

A Persistence Paradox

Fang Wu and Bernardo A. Huberman
HP Laboratories
Palo Alto, CA 94304

November 12, 2009

Abstract

It is widely believed that persistence in most endeavors is key to their success. By its very nature, persistence enhances both the quality of an outcome and its probability of success because people are willing to endure failures before achieving a desired goal. To test this hypothesis on a massive scale, we studied the production histories and success dynamics of 10 million videos uploaded to a popular video website. Our results reveal that while the average quality of submissions does increase with the number of uploads, the more frequently an individual uploads content the less likely it is that it will reach a popularity threshold. These paradoxical findings, which hold both at the aggregate and individual levels, throw light on the act of production in the attention economy.

Anyone trying to seek attention to a new idea, scientific result or media creation faces the challenge of having it attended to by a sizable audience. This age old problem has become exacerbated with the advent of the web, which allows content to be easily published by millions of people while increasing the competition that it faces for the attention of users.¹ Worse, attention over content is distributed in a highly skewed fashion, implying that while a few items receive a lot of attention, most receive a negligible amount [8, 11, 16]. And yet people persistently upload content to social media sites, hoping for the highly unlikely outcome of reaching a wide audience. As in other endeavors, in the attention economy [3, 4, 5, 6] persistence is construed by content producers as a key to success.

In environments such as **YouTube**, the notion of success is not measured by monetary value but by the amount of attention each video receives [9]. Although **YouTube** serves more than 100 million video downloads per day on its site, neither the individual producers nor the website extract direct revenue from this business. What they gain instead is peoples' attention, which could in principle be transformed into revenue through advertisement or some other efforts. If we view **YouTube** as a producer of videos that outsources its production to millions of individual producers, then the customer life time value of each producer should be estimated by the total amount of attention generated by his or her uploaded videos. And while the motivations of individuals to upload videos are more diverse than those of the service provider, they nevertheless enjoy the attention they receive regardless of the number of downloads. For example, there are those who upload a video in order to share it with a number of family members and friends, and while getting their attention to their video is their intent, popularity in terms of a very high number of downloads is not. On the other hand, there are many others who do cherish high attention income in the form of large number of downloads for their videos, including those who attempt to make substantial profits through the **YouTube** partner program [13].

We performed our study by analyzing data from **YouTube**, the most popular video website on the Internet where users can upload, view, and share

¹This is because information-rich regimes are characterized by keen competition for the user's attention [3, 12].

video clips for free. The website’s interface displays a view count number right next to each video title which measures how many times the video has been watched since it was uploaded. As a video is watched by more users, its view count grows accordingly. If a video has eventually been viewed a substantial number of times, it is promoted to YouTube’s front page, where it can reach a larger audience and receive even more attention.

Our data set contains 9,896,784 videos submitted by 579,470 producers as of April 30, 2008. For each video we obtained its date stamp, the uploader’s id, and its view count as of April 30, 2008. While typically videos get viewed several thousand times, the top ones can receive more than one million views. Fig. 1 plots the average and 99% quantile of view counts for videos submitted in each week. A declining trend is evident, showing that older videos tend to receive more view counts than newer videos, simply because they have been on the website longer. Thus, while the absolute view count offers some quantitative measure of the “attractiveness” of each video, it only makes sense to compare the view counts of two videos submitted not far apart in time.

We used a nonparametric procedure to remove the time dependence of view counts in our data set. We defined the *popularity* of each video to be the quantile of its view count among all videos submitted in the same week. For example, a video with popularity 0.7 receives more view counts than 70% of the videos submitted in the same week. By mapping each video’s absolute view count into this relative popularity measure, statements such as “Alice’s fifth video is more popular than her fourth video” become meaningful.² All of our analysis is based on this relative popularity measure instead of absolute view count.

It is well known that in many information-rich economies including YouTube, the amount of attention received by each item follows a skewed distribution with a long tail [16]. In particular, it has been established that on YouTube the top 1% of videos are viewed several hundred more times than the bottom 50% combined, and about fifty times more often than the bot-

²Strictly speaking, we have made the implicit assumption that videos submitted close in time evolve according to some “universal” pattern, so that their view count quantile is relatively stable. This has been verified empirically [14, 16].




Figure 1: (a) Average view count of videos uploaded in different weeks. The horizontal axis represents the number of weeks from January 1, 2006. The vertical axis marks the average view count of videos uploaded in the same week. All videos' view counts are measured on April 30, 2008. As can be seen, older videos tend to receive more view counts than newer videos, because they have been on YouTube for a longer time. (b) 99% quantile of videos uploaded in the same week, again exhibiting a declining trend.

tom 80% combined [7]. Thus, it is important for YouTube to identify the top contributors whose videos generate vastly more attention than all others.

In what follows, we will say that a video is a *hit* if its view count exceeds a popularity threshold. For ease of presentation we will concentrate on a particular value of the threshold, in our case 1%, but our results are valid for a wide range of threshold values (from 0.1% to 10%). We also define a producer to be a *hit producer* if at least one of her uploaded videos crosses the threshold.

By definition 1% videos in our data-set were hits. 7.5% of the 579,470 producers, or 43,508 producers, uploaded at least one hit video before April 30, 2008 and hence belong to the set of hit producers. We observed that about 21.0% of all hit producers earned the hit label with their first video. Given that only 5.9% of all videos are first submissions, this number is impressively large. What is even more impressive is that about 34.1% of the hit producers actually produced a hit in their first week (i.e. at least one of the videos they uploaded in their first week received top 1% popularity within that week). Fig. 2 plots the distribution of the index and week of the hit producers' first hit. Again, one can see from the figure that a considerable proportion of hit producers produced a hit early on, with only a few submissions within the first few weeks. In particular, more than half of the hit producers achieved a hit while uploading the first 5 videos, and more than half of those achieved a hit within 6 weeks.

We next studied how the popularity of submitted videos changes over time. One might think that a producer's latter videos would tend to find more favor than earlier ones, especially when one considers the many learning advantages producers obtain from past submissions: improved video creation skills, better knowledge about the audiences' taste, better marketing techniques, etc.³ If this were true, YouTube should value its more experienced contributors more than newcomers.

To test this hypothesis we measured the popularity dynamics of each contributor's videos at the aggregate level. We calculated the average hit ratio for all videos with submission index k (i.e. the k 'th uploaded video)

³For a counterexample see [1].




Figure 2: (a) Fraction of hit producers whose first hit video is their k 'th submission. (b) Fraction of hit producers whose first hit was submitted in the k 'th week since their first submission.

and plotted it as a function of k in Fig. 3(a). Surprisingly, the figure exhibits a clear decreasing trend, indicating that on average older contributors are less likely to create hits than new contributors. What is more surprising is that when we plot the average rating of videos with different submission index, an increasing trend is observed. As Fig. 3(b) clearly illustrates, producers on average receive higher ratings for their later videos, while getting decreasing hit ratios with increasing number of submissions. If the quality of videos increases, why should they receive less attention?

Before we attempt to resolve this paradox, we rule out some possible explanations. First, the declining trend in Fig. 3(a) is not caused by an ever increasing number of contributors and submissions. This is because the threshold we picked to define a “hit” is set at a fixed quantile which is independent of the total number of videos. Second, the trend is not a simple artifact of producer heterogeneity. To see this, consider two types of contributors: “professionals” who produce high-quality videos with high hit ratios, p_H , which are constant over time, and “amateurs” who produce poor quality videos with low hit ratios, p_L , that are also constant over time.




Figure 3: (a) Average hit ratio of videos with different submission index. A producer’s first video has index 1, second video has index 2, and so on. (b) Average rating of videos with different submission index.

Suppose professionals produce at a lower rate than amateurs, so that at any fixed time videos with a very large submission index are most likely to be submitted by amateurs, while videos with a small submission index may come from both amateurs and professionals. Such a model could account for the declining trend in Fig. 3(a) without assuming any time dependence in the hit ratios. To show that this is not the case however, in Fig. 4 we plot the empirical hit ratio for the first k videos from those who submitted no less than k videos, for $k = 5$ and 10. As can be seen, the declining trend still persists.

We now turn our focus from the service provider (**YouTube**) to the actual content producers. In many cases the goal of individual producers seems to align with **YouTube**, which is to attain a large level of attention for their videos.

According to the dictionary, the word “persist” is defined as “to go on resolutely or stubbornly in spite of opposition, importunity, or warning.” In our setting this corresponds to those producers who keep uploading despite the fact that their previous uploads never achieved the top 1% popularity.




Figure 4: Average hit ratio of videos from those who submitted no less than (a) 5 videos and (b) 10 videos, as a function of the submission index.

Formally, let us define a producer’s *persistence level* to be the number of failures he/she is willing to endure before the first hit. Thus, a producer with persistence level k ($k \geq 0$) would upload no less than k videos before coming to a stop, and would refrain from uploading the $(k + 1)$ ’th video if all of his/her first k attempts fail. We also say that user A is more/less persistent than user B if A’s persistence level is higher/lower than B.

Notice that there are two reasons why a producer’s persistence level is not a directly observable quantity within our data-set. First, our data ends on April 30, 2008. If a producer uploaded k videos before April 30, 2008 and all of them failed, we do not know whether she would still upload a $(k + 1)$ ’th video after April 30, 2008. Second, if a producer failed k times and then got a first success, all we can infer is that her persistence level must be no less than k . In fact, if a producer was lucky enough to have succeeded in the first video, we cannot make any inferences as to her persistence level.

Although it is not possible to measure the persistence level of each individual producer, one can measure the *conditional hit probability* or *hazard function* of the whole population, which is defined as the conditional probability that the k ’th attempt is a hit given the first $k - 1$ are not. For our




Figure 5: (a) The conditional hit probability of YouTube producers. The blue line is the conditional winning probability of a lottery with 0.01 winning probability. (b) The red line plots the success probability of YouTube producers as a function of their persistence level. The blue line plots the success probability of a producer who participates in a lottery with 0.01 winning probability for each draw, as a function of her persistence level.

data this amounts to

$$h(k) = \frac{\text{number of producers whose first } k-1 \text{ videos are not hits and whose } k\text{'th video is a hit}}{\text{number of producers whose first } k-1 \text{ videos are not hits and still uploaded a } k\text{'th video}}. \quad (1)$$

Fig. 5(a) plots the conditional hit probability for YouTube producers. We see that $h(k)$ declines with the submission index, indicating that later submissions are less likely to be hits. This result is somewhat surprising for one might expect the contributor to learn from past experiences and improve the chance to succeed. As the data shows, this is not the case.

A possible explanation of this surprising result could stem from the often implicit, but crucial assumption, that improvements along the way are what make persistence the cause of success. Indeed, if the same mistake is repeated over and over again, what is the point of being persistent? In order to test

whether success due to persistence with improvement is more than a matter of luck we compare our results with the benchmark case of a lottery with constant winning probability. While someone with no budget constraints can try to win a lottery as often as possible, the number of attempts will neither improve nor worsen the winning probability of the *next* lottery purchase. In this sense the lottery can be regarded as a game of “pure luck”, in which persistence does not play a role in determining winning probability.

In order to compare our results from YouTube with the benchmark case of a lottery, we plot in Fig. 5(a) the conditional winning probability of a lottery with winning probability 0.01, where 0.01 is the fraction of hit videos in our study of YouTube. Naturally the conditional winning probability for such a lottery is always 0.01, independent of past history. As can be seen in the figure, the conditional hit probability for YouTube is *worse* than the lottery from the second video thereafter.

More insight into this result can be obtained by calculating the empirical hit probability $p(k)$ of a producer with persistence level k , i.e. the probability that the producer succeeds at least once before he/she quits. It is not hard to show that for the lottery case $p(k) = 1 - 0.99^k$, while for YouTube it is given by $p(k) = 1 - \prod_{i=1}^k (1 - h(i))$, where $h(k)$ is estimated by Eq. (1). These two functions are plotted in Fig. 5(b). As can be seen, the empirical hit probability for a YouTube producer is lower than that of a lottery buyer regardless of the persistence level. In particular, while the success probability of a lottery buyer with persistence level 100 is 63.4%, that of a YouTube producer with the same persistence level is only 22.6%. Thus, the benefit of persistence on YouTube is much worse for content producers than that of a lottery.

We point out that while this analysis of empirical conditional probabilities at the aggregate level is standard in the literature, it is not without its flaws. In particular, under certain circumstances the decline in the conditional hit probability may be due to population heterogeneity and not because there is any underlying time dynamics at the individual level [2]. Thus, the analysis should only be treated as a benchmark model which assumes that the population under study is homogeneous at large. This well-known “ruse of heterogeneity” [17] can be resolved by analyzing our data at a refined indi-

	rating increases	rating decreases
popularity increases	64,131	67,663
popularity decreases	82,412	68,540

Table 1: The change in popularity and average rating for all contributors who have submitted at least two videos.

vidual level, as we do now.

We want to test if each individual contributor follows the same dynamics as the whole population, i.e. her later videos tend to receive less popularity than the earlier ones, even though their quality does improve. We start our analysis with the first two submissions. For each contributor who has submitted no less than two videos, we measured whether the popularity (measured by the quantile of view count) and average rating of her first two videos increase or decrease. Since there are four possible outcomes, we counted the number of contributors falling into each category. The result is listed in Table 1. As one can see, the majority of contributors falls into the lower-left quadrant, i.e. the popularity of the uploads decreases even though their quality improves.

In order to test the robustness of this result, we repeated the measurement for all 105 ordered pairs $\{(i, j) : 1 \leq i < j \leq 15\}$ in addition to the first two. We found out that the lower-left quadrant dominates in all 105 tables, and that the sum of the bottom row (popularity decreases) is always larger than the sum of the top row (popularity increases). Also, for those who submitted no less than $2n$ videos, we measured whether the median popularity of their first n videos (denoted by p_{m1}) is greater or less than the median popularity of her second n videos (denoted by p_{m2}). We found that the majority (i.e. more than half) of contributors have $p_{m1} > p_{m2}$, for all $n = 1, \dots, 10$. From these facts it is safe to conclude that if we randomly draw a contributor from the population, it is more likely that his later videos will receive a lower popularity than his earlier ones.

Next we studied how each producer’s chance of producing a hit evolves with time. In particular, we examined whether it took each contributor more or fewer attempts to achieve a second hit than a first one. Our data

set contained 43,508 hit producers. Suppose a producer achieved her first hit after n_1 attempts (i.e. her n_1 'th video is her first hit). Furthermore, assume that she would receive a second hit after n_2 attempts, where n_2 may or may not be directly observable in our data set. If we consider her uploads after n_1 there are three possibilities: (1) She uploaded no less than n_1 videos after the first hit and did not achieve a second hit with those n_1 videos, in which case we know that $n_2 > n_1$. (2) She achieved a second hit with no more than n_1 additional videos, in which case we know that $n_2 \leq n_1$. (3) She uploaded less than n_1 videos after the first hit and still had not achieved a second hit, in which case the available information is not enough for us to deduce the relationship between n_1 and n_2 . The number of hit producers falling into these three categories is 23,620, 9,146 and 10,742, respectively. We see that category 1 is more than twice as large as category 2 and larger than category 2 and 3 combined. Thus, even if all producers in category 3 end up achieving a second hit in no more than n_1 attempts (which is highly improbable), overall it takes the majority of contributors more attempts to achieve a second hit than the first hit, if a second hit is achieved at all. In other words, on average it is harder for each contributor to achieve a second hit than a first hit.

In this paper we have shown that in a pure attention economy [3] the more frequently an individual uploads content the less likely it is that it will reach a success threshold. In contrast to the extensive literature on attention at the individual level [10], this persistence paradox is a collective effect rather than a single individual one, since the success we measured is a competitive outcome. Furthermore, this paradoxical result is compounded by the fact that the average quality of submissions does increase with the number of uploads, with the likelihood of success less than that of playing a lottery.

The paradox itself does not seem to have at present a definitive explanation. A possible one is that as producers submit several videos over time their novelty, and hence their appeal to a wide audience, tends to decrease. As an example, the all time no. 1 YouTube video, “ Evolution of Dance” by judsonlaipply has achieved more than 122 million views as of July 2009. It happens to be the author’s first video. Encouraged by the huge success the author released two subsequent videos “EOD2 Info” and “Stop Pain Start

Dancing” both on the same subject “dancing”, but neither of them received more than 1/4 million views.

Finally there is the behavioral issue of why people persist in light of the small odds of making it in the charts. A possible answer is that as in many other instances, individuals overestimate the true probabilities of winning when they are small, thus increasing their efforts while their chances of becoming successful decrease over time. This is similar to the longshot anomalies observed and discussed in the context of gambling [15]. While in those cases plausible explanations have been suggested, we do not have enough behavioral data from content providers to elucidate their motivations.

References

- [1] Audia, Pino G., Edwin A. Locke and Ken G. Smith (2000) The Paradox of Success: An Archival and a Laboratory Study of Strategic Persistence Following Radical Environmental Change. *The Academy of Management Journal*, Vol. 43, No. 5, pp. 837–853.
- [2] Fader, Peter S. and Bruce G. S. Hardie (2009) Customer-Base Valuation in a Contractual Setting: The Perils of Ignoring Heterogeneity. Forthcoming in *Marketing Science*.
- [3] Falkinger, J. (2007) Attention Economies. *Journal of Economic Theory*, 133:266–294.
- [4] Franck, G. (1999) Science Communication: A Vanity Fair. *Science*, 286:53–55.
- [5] Goldhaber, M. H. (1998) Attention Economics and the Net. *First Monday*, 2.
- [6] Hirshleifer, David A. and S. H. Teoh (2006) Limited Investor Attention and Stock Market Misreactions to Accounting Information. Working Paper.

- [7] Hancock, Denis (2009) Putting the YouTube Long Tail in Perspective. *Wikinomics*, March 4, 2009.
- [8] Huberman, Bernardo A. (2001) *The Laws of the Web*. MIT Press.
- [9] Huberman, Bernardo A., Daniel Romero and Fang Wu (2009) Crowdsourcing, Attention and Productivity. Forthcoming in *Journal of Information Sciences*.
- [10] Kahneman, D. (1973) *Attention and Effort*, Prentice-Hall.
- [11] Klamer, A. and H. P. Van Dalen (2002) Attention and the Art of Scientific Publishing. *Journal of Economic Methodology*, 9:289–315.
- [12] Simon, H. A. (1971) Designing Organizations for an Information Rich World. In M. Greenberger (ed.), *Computers, Communications and the Public Interest*, Johns Hopkins Press, pp. 38–52.
- [13] Stelter, B. (2008) YouTube Videos Pull in Real Money. *New York Times*, December 10, 2008.
- [14] Szabo, Gabor and Bernardo A. Huberman (2009) Predicting the Popularity of Online Content. Forthcoming in *Communications of the ACM*.
- [15] Thaler, Richard H. and William T. Ziemba (1988) Parimutuel Betting Markets: Race-tracks and Lotteries. *Journal of Economic Perspectives*, Vol. 2, No. 2, pp. 161–174.
- [16] Wu, Fang and Bernardo A. Huberman (2007) Novelty and Collective Attention. *Proceedings of the National Academy of Sciences*, 104:17599–17601.
- [17] Vaupel, J. W. and A. I. Yashin (1985) Heterogeneity’s Ruses: Some Surprising Effects of Selection on Population Dynamics. *The American Statistician*, 39, 176–185.