

WIKIPEDIA, FREE KNOWLEDGE AND PEACE*

■ JIMMY WALES

I was asked to speak on the topic of peace and Wikipedia, which is a pretty big topic, and I've been fortunate enough in the past few days to have conversations over lunch and dinner with several people here and it became apparent that there's a lot of interest in how Wikipedia actually works, so I'm going to cover that as a base for about ten minutes to really explain, as quickly as I can, some of the mechanics of how Wikipedia actually functions. Once that base is covered I'll talk about our neutrality policy in particular, and how I think that that contributes to, and can contribute to, peace.

The original vision for Wikipedia is for all of us to imagine a world in which every single person on the planet is given free access to the sum of all human knowledge, so every piece of this vision statement is important. When we talk about every single person on the planet the project is inherently global in scope and – in vision, anyway – reaches far beyond just the people who have access to the Internet. When I say “free access” I mean something very particular by that, which, as was just discussed in the previous lecture, has more to do with open sharing than it has to do with cost. An old joke from the open source software community is, “we mean free as in speech, not free as in beer”. Free beer is great, people love to get a free beer, but we're talking about something more fundamental than just cost, we're talking about the licensing condition that allows anyone to copy, modify, and redistribute modified versions of all of our work and people can do that commercially or non commercially, which gives rise to a lot of interesting side-projects and things that go on. Finally, when we talk about the sum of all human knowledge, this is really, in a way, a limitation on what Wikipedia is. Wikipedia is not everything in the world, Wikipedia is not a textbook, Wikipedia is not YouTube: I love funny cat videos as much as anybody does but they don't belong in Wikipedia. What we have is an essentialized summary of human knowledge with the depth of material depending on the context.

How does Wikipedia work? Although Wikipedia is an open community, Wikipedia is not an anarchy. We have several hundred pages of editorial rules and guidelines and one of the amusing things people will often say to

* Transcript unrevised by author.

me is, has your community ever thought of ____ (and you can fill in what you want there) and the answer is, probably yes, probably not only have they thought of it, they've had a 30-page discussion about the names of rivers in Poland, or whatever it might be. It's an incredibly verbal community who think really hard about editorial policies and guidelines. The way this all works is that we have a complex social structure within our community, which I will try to explain very briefly. What we have is a curious but workable mix of several things: consensus, and by consensus what I mean is, whenever there is some disagreement about what an entry in Wikipedia should say we try, over and over and over again, on the discussion page, to find a formulation that will have very wide support, that even supporters and opponents will be able to agree with and I'll go into that in more detail when I talk about the neutrality policy. We have some elements of democracy: sometimes we do have informal votes on things, sometimes there is no simple way to compromise. If the entry in question is the Eiffel Tower in Paris, everybody agrees that we need to have a picture of the Eiffel Tower, but there's maybe disagreement about which picture to use and we'll have a little straw poll on it and see what happens. Those votes are typically not binding, but they do help to generate consensus. People may say, well I think this picture is better than that one but I see that 80% of people disagree with me so I'll drop the issue and let it go. Also, we have administrators who are elected by the community but, once elected, it's very, very difficult for someone to lose their status as an administrator. In a certain sense, there's a bit of aristocracy, these are the people in the community who've gained a certain status, a social status, a technical status as administrators, and they need a certain amount of independence to do their work and therefore they can't be unelected. They can be removed as administrators, it does happen from time to time, but it's quite rare. And then we have a little bit of monarchy, and this is my role in the community. I always say that Americans have a very hard time understanding how Wikipedia functions, as compared to the British, who do understand the idea of layers and layers of protocol and things that really rest almost as much on tradition and custom as anything else, and the idea of a monarch like me who mainly just waves at crowds is something that the Brits very much do understand.

The first thing to know is that anyone can edit Wikipedia so, within this complex social structure, the first thing to know is that for 99.9% – I don't know the exact number, it's actually higher than that – of the pages in Wikipedia anyone can edit anything at any time without even logging in. You don't even have to log in, you just click "edit", you make a change and you save and it goes live immediately. What happens if you don't log in is

that your IP address – the address identifying your computer – is shown, we can see that the edit was made by someone who wasn't logged in, we can see your IP address, and then once that edit is made it goes to the recent changes page – every change to Wikipedia is posted on the recent changes page – and there's a large group of people who spend their time – it's almost like a videogame for them – monitoring the recent changes page and looking at all the changes that come in. There are various tools that they use, robots that are looking for certain types of changes that are easy to flag as problematic so, for example, if a page goes from 20,000 characters down to 3 characters, you probably know that's not a good edit, somebody just blanked a page, that's probably wrong, a human needs to look at it and see what's going on. There are edit filters that look for curse words. Sometimes it can be appropriate in Wikipedia to have a curse word, you know, if a famous person has said a famous statement that includes a curse word it wouldn't be absolutely banned in Wikipedia but usually, when somebody puts in a curse word, they're up to no good and so that gets flagged and humans look at that more quickly. Additionally, every editor who is logged in has their personal watchlist. Most of the experienced editors have it set so that anything they ever edit goes on to their watchlist, so once you've participated you can see on your watchlist, the next day when you log in, what's changed in the areas you're interested in since yesterday. That's a slower filter, it may take several hours before somebody sees it on their watchlist, and that is the way a lot of ongoing discussions take place.

Anonymous edits are viewed, quite rightly, with suspicion. People say, this is someone we don't know, this is somebody who didn't even bother to log in, so if people want to gain a reputation in the community they log in, they gain a stable identity, they have an identity, everyone can see what they're doing, it's all very transparent and open, and there are certain pages that are semi-protected, we call it, due to ongoing troubles. An example of this would be a page like "George W. Bush" which, if you leave completely open so that people can edit without logging in it's exhausting, the number of people who come by and insert curse words or blank the page or do bad things. I don't mean to single out George W. Bush, this is true for a lot of very prominent people, he is just particularly prominent, particularly controversial. There are a lot of pages that end up semi-protected and in order to edit those pages you have to log in, have an account, and you have to have an account for four days. Once you've had an account for four days, then you're allowed to edit semi-protected pages. That's a very low barrier to entry. The idea here is that we still want to have this open community where people can participate, it's something like a cooling off period, if you

just come by and you're angry and you want to insult the President you have to wait four days. As it turns out, this eliminates 99% of that kind of behaviour. After four days people feel a little bit better and they realise, if they just log in and insult the President, they're just going to be banned and they'll have to start over again. This really does help a lot. And then we have full protection. This is when there's a big argument that's broken out between established editors, emotions are running high and there's a decision made to say we need a 24-hour cooling off period, we need to say, everybody step back, go to the discussion page, work out your difficulty and then come back and edit the article. The only people who can edit an article while it's under full protection are those administrators who are elected by the community but, as a tradition and custom, they don't edit those pages. This is mainly just to calm things during a big fight. The administrators are elected by the community, they do have additional powers, they can protect the pages, they can block people from editing, but their powers are socially limited. They are police, not judges, and so their actions can be reviewed. If someone is blocked in Wikipedia there is an avenue for appeal. You can say, I was blocked unfairly, I was blocked for disagreeing with someone, and this is where it becomes important that the administrators are something of an aristocracy, they don't have to bow to immediate popular will, but they do have checks and balances, there is an ability for administrators to lose their administration. How does that happen? We have the arbitration committee. This is a committee that is elected by the community and then appointed by me. The arbitration committee can issue binding decisions and in many cases they take on very difficult and thorny cases where there's been a longstanding, ongoing controversy and good editors haven't been able to resolve the conflict. After the arbitration committee makes a decision the only possible appeal is to me, and that is an avenue that is open to people, but I only review things for significant injustice or abuse of authority, I don't view it as appropriate for me to overturn the arbitration committee just because I disagree. I don't think that's the right role for a final appeal. I'm here as a final check and balance to say, if the arbitration committee decides to go crazy and starts banning people for reasons that don't make any sense, I'm there to say no, something has to give. My role in the community is, I would say, constitutionally ambiguous, and I would like to keep it that way. The rules say I can set policy by *fiat*, but whether such policy would be popular and accepted is an open question, so I avoid it, much like the Queen of England who, in theory, has a lot of power, but avoids using it because probably, if she did, she would lose her job. I try to avoid doing things unless absolutely necessary. The last time

I did was in July 2010 and that policy held firm, we had a policy on child protection where we really wanted to be quite firm about issues surrounding advocacy of paedophilia and there was some debate and discussion in the community as to whether it was policy or just tradition and that very much stuck within the community and now has been ratified.

This system is very complex, obviously, and it has a lot of checks and balances within it and unfortunately, of course, with any system as complex as this there's also some bureaucracy and so one of the questions that arises for people is what is the quality that results from this process. Now that you've heard a little bit about how we do it, I hope you feel a little bit reassured that it isn't just randomly anyone can do anything at any time, and the best academic studies of quality that we have suggest that the quality is similar to, or in some cases, better than Britannica. There was an academic study that showed the number of errors in Wikipedia versus Britannica is quite similar. But the truth is, and this is important to understand, this is not because we are great, this is because Britannica was never as good as people thought it was! The truth is, doing good quality reference material for the general public is a very hard task. We do make a lot of errors, and sometimes we get beaten up in the press for those errors and sometimes that's fair and sometimes that's not fair, but we are aware that there's always room for improvement and quality.

I want to move now into talking about the subject of Wikipedia and peace. I do think that we have, at least potentially, a significant role to play here, and it's something that I very much desire for Wikipedia as a mediating and moderating influence on the discourse of the Internet. Wikipedia now reaches nearly 500 million people every month. The last numbers I have, last month I think it was 481 million people read at least one thing in Wikipedia, and we are the most linguistically diverse website, we are in many, many languages. Just to give you some of the scope of that, English, German, French and Dutch all have over one million articles. We have six more languages that have over 500,000, we have 39 languages with over 100,000 articles and we have 108 that have over 10,000 articles. Ten thousand articles is still not a comprehensive encyclopaedia, but it's important to remember that for many of these languages – a good example would be Swahili, which has around 30,000 articles – there has never existed an encyclopaedia as we understand it. They've never had an encyclopaedia in their language and so this is the first encyclopaedia for many, many people all around the world. The growth of Wikipedia in the developing world is something that I'm extremely excited about and one of the things that I've been trying to communicate to people in the past few months is something

that I think a lot of people don't realise yet, which is how quickly Internet access is growing now in the developing world. Just to take a look at one country, Nigeria, what we've seen is that in 2000, 0.1% of Nigerians were online. By 2006 that had exploded, you might think, to 3.1%. In 2009 it was 16.1% and in 2011, the latest numbers I have access to, 29% of Nigerians were on the Internet. This is an incredibly important transformative change to society, in the same way that it was an important and transformative change to the wealthy societies of the world, people are able to get online. And what we're seeing, in the Yoruba Wikipedia, which is one of the major languages of Nigeria, just as Internet access has nearly doubled in the last year, the size of Yoruba Wikipedia has doubled in the last year, as people come online they become very interested and very eager to work in their own language. This is true even of people who are part of the diaspora. In many cases, talking to people who are in the diaspora, who are working in wealthy countries, they still have a very strong passion for their mother tongue but in the days when no one back home had Internet access there wasn't much inspiration to write in Yoruba Wikipedia. If you're a professor and you're working in Europe you might feel, well, if I write in Yoruba no one will read it, because they're not on the Internet yet. This is increasingly not true. Increasingly, if you're a professor out in the world somewhere, you can write in your home language and people back home will be able to read it, and so it's a very important mechanism to allow people to share their knowledge with people back home.

Some of Wikipedia's policies are very important. One of the most important ones is assume good faith. The idea here is that if you see someone doing something wrong in Wikipedia it usually isn't because they're being malicious, it's because they're trying to do the right thing in a wrong way. This is a very optimistic view of people but it turns out to be empirically justified. In most of the cases where you see someone come into Wikipedia and they're doing something that you think is bad if you challenge them and say, why are you doing that, this isn't the right thing to do, they'll reply and you'll have a conversation with them and they turn into much better contributors. We can't be super naïve. Obviously some people are just malicious people or crazy people or whatever and we do have to block people, we do have to deal with that, but by and large it turns out that most people are basically good and they're more than happy to contribute in a positive way.

Perhaps the most important policy of Wikipedia is what we call "NPOV" in our bit of jargon, "neutral point of view". This is one of the earliest policies of Wikipedia and it's one that I set down at the beginning of the project and said that NPOV is non negotiable. The question of

whether or not Wikipedia should strive to be as neutral as possible is a question we don't discuss or debate within the community, it's just part of the context of how we operate and who we are. What do we mean by this, by "neutral point of view"? Obviously the questions of bias and neutrality are, in many ways, very deep philosophical questions and we try to be as sophisticated as we can be about that. Insofar as possible, in any controversy, Wikipedia should not take sides but should fairly describe all significant viewpoints so that someone new to the subject can actually understand what is the debate about. Rather than that you come to what we think is the correct conclusion, it is more important to us that you really understand what all sides have said and why they say what they say. We try to maintain a subtle but philosophically sophisticated policy and so it precludes things and this is written directly into the policy. It precludes, for example, equating mainstream views with extreme minority viewpoints, so our view of neutrality is not to say, "some say the moon is made of rocks and some say the moon is made of cheese, who knows". We don't do that. We say, look, we want all significant viewpoints and there's a lot of discussion and debate on what constitutes a significant viewpoint, when does a minority viewpoint warrant equal treatment alongside another, and when does a minority viewpoint deserve to be simply ignored or simply moved out to a separate entry. These are very complex questions and there is no simple rule or simple answer. My view of this is that the only way to come to even an approximately sensible answer is through is an ongoing discussion, dialogue and debate, carried on in good faith and with thoughtful reason. For many controversies our mass media have been more inflammatory than informative, with more of an interest in encouraging combative discourse than encouraging a thoughtful dialogue and I hope that we try really hard to be an antidote to that.

I'll just tell one story, one example that I think is telling. I was in Taiwan and we had a group of students who were showing me around and the young man they had assigned to be my driver – who is a very active Wikipedia editor – told me that he had been raised in a very nationalist, Taiwanese household and he had been raised to think that the mainland Chinese were all completely brainwashed and that they knew nothing about history. He started to work in Wikipedia and working alongside people from the mainland and working on entries like the history of Taiwan and he said to me, "you know, there's a lot of really nice, really smart people. I still think they're wrong about a lot of things but I can kind of see where they are coming from". And so he had made that transition of thinking of them as mindless robots to thinking of them as human beings with whom

he had a disagreement, which I think is a huge step forward; in many parts of the world that is an incredible step forward that will help to further the cause of peace.

I'm going to close now. I know I'm slightly over time but as I'm quoting *Pacem in Terris* I think it's probably ok.

As we know from experience, men frequently differ widely in knowledge, virtue, intelligence and wealth, but that is no valid argument in favor of a system whereby those who are in a position of superiority impose their will arbitrarily on others. On the contrary, such men have a greater share in the common responsibility to help others to reach perfection by their mutual efforts (PT §87).

This is the spirit that we absolutely see in Wikipedia: the best academics who participate in Wikipedia understand that many of the people they are dealing with will have less knowledge, will be less reasonable. However, when they are filled with this spirit of sharing and love of knowledge, they are very patient and they realise it's important to work with everyone to come to a mutual understanding.

Thank you.