Veterans of this group may recall that a long time ago, I said that I would write an expository article entitled "Forcing for dummies" and post it here. Well, here it is at last. I'm still not totally happy with it, but I figured if I postponed any longer I might never post it.

Logicians may wish to skip right to the end where I have some questions that I don't know the answer to.

I am deeply indebted to Matthew Wiener and Andreas Blass, who have patiently answered my many stupid questions on the subject. The book I have studied most is Kunen's _Set_Theory_.

I apologize that this article is so long; I have tried to be concise but the subject is complex.

1. Models of ZFC

Perhaps the most famous theorem proved using forcing is the independence of the continuum hypothesis (CH). More precisely, what Cohen proved was that if ZFC is consistent, then CH is not a theorem of ZFC. Goedel had already proved that if ZFC is consistent then ~CH, the negation of CH, is not a theorem of ZFC (using his concept of "constructible sets").

Just how does one go about proving such a result? At a very high level, the structure of the proof is what you would expect: one writes down a very precise statement of what the axioms of set theory are (ZFC) and what ~CH is, and then one constructs a mathematical structure that satisfies both ZFC and ~CH. This structure is said to be a "model" of the axioms. Although the term "model" is not often seen in mathematics outside of formal logic, it is actually a familiar concept. For example, in group theory, a "model of the group-theoretic axioms" is just a group, i.e., a set with a binary operation satisfying blah blah blah.

Analogously, we could invent a term---say, "universe"---to mean "a structure that is a model of ZFC." Then we could begin a course in set theory with a definition like, "A `universe' is a set M together with a binary relation E satisfying..." followed by a long list of axioms such as

> If x and y are distinct elements of M then either there exists z in M such that z E x but not z E y, or there exists z in M such that z E y but not z E x. (This is the "Axiom of Extensionality.")

[Note: It so happens that the term "universe" is not a standard one; for some reason there is standard name for the *axioms* but no standard name for the *structures that satisfy the axioms*, even though in the rest of mathematics it's the other way around. But the point is still valid.]

Early on in a course in group theory, one gives lots of examples of groups. Here we encounter the first awkward feature of ZFC; strictly speaking, nobody has ever exhibited even *one* model of ZFC (without using assumptions that are not universally accepted among all mathematicians)! Moreover, Goedel's second incompleteness theorem essentially says that nobody ever will.

Fortunately, this is not as nasty a problem as it might seem at first. For example, one object that is almost a model of ZFC is V, the class of all sets. If we take M = V and we take E to mean "is a member of," then we see that the axiom of extensionality simply says that two sets are equal if and only if they contain the same elements---a manifestly true statement. The rest of the axioms of ZFC are similarly self-evident when M = V. The catch, of course, is that V is too large to actually be a *set*, and we stipulated that M had to be a set.

Furthermore, Goedel incompleteness notwithstanding, nobody has ever proved that there *aren't* any models of ZFC, so if we are bothered by this point, we can treat the existence of models of ZFC like any other unproved hypothesis, such as the Riemann hypothesis or P != NP. That is, we can assume it freely as long as we remember to put a conditional clause in our theorems. So from now on, let us assume that models of ZFC exist.

A second possible stumbling block is that some people initially approach axiomatic set theory with the expectation that we have to suspend all our previous preconceptions about sets. It may therefore seem surprising that I suggested starting the discussion by saying, "A `universe' is a *set* together with..." Wait a minute---what's a set? Isn't it circular to define axioms for sets in terms of sets?

Different people give different answers to this question. My recommendation is to approach the study of ZFC like the study of any other mathematical subject. In particular, you should use your ordinary, "naive" mathematical knowledge of sets to understand the sentence, "A `universe' is a set together with..." Do not harbor expectations, at least not at first, that ZFC will tell you what sets "really are" or which mathematical practices are "really rigorous." If it helps, think of ZFC as axioms for "universe theory" not "set theory"; then the apparent circularity goes away.

Next point: In group theory, the deepest theorems are not theorems about *all* groups, but about special classes of groups (finite simple groups, compact Lie groups, etc.). Similarly, in the theory of forcing, we will be interested in models of ZFC with some additional properties. The first property we will require is that the elements of M be "well-founded sets." This means sets that are built up from the empty set recursively, using operations such as union, pairing, powerset, etc. Thus the empty set {} is well-founded, as are {{}} and {{}, {{}}, {{{}}}, ...}. This requirement is not really crucial; as in the rest of mathematics, it doesn't really matter what the "atoms" out of which your structures are built are, so long as they have the properties you want them to have. In fact, it might seem rather confusing that not only is M a set, but now the elements of M (and the elements of the elements of M, and...) are also sets! However, well-foundedness turns out to be very convenient.

The second requirement is a requirement on the relation E. We will require that x E y if and only if x is a member of y. (Remember, x and y, being elements of M, are [well-founded] sets, and hence it makes sense to talk about x being a member of y.) This might again seem like a confusing requirement. Naively, we might guess that to construct "weird" models of ZFC in which ~CH holds, we might have to let E be some kind of weird relation that bears no resemblance whatsoever to ordinary set membership. Once we require E to coincide with ordinary set membership, we might wonder whether there is any model of ZFC other than V itself (ignoring temporarily the fact that V is not a set). Surprisingly, it turns out that there is still plenty of "wiggle room" left after this requirement.

The third requirement is that M be "transitive," i.e., that every member of M is actually a subset of M. If you have worked with ordinal numbers before, you will not be too surprised that transitivity turns out to be a very useful condition. Otherwise, it might seem unmotivated; however, some motivation will become apparent later.

The final requirement is perhaps the most important (from the point of view of forcing), and that is that M be a *countable* set. If you have studied some logic then you may remember the Loewenheim-Skolem theorem, which says that if a countable set of axioms has any models at all, then it has a model which is a countable set. Otherwise, the countability condition may make you even more incredulous that any such M could exist. Since our main purpose here is to discuss forcing and I don't want to spend any more time on preliminaries, I will simply have to ask you to take it on faith that a very rich array of examples remains even after the countability condition is imposed.

We use the standard abbreviation "c.t.m." (countable transitive model) for models of ZFC that satisfy the four extra requirements above.

2. What a c.t.m. satisfying ~CH looks like

What do c.t.m.'s look like? It turns out that every c.t.m. M contains all the natural numbers, if we adopt the understanding that 0 is the empty set, 1 = {{}}, 2 = {{}, {{}}}, and in general n is the set {0, 1, ..., n-1}. It also contains omega, the infinite set of all the natural numbers. What about the powerset 2^omega (the set of all subsets of omega)? This is easy to get confused about, because on the one hand 2^omega is uncountable, and since M is transitive, if 2^omega were a member of M it would also be a subset of M, but since M is countable it can't have an uncountable subset. On the other hand, ZFC contains a Powerset Axiom which says that if x is in M then so is the powerset of x. So what's going on?

Let us look carefully at what the Powerset Axiom really states. It says that for every x in M, there exists a y in M with the following property: if z is a member of M such that every w in M satisfying w E z also satisfies w E x, then z E y. We see now that even if E is interpreted as membership (requirement 2 in the previous section), it does not follow that y is the set of *all* subsets of z. First of all, it is not even clear that the z's in this axiom are actually subsets of x; the axiom doesn't require that *every w* satisfying w E z also satisfies w E x; it just requires that *every w in M* satisfying w E z satisfies w E x. Now, it turns out that because M is transitive, the z's *are* in fact subsets of x. What "goes wrong" is the rest of the axiom: y does not contain *every subset of x*; it only contains *those subsets of x that are in M*. So it is perfectly possible that this "powerset" of x is countable.

What should we call y? Calling it the "powerset of x" is potentially confusing; I prefer to reserve this term for the actual set of all subsets of x. The usual jargon is to call y "the powerset of x in M."

As an exercise in understanding this concept, consider Cantor's famous theorem that the powerset of omega is uncountable. Cantor's proof can be mimicked using the axioms of ZFC to produce a formal theorem of ZFC. This yields: "The powerset of omega in M is uncountable in M." In order to see more clearly what this is saying, let us expand this out more fully to obtain: "There is no bijection in M between omega and the powerset of omega in M," where a bijection is a certain set (in fact, a certain set of ordered pairs, where an "ordered pair" <x,y> is defined set-theoretically as {{x}, {x,y}}). So even though the powerset of omega in M is countable, and we can construct a bijection in the "real world" between omega and the powerset of omega in M, it turns out that this bijection is *not a member of M*. There is therefore no contradiction between being "uncountable in M" and being countable (this is known as "Skolem's paradox").

Once you grasp this point that appending "in M" is crucial and can change the meaning of a term or a sentence dramatically, you may start to worry about all kinds of things in the preceding paragraphs. For example, shouldn't we distinguish between "omega" and "omega in M"? This is a legitimate worry. Fortunately, the transitivity of M implies that a lot of things, including "is a subset of," "is a function," "omega," and other basic concepts, are *absolute*, meaning that they mean the same whether or not you append "in M." A complete treatment of forcing must necessarily include a careful discussion of which concepts are absolute and which are not. However, this is a rather tedious affair, so we will gloss over it.

Instead, we will simply warn the reader when something is not absolute.

We can now present the outlines of how one might go about constructing a model for ZFC + ~CH. Start with any c.t.m. M. By the elementary theory of cardinals, there is a set k in M that is omega_2 in M (omega_2, or aleph_2 if you prefer, is the second cardinal larger than omega). Of course k is countable, but there is no bijection *in M* between k and omega. So now let us construct a function f from the Cartesian product k x omega into 2. This may be interpreted as a sequence of k functions from omega into 2. We can easily arrange that these k functions are all distinct. Now if f is already in M, then M satisfies ~CH! The reason is that functions from omega into 2 can be identified with subsets of omega, and f therefore shows us that the powerset of omega in M must be at least omega_2 in M. If f is missing from M, then we simply insert it into M to get a bigger set M' that satisfies ~CH. Q.E.D., right?

If this sounds too easy, you're right. There are many problems with this naive idea. First of all, just adding an arbitrary set out of the blue to M isn't necessarily going to result in something else that satisfies ZFC. We can't take just any group and adjoin an arbitrary extra element and hope that the result will still be a group, so why should we expect that this will work with models of ZFC? Where, for instance, is the powerset of f in M'? We have not done anything to ensure that it exists.

Nevertheless, this basic idea, of adding f (along with some other associated sets that we need to keep ZFC happy) to M to get a larger c.t.m. M' that contains some sets that were "missing" from M, turns out to work, if we are sufficiently careful and clever about the details. Specifically, we have the following.

Fundamental Theorem. If M is a c.t.m., P is an element of M [satisfying certain technical conditions], and G is a subset of P [satisfying certain technical conditions], then the set M[G], obtained by adjoining G [plus some other auxiliary sets] to M, is also a c.t.m.

The bracketed omissions will be filled in and explained shortly, but let me first remark that the Fundamental Theorem is very powerful, because the technical conditions on G and P are fairly mild, leaving tremendous scope for creativity. By cleverly choosing G and P one can construct M[G]'s that satisfy all kinds of set-theoretic statements. Even today, nearly all of the hundreds of known independence results in logic are ultimately based in one way or another on this one fundamental construction.

The reader may be puzzled that the word "forcing" does not appear in the Fundamental Theorem. Forcing, as we shall see later, comes in when one tries to prove that M[G] satisfies various statements (or, in the usual jargon, that various statements are "true in M[G]"). Forcing is needed to prove the Fundamental Theorem, since we have to prove that each axiom of ZFC is true in M[G]. Forcing is also needed to prove that ~CH is true in M[G] (for a suitable G).

3. Names

We begin by describing the "auxiliary sets" that, along with G, comprise M[G]. A crucial idea that pervades the theory of forcing is that some facts about G (or more generally M[G]) depend on the particular G in question, while others are "general" facts that are true for all G. Consider the analogy of field extensions: If F is a field and we want to adjoin a new element X to the field, then no matter what X is, we know that there is also going to be an element X^2 + 1 in F(X); on the other hand, whether X^2 + 1 = 0 depends on the particular X chosen.

In the case of fields there is no pressing need to have different notations for a "general" X and a "particular" X because the situation is sufficiently simple that no confusion arises if we use "X" for everything. Things are more complicated in ZFC, and it is important to have different notations for the "general" and "particular" cases. By analogy with the way natural language mirrors certain general features of the real world, we will define "names" for the elements of M[G]. The names will be defined without reference to G; however, the *values* of the names---the sets that the names are names *of*---will depend on G. The precise definition of a name (or more precisely, a "P-name" since it depends on P) is:

A set x is a P-name if (and only if) all its members are ordered pairs <y,p> where y is a P-name and p is in P.

If x is a P-name then we also define the domain dom(x) by

dom(x) = {y : <y,p> is in x for some p in P}

The definition of a P-name is rather tricky so let us examine it carefully. The first subtlety is that it appears to be circular. However, this apparent circularity is really just a shorthand for a recursive definition. For example, the empty set is a P-name, because all the members of the empty set certainly satisfy the required property. Then, once we know that the empty set is a P-name, we can build other sets that are P-names using the definition.

A P-name x is supposed to name some set X that is related to G, but x is designed for discussion of facts that are true "in general" for all G, so the members of x (or more precisely, of dom(x)) are not, as we might at first expect, the names of all the elements of X---that's "too specific." Rather, the members of x are the names of everything that is *potentially* in X. Now, if I happen to know exactly which particular G I'm interested in, then I should be able to figure out *exactly* which set x names. To this end, each member of x is "tagged" with an element of P. If I know G, then I can go through each member of x, and whenever the "tag" is actually in G, then I transform the potential membership into actual membership, whereas if the tag is not in G, then I discard the potential member. Formally, the "value of a P-name x with respect to G," written val(x,G), is defined to be

val(x,G) = {val(y,G) : <y,p> is in x for some p in G}

This is again a recursive definition. The value of the empty set is always the empty set, and we can build up other values from this.

One thing that may still seem obscure in this definition is the role of P. Couldn't we simply replace "P" by "M" everywhere? Why do we need P? It is true that we could define names and their values without reference to P, but there are two reasons for including it: (1) to prove the main theorems of forcing, we will need to impose a certain structure on P, which we cannot in general impose on M itself, and (2) after we finish building the machinery and want to apply it, having the extra degree of freedom of choosing a good P will be very valuable in constructing c.t.m.'s with special properties.

The names construction is quite powerful. For example, let us, now and throughout the rest of this article, require that P always contain a distinguished element called "1" and let us also require that G always contain 1. (We previously used "1" for the set {{}} but there should be no confusion because we will have no further occasion to mention {{}}.) Then there is a name---in fact, a "canonical" name---for every element X of M. This can be proved by recursion: if every element Y of X has a name Y*, then we can formulate the name X* = {<Y*,1> : all Y in X}. Then since 1 is always in G, it follows that val(X*,G) = X. (Furthermore, this construction turns out to be absolute, so X* is in M.) From this we can deduce that G has a name too, the "diagonal" name G* = {<X*,X> : all X in P}. Chasing the definition, we see that the value of this name is the set of values of X* as X runs over all elements of G, i.e., val(G*,G) = G.

We can now fill in one blank in the Fundamental Theorem. Given M, P, and G, the formal definition of M[G] is:

```
  M[G] = {val(x,G) : x is a P-name in M}.
```

We would like to claim is that M[G] is a c.t.m. if M is. However, this turns out to be false without further conditions on P and G. Why this is so is the topic of the next section.

4. A naive attempt to prove that ZFC is true in M[G]

The ideas presented so far, while not sufficient to prove that M[G] satisfies ZFC, do buy us *some* of the axioms of ZFC. In particular, M[G] satisfies the Pairing and Union axioms of ZFC. I didn't define what these axioms are, but they are easy to state. For example, Pairing says that for every X and Y in M[G] there exists Z in M[G] such that X E Z and Y E Z. To see this, suppose that X and Y are in M[G]. By definition of M[G], X = val(x,G) and Y = val(y,G) for some P-names x and y in M. Now set z = {<x,1>, <y,1>}. This is a P-name, and it is also in M (the latter fact requires proof, but it is essentially because M satisfies Pairing). So Z = val(z,G) = {X,Y} is in M[G] and it has the desired property (remember that "E" denotes membership). Pairing is so easy that we don't even need most of the power of the concept of names to prove that it is true in M[G].

Trouble arises, however, with other axioms of ZFC, such as the Powerset axiom. Suppose X is in M[G], and let x be a P-name in M such that val(x,G) = X. We need to find a P-name y in M such that Y = val(y,G) contains every subset of X that lies in M[G]. The following definition is a natural try:

```
   y = {<z,1> : z is a P-name in M and dom(z) is a subset
                of dom(x)}
```

Now suppose Z is a subset of X in M[G]. There must exist a P-name z in M such that Z = val(z,G). We would love to be able to say that dom(z) must be a subset of dom(x), because val(z,G) would then be in Y, and we would be done. But unfortunately, this need not be true. The problem is that z might contain a lot of "potential" elements that are not in dom(x) but that "disappear" when you take the value, because they are tagged with elements of P that are not in G. So it seems we're stuck.

All is not lost, though. Different P-names in M may actually have the same value; although an arbitrary z that names Z might not have a domain contained in dom(x), perhaps a careful choice of z will.

But how do we find such a z? Answering this question leads us (finally!) to the concept of forcing, to which we now turn.

5. Forcing

To make progress we need to investigate further the idea that some facts about G are specific to G while other facts are "true in general." For example, whether or not the empty set is a member of G is something that is very specific to G. In contrast, consider the following statement: If the empty set is a member of G, then G is nonempty. This is something that can be said about G *without* knowing anything specific about G. For a slightly less trivial example, we can also say that if omega is a member of G, then the union of all the members of G is an infinite set. More generally, without any knowledge of facts that are particular to G, it is possible to construct quite sophisticated statements of the form, "*If* G contains a certain element p, *then* such-and-such must be true." Let us now define forcing:

> Let p be an element of P and let phi be any sentence whose nouns are P-names. Then we say that "p ||- phi" (read: "p forces phi") if, for any subset G of P [satisfying technical conditions], p being a member of G implies that phi is true in M[G] (i.e., phi is true when the names are replaced by their values).

This definition requires some explanation. First, it is not actually the true definition of forcing, unless the technical conditions are inserted (we will return to this later; ignore it for now). The main part of the definition is best clarified by example. Write 0 for the empty set, and let G* be the canonical name of G. Let phi be the sentence, "G* is nonempty." Then "0 ||- phi" states that for any G, 0 being a member of G implies the following statement: "G is nonempty" (I have replaced, in phi, the name G* by its value, G). This is true. On the other hand, if omega* is the canonical name of omega, then we would not expect "0 ||- omega* is in G*" to be true; that is, we would not expect the fact that 0 is in G to imply that omega is in G, because some G's might contain 0 and not omega.

How does this concept help us? Let us return to the problem in the previous section, of proving that M[G] satisfies the Powerset Axiom. With the same notation that was used there, suppose we now define z' as follows:

> z' = {<q,r> : q is in dom(x) and r ||- "q is in z"}.

By definition, dom(z') is a subset of dom(x). What is val(z',G)? Well, this is determined by which "tags" are in G. But note that by the definition of forcing, if r is in G, then val(q,G) is in val(z,G) = Z. Therefore, val(z',G) is a subset of Z.

This is progress, but what we really want is for val(z',G) to be actually equal to Z, and this is not at all clear. What we need in order to prove val(z',G) = Z is that for every Q in Z, there is a P-name q of Q in dom(x) and a P-name r in G such that r ||- "q is in z." Finding a suitable q turns out not to be that hard, but the existence of r is not so clear; what we need is a theorem that for every statement phi (such as "Q is in Z") that is actually true in M[G] for a particular G, there is some element r of that particular G such that merely knowing that r is in G is enough to deduce that phi is true of G. This doesn't seem plausible. In general, knowing just one element of G doesn't tell us very much about G.

However, we can fix this problem if we restrict G and P to have a certain structure, so that some elements r of G are "more informative" than others, i.e., knowing that r is in G tells us more about G than knowing that some other element p is in G. The simplest way to achieve this is to impose a partial ordering < on P (with "1" being the unique maximal element) and to require G to be a *filter*, i.e., for G to have the following two properties:

1. If r is in G and r <= p then p is also in G.
2. For every p and q in G there exists r in G such that r <= p and r <= q.

The first of these conditions ensures that knowing that some element r is in G tells us not only that r is in G but that every element >= r is also in G. Thus, r is "more informative" than p if r <= p. The second condition may be thought of as capturing the idea that for any two elements p and q, there is some element r that is at least as informative as p and q put together. If we are going to insist on looking at only single elements of G to get information about G, then this is clearly a useful condition to impose.

If we now revise our definition of ||- by inserting the technical conditions that P be a partially ordered set and G be a filter, then this gives us more hope of being able to find r in G such that r ||- "Q is in Z." Unfortunately, even the condition of being a filter turns out not to be quite good enough. Intuitively, the reason is that since P may be infinite, there is no guarantee that, for an arbitrary statement phi that asserts something about G, there will exist an r in G that will give us enough information to decide phi. We might try finding smaller and smaller r's in the partial order that give us more and more information but we might never find one that tells us what we need to know.

The following condition on G solves this problem. Define a subset D of P to be "dense" if for every p in P there exists q in D such that q <= p. A filter G is "generic" if it intersects every dense subset.

It can now be proved that if we require P to be a partially ordered set with unique maximal element 1 and we define "p ||- phi" to mean that for all generic filters G in P, phi is true in M[G], then the following is true:

Forcing Fact 1. Fix a particular generic filter G of P. Then for any phi that is true in M[G], there exists r in G such that r ||- phi.

If we now go back to our attempt to prove the Powerset Axiom, we see that Forcing Fact 1 is just what we need to prove that val(z',G) = Z. The only thing left to prove is that z' is actually in M. For this, one needs the following:

Forcing Fact 2. ||- is definable in M.

I won't say precisely what Forcing Fact 2 means, but intuitively it asserts that even when we add the side conditions (partial order, generic filter), the definition of ||- is still sufficiently "general" that it does not require any specific knowledge of G; therefore, sets like z' that are defined using ||- still lie inside M.

Proving Forcing Facts 1 and 2 is not easy, even after all the above definitions (which themselves are not easy to come up with!) have been provided. Here I must refer you to a Kunen's book for details. The work required to prove them, however, pays off great dividends. Not only Powerset, but all the other axioms of ZFC, can be proved to be true in M[G], using just the two Forcing Facts (without ever having to delve into the proof of the Facts themselves), and following a similar pattern to our proof of Powerset. The Fundamental Theorem is therefore proved! For convenience, we restate it with all the blanks filled in.

Fundamental Theorem. If M is a c.t.m., P is a partially ordered set in M with unique maximum element 1, and G is a nonempty generic filter of P, then the set M[G] = {val(x,G) : x is a P-name in M} is also a c.t.m.

6. ~CH revisited

Let's see how to apply the Fundamental Theorem to the continuum hypothesis. One's first thought might be to take P to be the set of all ordered pairs <x,y> with x in k x omega and y in 2, hoping that G can be chosen to be a function from k x omega to 2. However, this does not quite work because the members of a function are ordered pairs, and no particular ordered pair is "more informative" about the function than any other ordered pair.

On the other hand, large *subsets* of ordered pairs carry more information than smaller subsets, so this is a hint that we should "pass to the powerset." More precisely, let P be the set of all "finite partial functions" from k x omega into 2, i.e., functions from a finite subset of k x omega into 2. Partially order these functions by reverse inclusion, i.e., f <= g if the set of all ordered pairs in g is contained in the set of all ordered pairs of f. Now it is not hard to see that for each y in k x omega, the set D_y of all elements of P that are defined on y (i.e., that include y as part of their domain) is dense. Therefore, if G is a generic filter in P, G must intersect every such D_y, and the union of all the elements of G will be a *total* function from k x omega to 2. The Fundamental Theorem tells us that we can successfully "insert" this "missing" function into M if it wasn't there already. Furthermore, the fact that G is generic gives us automatically, as a bonus, the fact that each of the k functions from omega to 2 are distinct! (See if you can prove this, by constructing a suitable dense subset in P. You can look up the answer in Kunen.) There is no need to cleverly choose the "right" generic filter; any generic filter will do in this case.

Are we done? Well, almost...there are two more points to take care of. First of all, we haven't yet established that generic filters exist. This turns out not to be difficult to prove, but it does hinge on the countability of M (this is where the countability of M turns out to be important). Essentially, because of countability, one can enumerate all the dense subsets, and pick one from each; the filter generated by all these elements will be generic.

The second point is more substantial. It is conceivable to me that Cohen could have devised forcing, proved the Fundamental Theorem, and still failed to prove anything new about the continuum hypothesis! The reason is that nothing in the general framework of forcing guarantees that in M[G], the cardinal k will still equal omega_2. Being a cardinal is not an absolute property, so large cardinals can "collapse" into lower ordinals when you pass to M[G]. The reason is that the auxiliary sets in M[G] might include a bijection between k and a smaller cardinal. Luckily for Cohen, though, it can be proved that cardinal collapse does not occur for the particular P described above. See Kunen for details.

7. Epilogue

One reason I'm still not totally happy with this exposition is that the definitions "p ||- phi" and "G is a generic filter" still seem to be pulled out of a hat, even though one can get some a posteriori intuition about why each part of the definition is important. There do exist alternative approaches to forcing, e.g., using "Boolean-valued models." This is somewhat more intuitive because the idea that some statements require more "particular information about G" to be decided is encoded "directly," i.e., we take the set of all statements and impose a partial order on this set. The concept of a generic filter becomes simpler (essentially, it becomes an "ultrafilter" in a complete Boolean algebra). However, with this approach one needs to hit on the idea of embedding partial orders in a complete Boolean algebra; I'm not sure how to motivate this. Question for the logicians: are there nice applications of forcing where P is a complete Boolean algebra and no embedding is involved?

Another question for the logicians: Say we drop the word "generic" from the definition of forcing, and we say that "p pseudoforces phi" if for every filter G in P, p being in G implies that phi is true in M[G]. Is pseudoforcing definable in M? At first I thought it might be, because it doesn't seem that one needs to know anything about any particular G to decide whether p pseudoforces phi, but Andreas Blass thinks the answer is probably no.