certain finite initial segment of choices that is uninhibited for M'. By this he refers to the fact that we are building in the opposite direction from the usual, so that now species corresponding to descendants are constructed before the species corresponding to their immediate ascendant. The well-ordered species $T_{sn_1\dots n_r\nu}$ to be added in the infinite sum with $T_{sn_1\dots n_r}$ as top are now determined as follows:

- If $F_{sn_1\dots n_r\nu}$ is inadmissible, we take a primitive species consisting of ν as a null element.
- If $F_{sn_1\dots n_r\nu}$ is admissible and moreover is an element of μ_1, we take a primitive species consisting of ν as a full element.
- In the remaining case, where $F_{sn_1\dots n_r\nu}$ is admissible but is not an element of μ_1, that is, is unsecured, we appeal to our backward induction hypothesis, by which we already have a well-ordering $T_{sn_1\dots n_r\nu}$ of the little bar of the subtree dominated by $F_{sn_1\dots n_r\nu}$; this $T_{sn_1\dots n_r\nu}$ will be the summand.

One sees how the construction of the well-ordering depends on the fact that μ_1 is a decidable species, which in turn was guaranteed by the decidability of the bar.

The primitive species of $T_{sn_1\dots n_r}$ that have been predicated full correspond to the subspreads of M determined by the F_α.

The primitive species of $T_{sn_1...n_r}$ that have been predicated null do not correspond to anything in M, as they correspond to inhibited sequences, to which the correlation law does not apply. Brouwer (later) called elements that are inhibited while their immediate ascendant is not, 'stops'. The rationale for including the stops in the construction of T is that it guarantees that the generating operation of the second kind employed will indeed have a well-ordered species to add for every ν.

Taken together, the full and null primitive species of $T_{sn_1...n_r}$ correspond to the species τ of secured elements, except that τ also includes extensions of inhibited sequences.

The well-ordering of the thin bar can be further described as follows. Consider any two elements of the thin bar $a = \langle a_0, \ldots, a_m \rangle$ and $b = \langle b_0, \ldots, b_n \rangle$. Because the bar is thin, neither element is a descendant of the other; so even when $m \neq n$, one can indicate the position at which they first differ, say $a_i \neq b_i$. Then T is constructed such that, if $a_i < b_i$ then $a \prec b$, and, conversely, if $b_i < a_i$ then $b \prec a$. So $\langle 2, 3, 4 \rangle$ comes after $\langle 2, 1 \rangle$, and before $\langle 5 \rangle$ (assuming that these sequences are all in a given thin bar). Thus, T is a lexicographical ordering.

4.3 Some comments on Brouwer's proof

1. Each of the well-orderings $T_{sn_1...n_r}$ is constructed strictly upwards; hence, to arrive at them, one never first needs a well-ordering corresponding to a node above $F_{sn_1...n_r}$. As a node's having the well-ordering property implies that it is securable, this suggests that it cannot be essential first to have a canonical proof of the securability of a node higher up in the tree. But this is exactly the direction of argument in a ζ-inference; so ζ-inferences seem redundant. If one proves this independently, this fact can be exploited to prove the bar theorem in a slightly different way. This was done by Brouwer in the 1924 version, which proceeds, in effect, by removing all the ζ-inferences from the canonical proofs (this is done explicitly in an elucidatory companion paper of the same year [18], unfortunately not included in [104], and in the 1954 proof (but see remark 3 below)). The inductive definition of the elements in the thin bar then is a copy of the inductive definition of the resulting canonical proof.

2. While eliminating the ζ-inferences simplifies the proof, the 1927 proof shows that it is not essential to do so, and it is therefore not this eliminability that is the main idea behind the bar theorem. Indeed, in the course of a canonical proof, the securability of a certain $F_{sn_1...n_r}$ may be established more than once, and if it is, there are equally many different proofs of the well-ordering property of that element; but it doesn't matter if there are redundant inferences, in two senses.

It doesn't matter in the sense that in our effort to construct a well-ordering $T_{sn_1...n_r}$, for each ν we need only one well-ordering $T_{sn_1...n_r\nu}$; whether this is obtained from the first proof of the securability of $F_{sn_1...n_r\nu}$, or from a later one, is immaterial. Any one will do.

But as Brouwer mentions in footnote 9 of the 1927 paper, it also doesn't matter in the stronger sense that the well-orderings obtained on the basis of different canonical proofs in fact all come out the same; his term is that they are all 'generation-equivalent'. This is defined as follows [22, pp.452–453]:

- Two well-ordered primitive species F' and F'' are generation-equivalent if they consist of the same element and moreover this element is either full for both or null for both.
- Two well-ordered non-primitive species F' and F'' are generation-equivalent if for arbitrary ν the constructional underspecies F'_ν and F''_ν either both fail to exist, or both exist and are generation-equivalent. [translation taken from [67, p.461]]

(Note that because our species T is generated in the second generating operation, the constructional underspecies T_ν exists for every ν.) What it means for two well-orderings to be generation-equivalent is that although they are different intensionally, as they are defined on the basis of different canonical proofs, they are all built up in exactly the same way. If you would draw them on paper, their pictures would look the same. This is shown by the following induction hinted at by Brouwer.

induction basis

Let $T^k_{sn_1...n_r}$ be a primitive species, corresponding to the canonical proof $k_{sn_1...n_r}$ of the securability of $F_{sn_1...n_r}$.

Assume that there is a different canonical proof of the securability of $F_{sn_1...n_r}$, $m_{sn_1...n_r}$, and let $T^m_{sn_1...n_r}$ be the primitive species corresponding to that.

Now the two species $T^k_{sn_1...n_r}$ and $T^m_{sn_1...n_r}$ are generation-equivalent, because in both cases, the element $F_{sn_1...n_r}$ is the same, and only that determines whether the unique element of $T^k_{sn_1...n_r}$ and $T^m_{sn_1...n_r}$ will be a full or a null element.

induction step

Let $T^k_{sn_1...n_r}$ be a species, built up by adding, in the second generating operation, the species $T^k_{sn_1...n_r\nu}$. The species $T^k_{sn_1...n_r}$ corresponds to the canonical proof $k_{sn_1...n_r}$ of the securability of $F_{sn_1...n_r}$.

Assume that there is a different canonical proof of that, $m_{sn_1...n_r}$, yielding the species $T^m_{sn_1...n_r}$. This species is built up by adding, in the second generating operation, the species $T^m_{sn_1...n_r\nu}$.

By induction hypothesis, for every ν, $T^k_{sn_1...n_r\nu}$ and $T^m_{sn_1...n_r\nu}$ are generation-equivalent: they correspond to different canonical proofs of the securability of the same element $F_{sn_1...n_r\nu}$.

But then, by the definition of generation-equivalence, $T^k_{sn_1\ldots n_r}$ and $T^m_{sn_1\ldots n_r}$ are generation-equivalent.

3. That the bar has to be decidable for the bar theorem to be true was brought to light by Kleene [95, pp.87–88]. Van Dalen gives the following simplified version [37, p.102]. Consider the universal spread and a decidable predicate $A(x)$ for which as yet we have neither a proof of $\forall x A(x)$ nor of $\neg\forall x A(x)$. Now define the species B as follows:

$$\langle\rangle \in B \Leftrightarrow \neg\forall x A(x)$$
$$\langle n\rangle \in B \Leftrightarrow A(n)$$

It will be immaterial what other elements B might contain. Every path α through $\langle\rangle$ hits B, so B is a bar. This is seen as follows. For any particular $\alpha(0)$—the first value on the path α—we can always find out whether A holds of it or not, for it was given that $A(x)$ is decidable. If it holds, then $\langle\alpha(0)\rangle \in B$; if it doesn't, then we have a counterexample to $\forall x A(x)$, therefore $\neg\forall x A(x)$, and $\langle\rangle \in B$.

B is, at present, not a decidable bar. For if it were, then, in particular, $\langle\rangle \in B \vee \langle\rangle \notin B$. Combining this with the intuitionistically valid $\neg\neg\forall x A(x) \rightarrow \forall x\neg\neg A(x)$ and with $\forall x\neg\neg A(x) \rightarrow \forall x A(x)$, valid because A is decidable, we obtain $\neg\forall x A(x) \vee \forall x A(x)$; but this contradicts our hypothesis that we do not yet have a proof of the latter.

Now assume that the bar theorem holds for this B, so that it contains a well-ordered thin bar B'; being a well-ordered species, B' is decidable. It follows that, in particular, $\langle\rangle \in B' \vee \langle\rangle \notin B'$. But then we get, by the same reasoning as above, $\neg\forall x A(x) \vee \forall x A(x)$, contradiction. Thus, we have a weak counterexample to the bar theorem for bars that are not decidable: it is not shown that it is false, but it is shown that it cannot be true as long as there are such predicates A satisfying the conditions mentioned.

I mentioned that the decidability of the bar is implicit in the premises of the 1924 and 1927 proofs of the bar theorem (because of the continuity principle), but it is neither explicit nor implicit in the proof from 1954. The countexample shows, therefore, that the latter proof must be incorrect. For further discussion, see [107] and [49, 3.4].

4. Various authors—Heyting, Kleene, Troelstra, Dummett—agree with Brouwer's remark on securability in his footnote 7, quoted on p.50 above. They choose to adopt that in the form of the axiom schema BI_D for bars in the universal spread (one axiom for each specific bar). It has the form of an implication, where the antecedent consists of a conjunction of four conditions. For

clarity, I will write each conjunct on a new line:

$$
\begin{aligned}
BI_D \quad & \forall\alpha\exists x(\bar{\alpha}x \in B)\wedge && (4.1)\\
& \forall n(n \in B \vee n \notin B)\wedge && (4.2)\\
& \forall n(n \in B \rightarrow n \in Q)\wedge && (4.3)\\
& \forall n(\forall y(n * y \in Q) \rightarrow n \in Q) \rightarrow && (4.4)\\
& \langle\rangle \in Q && (4.5)
\end{aligned}
$$

(4.1) expresses that B is a bar (not necessarily thin). (4.2) adds to this that B is decidable—hence the 'D' in BI_D. Let Q be the species of securable sequences; then (4.3) specfies that whenever an element is in the bar, it is securable. In (4.4), $n*y$ means the element n extended by one choice y. For example, $\langle 1,2,3\rangle * 4$ is $\langle 1,2,3,4\rangle$. (4.4) translates Brouwer's stipulation that securability is the property that, whenever it holds for all immediate descendants of an element, holds for that element itself. The conclusion drawn from (4.1)–(4.4) is that the root of the tree is securable (4.5).

As explained after the quote from Brouwer's footnote on p.51, BI_D immediately implies the well-ordering property for any $F_{sn_1\ldots n_r}$ in σ, and in particular for the root $\langle\rangle$; so if one finds BI_D evident, one can leave aside the long argument based on analysis of proofs into canonical proofs. Of particular interest is Kleene's proof that BI_D is independent of the other intuitionistic principles as he formalized them [95, p.113]. This means that if one wishes to prove the validity of the schema, one has to adopt a new principle in its place to prove it from; and Kleene remarks, 'We are unconvinced that any known substitute is more fundamental and intuitive' [95, p.51]. But as we have seen, the substitute of the analysis in terms of canonical proofs is more fundamental in the sense that it makes the role of intentionality in proofs explicit. Perhaps one should consider the long proof first of all an explication of the principle in the footnote [75, p.45]. Indeed, it has been shown by Martino and Giaretta that Brouwer's claim that any proof of the existence of a bar can be analysed into his three elementary inferences is, if one accepts the continuity principle, equivalent to BI_D [107].[33] Brouwer sometimes wondered if the basic relations on which the elementary inferences are based couldn't be reduced to even more basic relations [34, p.13]. This is still an open question (see p.65).

4.4 The fan theorem

From the bar theorem, Brouwer proved the fan theorem. In turn, the fan theorem is used to prove that all total functions on the continuum (intuitionistically conceived) are continuous, and uniformly continuous at that. The importance of this result was explained at the beginning of this chapter.

Recall that a spread M is a finitary tree or fan if each node in it has only finitely many immediate descendants. (According to the definition of a spread, it always has at least one.) Intuitionistically, 'each node has only finitely many descendants' means 'for each node (in the underlying tree) we can determine a number k such that no choice greater than k is admissible at that node'. If we cannot do this, then we don't know that M is a fan. Brouwer states the fan theorem as follows:

> If with each element e of a finitary spread M a natural number β_e is associated, a natural number z can be specified such that β_e is completely determined by the first z choices generating e. [21, p.462]

The point is that, while β_e in general will depend on finite elements in the underlying tree and the correlation law, it does not depend on any infinite choice sequence.

One can prove that the unit continuum, i.e. the closed interval $[0, 1]$, can be represented by a fan, for example the fan J we saw earlier, p.43.[34] (To do this, one has to show that every element of the fan falls within that interval and that, conversely, every element of that interval coincides with an element of the fan. An open interval cannot be represented by a fan as in such an interval there is no leftmost element and no rightmost element.) From this, Brouwer proved the uniform continuity theorem: a total function from the closed interval $[0, 1]$ to $\mathbb{R}$ is uniformly continuous on $[0, 1]$, that is, in a standard formulation,

$$\forall\epsilon\exists\delta\forall x_1\forall x_2(|x_1 - x_2| < \delta \rightarrow |f(x_1) - f(x_2)| < \epsilon)$$

for positive δ, ϵ and $x_1, x_2 \in [0, 1]$.

An immediate consequence (a matter of manipulating the quantifiers in front) of the uniform continuity theorem is the continuity theorem, which is not stated by Brouwer: a total function from the closed interval $[0, 1]$ to $\mathbb{R}$ is continuous on $[0, 1]$, that is, again in a standard formulation,

$$\forall\epsilon\forall x_1\exists\delta\forall x_2(|x_1 - x_2| < \delta \rightarrow |f(x_1) - f(x_2)| < \epsilon)$$

for positive δ, ϵ and $x_1, x_2 \in [0, 1]$

One can see from the order of the quantifiers why the uniform continuity theorem is a stronger result than the continuity theorem: for a given ϵ, uniform continuity demands that the same δ work for all x_1 simultaneously, whereas for ordinary continuity, δ may vary with each x_1. It is therefore plausible that uniform continuity should require knowledge of the structure of the bar whereas ordinary continuity does not.[35] The uniform continuity theorem is a much more powerful weapon than the continuity theorem; this will have added to Brouwer's pride in having established it.

An important consequence of the continuity theorem[36] is the unsplittability of the unit continuum: suppose $[0, 1] = A \cup B$ and $A \cap B = \emptyset$, then f defined by

$$f(x) = \begin{cases} 0 & \text{if } x \in A \\ 1 & \text{if } x \in B \end{cases}$$

is total and therefore, by the continuity theorem, continuous. But then f must be constant, so either $[0, 1] = A$ or $[0, 1] = B$.

These results for the unit continuum—continuity, uniform continuity, and unsplittability—generalize to the whole continuum. A remarkable consequence is that it is false that every real number is either rational or irrational. For if it were, we could obtain a splitting of the continuum by assigning 0 to rational, and 1 to irrational real numbers.[37] This is a strong counterexample to one form of the principle of excluded middle:

$$\neg\forall x \in \mathbb{R}(x \in \mathbb{Q} \vee x \notin \mathbb{Q})$$

Note that this does not mean that it is contradictory to hold of any particular real number c that it is rational or irrational; but it does mean that it is contradictory to hold it for all real numbers simultaneously.

The fan theorem, from which Brouwer derived these results, holds for the special kind of spread that fans are, but not for spreads in general. Here is a counterexample. Consider the universal spread and define a function f on it by $f(\alpha) = \alpha(\alpha(0))$, that is, f assigns to α the value of its $\alpha(0)$-th element. As α is an element of the universal spread, this means that any arbitrary choice for $\alpha(0)$ is admitted. But then there can be no upper bound on $\alpha(0)$, and hence not on the length of the segment of any α that one has to know before one can determine $f(\alpha)$. In this case there is no z as claimed by the fan theorem.

Before looking at how the fan theorem follows from the bar theorem, let us first, as we did above for the bar theorem, see why a certain simple-minded approach doesn't work (compare [52], the general tenet of which however I do not accept). Start at the root of the underlying tree and try all admissible paths of length 1. This can be done as in a fan we know that there are only finitely many. Put all paths that just hit the bar aside; of the remaining ones, now try their admissible immediate extensions. In other words, try all admissible paths of length 2 such that their predecessor hasn't hit the bar already. Of these there are likewise only finitely many. Put all paths that just hit the bar aside, etc. Keep repeating the process; as every infinite path will at some point hit the bar, the process will end. The length at which all paths of that length are put aside is the maximum length a path can have in this fan when it hits the bar. This is the z we were looking for.

However, this reasoning is circular. To know that the process will end means, intuitionistically, that we can determine an upper bound on the length of paths to the thin bar. But that we can is precisely what we are trying to prove.

This is where the bar theorem comes in. As the fan is barred by a decidable bar, the bar theorem applies, and yields a well-ordered thin bar. That well-ordering is constructed using generating operations of the second kind. This operation adds together infinitely many well-ordered species $T_{sn_1\dots n_r\nu}$ that are determined by the rules we saw on p.55. In particular, $T_{sn_1\dots n_r\nu}$ will be a null element exactly if choosing ν does not yield an admissible extension of $F_{sn_1\dots n_r}$.

Now, given that M is a fan, we know that there are always only finitely many admissible extensions. That is, for each $F_{sn_1\dots n_r}$ we can indicate a number k such that for all $\nu > k$, $F_{sn_1\dots n_r\nu}$ is inadmissible. But then for all $\nu > k$, $T_{sn_1\dots n_r\nu}$ will be a null element. If we go through the $T_{sn_1\dots n_r\nu}$, starting at

$T_{sn_1\ldots n_r0}$, then all the elements among them that are not null will have been reached, at the latest, by the time we have arrived at $T_{sn_1\ldots n_rk}$. Let these elements form the content of a species $Q_{sn_1\ldots n_r}$. It can be constructed as a well-ordered species in parallel with $T_{sn_1\ldots n_r}$, this time using a generating operation of the first kind. When $T_{sn_1\ldots n_r0}$ to $T_{sn_1\ldots n_rk}$ have been determined, run through them to determine the summands whose sum will be $Q_{sn_1\ldots n_r}$, as follows:

- If $T_{sn_1\ldots n_r\nu}$ is a primitive species consisting of ν as a null element, and therefore corresponds to an inadmissible element and determines no element of M, we skip it.
- If $T_{sn_1\ldots n_r\nu}$ is a primitive species consisting of ν as a full element, we put a primitive species $Q_{sn_1\ldots n_r\nu}$ consisting of ν as a full element next in the list of summands.
- In the remaining case, where $T_{sn_1\ldots n_r\nu}$ is not a primitive species, we appeal to our backward induction hypothesis, by which we already have a well-ordered species $Q_{sn_1\ldots n_r\nu}$; we put this species next in the list of summands.

Adding the summands on the finite list in a generating operation of the first kind, we obtain $Q_{sn_1\ldots n_r}$. It has only full elements and no null ones. We constructed it using only generating operations of the first kind, so the species of its elements must be finite, as is shown by an easy induction. Because of the way it is constructed, $Q_{sn_1\ldots n_r}$ determines a well-ordering of the species of full elements of $T_{sn_1\ldots n_r}$, which therefore is also finite.

In particular, the species of full elements of $T_{\langle\rangle}$—the species, that is, of all nodes in the thin bar μ_1—is well-ordered and finite. In that case, a natural number z can be indicated such that the maximum length of a path from the root to the bar is z: just run through the finite well-ordering of full elements of $T_{\langle\rangle}$ and keep track of the deepest node found so far. As it takes at most z choices to hit the bar from the root, the natural number β_e assigned to an element e of M is completely determined by the first z choices generating e. This proves the fan theorem.

From a classical point of view, one proves the same much quicker, from König's lemma:[38]

> If a fan contains infinitely many nodes, it contains an infinite path

Taking the contraposition gives

> If a fan contains only finite paths, it contains finitely many nodes
> (and hence there is an upper bound on the length of the paths)

Note that a spread that has only finite paths can still contain infinitely many nodes: not in depth, but in width. But such a spread is not a fan. As a bar in a fan cuts off all infinite paths at some point, it determines a fan having only finite paths, and if we then apply the contraposition of König's lemma to that, we obtain a classical version of the fan theorem.

But this version does not have the same strong content as the intuitionistic one, for the latter provides us with a construction of an upper bound. The classical version merely says that there is such a bound, without further informing us what this bound is. This is why the move of contraposing König's lemma is not intuitionistically valid: it introduces an existential statement without supplying a construction to find a number that is a witness to it. But then we have, on the intuitionistic interpretation of logic, no right to hold that statement true.

It might be instructive to see why the classical proof of König's lemma itself is not intuitionistically valid; the reason why also shows that there is little reason to suppose that König's lemma is intuitionistically true at all.

The proof has a certain constructive flavour, as it defines an infinite path α through the fan inductively [134, p.8]. We let the path start at the root, which has, by hypothesis, infinitely many descendants, and set $\alpha(0) = \langle\rangle$. This is the induction basis. The induction step is based on the observation that, of the finite number of immediate descendants that a node $\alpha(n)$ in the fan has, at least one must have infinitely many descendants. For if none had, then the fan couldn't contain infinitely many nodes, as it does, by the hypothesis of the theorem. Pick such an immediate descendant for $\alpha(n+1)$, the next node on the infinite path α. By induction from the root down, then, we have defined $\alpha(n)$ for all values of n, and this determines an infinite path, as was asked for.

This proof, however, is not really constructive. It employs the principle of the excluded middle in the form 'Each immediate descendant either has finitely many descendants, or it has infinitely many'. But it is not effectively decidable which is the case, and therefore it is not intuitionistically true. Obviously, a trial-and-error search for an immediate descendant of $\alpha(n)$ that has infinitely many descendants is out of the question; and neither do we have another general method to determine one. In a specific fan of course the spread law may be such that from inspecting it we can, at any given node, effectively pick out such an immediate descendant as required; but there is no reason to suppose that this is the case for every fan.

König's lemma was proved in 1926, two years after Brouwer's first proof of the fan theorem. Historically, the two results seem to be unrelated. It is of some interest that here we have a theorem of which the intuitionistic proof preceded the classical one.