WELL-ORDERED SETS, ORDINALITIES AND THE AXIOM OF CHOICE

PETE L. CLARK

1. The Calculus of Ordinalities

1.1. **Well-ordered sets and ordinalities.**

The discussion of cardinalities in Chapter 2 suggests that the most interesting thing about them is their order relation, namely that any set of cardinalities forms a well-ordered set. So in this section we shall embark upon a systematic study of well-ordered sets. Remarkably, we will see that the problem of classifying sets up to bijection is literally contained in the problem of classifying well-ordered sets up to order-isomorphism.

Exercise 1.1.1: Show that for a linearly ordered set X, TFAE:
(i) X satisfies the descending chain condition: there are no infinite strictly descending sequences $x_1 > x_2 > \ldots$ in X.
(ii) X is well-ordered.

We need the notion of "equivalence" of of well-ordered sets. A mapping $f : S \to T$ between partially ordered sets is **order preserving** (or an **order homomorphism**) if $s_1 \leq s_2$ in S implies $f(s_1) \leq f(s_2)$ in T.

Exercise 1.1.2: Let $f : S \to T$ and $g : T \to U$ be order homomorphisms of partially ordered sets.
a) Show that $g \circ f : S \to U$ is an order homomorphism.
b) Note that the identity map from a partially ordered set to itself is an order homomorphism.
(It follows that there is a **category** whose objects are partially ordered sets and whose morphisms are order homomorphisms.)

An **order isomorphism** between posets is a mapping f which is order preserving, bijective, and whose inverse f^{-1} is order preserving. (This is the general – i.e., categorical – definition of isomorphism of structures.)

Exercise 1.1.3: Give an example of an order preserving bijection f such that f^{-1} is not order preserving.

However:

Lemma 1. *An order-preserving bijection whose domain is a totally ordered set is an order isomorphism.*

Exercise 1.1.4: Prove Lemma 1.

Lemma 2. *(Rigidity Lemma) Let S and T be well-ordered sets and $f_1,\ f_2 : S \to T$ two order isomorphisms. Then $f_1 = f_2$.*

Proof: Let f_1 and f_2 be two order isomorphisms between the well-ordered sets S and T, which we may certainly assume are nonempty. Consider S_2, the set of elements s of S such that $f_1(s) \neq f_2(s)$, and let $S_1 = S \setminus S_2$. Since the least element of S must get mapped to the least element of T by any order-preserving map, S_1 is nonempty; put $T_1 = f_1(S_1) = f_2(S_1)$. Supposing that S_2 is nonempty, let s_2 be its least element. Then $f_1(s_2)$ and $f_2(s_2)$ are both characterized by being the least element of $T \setminus T_1$, so they must be equal, a contradiction.

Exercise 1.1.5: Let S be a partially ordered set.
a) Show that the order isomorphisms $f : S \to S$ form a group, the **order automorphism group** $\mathrm{Aut}(S)$ of S. (The same holds for any object in any category.)
b) Notice that Lemma 2 implies that the automorphism group of a well-ordered set is the trivial group.[1]
c) Suppose S is linearly ordered and f is an order automorphism of S such that for some positive integer n we have $f^n = \mathrm{Id}_S$, the identity map. Show that $f = \mathrm{Id}_S$. (Thus the automorphism group of a linearly ordered set is **torsionfree**.)
d) For any infinite cardinality κ, find a linearly ordered set S with $|\mathrm{Aut}(S)| \geq \kappa$. Can we always ensure equality?
e)** Show that every group G is (isomorphic to) the automorphism group of some partially ordered set.

Let us define an **ordinality** to be an order-isomorphism class of well-ordered sets, and write $o(X)$ for the order-isomorphism class of X. The intentionally graceless terminology will be cleaned up later on. Since two-order isomorphic sets are equipotent, we can associate to every ordinality α an "underlying" cardinality $|\alpha|$: $|o(X)| = |X|$. It is natural to expect that the classification of ordinalities will be somewhat richer than the classification of cardinalities – in general, endowing a set with extra structure leads to a richer classification – but the reader new to the subject may be (we hope, pleasantly) surprised at how much richer the theory becomes.

From the perspective of forming "isomorphism classes" (a notion the ontological details of which we have not found it profitable to investigate too closely) ordinalities have a distinct advantage over cardinalities: according to the Rigidity Lemma, any two representatives of the same ordinality are *uniquely* (hence canonically!) isomorphic. This in turn raises the hope that we can write down a *canonical* representative of each ordinality. This hope has indeed been realized, by von Neumann, as we shall see later on: the canonical representatives will be called "ordinals." While we are alluding to later developments, let us mention that just as we can associate a cardinality to each ordinality, we can also – and this is much more profound – associate an ordinality $o(\kappa)$ to each cardinality κ. This assignment is *one-to-one*, and this allows us to give a canonical representative to each cardinality, a "cardinal." At least at the current level of discussion, there is no purely mathematical advantage to the passage from cardinalities to cardinals, but it has a

[1]One says that a structure is **rigid** if it has no nontrivial automorphisms.

striking ontological[2] consequence, namely that, up to isomorphism, we may develop all of set theory in the context of "pure sets", i.e., sets whose elements (and whose elements' elements, and ...) are themselves sets.

But first let us give some basic examples of ordinalities and ways to construct new ordinalities from preexisting ones.

1.2. **Algebra of ordinalities.**

Example 1.2.1: Trivially the empty set is well-ordered, as is any set of cardinality one. These sets, and only these sets, have unique well-orderings.

Example 1.2.2: Our "standard" example $[n]$ of the cardinality n comes with a well-ordering. Moreover, on a finite set, the concepts of well-ordering and linear ordering coincide, and it is clear that there is up to order isomorphism a unique linear ordering on $[n]$. Informally, given any two orderings on an n element set, we define an order-preserving bijection by pairing up the least elements, then the second-least elements, and so forth. (For a formal proof, use induction.)

Example 1.2.3: The usual ordering on $\mathbb{N}$ is a well-ordering. Notice that this is isomorphic to the ordering on $\{n \in \mathbb{Z} \mid n \geq n_0\}$ for any $n_0 \in \mathbb{Z}$. As is traditional, we write ω for the ordinality of $\mathbb{N}$.

Exercise 1.2.4: For any ordering $\leq$ on a set X, we have the opposite ordering $\leq'$, defined by $x \leq' y$ iff $y \leq x$.
a) If $\leq$ is a linear ordering, so is $\leq'$.
b) If both $\leq$ and $\leq'$ are well-orderings, then X is finite.

For a partially ordered set X, we can define a new partially ordered set $X^+ := X \cup \{\infty\}$ by adjoining some new element ∞ and decreeing $x \leq \infty$ for all $x \in X$.

Proposition 3. *If X is well-ordered, so is X^+.*

Proof: Let Y be a nonempty subset of X^+. Certainly there is a least element if $|Y| = 1$; otherwise, Y contains an element other than ∞, so that $Y \cap X$ is nonempty, and its least element will be the least element of Y.

If X and Y are order-isomorphic, so too are X^+ and Y^+, so the passage from X to X^+ may be viewed as an operation on ordinalities. We denote $o(X^+)$ by $o(X) + 1$, the **successor ordinality** of $o(X)$.

Note that all the finite ordinalities are formed from the empty ordinality 0 by iterated successorship. However, not every ordinality is of the form $o' + 1$, e.g. ω is clearly not: it lacks a maximal element. (On the other hand, it is obtained from *all* the finite ordinalities in a way that we will come back to shortly.) We will say that an ordinality o is a **successor ordinality** if it is of the form $o' + 1$ for some ordinality o' and a **limit ordinality** otherwise. Thus 0 and ω are limit ordinals.

[2]I restrain myself from writing "ontological" (i.e., with quotation marks), being like most contemporary mathematicians alarmed by statements about the reality of mathematical objects.

Example 1.2.6: The successor operation allows us to construct from ω the new ordinals $\omega + 1$, $\omega + 2 := (\omega + 1) + 1$, and for all $n \in \mathbb{Z}^+$, $\omega + n := (\omega + (n-1)) + 1$: these are all distinct ordinals.

Definition: For partially ordered sets $(S_1, \leq_1)$ and $(S_2, \leq_2)$, we define $S_1 + S_2$ to be the set $S_1 \coprod S_2$ with $s \leq t$ if either of the following holds:
(i) For $i = 1$ or 2, s and t are both in S_i and $s \leq_i t$;
(ii) $s \in S_1$ and $s \in S_2$.

Informally, we may think of $S_1 + S_2$ as "S_1 followed by S_2."

Proposition 4. *If S_1 and S_2 are linearly ordered (resp. well-ordered), so is $S_1 + S_2$.*

Exercise 1.2.5: Prove Proposition 4.

Again the operation is well-defined on ordinalities, so we may speak of the **ordinal sum** $o + o'$. By taking $S_2 = [1]$, we recover the successor ordinality: $o + [1] = o + 1$.

Example 1.2.6: The ordinality $2\omega := \omega + \omega$ is the class of a well-ordered set which contains one copy of the natural numbers followed by another. Proceeding inductively, we have $n\omega := (n-1)\omega + \omega$, with a similar description.

Tournant dangereuse: We can also form the ordinal sum $1 + \omega$, which amounts to adjoining to the natural numbers a smallest element. But this is still order-isomorphic to the natural numbers: $1 + \omega = \omega$. In fact the identity $1 + o = o$ holds for every infinite ordinality (as will be clear later on). In particular $1 + \omega \neq \omega + 1$, so beware: the ordinal sum is not commutative! (To my knowledge it is the only non-commutative operation in all of mathematics which is invariably denoted by "+".) It is however immediately seen to be associative.

The notation 2ω suggests that we should have an ordinal product, and indeed we do:

Definition: For posets $(S_1, \leq_1)$ and $(S_2, \leq_2)$ we define the **lexicographic product** to be the Cartesian product $S_1 \times S_2$ endowed with the relation $(s_1, s_2) \leq (t_1, t_2)$ if(f) either $s_1 \leq t_1$ or $s_1 = t_1$ and $s_2 \leq t_2$. If the reasoning behind the nomenclature is unclear, I suggest you look up "lexicographic" in the dictionary.[3]

Proposition 5. *If S_1 and S_2 are linearly ordered (resp. well-ordered), so is $S_1 \times S_2$.*

Exercise 1.2.7: Prove Proposition 5.

As usual this is well-defined on ordinalities so leads to the **ordinal product** $o \cdot o'$.

Example 1.2.8: For any well-ordered set X, $[2] \cdot X$ gives us one copy $\{(1, x) \mid x \in X\}$ followed by another copy $\{(2, x) \mid x \in X\}$, so we have a natural isomorphism of $[2] \cdot X$ with $X + X$ and hence $2 \cdot o = o + o$. (Similarly for $3o$ and so forth.) This time we should be prepared for the failure of commutativity: $\omega \cdot n$ is isomorphic to ω. This allows us to write down $\omega^2 := \omega \times \omega$, which we visualize by starting with the positive integers and then "blowing up" each positive integer to give a whole

[3] Ha ha.

order isomorphic copy of the positive integers again. Repeating this operation gives $\omega^3 = \omega^2 \cdot \omega$, and so forth. Altogether this allows us to write down ordinalities of the form $P(\omega) = a_n\omega^n + \ldots + a_1\omega + a_0$ with $a_i \in \mathbb{N}$, i.e., polynomials in ω with natural number coefficients. It is in fact the case that (i) distinct polynomials $P \neq Q \in \mathbb{N}[T]$ give rise to distinct ordinalities $P(\omega) \neq Q(\omega)$; and (ii) any ordinality formed from $[n]$ and ω by finitely many sums and products is equal to some $P(\omega)$ – even when we add/multiply in "the wrong order", e.g. $\omega * 7 * \omega^2 * 4 + \omega * 3 + 11 = \omega^3 + \omega + 11$ – but we will wait until we know more about the ordering of ordinalities to try to establish these facts.

Example 1.2.9: Let $\alpha_1 = o(X_1), \ldots, \alpha_n = o(X_n)$ be ordinalities.
a) Show that $\alpha_1 \times (\alpha_2 \times \alpha_3)$ and $(\alpha_1 \times \alpha_2) \times \alpha_3$ are each order isomorphic to the set $X_1 \times X_2 \times X_3$ endowed with the ordering $(x_1, x_2, x_3) \leq (y_1, y_2, y_3)$ if $x_1 < y_1$ or $(x_1 = y_1$ and $(x_2 < y_2$ or $(x_2 = y_2$ and $x_3 \leq y_3)))$. In particular ordinal multiplication is associative.
b) Give an explicit definition of the product well-ordering on $X_1 \times \ldots \times X_n$, another "lexicographic ordering."

In fact, we also have a way to exponentiate ordinalities: let $\alpha = o(X)$ and $\beta = o(Y)$. Then by α^β we mean the order isomorphism class of the set $Z = Z(Y, X)$ of all functions $f : Y \to X$ with $f(y) = 0_X$ (0_X denotes the minimal element of X) for all but finitely many $y \in Y$, ordered by $f_1 \leq f_2$ if $f_1 = f_2$ or, for the greatest element $y \in Y$ such that $f_1(y) \neq f_2(y)$ we have $f_1(y) < f_2(y)$.

Some helpful terminology: one has the zero function, which is 0 for all values. For every other $f \in W$, we define its **degree** $y_{\deg}$ to be the largest $y \in Y$ such that $f(y) \neq 0$ and its **leading coefficient** $x_l := f(y_{\deg})$.

Proposition 6. *For ordinalities α and β, α^β is an ordinality.*

Proof: Let Z be the set of finitely nonzero functions $f : Y \to X$ as above, and let $W \subset Z$ be a nonempty subset. We may assume 0 is not in W, since the zero function is the minimal element of all of Z. Thus the set of degrees of all elements of W is nonempty, and we may choose an element of minimal degree y_1; moreover, among all elements of minimal degree we may choose one with minimal leading coefficient x_1, say f_1. Suppose f_1 is not the minimal element of W, i.e., there exists $f' \in W_2$ with $f' < f_1$. Any such f' has the same degree and leading coefficient as f_1, so the last value y' at which f' and f_1 differ must be less than y_1. Since f_1 is nonzero at all such y' and f_1 is finitely nonzero, the set of all such y' is finite and thus has a *maximal* element y_2. Among all f' with $f'(y_2) < f(y_2)$ and $f'(y) = f(y)$ for all $y > y_2$, choose one with $x_2 = f'(y_2)$ minimal and call it f_2. If f_2 is not minimal, we may continue in this way, and indeed get a sequence of elements $f_1 > f_2 > f_3 \ldots$ as well as a descending chain $y_1 > y_2 > \ldots$. Since Y is well-ordered, this descending chain must terminate at some point, meaning that at some point we find a minimal element f_n of W.

Example 1.2.10: The ordinality ω^ω is the set of all finitely nonzero functions $f : \mathbb{N} \to \mathbb{N}$. At least formally, we can identify such functions as polynomials in ω with $\mathbb{N}$-coefficients: $P_f(\omega) = \sum_{n \in \mathbb{N}} f(n)\omega^n$. The well-ordering makes $P_f < P_g$ if the at the largest n for which $f(n) \neq g(n)$ we have $f(n) < g(n)$, e.g.

$\omega^3 + 2\omega^2 + 1 > \omega^3 + \omega^2 + \omega + 100$.

It is hard to ignore the following observation: ω^ω puts a natural well-ordering relation on all the ordinalities we had already defined. This makes us look back and see that the same seems to be the case for all ordinalities: e.g. ω itself is order isomorphic to the set of all the finite ordinalities $[n]$ with the obvious order relation between them. Now that we see the suggested order relation on the ordinalities of the form $P(\omega)$ one can check that this is the case for them as well: e.g. ω^2 can be realized as the set of all linear polynomials $\{a\omega + b \mid a, b \in \mathbb{N}\}$.

This suggests the following line of inquiry:

(i) Define a natural ordering on ordinalities (as we did for cardinalities).
(ii) Show that this ordering *well-orders* any set of ordinalities.

Exercise 1.2.11: Let α and β be ordinalities.
a) Show that $0^\beta = 0$, $1^\beta = 1$, $\alpha^0 = 1$, $\alpha^1 = \alpha$.
b) Show that the correspondence between finite ordinals and natural numbers respects exponentiation.
c) For an ordinal α, the symbol α^n now has two possible meanings: exponentiation and iterated multiplication. Show that the two ordinalities are equal. (The proof requires you to surmount a small left-to-right lexicographic difficulty.) In particular $|\alpha^n| = |\alpha|^n = |\alpha|$.
d) For any infinite ordinal β, show that $|\alpha^\beta| = \max(|\alpha|, |\beta|)$.

Tournant dangereuse: In particular, it is generally *not* the case that $|\alpha^\beta| = |\alpha|^{|\beta|}$: e.g. 2^ω and ω^ω are both countable ordinalities. In fact, we have not yet seen any uncountable well-ordered sets, and one cannot construct an uncountable ordinal from ω by any finite iteration of the ordinal operations we have described (nor by a countable iteration either, although we have not yet made formal sense of that). This leads us to wonder: are there any uncountable ordinalities?

1.3. **Ordering ordinalities.** Let S_1 and S_2 be two well-ordered sets. In analogy with our operation $\leq$ on sets, it would seem natural to define $S_1 \leq S_2$ if there exists an order-preserving injection $S_1 \to S_2$. This gives a relation $\leq$ on ordinalities which is clearly symmetric and transitive.

However, this is *not* the most useful definition of $\leq$ for well-ordered sets, since it gives up the rigidity property. In particular, recall Dedekind's characterization of infinite sets as those which are in bijection with a proper subset of themselves, or, equivalently, those which *inject* into a proper subset of themselves. With the above definition, this will still occur for infinite ordinalities: for instance, we can inject ω properly into itself just by taking $\mathbb{N} \to \mathbb{N}$, $n \mapsto n+1$. Even if we require the least elements to be preserved, then we can still inject $\mathbb{N}$ into any infinite subset of itself containing 0.

So as a sort of mild *deus ex machina* we will work with a more sophisticated

order relation. First, for a linearly ordered set S and $s \in S$, we define

$$I(s) = \{t \in S \mid t < s\},$$

an **initial segment** of S. Note that every initial segment is a proper subset. Let us also define

$$I[s] = \{t \in S \mid t \leq s\}.$$

Now, given linearly ordered sets S and T, we define $S < T$ if there exists an order-preserving embedding $f : S \to T$ such that $f(S)$ is an initial segment of T (say, an **initial embedding**). We define $S \leq T$ if $S < T$ or $S \cong T$.

Exercise 1.3.1: Let $f : S_1 \to S_2$ and $g : T_1 \to T_2$ be order isomorphisms of linearly ordered sets.
a) Suppose $s \in S_1$. Show that $f(I(s)) = I(f(s))$ and $f(I[s])) = I(f[s])$.
b) Suppose that $S_1 < T_1$ (resp. $S_1 \leq T_1$). Show that $S_2 < T_2$) (resp. $S_2 \leq T_2$).
c) Deduce that $<$ and $\leq$ give well-defined relations on any set $\mathcal{F}$ of ordinalities.

Exercise 1.3.2: a) Show that if $i : X \to Y$ and $j : Y \to Z$ are initial embeddings of linearly ordered sets, then $j \circ i : X \to Z$ is an initial embedding.
b) Deduce that the relation $<$ on any set of ordinalities is transitive.

Definition: In a partially ordered set X, a subset Z is an **order ideal** if for all $z \in Z$ and $x \in X$, if $x < z$ then $x \in Z$. For example, the empty set and X itself are always order ideals. We say that X is an **improper** order ideal of itself, and all other order ideals are **proper**. For instance, $I[s]$ is an order ideal, which may or may not be an initial segment.

Lemma 7. *("Principal ideal lemma") Any proper order ideal in a well-ordered set is an initial segment.*

Proof: Let Z be a proper order ideal in X, and s the least element of $X \setminus Z$. Then a moment's thought gives $Z = I(s)$.

The following is a key result:

Theorem 8. *(Ordinal trichotomy) For any two ordinalities $\alpha = o(X)$ and $\beta = o(Y)$, exactly one of the following holds: $\alpha < \beta$, $\alpha = \beta$, $\beta < \alpha$.*

Corollary 9. *Any set of ordinalities is linearly ordered under $\leq$.*

Exercise 1.3.3: Deduce Corollary 9 from Theorem 8. Is the Corollary equivalent to the Theorem?

Proof of Theorem 8: Part of the assertion is that no well-ordered set X is order isomorphic to any initial segment $I(s)$ in X (we would then have both $o(I(s)) < o(X)$ and $o(I(s)) = o(X)$); let us establish this first. Suppose to the contrary that $\iota : X \to X$ is an order embedding whose image is an initial segment $I(s)$. Then the set of x for which $\iota(x) \neq x$ is nonempty (otherwise ι would be the identity map, and no linearly ordered set is *equal* to any of its initial segments), so let x be the least such element. Then, since ι restricted to $I(x)$ is the identity map, $\iota(I(x)) = I(x)$,

so we cannot have $\iota(x) < x$ – that would contradict the injectivity of ι – so it must be the case that $\iota(x) > x$. But since $\iota(X)$ is an initial segment, this means that x is in the image of ι, which is seen to be impossible.

Now if $\alpha < \beta$ and $\beta < \alpha$ then we have initial embeddings $i : X \to Y$ and $j : Y \to X$. By Exercise 1.3.2 their composite $j \circ i : X \to X$ is an initial embedding, which we have just seen is impossible. It remains to show that if $\alpha \neq \beta$ there is either initial embedding from X to Y or vice versa. We may assume that X is nonempty. Let us try to build an initial embedding from X into Y. A little thought convinces us that we have no choices to make: suppose we have already defined an initial embedding f on a segment $I(s)$ of X. Then we *must* define $f(s)$ to be the least element of $Y \setminus f(I(s))$, and we *can* define it this way exactly when $f(I(s)) \neq Y$. If however $f(I(s)) = Y$, then we see that f^{-1} gives an initial embedding from Y to X. So assume Y is not isomorphic to an initial segment of X, and let Z be the set of x in X such that there exists an initial embedding from $I(z)$ to Y. It is immediate to see that Z is an order ideal, so by Lemma 7 we have either $Z = I(x)$ or $Z = X$. In the former case we have an initial embedding from $I(z)$ to Y, and as above, the only we could not extend it to x is if it is surjective, and then we are done as above. So we can extend the initial embedding to $I[x]$, which – again by Lemma 7 is either an initial segment (in which case we have a contradiction) or $I[x] = X$, in which case we are done. The last case is that $Z = X$ has no maximal element, but then we have $X = \bigcup_{x \in X} I(x)$ and a uniquely defined initial embedding ι on each $I(x)$. So altogether we have a map on all of X whose image $f(X)$, as a union of initial segments, is an order ideal. Applying Lemma 7 yet again, we either have $f(X) = Y$ – in which case f is an order isomorphism – or $f(X)$ is an initial segment of Y, in which case $X < Y$: done.

Exercise 1.3.4: Let α and β be ordinalities. Show that if $|\alpha| > |\beta|$, then $\alpha > \beta$. (Of course the converse does not hold: there are many countable ordinalities.)

Corollary 10. *Any set $\mathcal{F}$ of ordinalities is* well-ordered *with respect to $\leq$.*

Proof: Using Exercise 1.1.1, it suffices to prove that there is no infinite descending chain in $\mathcal{F} = \{o_\alpha\}_{\alpha \in I}$. So, seeking a contradiction, suppose that we have a sequence of well-ordered sets S_1, $S_2 = I(s_1)$ for $s_1 \in S_1$, $S_3 = I(s_2)$,...,$S_{n+1} = I(s_n)$ for $s_n \in S_n$,.... But all the S_n's live inside S_1 and we have produced an infinite descending chain $s_1 > s_2 > s_3 > \ldots > s_n > \ldots$ inside the well-ordered set S_1, a contradiction.

Thus any set $\mathcal{F}$ of ordinalities itself generates an ordinality $o(\mathcal{F})$, the ordinality of the well-ordering that we have just defined on $\mathcal{F}$!

Now: for any ordinality o, it makes sense to consider the set $I(o)$ of ordinalities $\{o' \mid o' < o\}$: indeed, these are well-orderings on a set of cardinality at most the cardinality of o, so there are at most $2^{|o| \times |o|}$ such well-orderings. Similarly, define

$$I[o] = \{o' \mid o' \leq o\}.$$

Corollary 11. *$I(o)$ is order-isomorphic to o itself.*

Proof: We shall define an order-isomorphism $f : I(o) \to o$. Namely, each $o' \in I(o)$ is given by an initial segment $I(y)$ of o, so define $f(o') = y$. That this is an order isomorphism is essentially a tautology which we leave for the reader to unwind.

1.4. **The Burali-Forti "Paradox".**

Do the ordinalities form a set? As we have so far managed to construct only countably many of them, it seems conceivable that they might. However, Burali-Forti famously observed that the assumption that there is a set of all ordinalities leads to a paradox. Namely, suppose $\mathbb{O}$ is a set whose elements are the ordinalities. Then by Corollary 10, $\mathbb{O}$ is itself well-ordered under our initial embedding relation $\leq$, so that the ordinality $o = o(\mathbb{O})$ would itself be a member of $\mathbb{O}$.

This is already curious: it is tantamount to saying that $\mathbb{O}$ is an element of itself, but notice that we are not necessarily committed to this: $(\mathbb{O}, \leq)$ is order isomorphic to one of its members, but maybe it is not *the same* set. (Anyway, is $o \in o$ paradoxical, or just strange?) Thankfully the paradox does not depend upon these ontological questions, but is rather the following: if $o \in \mathbb{O}$, then consider the initial segment $I(o)$ of $\mathbb{O}$: we have $\mathbb{O} \cong o \cong I(o)$, but this means that $\mathbb{O}$ is order-isomorphic to one of its initial segments, in contradiction to the Ordinal Trichotomy Theorem (Theorem 8).

Just as the proof of Cantor's *paradox* (i.e., that the cardinalities do not form a set) can be immediately adapted to yield a profound and useful *theorem* – if S is a set, there is no surjection $S \to 2^S$, so that $2^{|S|} > |S|$ – in turn the proof of the Burali-Forti paradox immediately gives the following result, which we have so far been unable to establish:

Theorem 12. *(Burali-Forti's Theorem) For any cardinal κ, the set $\mathcal{O}_\kappa$ of ordinalities o with $|o| \leq \kappa$ has cardinality greater than κ.*

Proof: Indeed, $\mathcal{O}_\kappa$ is, like any set of ordinalities, well-ordered under our relation $\leq$, so if it had cardinality at most κ it would contain its own ordinal isomorphism class o as a member and hence be isomorphic to its initial segment $I(o)$ as above.

So in particular there are uncountable ordinalities. There is therefore a *least* uncountable ordinality, traditionally denoted ω_1. This least uncountable ordinality is a truly remarkable mathematical object: mere contemplation of it is fascinating and a little dizzying. For instance, the minimality property implies that all of its initial segments are countable, so it is not only very large as a set, but it is extremely difficult to traverse: for any point $x \in \omega_1$, the set of elements less than x is countable whereas the set of elements greater than x is uncountable! (This makes Zeno's Paradox look like kid stuff.) In particular it has no largest element so is a limit ordinal.[4]

On the other hand its successor ω_1^+ is also of interest.

Exercise 1.4.1 (Order topology): Let S be a totally ordered set. We endow S with the **order topology**, which is the topology generated by by infinite rays of the form

$$(a, \infty) = \{s \in S \mid a < s\}$$

and

$$(-\infty, b) = \{s \in S \mid s < b\}.$$

[4]In fact this only begins to express ω_1's "inaccessibility from the left"; the correct concept, that of **cofinality**, will be discussed later.

Equivalently, the open intervals $(a,b) = (a,\infty) \cap (-\infty,b)$ together with the above rays and $X = (-\infty,\infty)$[5] form a basis for the topology. A topological space which arises (up to homeomorphism, of course) from this construction is called a **linearly ordered space**.
a) Show that the order topology on an ordinal o is discrete iff $o \leq \omega$. What is the order topology on $\omega + 1$? On 2ω?
b) Show that order topologies are Hausdorff.
c) Show that an ordinality is compact iff it is a successor ordinality. In particular $I[o]$ is the one-point compactification of $I(o) \cong o$; deduce that the order topology on an ordinality is Tychonoff.
d)* Show that, in fact, any linearly ordered space is normal, and moreover all subspaces are normal.
e) A subset Y of a linearly ordered set X can be endowed with two topologies: the subspace topology, and the order topology for the ordering on X restricted to Y. Show that the subspace topology is always finer than the order topology; by contemplating $X = \mathbb{R}$, $Y = \{-1\} \cup \{\frac{1}{n}\}_{n \in \mathbb{Z}^+}$ show that the two topologies need not coincide.
f) Show that it may happen that a subspace of a linearly ordered space need not be a linearly ordered space (i.e., there may be *no* ordering inducing the subspace topology). Suggestion: take $X = \mathbb{R}$, $Y = \{-1\} \cup (0,1)$. One therefore has the notion of a **generalized order space**, which is a space homeomorphic to a subspace of a linearly ordered space. Show that no real manifold of dimension greater than one is a generalized order space.
g) Let X be a well-ordered set and Y a nonempty subset. Show that the embedding $Y \to X$ may be viewed as a *net* on X, indexed by the (nonempty well-ordered, hence directed) set Y. Show that for any ordinality o the net $I(o)$ in $I[o]$ converges to o.

Exercise 1.4.2: Let $\mathcal{F}$ be a set of ordinalities. As we have seen, $\mathcal{F}$ is well-ordered under our initial embedding relation $<$ so gives rise to an ordinality $o(\mathcal{F})$. In fact there is another way to attach an ordinality to $\mathcal{F}$.
a) Show that there is a least ordinality s such that $\alpha \leq s$ for all $\alpha \in \mathcal{F}$. (Write $\alpha = o(X_\alpha)$, apply the Burali-Forti theorem to $|2^{\coprod_{\alpha \in \mathcal{F}} X_\alpha}|$, and use Exercise 1.3.4.) We call this s the **ordinal supremum** of the ordinalities in $\mathcal{F}$.
b) Show that an ordinality is a limit ordinality iff it is the supremum of all smaller ordinalities.
c) Recall that a subset T of a partially ordered set S is **cofinal** if for all $s \in S$ there exists $t \in T$ such that $s \leq t$. Let α be a limit ordinality, and $\mathcal{F}$ a subset of $I(\alpha)$. Show that $\mathcal{F}$ is cofinal iff $\alpha = \sup \mathcal{F}$.
d) For any ordinality α, we define the **cofinality** $\mathrm{cf}(\alpha)$ to be the minimal ordinality of a cofinal subset $\mathcal{F}$ of $I(\alpha)$. E.g., an ordinality is a successor ordinality iff it has cofinality 1. Show that $\mathrm{cf}(\omega) = \omega$ and $\mathrm{cf}(\omega_1) = \mathrm{cf}(\omega_1)$. What is $\mathrm{cf}(\omega^2)$?
e*) An ordinality is said to be **regular** if it is equal to its own cofinality. Show that for every cardinality κ, there exists a regular ordinality o with $|o| > \kappa$.
g) (For D. Lorenzini) For a cardinality κ, let o be a regular ordinality with $|o| > \kappa$.

[5]This calculus-style interval notation is horrible when S has a maximal or minimal element, since it – quite incorrectly! – seems to indicate that these elements "$\pm\infty$" should be excluded. We will not use the notation enough to have a chance to get tripped up, but beware.

Show that any linearly ordered subset of cardinality at most κ has an upper bound in o, but $I(\kappa)$ does not have a maximal element.[6]

1.5. **Von Neumann ordinals.**

Here we weish to report on an idea of von Neumann, which uses the relation $I(o) \cong o$ to define a canonical well-ordered set with any given ordinality. The construction is often informally defined as follows: "we inductively define o to be the set of all ordinals less than o." Unfortunately this definition is circular, and not for reasons relating to the induction process: step back and see that it is circular in the most obvious sense of using the quantity it purports to define!

However, it is quite corrigible: rather than building ordinals out of nothing, we consider the construction as taking as input a well-ordered set S and returning an order-isomorphic well-ordered set $vo(S)$, the **von Neumann ordinal** of S. The only property that we wish it to have is the following: if S and T are order-isomorphic sets, we want $vo(S)$ and $vo(T)$ to be not just order-isomorphic but *equal*. Let us be a bit formal and write down some axioms:

(VN1) For all well-ordered sets S, we have $vo(S) \cong S$.
(VN2) For well-ordered S and T, $S \cong T \implies vo(S) = vo(T)$.

Consider the following two additional axioms:

(VN3) $vo(\emptyset) = \emptyset$.
(VN4) For $S \neq \emptyset$, $vo(S) = \{vo(S') \mid S' < S\}$.

The third axiom is more than reasonable: it is forced upon us, by the fact that there is a unique empty well-ordered set. The fourth axiom is just expressing the order-isomorphism $I(o) \cong o$ in terms of von Neumann ordinals. Now the point is that these axioms determine all the von Neumann ordinals:

Theorem 13. *(von Neumann) There is a unique correspondence $S \mapsto vo(S)$ satisfying (VN1) and (VN2).*

Before proving this theorem, let's play around with the axioms by discussing their consequences for finite ordinals. We know that $vo(\emptyset) = \emptyset = [0]$. What is $vo([1])$? Well, it is supposed to be the set of von Neumann ordinals strictly less than it. There is in all of creation exactly one well-ordered set which is strictly less than $[1]$: it is $\emptyset$. So the axioms imply

$$vo([1]) = \{\emptyset\}.$$

How about $vo([2])$? The axioms easily yield:

$$vo([2]) = \{vo[0], vo[1]\} = \{\emptyset, \{\emptyset\}\}.$$

Similarly, for any finite number n, the axioms give:

$$v0([n]) = \{vo[0], vo[1], \ldots, vo[n-1]\},$$

[6]This shows that one must allow chains of arbitrary cardinalities, and not simply ascending sequences, in order for Zorn's Lemma to hold.

or in other words,

$$vo([n]) = \{vo[n-1], \{vo[n-1]\}\}.$$

More interestingly, the axioms tell us that the von Neumann ordinal ω is precisely the set of all the von Neumann numbers attached to the natural numbers. And we can track this construction "by hand" up through the von Neumann ordinals of 2ω, ω^2, ω^ω and so forth. But how do we know the construction works (i.e., gives a unique answer) for every ordinality?

The answer is simple: by induction. We have seen that the axioms imply that at least for sufficiently small ordinalities there is a unique assignment $S \mapsto vo(S)$. If the construction does not always work, there will be a smallest ordinality o for which it fails. But this cannot be, since it is clear how to define $vo(o)$ given definitions of all von Neumann ordinals of ordinalities less than o: indeed, (VN4) tells us exactly how to do this.

This construction is an instance of **transfinite induction**. This is the extension to general well-ordered sets of the principle of complete induction for the natural numbers: if S is a well-ordered set and T is a subset which is (i) nonempty and (ii) for all $sinS$, if the order ideal $I(s)$ is contained in T, then s is in T; then T must in fact be all of S. We trust the proof is clear.

Note that transfinite induction generalizes the principle of *complete* induction, not the principle of mathematical induction which says that if 0 is in S and $n \in S \implies n+1 \in S$, then $S = \mathbb{N}$. This principle is not valid for any ordinality larger than ω, since indeed ω is (canonically) an initial segment of every larger ordinality and the usual axioms of induction are satisfies for ω itself. All this is to say that in most applications of transfinite induction one must distinguish between the case of successor ordinals and the case of limit ordinals. For example:

Exercise 1.5.1: Show that for any well-ordered set S, $vo(S^+) = \{vo(S), \{vo(S)\}\}$.

We should remark that this is not a foundationalist treatment of von Neumann ordinals. It would also be possible to define a von Neumann ordinal as a certain type of set, using the following exercise.

Exercise 1.5.2: Show that a set S is a von Neumann ordinal iff:
(i) if $x \in S$ implies $x \subset S$;
(ii) the relation $\subset$ is a well-ordering on elements of S.

For the rest of these notes we will drop the term "ordinality" in favor of "ordinal." The reader who wants an ordinal to be something in particular can thus take it to be a von Neumann ordinal. This convention has to my knowledge no real mathematical advantage, but it has some very convenient notational consequences, as for instance the following definition of "cardinal."

1.6. **A definition of cardinals.** Here we allow ourselves the following result, which we will discuss in more detail later on.

Theorem 14. *(Well-ordering theorem) Assuming the Axiom of Choice, every set S can be well-ordered.*

We can use this theorem ("theorem"?) to reduce the theory of cardinalities to a special case of the theory of ordinalities, and thus, we can give a concrete definition of cardinal numbers in terms of Von Neumann's ordinal numbers.

Namely, for any set S, we define its cardinal $|S|$ to be the smallest von Neumann ordinal o such that o is equivalent to (i.e., in bijection with) S.

In particular, we find that the finite cardinals and the finite ordinals are the same: we have changed our standard n element set from $[1, n]$ to the von Neumann ordinal n, so for instance $3 = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}$. On purely mathematical grounds, this is not very exciting. However, if you like, we can replace our previous attitude to what the set $[n] = \{1, \ldots, n\}$ "really is" (which was, essentially, "Why are you bothering me with such silly questions?") by saying that, in case anyone asks (we may still hope that they do not ask), we identify the non-negative integer n with its von Neumann ordinal. Again, this is not to say that we have discovered what 3 really is. Rather, we noticed that a set with three elements exists in the context of **pure set theory**, i.e., we do not have to know that there exist 3 objects in some box somewhere that we are basing our definition of 3 on (like the definition of a meter used to be based upon an actual meter stick kept by the Bureau of Standards). In truth 3 is not a very problematic number, but consider instead $n = 10^{10^{10^{10}}}$; the fact that n is (perhaps) greater than the number of distinct particles in the universe is, in our account, no obstacle to the existence of sets with n elements.

Let's not overstate the significance of this for finite sets: with anything like a mainstream opinion on mathematical objects[7] this is completely obvious: we could also have defined 0 as $\emptyset$ and n as $\{n - 1\}$, or in infinitely many other ways. It becomes more interesting for infinite sets, though.

That is, we can construct a theory of sets without *individuals* – in which we never have to say what we mean by an "object" as an element of a set, because the only elements of a set are other sets, which ultimately, when broken up enough (but possibly infinitely many) times, are lots and lots of braces around the empty set. This is nice to know, most of all because it means that in practice we don't have to worry one bit about what the elements of are sets are: we can take them to be whatever we want, because each set is equivalent (bijective) to a *pure set*. If you would like (as I would) to take a primarily Bourbakistic view of mathematical structure – i.e., that the component parts of any mathematical object are of no importance whatsoever, and that mathematical objects matter only as they relate to each other – then this is very comforting.

Coming back to the mathematics, we see then that any set of cardinals is in particular a set of ordinals, and the notion of $<$ on cardinals induced in this way is the same as the one we defined before. That is, if α and β are von Neumann cardinals, then $\alpha < \beta$ holds in the sense of ordinals iff there exists an injection from α to β but not an injection from β to α.

[7]The only contemporary mathematician I know who would have problems with this is Doron Zeilberger.

Exercise 1.6.1: Convince yourself that this is true.

Thus we have now, at last, proved the Second Fundamental Theorem of Set Theory, modulo our discussion of Theorem 14.

2. The Axiom of Choice and some of its equivalents

2.1. Introducing the Axiom of Choice.

Now we come clean. Many of the results of Chapter II rely on the following "fact":

Fact 15. *(Axiom of Choice (AC)): For any nonempty family I of nonempty sets S_i, the product $\prod_{i \in I} S_i$ is nonempty.*

Remark: In other words, any product of nonzero cardinalities is itself nonzero. This is the version of the axiom of choice favored by Bertrand Russell, who called it the "multiplicative axiom." Aesthetically speaking, I like it as well, because it seems so simple and self-evident.

Exercise 2.1: Show that if (AC) holds for all families of pairwise disjoint sets S_i, it holds for all nonempty families of nonempty sets.

However, in applications it is often more convenient to use the following reformulation of (AC) which spells out the connection with "choice".

(AC$'$): If S is a set and $I = \{S_i\}$ is a nonempty family of nonempty subsets of S, then there exists a **choice function**, i.e., a function $f : I \to S$ such that for all $i \in I$, $f(S_i) \in S_i$.

Let us verify the equivalence of (AC) and (AC$'$).

(AC) $\implies$ (AC$'$): By (AC), $\mathcal{S} = \prod_{i \in I} S_i$ is nonempty, and an element f of $\mathcal{S}$ is precisely an assignment to each $i \in I$ of an element $f(i) \in S_i \subset S$. Thus f determines a choice function $f : I \to S$.

(AC$'$) $\implies$ (AC): Let $I = \{S_i\}$ be a nonempty family of nonempty sets. Put $S = \bigcup_{i \in I} S_i$. Let $f : I \to S$ be a choice function: for all $i \in I$, $f(S_i) \in S_i$. Thus $\{f(i)\}_{i \in I} \in \prod_{i \in I} S_i$.

The issue here is that if I is infinite we are making infinitely many choices – possibly with no coherence or defining rule to them – so that to give a choice function f is in general to give an infinite amount of information. Have any of us in our daily lives ever made infinitely many independent choices? Probably not. So the worry that making such a collection of choices is not possible is not absurd and should be taken with some seriousness.

Thus the nomenclature *Axiom* of Choice: we are, in fact, asserting some feeling about how infinite sets behave, i.e., we are doing exactly the sort of thing we had earlier averred to try to avoid. However, in favor of assuming AC, we can say: (i) it is a fairly basic and reasonable axiom – if we accept it we do not, e.g., feel the need to justify it in terms of something simpler; and (ii) we are committed to it, because most of the results we presented in Chapter II would not be true without

it, nor would a great deal of the results of mainstream mathematics.

Every student of mathematics should be aware of some of the "facts" that are equivalent to AC. The most important two are as follows:

Fact 16. *(Zorn's Lemma) Let S be a partially ordered set. Suppose that every chain C – i.e., a totally ordered subset of S – has an upper bound in S. Then S has a maximal element.*

Theorem 17. *The axiom of choice (AC), Zorn's Lemma (ZL), and the Well-Ordering Theorem (WOT) are all equivalent to each other.*

Remark: The fact that we are asserting the logical equivalence of an axiom, a lemma and a theorem is an amusing historical accident: according to the theorem they are all on the same logical footing.

WOT $\implies$ AC: It is enough to show WOT $\implies$ AC$'$, which is easy: let $\{S_i\}_{i \in I}$ be a nonempty family of nonempty subsets of a set S. Well-order S. Then we may define a choice function $f : I \to S$ by mapping i to the least element of S_i.

AC $\implies$ ZL: Strangely enough, this proof will use transfinite induction (so that one might initially think WOT would be involved, but this is absolutely not the case). Namely, suppose that S is a poset in which each chain C contains an upper bound, but there is no maximal element. Then we can define, for every ordinal o, a subset $C_0 \subset S$ order-isomorphic to o, in such a way that if $o' < o$, $C_{o'} \subset C_o$. Indeed we define $C_\emptyset = \emptyset$, of course. Assume that for all $o' < o$ we have defined $C_{o'}$. If o is a limit ordinal then we define $C_o := \bigcup_{o'<o} C_{o'}$. Then necessarily C_0 is order-isomorphic to o: that's how limit ordinals work. If $o = o' + 1$, then we have $C_{o'}$ which is assumed not to be maximal, so we choose an element x of $S \setminus C_{o'}$ and define $x_o := x$. Thus we have inside of S well-ordered sets of all possible order-isomorphism types. This is clearly absurd: the collection $o(|S|)$ of ordinals of cardinality $|S|$ is an ordinal of cardinality greater than the cardinality of S, and $o(|S|) \hookrightarrow S$ is impossible.

But where did we use AC? Well, we definitely made some choices, one for each non-successor ordinal. To really nail things down we should cast our choices in the framework of a choice function. Suppose we choose, for each well-ordered subset W of X, an element $x_W \in X \setminus W$ which is an upper bound for W. (This is easily phrased in terms of a choice function.) We might worry for a second that in the above construction there was some compatibility condition imposed on our choices, but this is not in fact the case: at stage o, any upper bound x for C_o in $S \setminus C_o$ will do to give us $C_{o+1} := C_o \cup \{x\}$. This completes the proof.

Remark: Note that we showed something (apparently) slightly stronger: namely, that if every well-ordered subset of a poset S has an upper bound in S, then S has a maximal element. This is mildly interesting but apparently useless in practice.

ZL $\implies$ WOT: Let X be a non-empty set, and let $\mathcal{A}$ be the collection of pairs $(A, \leq)$ where $A \subset X$ and $\leq$ is a well-ordering on A. We define a relation $<$ on $\mathcal{A}$: $x < y$ iff x is equal to an initial segment of y. It is immediate that $<$ is a strict partial ordering on $\mathcal{A}$. Now for each chain $C \subset \mathcal{A}$, we can define x_C to be the union

of the elements of C, with the induced relation. x_C is itself well-ordered with the induced relation: indeed, suppose Y is a nonempty subset of x_C which is not well-ordered. Then Y contains an infinite descending chain $p_1 > p_2 > \ldots > p_n > \ldots$. But taking an element $y \in C$ such that $p_1 \in y$, this chain lives entirely inside y (since otherwise $p_n \in y'$ for $y' > y$ and then y is an initial segment of y', so $p_n \in y'$, $p_n < p_1$ implies $p_n \in y$), a contradiction.

Therefore applying Zorn's Lemma we are entitled to a maximal element $(M, \leq_M)$ of $\mathcal{A}$. It remains to see that $M = X$. If not, take $x \in X \setminus M$; adjoining x to $(M, \leq_M)$ as the maximum element we get a strictly larger well-ordering, a contradiction.

Remark: In the proof of AC $\implies$ ZL we made good advantage of our theory of ordinal arithmetic. It is possible to prove this implication (or even the direct implication AC $\implies$ ZL) directly, but this essentially requires proving some of our lemmata on well-ordered sets on the fly.

2.2. **Some equivalents and consequences of the Axiom of Choice.** Although disbelieving AC is a tenable position, mainstream mathematics makes this position slightly unpleasant, because Zorn's Lemma is used to prove many quite basic results. One can ask which of these uses are "essential." The strongest possible case is if the result we prove using ZL can itself be shown to imply ZL or AC. Here are some samples of these results:

Fact 18. *For any infinite set A, $|A| = |A \times A|$.*

Fact 19. *For sets A and B, there is an injection $A \hookrightarrow B$ or an injection $B \hookrightarrow A$.*

Fact 20. *Every surjective map of sets has a section.*

Fact 21. *For any field k, every k-vector space V has a basis.*

Fact 22. *Every proper ideal in a commutative ring is contained in a maximal proper ideal.*

Fact 23. *The product of any number of compact spaces is itself compact.*

Even more commonly one finds that one can make a proof work using Zorn's Lemma but it is not clear how to make it work without it. In other words, many statements seem to require AC even if they are not equivalent to it. As a simple example, try to give an explicit well-ordering of $\mathbb{R}$. Did you succeed? In a precise formal sense this is impossible. But it is intuitively clear (and also true!) that being able to well-order a set S of any given infinite cardinality is not going to tell us that we can well-order sets of all cardinalities (and in particular, how to well-order 2^S), so the existence of a well-ordering of the continuum is not equivalent to AC.

Formally, speaking one says that a statement *requires* AC if one cannot prove that statement in the Zermelo-Fraenkel axiomation of set theory (ZF) which excludes AC. (The Zermelo-Fraenkel axiomatization of set theory including the axiom of choice is abbreviated ZFC; ZFC is essentially the "standard model" for sets.) If on the other hand a statement requires AC in this sense but one cannot deduce AC from ZF and this statement, we will say that the statement *merely requires* AC. There are lots of statements that merely require AC:[8]

[8]This list was compiled with the help of the Wikipedia page on the Axiom of Choice.

Theorem 24. *The following facts merely require AC:*
a) The countable union of countable sets is countable.
b) An infinite set is Dedekind infinite.
c) There exists a non(-Lebesgue-)measurable subset of $\mathbb{R}$.
d) The Banach-Tarski paradox.
e) Every field has an algebraic closure.
f) Every field extension has a relative transcendence basis.
g) Every Boolean algebra contains a prime ideal (BPIT).
h) Every Boolean algebra is isomorphic to a Boolean algebra of sets (Stone representation theorem).
i) Every subgroup of a free group is free.
j) The Hahn-Banach theorem (on extension of linear functionals), the open mapping theorem, the closed graph theorem, the Banach-Alaoglu theorem.
k) The Baire category theorem.
l) The existence of a Stone-Cech compactification of every completely regular space.

Needless to say the web of implications among all these important theorems is a much more complicated picture; for instance, it turns out that the BPIT is an interesting intermediate point (e.g. Tychonoff's theorem for Hausdorff spaces is equivalent to BPIT). Much contemporary mathematics is involved in working out the various dependencies.

In summary, if your beliefs about sets are the same as the standard ones except that you do not admit any form of AC (again, exactly what this means is something that we have not spelled out), then you will find that there is an amazing array of mathematical theorems that you will not be able to prove. If instead of being entirely agnostic about AC you believe a strong enough condemnation of it (i.e., you believe one of the many axioms which is independent of ZF and contradicts AC), then you will be able to prove false some of the results in standard mathematics.

Notable here is the existence of a relatively mild denial of AC which allows most familiar analytic results to remain true but implies that every subset $\mathbb{R}$ is Lebesgue measurable. There are analysts who advocate the use of this axiom, noting that it simplifies the theory: in proving Fubini-type theorems on integrals over product measure spaces, one has to verify that the measurability of the given functions implies the measurability of certain auxiliary functions, a verification which is tedious and unpleasant (and nontrivial). Like most people who lost an hour of their lives somewhere in their early 20's sitting through the proof of Fubini's theorem, I have some sympathy for this position.

What should your attitude be towards AC? You will, of course, have to decide for yourself, although again a sincere agnosticism or disbelief could lead you to state and prove different theorems. My own take on AC (which is rather standard to the extent of coming uncomfortably close to parroting the corresponding paragraph in Kaplansky's book, but it is nevertheless how I feel) is a sort of middle-ground: when you use a result which requires (merely or otherwise) AC, you should acknowledge this – not necessarily with a large fanfare; if you used Zorn's Lemma somewhere it is plausible that your result requires AC, whether or not the set theorists have proven its independence from ZF – and take mental note: it means

that there is some obstacle to making your result explicit in full generality. Now if you are working in some fairly concrete area of mathematics (like number theory), perhaps there are some interesting special cases of your general result which you might be able to make explicit with a different and more perspicuous argument. In general when you prove a theorem asserting the existence of an object, it is good to know whether or not you can actually *construct*, in some algorithmic sense, such an object. The advent of computers has done wonders for constructive mathematics, a philosophy which only 50 years ago looked rather eccentric. You've proven the existence of a genus one curve with certain properties, have you? Well, can you program a computer to find one? (I'm afraid I can't, usually.)

One thing that the majority of working mathematicians would probably agree with is that while uncountable sets exist in the sense of convenience and noncontradiction, they do not exist in the same *visceral* sense of things that you can get your computer to spit out. Twenty-first century mathematics is at the same time more abstract *and* more concrete than mathematics one hundred years before.