On Godwin's Law: A Statistical Analysis on the Distribution of Nazi Analogies in Online Discussion

Alex Miller
alex@miller.im

November 2013

Abstract

Godwin's Law is a well known internet adage about the prevalence of analogies to Nazi Germany in online discussions. Though Godwin's Law can be shown to be mathematically true by a triviallly universal proof, this does not account for the perceived frequency of Nazi analogies specifically by Godwin and others. The undertaking of this research is to study the empirical distribution of online Nazi-comparisons by examining the comments section of major online news outlets. It is shown that number of comments until the first Nazi-comparison does not follow a geometric distribution (as presumed by the trivial theory), but rather demonstrates power-law characteristics above a specified threshold. The dataset in this study is shown to resemble a type II Pareto (Lomax) distribution, though ways to improve this model are suggested in conclusion.

1 Introduction

The social dynamics of online discussion fora have received increased media and research attention in the past year. Perhaps the most well-known topic in recent discourse is the University of Wisconson-Madison study [1] that prompted Popular Science to remove its online discussion boards due to their effect on readers' perception of their articles [2]. Results such as this underscore the importance of understanding discussion dynamics not simply from a statistical perspective, but also a social one.

One of the first and oldest observations about social behavior in online discussion threads was formulated by Mike Godwin in 1990. In what has become known as "Godwin's Law of Nazi Analogies", Godwin asserts: "*As an online discussion grows longer, the probability of a comparison involving Nazis or Hitler approaches one*" [3]. Despite its tongue-in-cheek characterization as a "law", it is still surprising to find very little academic literature on this subject. The current research attempts to formally study Godwin's Law and its properties in real online discussion fora.

Initially, Godwin's assertion was less a "law" than an informal observation about the prevalence of analogies to Nazi Germany in Usenet discussion groups. Nonetheless, Godwin's Law anectodally appears to hold true universally. Indeed, as discussed in [3], a key aspect of Godwin's Law is that it appears to hold true regardless of the topic of conversation (esp. when references to Nazi Germany would be unexpected).

Copyright is held by the author/owner. To copy, redistribute, or republish this work, in part or in whole, requires prior specific permission and/or a fee.

As recalled in [4], Godwin's intention behind his eponymous law was to to invent a meme that would neutralize the "Nazi-comparison meme". Quoting from Godwin himself (in reference to his law):

> *Although deliberately framed as if it were a law of nature or of mathematics, its purpose has always been rhetorical and pedagogical: I wanted folks who glibly compared someone else to Hitler or to Nazis to think a bit harder about the Holocaust.*

Godwin's goal is a worthy one. Though this study may betray his intention by studying the law mathematically, it may also contribute to Godwin's goal by laying the groundwork for future research. Presently, no literature exists on the prevalence of Nazi-comparisons in online dicussion. This prevents even the most essential questions about this subject from being studied; namely, is Godwin's Law succeeding as a "counter-meme"? Is the incidence of Nazi-comparisons decreasing over time? Is it inversely related to awareness about Godwin's law itself?

The current research does not address these questions, but rather the more basic ones of: (1) How accurate is Godwin's Law (quantitatively) in real discussion fora? And (2) what is the distribution of Nazi-comparisons in online discussion?

2 Background

2.1 Formalization & Terminology

Let us first formalize Godwin's Law in the language of mathematical precision. Though the statement, "As an online discussion grows longer" can be interpeted as a reference to an increasing quantity of words, characters, or digital memory correspoding to an online discussion, we will instead refer a thread's "length" as the integral number of entries or comments in the discsussion. Furthermore, when referring to an entry that makes a "comparison involving Nazis or Hitler", we will euphmistically call this a "Godwin positive" entry or a "Godwin match". Let us all refer to the number of comments in a thread until the first Godwin match as the thread's "Godwin length".

2.2 Triviality

Presumably, Godwin's Law is trivially true. Let l be the length of arbitrary thread and assume the probability that a thread participant makes a Godwin positive comment is greater than or equal to some $p > 0$ at all times. We can then think of each new comment in the thread as a Bernoulli trial. The distribution of Godwin length will be bounded below by a geometric distribution with parameter p. Let $F(l)$ be the cumulative distribution of the first Godwin positive comment in an arbitrary thread and let $G(l) = 1 - (1 - p)^l$ be the theoretical lower-bound geometric distribution. Then Godwin's Law may be proven mathematically using the squeeze theorem:

$$1 \geq F(l) \geq G(l)$$

$$\lim_{l \to \infty} G(l) = 1$$
$$\Rightarrow \lim_{l \to \infty} F(l) = 1$$

Of course, nothing about this argument has been unique to Godwin's Law. The same logic could apply be applied to to any arbitrary topic, so long as we assume that the probability of it being mentioned is always greater than zero.

Indeed, Godwin's Law has been criticized for this very reason. French blogger Brogol [5] points out the apparent absurdity of Godwin's Law (using roughly the same logic above) by declaring his own law:

> *As an online discussion grows longer, the probability of finding a comparison involving platypi approaches 1*

Clearly, this observation has the potential to diminish the appeal Godwin's Law as an interesting subject. Thus, the over-arching goal of this research is not to simply "prove" Godwin's Law, but rather demonstrate empirically that references to Nazi Germany in online discussion are more common than those to an arbitrary subject (such as platypi). This task is not fully accomplished in the present study, though the findings here are constructive in informing future inquiries into the subject.

2.3 Related Literature

Much of the literature about Godwin's Law is found in popular media outlets. These articles are typically either a criticism of the universal "ban" that Godwin's Law pronounces on mentioning Nazis ([6], Salon.com) or just commentary on the the prevalence of Nazi-analogies in general discourse ([7], Reason.com).

Though there is apparently no strictly academic literature on the subject of Godwin's Law, there are numerous studies on the dynamics of online discussions. In [8], Mishne & Glance undertook the task of studying comment sections in the whole blogosphere. Their results demonstrated an approximately power law distribution in the overall length of weblog discussion threads. In both [9] and [10], the authors studied the mechanisms by which the distribution of various blog metrics (e.g., length, number of authors, post in-degree, thread depth) arise. These results can be very instructive for explaining discussion meta-data, but less so for analyzing the contents of online discussion.

3 Methodology

The data for this study were extracted from the comment sections following online news and commentary articles. The sites used were www.cnn.com, www.npr.org, and abcnews.go.com. All three domains are home to major national news organizations, each with thousands of news articles and accompanying comment sections. Each website has daily stories on a variety of subjects and significant readership (and consequent discussion participation).

To gather the data, each site was first crawled (using a third-party spider) to generate a list of URIs with discussion sections. Each URI was then processed that scraped both the webpage content and the comment section (full description of this process will appear in a forthcoming article). Comments were sorted chronologically and searched for matches from the list of keywords below. This list was generated by considering the most salient subjects pertaining to Nazi Germany and their most common misspellings (identified by data from http://spellweb.com).

- nazi
- nazis
- natzi
- nazie
- nazzi
- adolf
- hitler
- hilter
- hittler
- hitlar
- weimar
- gestapo
- gistapo
- himmler
- goebbels
- goebells
- goebels
- fuehrer
- fuhrer
- third reich
- 3rd reich
- third riech
- 3rd riech
- third rike
- 3rd rike
- holocaust
- concentration camps
- concentration camp
- mein kampf
- auschwitz
- dachau
- yellow badge
- yellow patch
- yellow badges
- yellow patches
- fuhrer
- furher
- holocost
- holocoust
- halocaust
- holocuast
- holacaust
- holocause
- hollocost
- holicost
- hallocaust
- holcaust
- holocast
- holocust
- holacuist
- hulocost
- haulocaust
- haoulocaust
- holucaust
- holocauset
- hholocaust
- holicaust
- aushwitz
- auchwitz
- auschwits
- auschwiz
- auschwtiz
- auswitz
- aushcwitz
- auchwats
- auschzites
- auchwtiz
- auschowitz

Each URI was processed and assigned a corresponding data-vector with the following information:

- Thread Length: Total number of comments in thread at time of observation
- Godwin Match Index: The index of the first comment containing any keyword(:=0 if no match found)
- Match Context: 100 (±50) character comment context around keyword
- Keywords found in page body: boolean value, TRUE/FALSE

TRUE was applied to the last entry for URIs for which any one of the keywords was found outside of the page's discussion section; these URIs were excluded from the final dataset.

4 Data

A total of 20175 URLs were crawled between the three sites. Of these, 16043 URLs contained discussions with ≥ 1 comment. Of this subset, 4716 threads contained a Godwin match. Below is a summary of the this study's dataset.

Table 1: Summary Data

Total URLs Processed	20175
Threads with ≥ 1 comment	16043
Total comments processed	10593454
Average Thread Length	659
Threads with ≥ 1 Godwin match	4716
Average Godwin Length	213

Table 2: Data by Source

	www.cnn.com	www.npr.org	abcnews.go.com
Threads With ≥ 1 Comment	8468	6126	1507
Average Thread Length	4024	1507	555
Threads With ≥ 1 Godwin Match	3996	555	127
Average Godwin Length	236	73	138

For reference, Figure 1 shows the distribution of overall thread length (in log-log scale) with the best-fit power law coefficient (in red). This distribution closely resemebles that found by Mishne & Glance [8]. Both appear to approximiately follow a power law, with slightly diminished quanitities of small values.

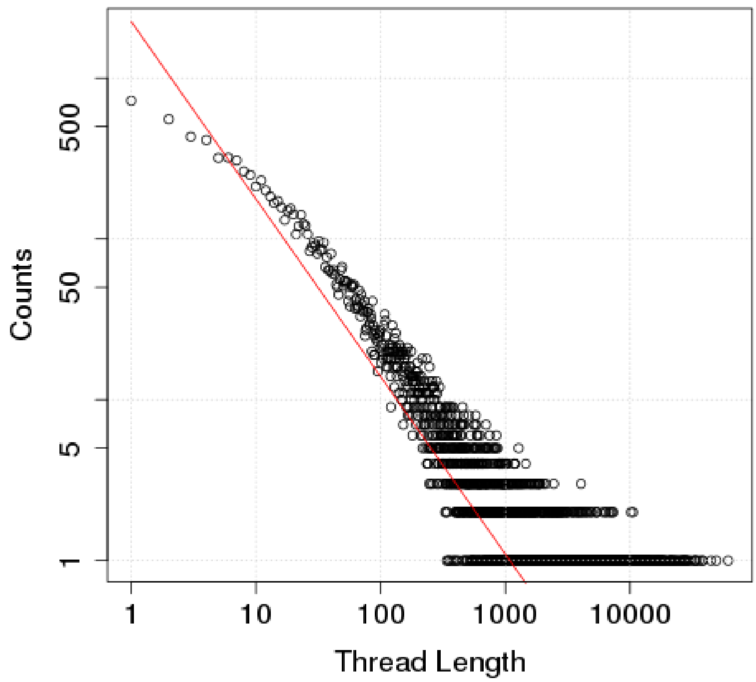

Figure 1: Distribution of thread length in log-log scale with best-fit power law coefficient

Figures 2 and 3 represent the data corresponding to the subset of with Godwin matches. Recall that a thread's Godwin length is the number of comments until the first Godwin positive entry.

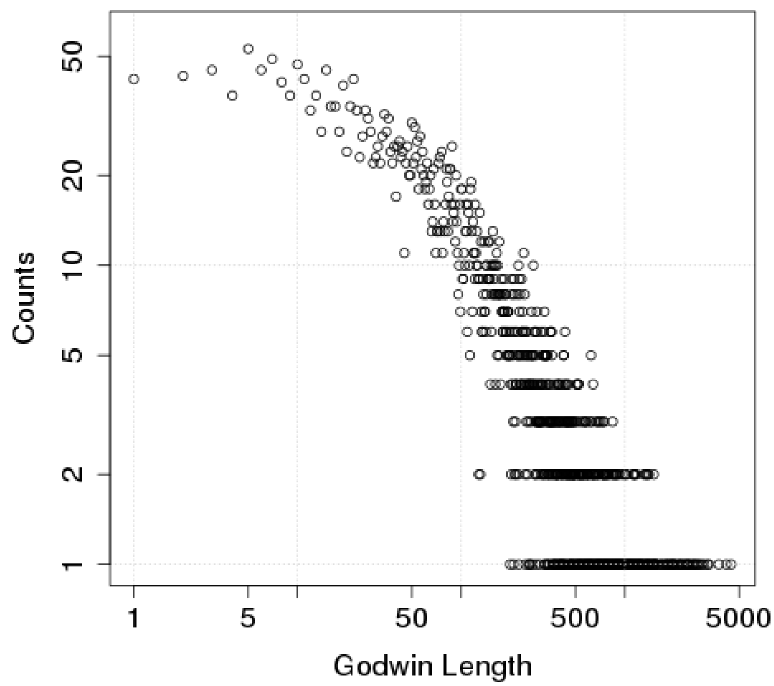

Figure 2: Distribution of Godwin length in log-log scale

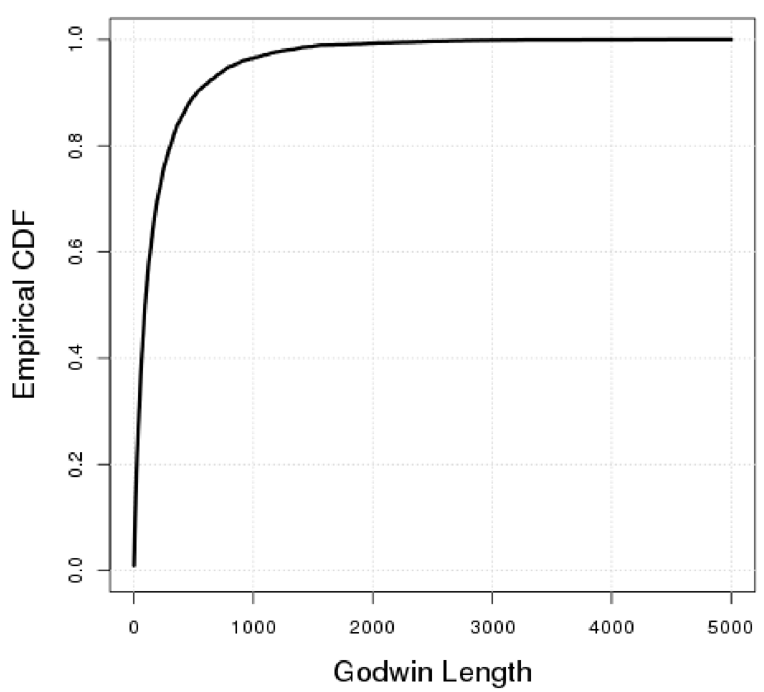

Figure 3: Empirical cumulative distribution of Godwin length

Table 3: Summary of the Inverse Empirical Godwin Length CDF

l	$\mathrm{ECDF}^{-1}(l)$
.1	11
.2	26
.3	44
.4	66
.5	94
.6	138
.7	198
.8	307
.9	526
.95	817
.99	1675

5 Analysis

Upon first glance at Figure 2, the distribution appears to resemble that of a geometric/exponential random variable. However, this is seen to not be the case upon further anlayis. See red CCDF line in figures 4 and 5 for the best fit geometric distribution. These figures reveal the dataset evidentaly has a heavy-tailed distribution.

Though the data clearly deviates from a strict power law or Pareto distribution, its tail does appear to follow a power law above a certain minimum value. Choosing an x_{min} of 1000, and fitting the data using maximum-likelihood method yields a theoretical power law distribution with parameter $\alpha = 3.56$. The KS statistic between this theoretical power distribution and the empirical distribution is 0.05689 with corresponding p-value of 0.656.

However, we can model the head of the distribution and account for its power-law tail by considering the Lomax (Parteo type II) distribution:

$$\text{PDF:} \quad \frac{\alpha\lambda^{\alpha}}{(x+\lambda)^{\alpha+1}}$$

$$\text{CCDF:} \quad \left[1 + \frac{x}{\lambda}\right]^{-\alpha}$$

Using the MLE method, the estimated parameters for the data's best-fit Lomax distribution are $\alpha = 2.080$, $\lambda = 246.866$. One can assess the accuracy of this distribution by comparing CCDFs. The green line Figure 4 (on log-log scale) represents the best-fit Lomax CCDF.

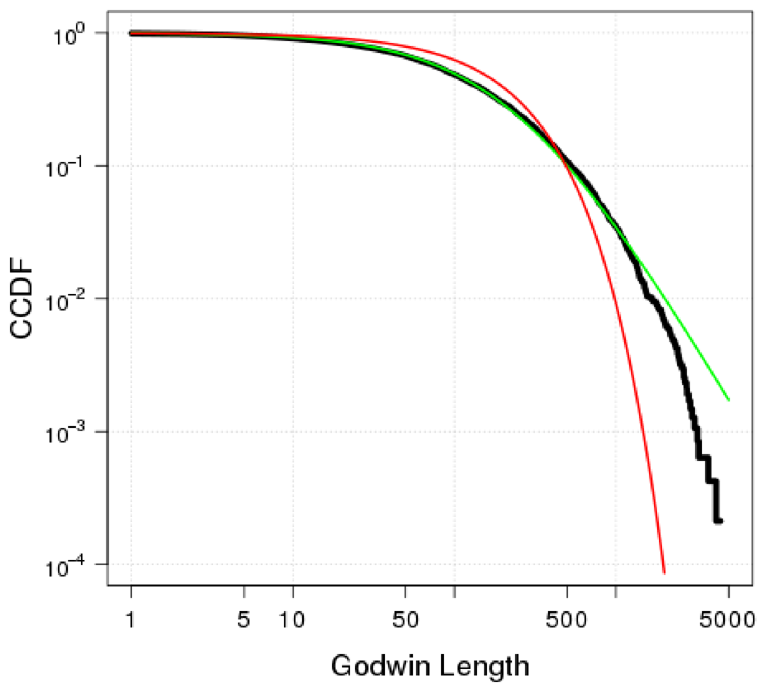

Figure 4: CCDF functions in log-log scale. Legend:
BLACK: experimentally measured Godwin Length distribution;
RED: best fit exponential (i.e., geometric) distribution;
GREEN: Best fit Pareto distribution (type II)

The semi-log plot in Figure 5 better represents the tail behavior of the distributions. As can be seen, the data's distribution is heavy-tailed, though less so than the theorietical Lomax curve.

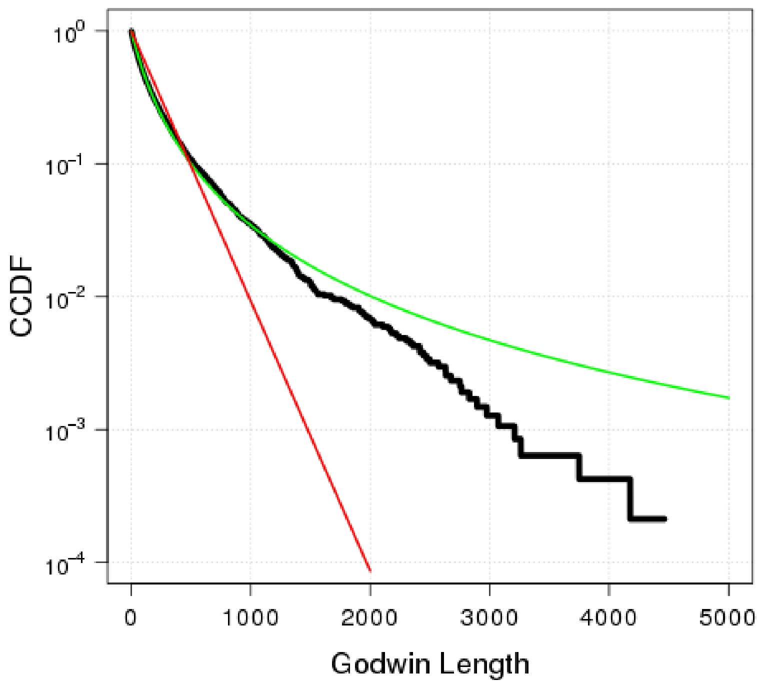

Figure 5: CCDF functions in semi-log scale. Legend:
BLACK: experimentally measured Godwin Length distribution;
RED: best fit exponential (i.e., geometric) distribution;
GREEN: Best fit Pareto distribution (type II)

Despite this deviant tail-behavior, the χ^2 goodness-of-fit test returns a p-value of 0.38. This allows us to assume the null hypothesis at the critical value of .05.

6 Conclusions & Future Work

This research has shown that the Godwin length of the observed data is not accurately modeled as a geometric random variable. Rather, power-law behavior was observed for values ≥ 1000 and the entire sample set was plausibly modeled by the Lomax distribution.

Strictly speaking, the Lomax distribution should be used for continuous random variables. A straightforward way to to extend the current analysis is to start with a discrete probably distribution, such as the zeta or Zipf. However, this approach will still need account for the non-power law behavior for small values. Cristelli et al. [11] have studied Zipf distributions with the same general skewness as that in this study and propose a correction factor for small sample values. This approach is promising for both modeling and explaining the power and non-power law behavior of the Godwin length distribution.

The data sources themselves may also deserve more scrutiny in future analyses. The nature of national news and commentary websites (such as those used here) ensures a diversity of article (and corresponding discussion) topics, though it is expected that the content of the articles in this study was more political in nature than "general" online content. It is possible that the prevalence of Godwin positive comments is higher in politically-themed discussions than general discussions. If this question were studied, the current research would be a very useful basis for comparison.

Also not taken into consideration in this study was the effect of a thread's topology on the applicability of Godwin's Law. The Disqus commenting platform employs a "threaded" commenting system (cf. [12]), which was entirely ignored in the current analysis in favor of purely chronological ordering. Many questions could be asked about the incidence of Godwin matches among various other types of thread characteristics than simply length (e.g., comment depth, degree, etc.). Studying the distribution of "Godwin time" (the amount of time passed between a thread's first comment and its first Godwin match) instead of Godwin length may also be a more accurate representation of the data given the chronological ordering.

As indicated by the many questions raised above, it is clear that many opportunities exist for further study on Godwin's Law. Being the first of its kind, this study will serve as a basis for future research on the topic for the author and will hopefully be the same for others.

7 References

1. Anderson, A. A., Brossard, D., Scheufele, D. A., Xenos, M. A. and Ladwig, P. (2013), The "Nasty Effect:" Online Incivility and Risk Perceptions of Emerging Technologies. Journal of Computer-Mediated Communication. doi: 10.1111/jcc4.12009

2. LaBarre, Suzanne. "Why We're Shutting Off Our Comments." Web log post. Popular Science. 24 Sept. 2013. Web. http://www.popsci.com/science/article/2013-09/why-were-shutting-our-comments.

3. Godwin, Mike. "Meme, Counter-meme." Wired. Oct. 1994. Web. http://www.wired.com/wired/archive/2.10/godwin.if_pr.html.

4. Godwin, Mike. "I Seem To Be A Verb: 18 Years of Godwin's Law." Jewcy. 30 Apr. 2008. Web. http://www.jewcy.com/arts-and-culture/i_seem_be_verb_18_years_godwins_law.

5. Brogol. "En Finir Avec Le Point Godwin." La Politeia. 2 Oct. 2010. Web. http://lapoliteia.com/en-finir-avec-le-point-godwin-critique-de-la-loi-de-godwin/.

6. Greenwald, Glenn. "The Odiousness of the Distorted Godwin's Law." Salon. 1 July 2010. Web. http://www.salon.com/2010/07/01/godwin/.

7. Weigel, David. "Hands Off Hitler!" Reason.com. 14 July 2005. Web. http://reason.com/archives/2005/07/14/hands-off-hitler.

8. Mishne, Gilad, and Natalie Glance. "Leave a Reply: An Analysis of Weblog Comments." (2006). Web.http://www.ambuehler.ethz.ch/CDstore/www2006/www.blogpulse.com/www2006-workshop/papers/wwe2006-blogcomments.pdf

9. Kumar, Ravi, Mohammad Mahdian, and Mary McGlohon. "Dynamics of Conversations." Mahdian.org. Web. http://www.mahdian.org/threads.pdf.

10. Wang, Chunyan, Mao Ye, and Bernardo A. Huberman. "From User Comments to On-line Conversations." HP.com. 2012. Web. http://www.hpl.hp.com/research/scl/papers/comments/comments.pdf

11. Cristelli, Matthieu, Luciano Pietronero, and Michael Batty. "There Is More than a Power Law in Zipf." Scientific Reports 2 (2012). DOI:10.1038/srep00812. http://www.nature.com/srep/2012/121108/srep00812/full/srep00812.html

12. Gómez, Vincenc, Hilbert J. Kappen, Nelly Litvak, and Andreas Kaltenbrunner. "Modeling the Structure and Evolution of Online Discussion Cascades." 26 July 2012. http://www.quantware.ups-tlse.fr/complexnetworks2012/slides/kaltenbrunner.pdf