Mind Your ∀'s and ∃'s

Stephen M. Walk

Stephen Walk (smwalk@stcloudstate.edu) earned both his Bachelor's and Master's degrees from the University of Northern Iowa, the latter under the direction of Michael Millar. His Ph.D., supervised by Peter Cholak, is from the University of Notre Dame. He has taught at St. Cloud State University in Minnesota since 1999, and his mathematical interests include computability theory and many-valued logic.

I'd like to introduce you to a mathematical structure that looks like a vector space, walks like a vector space, and quacks like a vector space, but most definitely is *not* a vector space. I should probably come up with a special name for this structure (*wector space*? *deflector space*? *Hannibal Lecter space*?), but for now I'll just call it $\mathcal{W}$.

The underlying set of $\mathcal{W}$ consists of two "copies" of $\mathbb{R}$, plus one extra element denoted $\perp$. The two copies of $\mathbb{R}$ are conceptualized as follows:

- the set of *pro-reals*, $\{\overline{a} : a \in \mathbb{R}\}$, and
- the set of *anti-reals*, $\{\underline{a} : a \in \mathbb{R}\}$.

Scalar multiplication for $\mathcal{W}$ (with $\mathbb{R}$ the set of scalars) is easily defined: If c is any scalar, then $c\overline{a}$ is defined to be $\overline{ca}$ and $c\underline{a}$ is defined to be $\underline{ca}$. As for addition, the sum of two pro-reals is given by the natural formula $\overline{a} + \overline{b} = \overline{a+b}$ and the sum of two anti-reals by $\underline{a} + \underline{b} = \underline{a+b}$.

What about the sum of a pro-real and an anti-real? The terminology is meant to suggest a dichotomy like that of matter and antimatter, where contact results in annihilation. The same idea motivates the addition of a pro-real to an anti-real in either order; the sum will be represented by the symbol "$\perp$," which we call the *obliterator* (since "annihilator" is already taken as a mathematical term).

We complete the definitions of addition and scalar multiplication to include the obliterator: For any $\mathbf{u}$ in $\mathcal{W}$ and any scalar c,

$$\perp + \mathbf{u} = \mathbf{u} + \perp = \perp \quad \text{and} \quad c\perp = \perp.$$

Now, $\mathcal{W}$ is not a vector space, but it is not far off. It looks like a vector space: Surely, as most undergraduate linear algebra students can attest, any structure that has both an addition and a scalar multiplication operation and is closed with respect to both of them is probably a vector space. It walks like a vector space: Of the ten usual vector space axioms, $\mathcal{W}$ is easily seen to satisfy eight. It quacks like a vector space—and this is the key observation here: As we'll see, if we are not careful about how we state or understand the other two vector space axioms, $\mathcal{W}$ *seems* to satisfy *all ten*, and thus a careless treatment of the vector space axioms can lead us to believe that certain structures are vector spaces when really they aren't. So $\mathcal{W}$ is a concrete example of why we (and our students) *do* need to be careful in our statement and understanding of axioms and, especially, in our use of quantifiers ("for all $\mathbf{u}$" and "there exists a $\mathbf{v}$").

The switch

I recently taught Elementary Linear Algebra for the first time, fortunate to have received a tip from a colleague and to have adopted an exceptionally good text. It was something of a surprise toward the middle of the semester when I realized that, of all things, the *vector space axioms* had been stated incorrectly in the text.

There is no need to review the vector space axioms here, except to note that they are usually presented as a list of ten, or a list of eight if closure under each operation is not listed explicitly as an axiom, and that most of them involve universal ("for all"/"for every"/"for any") quantification over a set V of vectors and a set of scalars, which we will take to be $\mathbb{R}$. I expected to see them stated as, for example, "For all $\mathbf{u}$ and $\mathbf{v}$ in V, $\mathbf{u}+\mathbf{v}=\mathbf{v}+\mathbf{u}$"; and "For all $\mathbf{u}$ in V and all scalars a and b, $a(b\mathbf{u})=(ab)\mathbf{u}$." This was what nature and the author of our text intended.

Instead, in what had almost certainly been an instance of well-meaning but misguided editorial meddling, someone had apparently decided that it was distracting to the eye to have all those *for all*s at the heads of *all* of the axioms, so they had been stripped away and replaced by what amounted to a blanket universal quantifier at the beginning of the axiom list. It was something to the effect that "The following statements have to be true for all elements $\mathbf{u}$, $\mathbf{v}$, and $\mathbf{w}$ in V and all scalars a, b, and c."

There would be absolutely no problem, pedagogically or mathematically, with this migration to the beginning if *all* of the axioms were meant to be universally quantified. It is a matter of basic logic that "for all" distributes over "and": Instead of, for instance,

$$\forall\mathbf{u}\, P_1(\mathbf{u}) \wedge \forall\mathbf{u}\, P_2(\mathbf{u}) \wedge \cdots \wedge \forall\mathbf{u}\, P_{10}(\mathbf{u}),$$

we can pull all the universal quantifiers to the front of the sentence to get

$$\forall\mathbf{u}\big[P_1(\mathbf{u}) \wedge P_2(\mathbf{u}) \wedge \cdots \wedge P_{10}(\mathbf{u})\big],$$

a sentence that is precisely equivalent to the one we started with—even if there is more than one variable $\mathbf{u}$ being quantified, and even if we are working in what logicians refer to as *many-sorted* logic, with different "sorts" of variables representing different sets such as vectors and scalars, which is really necessary for a consideration of the vector space axioms.

But there was a problem with that text's exposition, pedagogically and especially mathematically, because one of the vector space axioms is *not* meant to be universally quantified. This is the "identity axiom"—that is,

There is a designated element $\mathbf{0}$ such that for every $\mathbf{u}$ in V, $\mathbf{u}+\mathbf{0}=\mathbf{u}$,

which I'll denote here as Axiom `Zero`. Note that there are two quantifiers in Axiom `Zero`—a "there exists" and a "for every." The name $\mathbf{0}$ is a convenience; naming the element in the axiom is logically unnecessary, and the axiom is properly rendered in symbols as follows:

$$\exists\mathbf{x}\forall\mathbf{u}\,[\mathbf{u}+\mathbf{x}=\mathbf{u}].$$

In the text I was using, by contrast, the identity axiom was written as follows:

There is an element $\mathbf{0}$ such that $\mathbf{u}+\mathbf{0}=\mathbf{u}$,

and the reader has no doubt anticipated the problem before I state it: Since $\mathbf{u}$ had been quantified under the blanket universal quantifier at the beginning of the list ("The following have to be true for all $\mathbf{u}\ldots$"), this axiom was really being presented as

For every $\mathbf{u}$ there is an element $\mathbf{0}$ such that $\mathbf{u} + \mathbf{0} = \mathbf{u}$,

which when rendered in symbols becomes

$$\forall \mathbf{u}\, \exists \mathbf{x}\, [\mathbf{u} + \mathbf{x} = \mathbf{u}],$$

which is clearly different—*vastly* different—from what was intended.

I write "vastly" because it seems an obvious and vast difference to any experienced mathematician. To an Elementary Linear Algebra student, however, the distinction is not necessarily clear. The problem, I told the class, was that whereas the correct Axiom `Zero` posited the existence of *some* element that worked as a zero element for *every* $\mathbf{u}$ (here a *zero element for* $\mathbf{u}$ is, as you'd expect, an $\mathbf{x}$ such that $\mathbf{u} + \mathbf{x} = \mathbf{u}$), the altered version of the axiom said only that, for every $\mathbf{u}$, there was some zero element for it. In the altered version, which I'll call Axiom `Zero`*, there was no claim that the *same* zero element worked for *all* $\mathbf{u}$! Perhaps distinct $\mathbf{u}$ and $\mathbf{v}$ could have distinct zero elements $\mathbf{0}_\mathbf{u}$ and $\mathbf{0}_\mathbf{v}$, and $\mathbf{u} + \mathbf{0}_\mathbf{v}$ might *not* be equal to $\mathbf{u}$, nor $\mathbf{v} + \mathbf{0}_\mathbf{u}$ equal to $\mathbf{v}$.

I trotted out all my examples, from years of teaching logic in our Discrete Mathematics course, of how "$\forall x \exists y$" is different from "$\exists y \forall x$". My favorite example at the time (this was before Paul Stockmeyer's discussion [**4**] of Abraham Lincoln's famous quote about fooling all of the people some of the time) was from the Arthur Andersen criminal case in the summer of 2002: Judge Melinda Harmon had instructed the jurors that a conviction of the corporation did not have to be based on the standard "There is an executive y such that for every juror x, x thinks y is guilty" but, instead, on the clearly weaker standard "For every juror x there is an executive y such that x thinks y is guilty". Some students understood. Some would come to my office later. But the fundamental barrier to accepting the distinction in the current context, even among those who could see the logic, was the very idea that there could be different zero elements for different "vectors." "How could you have more than one zero in a vector space?" asked one student, and the rest of the class agreed.

"Well, that's just it," I said. "*If* the mathematical structure V *is* a vector space, then *of course* there is only one zero element. The problem is that, conceivably, a mathematical structure could satisfy all ten of these axioms, including the faulty one, and not really be a vector space."

One of my brighter students then put me on the spot, which is why brighter students are a joy to have around, by asking, "But can that *really* happen if the other nine axioms are satisfied?"

And I was flummoxed. Because, *yes*, Axiom `Zero`* as presented, with the universal and the existential quantifiers switched, is not, in a vacuum, equivalent to the correct axiom. But the other nine axioms impose a great deal of algebraic structure on whatever object we're studying. Perhaps any mathematical structure that satisfied those nine, plus the imperfect Axiom `Zero`*, *must* satisfy the correct Axiom `Zero` and must therefore really be a vector space. Perhaps the "new" list of axioms really was an equivalent set of vector space axioms and wouldn't let any impostors in, that is, wouldn't be satisfied by anything that wasn't really a vector space. There are different axiom systems for group theory [**2**, Theorem 1.6] and the theory of Boolean algebras [**3**, Section 5.4], for example; how did I know that the axioms as stated weren't equivalent to the usual vector space axioms? Until I knew whether this was the case, I couldn't be sure whether the distinction I'd pointed out was vitally important, somewhat important, or merely pedantic, much like the Mad Hatter's pointing out (correctly) the difference between "I mean what I say" and "I say what I mean".

Off the top of my head, I couldn't settle the matter. But in the safety of my office after class, I discovered $\mathcal{W}$.

Axiom Inv

When you follow the line of inquiry I was following—trying to find a structure that satisfies nine of the original axioms as well as Axiom Zero* but not the real Axiom Zero—you discover a couple of things immediately.

First, you realize that Axiom Zero* is not only weaker than the intended Axiom Zero, it's superfluous. The axioms specifying that $1\mathbf{u} = \mathbf{u}$ and $(c + d)\mathbf{u} = c\mathbf{u} + d\mathbf{u}$ (both properly universally quantified, of course) guarantee that for every $\mathbf{u}$, $0\mathbf{u}$ is a zero element for $\mathbf{u}$. Thus, Axiom Zero* didn't add anything, and I was really trying to decide whether Axiom Zero could be deduced from the other nine axioms. If it had been well known that the vector space axioms were independent, the inquiry would have ended here, but the vector space axioms are *not* independent (a different text [**1**, Section 5.1, Exercise 30] even has an exercise in which students are led through a demonstration of the "commutativity of addition" axiom from the others), and the independence or dependence of Axiom Zero with respect to the other nine isn't standard material for textbooks of linear algebra or logic. So I was left to my own devices.

Second, whether you replace Axiom Zero with Axiom Zero* or delete it altogether, it doesn't take long to discover that one of the other nine has been compromised. I'm referring to the axiom, which I'll refer to here as Axiom Inv, that says that

For every $\mathbf{u}$, there is an element $-\mathbf{u}$ such that $\mathbf{u} + (-\mathbf{u}) = \mathbf{0}$.

This, if we have only Axiom Zero* instead of Axiom Zero, is nonsense. As we are not assuming the existence of an element that acts as a zero element for everything, any use of the symbol "$\mathbf{0}$" in this axiom is ambiguous. So Axiom Inv must be rewritten before we really begin, and we will see that the way we rewrite Axiom Inv is of paramount importance.

One possible reformulation of Axiom Inv is the following, which I'll call Axiom Inv*:

For every $\mathbf{u}$, there is an element $-\mathbf{u}$ such that $\mathbf{u} + (-\mathbf{u}) = \mathbf{0}_\mathbf{u}$, where $\mathbf{0}_\mathbf{u}$ is *some* element such that $\mathbf{u} + \mathbf{0}_\mathbf{u} = \mathbf{u}$.

In other words, where Axiom Zero* states that for any $\mathbf{u}$ there is some $\mathbf{0}_\mathbf{u}$ that acts as a zero element for $\mathbf{u}$ (though perhaps not for any other element of the structure), Axiom Inv* goes on to say that $\mathbf{u}$ must have an inverse with respect to *at least one* of these zero elements, perhaps more. (We do not yet have any reason to believe a single $\mathbf{u}$ cannot have two or three or even infinitely many such zero elements.) Note that this Axiom Inv* does not add anything; just as Axiom Zero* is superfluous because of the presence of $0\mathbf{u}$ for every $\mathbf{u}$, similarly Axiom Inv* is superfluous because of the presence of $(-1)\mathbf{u}$. Nevertheless, Axiom Inv* is a reasonable reformulation of Axiom Inv.

Suppose, then, that we accept Axiom Zero* as written, interpret Axiom Inv in this reasonable manner as Axiom Inv*, and keep the other eight vector space axioms as they are. Is it possible for an impostor to satisfy all ten of these axioms but not be an actual vector space? Yes, of course; $\mathcal{W}$ fits the bill.

The impostor

It is easy to verify that $\mathcal{W}$ satisfies all the axioms in which we're interested. Satisfaction of the eight unchanged axioms is left as an exercise. For Axiom Zero*, let $\mathbf{u}$ be any element of $\mathcal{W}$. If $\mathbf{u}$ is a pro-real, then $\overline{0}$ is a zero element for $\mathbf{u}$, in fact the only zero element for $\mathbf{u}$. Similarly, if $\mathbf{u}$ is an anti-real, then $\underline{0}$ is the only zero element for $\mathbf{u}$. More interesting is the fact that if $\mathbf{u}$ is $\perp$, then *every* element of $\mathcal{W}$ is a zero element for $\mathbf{u}$, since $\perp$ plus anything is still $\perp$.

Thus, $\mathcal{W}$ satisfies Axiom Zero*. Before we continue, we note that $\mathcal{W}$ displays the following counterintuitive properties:

- Different elements of $\mathcal{W}$ can have different zero elements.
- A single element of $\mathcal{W}$ can have more than one zero element. In fact, a single element can have infinitely many zero elements. It can even have *all* of the elements of $\mathcal{W}$ serve as its zero elements.
- Most to the point, there is no *single* element of $\mathcal{W}$ that acts as a zero element for *every* element of $\mathcal{W}$, since pro-reals have only $\overline{0}$ and anti-reals have only $\underline{0}$. Thus, $\mathcal{W}$ is not a vector space.

Now let's take care of Axiom Inv*. Let $\mathbf{u}$ be any element of $\mathcal{W}$. If $\mathbf{u}$ is a pro-real $\overline{a}$, then $\overline{0}$ is a (*the*, as noted above) zero element for $\mathbf{u}$; and we can add $\overline{-a}$ to $\mathbf{u}$ to get $\overline{0}$. Similarly, if $\mathbf{u}$ is an anti-real $\underline{a}$, then $\underline{0}$ is a (*the*!!) zero element for $\mathbf{u}$, and we have $\mathbf{u} + \underline{-a} = \underline{0}$.

This leaves only the case $\mathbf{u} = \perp$. In this case, *every* element of $\mathcal{W}$ is a zero element for $\mathbf{u}$; in particular, $\perp$ is a zero element for $\mathbf{u}$, so to verify that Axiom Inv* holds we need only to find an inverse element "$-\mathbf{u}$" such that $\mathbf{u} + (-\mathbf{u}) = \perp$, that is, such that $\perp + (-\mathbf{u}) = \perp$. This is not much of a challenge: every element of $\mathcal{W}$ can play the role of $-\mathbf{u}$! With respect to this zero element, $\mathbf{u}$ has *infinitely many* inverses—another counterintuitive property. But Axiom Inv* says nothing about intuition; it specifies that there has to be *some* zero element with respect to which $\mathbf{u}$ has *some* inverse, and we've just seen that this is true no matter what $\mathbf{u}$ is.

It is worth noting that $\mathbf{u} = \perp$ does not have any inverses at all with respect to any of its other zero elements. This fact will turn out to be significant, as we'll see later. It also helps to crystallize the role of $\perp$: in some ways, $\perp$ is to addition in $\mathcal{W}$ what 0 is to multiplication in $\mathbb{R}$.

Variations: Instead of having just one obliterator, $\perp$, we can expand $\mathcal{W}$ to a structure $\mathcal{W}'$ in which all of the scalar multiples of $\perp$ are different. Thus we have, essentially, three copies of $\mathbb{R}$. The definitions of the operations are changed in the following manner:

- for any pro-real $\overline{a}$ and any anti-real $\underline{b}$, $\overline{a} + \underline{b} = (a + b)\perp$;
- for any pro-real $\overline{a}$ and any obliterator $c\perp$, $\overline{a} + c\perp = (a + c)\perp$, and similarly for any anti-real $\underline{a}$ and any obliterator $c\perp$; and
- for any scalars c and d, $c(d\perp) = (cd)\perp$.

Then $\mathcal{W}'$, like $\mathcal{W}$, satisfies the ten axioms in question and is still not a vector space. This structure is slightly more complicated than $\mathcal{W}$ but seems less pathological: Every scalar multiple of $\perp$ now has only three zero elements, $\overline{0}$, $(0\perp)$, and $\underline{0}$, rather than infinitely many.

We should note that there is no shortage of examples like $\mathcal{W}$ and $\mathcal{W}'$. For instance, we can create a structure like $\mathcal{W}'$ with any number of copies of $\mathbb{R}$, or even with one copy of $\mathbb{R}$ for every element of an arbitrary index set. Further, if $\{V_\alpha\}_{\alpha\in I}$ is *any* pairwise disjoint collection of real vector spaces, then the set

$$\left(\bigsqcup_{\alpha \in I} V_\alpha \right) \sqcup \{\perp\}$$

can be made into an example like $\mathcal{W}$. (The simplest possible model of the altered axioms is just $\mathbb{R} \sqcup \{\perp\}$.) The key in each case is that the sum of elements from different V_α's—or from different copies of $\mathbb{R}$—is a scalar multiple of $\perp$.

Uniqueness

It's more than a little disturbing to think that $\mathcal{W}$ contains more than one zero element. After all, when we introduce the vector space axioms, or when we introduce the idea of *identity* in any mathematical structure, the first thing we prove to our students is that there can be only *one* identity. The fragmenting of our notion of "$\mathbf{0}$" into a pluralistic idea of "different $\mathbf{0}_\mathbf{u}$'s for different $\mathbf{u}$'s" is alien at first; hence my students' struggles with the idea. It also renders nonsensical the ideas of linear independence, basis, and dimension that are fundamental to linear algebra. (For example, we can "span" $\mathcal{W}$ with a two-element set $\{\overline{a}, \underline{a}\}$, but such a set is not "linearly independent" because *any* linear combination of the elements is equal to a zero element.) On a deeper level, this fragmenting is antithetical to the foundation of anything that we call "algebra," that foundation being the ability to solve equations by undoing operations. If we know that $x + 3 = 7$, then we add -3 to both sides of the equation, appeal to associativity to get $x + (3 + (-3)) = 7 + (-3)$, and *take it for granted* that the "zero element" we get by adding 3 to -3 is something that we can add to x and leave x unchanged. In other words, we usually take for granted our ability to perform subtraction with these objects, and it is this ability that is supposed to be built into the group axioms and thus into the vector space axioms. Having different zero elements for different objects in the structure wreaks havoc on this ability.

Worse, though, is the idea that a *single* $\mathbf{u}$ could have *multiple* zero elements, as $\perp$ does in $\mathcal{W}$. It seems much more reasonable that, if $\mathbf{u}$ insists on having its own zero element, it should just pick *one* and stick with it, rather than having a host of them. Even *two* zero elements seem to be too many.

The question presents itself, then: What if we alter Axiom `Zero`* to prescribe uniqueness? That is, we will still not say anything about a single $\mathbf{0}$ that works for every $\mathbf{u}$, but we will strengthen Axiom `Zero`* to, say, Axiom `Zero`$^{\text{unique}}$ as follows:

For every $\mathbf{u}$ there is one *and only one* $\mathbf{0}_\mathbf{u}$ such that $\mathbf{u} + \mathbf{0}_\mathbf{u} = \mathbf{u}$.

Now the quantifier ordering is just as bad as before, but the existence quantifier demands uniqueness. (This is really equivalent to sticking in a third quantifier—a universal "$\forall$" one—after the "there exists," but we won't go into that.) Note that Axiom `Zero`$^{\text{unique}}$ is not superfluous; other axioms tell us that $0\mathbf{u}$ is a zero element for $\mathbf{u}$, but they do not imply that it is the *only* one. Also, given Axiom `Zero`$^{\text{unique}}$, there is only one way to read Axiom `Inv`: "For every $\mathbf{u}$, there is some $-\mathbf{u}$ such that $\mathbf{u} + (-\mathbf{u}) = \mathbf{0}_\mathbf{u}$," where $\mathbf{0}_\mathbf{u}$ is the single zero element of $\mathbf{u}$ guaranteed by Axiom `Zero`$^{\text{unique}}$.

Suppose we combine these versions of Axiom `Zero` and Axiom `Inv` with the other eight of the usual vector space axioms. If V is a structure that satisfies all ten of these axioms, then how do elements of V relate to each other? For example, take elements $\mathbf{u}$ and $\mathbf{v}$ of V, with zero elements $\mathbf{0}_\mathbf{u}$ and $\mathbf{0}_\mathbf{v}$, respectively. Then we have

$$\begin{aligned} \mathbf{0}_\mathbf{u} + (\mathbf{u} + \mathbf{v}) &= (\mathbf{0}_\mathbf{u} + \mathbf{u}) + \mathbf{v} \\ &= \mathbf{u} + \mathbf{v}, \end{aligned}$$

and thus $\mathbf{0_u}$ acts as a zero element not only for $\mathbf{u}$ but also for $\mathbf{u}+\mathbf{v}$. In precisely the same manner we can show that $\mathbf{0_v}$ acts as a zero element for $\mathbf{u}+\mathbf{v}$ as well. But by our new Axiom $\texttt{Zero}^{\text{unique}}$, $\mathbf{u}+\mathbf{v}$ has exactly one zero element, and thus $\mathbf{0_u}$ and $\mathbf{0_v}$ are the same. Since $\mathbf{u}$ and $\mathbf{v}$ were arbitrary, it follows that all zero elements are the same—so that there is in fact *one* $\mathbf{0}$ that works as a zero element for *every* element of V!

Thus, we *can* replace Axiom `Zero` with Axiom $\texttt{Zero}^{\text{unique}}$ to obtain an equivalent set of axioms for vector spaces. The Mad Hatter would be most displeased! In this case, the quantifiers "for all" and "there exists" really *can* be switched with impunity (though, if it would make the Mad Hatter feel any better, I have to reiterate that there really is a big logical difference between "there exists an $\mathbf{x}$" and "there exists a *unique* $\mathbf{x}$"). Moreover, the only other axioms we used were commutativity and associativity of addition.

Another look at Axiom `Inv`

Now suppose, again, that we are using Axiom `Zero*`: "For every $\mathbf{u}$ there is a $\mathbf{0_u}$ such that $\mathbf{u}+\mathbf{0_u}=\mathbf{u}$." Recall that we had to rewrite Axiom `Inv` because of its reference to $\mathbf{0}$, and that we rewrote it to assert, for every $\mathbf{u}$, the existence of an inverse for $\mathbf{u}$ with respect to *some* zero element for $\mathbf{u}$.

There is, of course, another way we could choose to rewrite Axiom `Inv`—as Axiom $\texttt{Inv}^{\forall}$:

> For every $\mathbf{u}$ and for *every* zero element $\mathbf{0}_{\mathbf{u},\alpha}$ of $\mathbf{u}$, there is an element $-\mathbf{u}_\alpha$ such that $\mathbf{u}+(-\mathbf{u}_\alpha)=\mathbf{0}_{\mathbf{u},\alpha}$.

Thus, an element $\mathbf{u}$ (of a structure V satisfying this newest set of axioms, with Axiom `Zero*` and Axiom $\texttt{Inv}^{\forall}$) might have more than one zero element, but there must be an inverse for $\mathbf{u}$ with respect to *every* one of its zero elements. For instance, $\mathbf{0}_{\mathbf{u},1}$ and $\mathbf{0}_{\mathbf{u},2}$ might be zero elements for $\mathbf{u}$; by Axiom $\texttt{Inv}^{\forall}$, then, there would have to be elements $-\mathbf{u}_1$ and $-\mathbf{u}_2$ such that

$$\mathbf{u}+(-\mathbf{u}_1)=\mathbf{0}_{\mathbf{u},1} \quad \text{and} \quad \mathbf{u}+(-\mathbf{u}_2)=\mathbf{0}_{\mathbf{u},2}.$$

Just for the fun of it, since

$$\mathbf{u}=\mathbf{u}+\mathbf{0}_{\mathbf{u},1},$$

we can add $-\mathbf{u}_2$ to each side and find, appealing to commutativity and associativity along the way, that

$$\mathbf{u}+(-\mathbf{u}_2)=\big(\mathbf{u}+(-\mathbf{u}_2)\big)+\mathbf{0}_{\mathbf{u},1}, \quad \text{and thus}$$
$$\mathbf{0}_{\mathbf{u},2}=\mathbf{0}_{\mathbf{u},2}+\mathbf{0}_{\mathbf{u},1}.$$

This is cute enough as it stands. Note that we can't do what seems natural—"subtract" $\mathbf{0}_{\mathbf{u},2}$ from both sides and conclude that $\mathbf{0}=\mathbf{0}_{\mathbf{u},1}$—because so far we have no reason to believe either that $\mathbf{0}$ exists or that subtraction would work for these elements. We can, however, start over with

$$\mathbf{u}=\mathbf{u}+\mathbf{0}_{\mathbf{u},2}$$

and add $-\mathbf{u}_1$ to both sides to find that

$$\mathbf{0}_{\mathbf{u},1} = \mathbf{0}_{\mathbf{u},1} + \mathbf{0}_{\mathbf{u},2}.$$

Then $\mathbf{0}_{\mathbf{u},1}$ and $\mathbf{0}_{\mathbf{u},2}$ are both equal to $\mathbf{0}_{\mathbf{u},1} + \mathbf{0}_{\mathbf{u},2}$ and thus are equal to each other. That would be no big deal, except that $\mathbf{0}_{\mathbf{u},1}$ and $\mathbf{0}_{\mathbf{u},2}$ were really *arbitrary* zero elements of $\mathbf{u}$. We've shown that they are equal, so $\mathbf{u}$ has only *one* zero element. And now that we think about it, $\mathbf{u}$ was arbitrary too; so *every* element of V has exactly one zero element. But as we saw in the previous section, this implies that there is a true $\mathbf{0}$ and that V is a vector space.

Thus, we *can* replace Axiom `Zero` with Axiom `Zero`* if we rewrite Axiom `Inv` as Axiom `Inv`$^\forall$ rather than as Axiom `Inv`*, and the logical structure of Axiom `Inv`$^\forall$ is no more complex than that of the usual Axiom `Inv`. Pedagogically, though, it seems to be a step backwards. The usual Axiom `Inv` seems much more natural, because it refers to the single designated $\mathbf{0}$ postulated by the usual Axiom `Zero` rather than to a whole host of conceivably different zero elements that merely *turn out* to be the same.

The moral for students is that they are free to be lazy or careless with one axiom, but only if they are willing to make another more complicated.

Final thoughts

Was my distinction between Axiom `Zero` and Axiom `Zero`* only pedantic? Certainly not—the existence of $\mathcal{W}$ shows that it does matter how Axiom `Zero` is written. Was the distinction vitally important? Well, not *vitally*, since the addition of a well-chosen but more complicated version of Axiom `Inv` will result in an equivalent set of axioms. But the distinction is important.

First, there is an immediate pedagogical value in providing examples of objects that are *not* vector spaces as well as objects that *are*. Second, there is a broader mathematical value in providing students yet another example of how switching the order of quantifiers in a sentence changes the meaning of the sentence. Perhaps this line of inquiry can even foreshadow similar questions for later classes; for example, what difficulties might occur if the ring theory axioms were misstated?

Finally, there is a practical value for the instructor that I didn't anticipate: I now had a reason for *not* giving partial credit if Axiom `Zero` was stated incorrectly on the next exam, so students knew they had to get the axioms exactly right. That in itself was worth the search.

Acknowledgments. I am grateful to the referees and the editor for suggestions that resulted in definite improvements to the paper. In the last paragraph of Section 3, in particular, both the simplest possible example (which I somehow missed!) and the idea to generalize $\mathcal{W}$ and $\mathcal{W}'$ were suggested by one of the referees.

References

1. Howard Anton, *Elementary Linear Algebra*, (8th ed.), Wiley, 2000.
2. I. Martin Isaacs, *Algebra: A Graduate Course*, Brooks/Cole, 1994.
3. Elliott Mendelson, *Boolean Algebra and Switching Circuits*, Schaum's Outline Series in Mathematics, McGraw-Hill, 1970.
4. Paul K. Stockmeyer, What did Lincoln really mean?, *College Mathematics Journal*, **35** (2004), 103–104.