DOI: 10.18727/0722-6691/5040

The Deadline Flurry Formula

Felix Stoehr[1]

[1] ESO

When having to deliver work to a fixed deadline people often wait until the very last minute, in part, because they procrastinate. While procrastination has been studied extensively in the psychology literature, few direct measures of human behaviour leading up to a deadline exist. Here we use metadata from the ALMA proposal submission process over the last five years and find that collective human behaviour for submitting work before a deadline can be described spectacularly well by a simple "universal" law. We also analyse this behaviour as a function of several other factors, such as gender, age, proposal size, number of co-authors and the subsequent success of a submitted proposal.

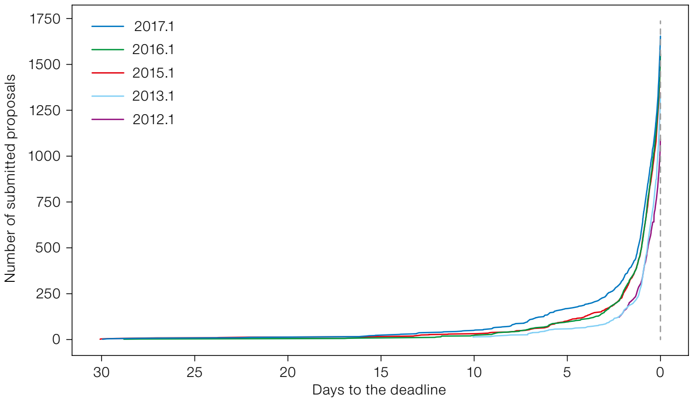

Figure 1. Evolution of the number of proposals submitted to ALMA as a function of the time remaining until the proposal deadline (since the first submission of the proposal) for the different ALMA proposal cycles (see legend). The vertical grey dashed line indicates the proposal deadline.

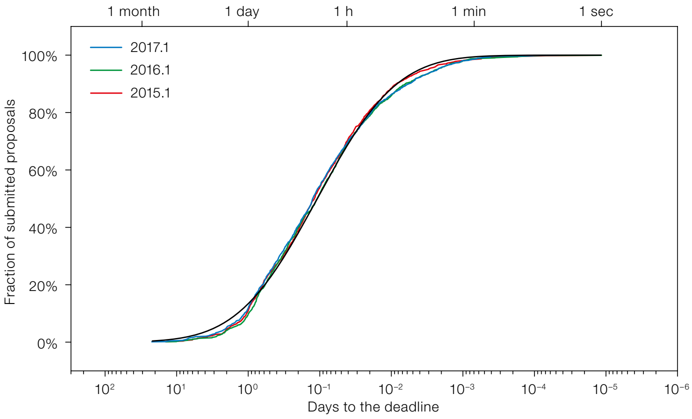

Figure 2. Fraction of submitted proposals as a function of log time to the deadline using the moments of last submission of the proposals for the last three ALMA proposal cycles (3–5). The black line shows the lognormal distribution with $\mu^* = 2\text{h}\ 36\text{m}$ and $\sigma^* = 0.1349$, explaining 99.87 % of the signal.

Introduction

It is deeply rooted in human nature that work which needs to be delivered by a fixed deadline is often delivered only at the very last moment. Together with factors like a large overall workload and the fact that the delivery itself is the culmination of a long period of work, one reason is certainly procrastination, i.e. the act of delaying a task that must be done. Procrastination was probably first mentioned by Hesiod around 700 BC (Evelyn-White, 1936). Steel (2006) compiled a large meta-study of the state of research on the nature of procrastination. Such psychological studies often focus on the causes (for example, task aversiveness, task delay, self-efficacy, active *vs.* passive procrastination) and the effects on performance. In most cases, these studies rely on questionnaires filled out by test persons (using, for example, the Tuckman Procrastination Scale [Tuckman, 1991]). The correlation between self-reported and actual procrastination, however, is relatively low (Tice & Baumeister, 1997).

In this study, we use the metadata from the Atacama Large Millimeter/Submillimeter Array (ALMA) observing proposal submission process since 2012. ALMA conducts essentially yearly calls for proposals, of one month's duration, with the particularity that users can modify their proposals as often as they wish before the deadline. The number of proposals received in the latest cycles is larger than that of any other single telescope proposal process worldwide.

The metadata are ideally suited to such an analysis. In particular the data are fully objective; the amount of work a Principal Investigator (PI) needs to deliver is fixed by the maximum length of the proposal of four pages (five pages in Cycle 5), which is strictly enforced; the proposals are received well distributed over all time zones; the proposal deadline is strict; additional metadata on the proposals are available (for example, whether the proposal is a student proposal and whether or not it was awarded observing time); and, finally, the incentive to work for a very significant amount of time on a proposal is very high. Indeed, with a single highly competitive call for ALMA proposals per year, the award of observing time can have an enormous positive impact on the career of an astronomer.

The precise timestamps of the proposal submissions not only allow us to measure the median time people submit before the deadline, but also allow us to study the entire cumulative proposal submission evolution. The only similar effort of which we are aware (Durakiewicz, 2016), is for proposal submission to the National Science Foundation (NSF).

Results

Figure 1 shows the cumulative number of submitted proposals for five ALMA proposal cycles as a function of the initial time they have been submitted to the ALMA Observing Tool. Given that proposal submission is a random process, one might have expected the curve to be a (cumulative) Gaussian distribution, where a few PIs submit early, a few very late and the bulk a certain time before the deadline. The actual distribution is, however, radically different. The submission rate is highest just before the deadline, with about 78 % of all proposals being submitted within the last 12 hours and 32 % submitted within the last hour.

Applying three transformations to these data — normalising, using the instant of the last time at which a PI has submitted a proposal, and using a log-time scale — we find a spectacular result (Figure 2). The proposal submission evolution curves fall nearly exactly on top of each other with a mean standard deviation from the average curve of only 0.02 %; however, note that the ultimate submission time was only available for the last three proposal cycles. More importantly, however, the evolution can be fitted by a simple Gaussian distribution. The black line shows the best-fit cumulative distribution function (CDF) of a Gaussian probability distribution function with mean μ^* = 2h 36m and standard deviation $\sigma^* = 0.1349$ using logarithmic time as random variable. Mathematically this function is the lognormal distribution with

$$\text{Lognormal CDF}\,(t|\mu^*,\sigma^*) = \frac{1}{2}\,erfc\left(-\frac{\ln(t)-\ln(\mu^*)}{\ln(\sigma^*)\sqrt{2}}\right)$$

where t denotes the time to the deadline in days and $\mu^* = e^{\mu}$ and $\sigma^* = e^{\sigma}$ (all dimensionless, expressed in units of days) are the median and the multiplicative standard deviation, respectively, and *erfc* is the complementary error function. μ^* and σ^* are also the mean and standard deviation of the normal distribution with $\chi = \ln(t)$. We will refer to this distribution as the deadline flurry formula (DFF).

This very simple model fits the average evolution of the three proposal cycles extremely well, explaining 99.87 % of the signal (coefficient of determination, r^2) and is valid over at least five decades in time to the deadline. Moreover, the DFF is “universal” in the sense that it can accurately describe the global evolution, can also describe the evolution of subsets of the data (see below) and requires only two parameters which have well known meaning.

Since the evolution of proposal submission over the three submission cycles is almost identical, the statistics suggest that the moment at which the possible improvement of a proposal in the remaining time is perceived not to be worth the effort is a very precise and distinct moment for all proposers, thus leading to the very small scatter. This finding is even more remarkable since there is a much larger scatter in the evolution when using the times of the first proposal submission. Even if PIs start earlier or later, due to external factors like attendance at ALMA proposal preparation community days, the moment at which more work is not considered worth it any more is an intrinsic value for every person. This can be seen as indirect support of theories of motivation, like the Time Motivational Theory (TMT; Steel & König, 2006).

Analysis of sub-samples

Our statistical sample is large enough that we can split it in various ways and study the submission behaviour of subsets; see Figures 3 and 4 as well as Table 1. We use bootstrapping to estimate the 95 % confidence interval of μ^* to be typically +/– 12 minutes, which allows us to assess the statistical significance of our results. We find that PIs from East Asia (EA) submit later (μ^* = 1h 52m) than their colleagues from Europe (EU) (2h 31m) or North America (NA) (3h 41m). As the proposal deadline, which is always set at 15:00 UT, corresponds to very different local times in the different regions (i.e., very late evening for EA, late afternoon for EU and (very) early morning for NA), our data do not allow us to determine whether or not the effect is due only to the different local deadline times or also to actual cultural differences. The fact that there is indeed an influence due to the local deadline time can clearly be seen in Figure 3. The short, flatter parts of the NA and EU curves correspond roughly to midnight until 09:00 local time.

In accord with earlier studies (Rotenstein et al., 2009; Kim & Seo, 2015), we find that PIs of proposals that were awarded observing time, and thus are those with the best performance (roughly 20 % of all proposals), typically submitted earlier (2h 57m). We also find that, in all cases, PI groups which submit the final version later also have started later (based on the time of first submission).

In the literature, stronger procrastination is found to be correlated with younger age (for example, Steel, 2006 and references therein); this coincides with our findings as proposals marked “student project” form the group that submits latest (1h 42m). However, it should be noted that the tag does not necessarily mean that the proposal was submitted by a PhD student as any proposal would bear that tag if the data are to be used for a student project. We use the Python package *sexmachine*[1] to estimate gender from first names; while the algorithm is less effective for EA we find similar fractions of female PIs among the different regions and so judge the subsequent analysis to be reflective of the missing or incorrectly identified population. We find that female PIs submit slightly later than their male colleagues, which is in contrast with earlier findings based on self-reporting of the test persons (for example, Mandap, 2016). Our finding is true globally, but also for each of the three regions separately, thus excluding any possible cultural influence. However, it should be noted that there may be multiple factors at play as, given the demographics of astronomers, the proportion of senior proposers among the male PIs is likely to be higher than that among female PIs.

Finally, we split the full sample into two halves by the amount of observing time requested, as well as by the number of co-authors, and find in both cases that the difference in μ^* between the two halves is larger than the 95 % confidence limit (see Table 1). It seems plausible that it takes more time to agree on and finalise a proposal if a larger number of co-authors is involved (2h 24m *vs.* 2h 52m) and that it takes more time to finish a complicated proposal asking for a lot of observing time (2h 27m *vs.* 2h 48m).

Durakiewicz (2016) studied the submission evolution of proposals to the NSF, finding a good fit using a modified hyperbolic function. The fitted function has only one fixed parameter, the length of the submission window *D* (days). For *D* = 30, as in our case, their function has an equivalent DFF μ^*-value of 21h 2m for the (final) submission, which is much larger than the 2h 36m of ALMA proposals (it is even larger than the value of ALMA’s first submissions) and thus not a good fit to

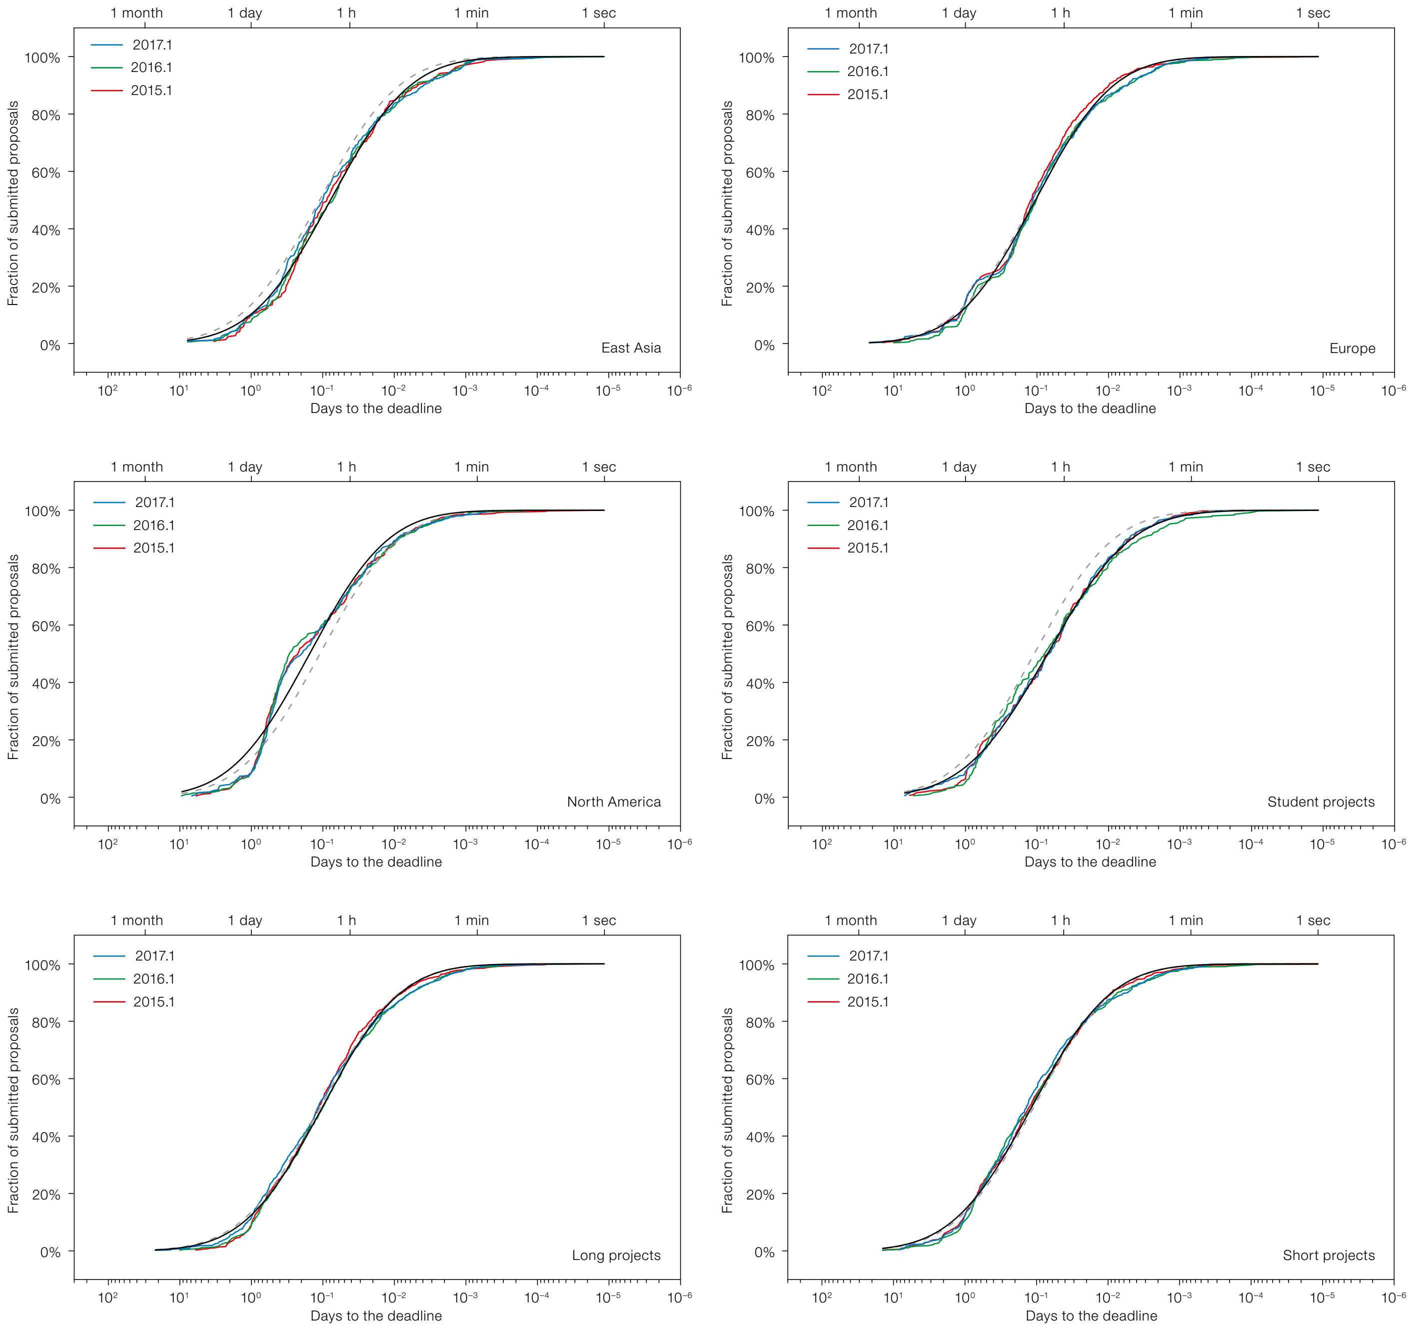

Figure 3. Distributions of the (last) proposal submission times of (from top left to bottom right) East Asia, Europe, North America, student projects, long projects and short projects. The best DFF fit is shown with a black solid line and, for reference, the best DFF fit for the global distribution is shown as a grey dashed line. In the North American PI plot, the effect of the early-morning proposal deadline in North America can be seen. This curve is responsible for the slight deviations of the global curve from the lognormal distribution.

our data. The very different timescales are probably due to the very different proposal processes. Indeed, NSF proposals are not submitted by the PIs directly but by the Sponsored Research office which also has to approve the requested budget.

Context

The study of proposal submission behaviour to a deadline in this work fits into a vast body of scientific quantities that follow a lognormal distribution: from the latency periods of diseases, to the amount of rainfall, the number of words in sentences, the age of marriage and the ratio of income to the size of people (Limpert, Staehl & Abbt, 2001).

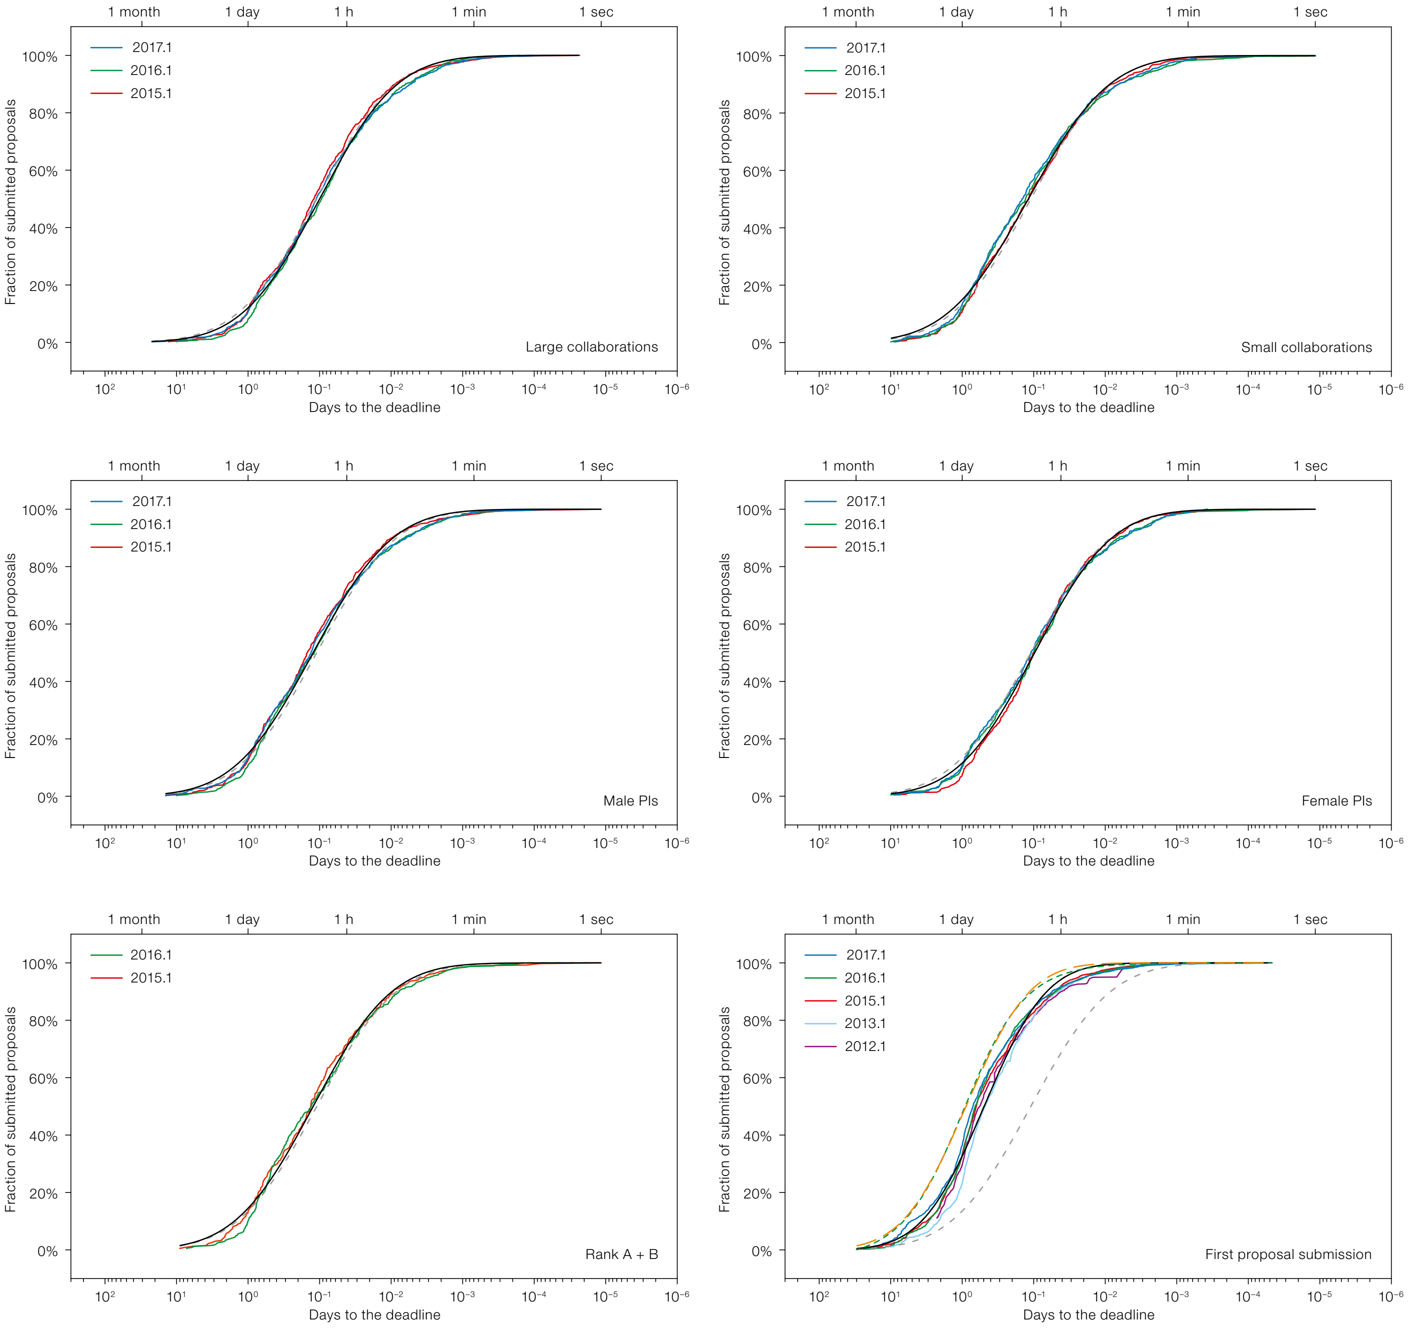

Figure 4. As Figure 3 for the sub groups: large collaborations; small collaborations; male PIs; female PIs; and successful proposals (rank A and B). The best DFF fit is shown with a black solid line and, for reference, the best DFF fit for the global distribution is shown as a grey dashed line. The lower right plot is the time of first proposal submission over all cycles; the modified hyperbolic with *D* = 30 days proposed by Durakiewicz (2016) has also been added (green dashed line) and fitted with a DFF (orange dashed line). In the plot of successful proposals, only two cycles are shown, as the accepted proposals of the last cycle (2017.1) were not known at the time of writing.

Lognormal distributions arise naturally in random processes when the random variable is a product of random variables; while the sum of the faces when rolling a number of dice repeatedly results in a normal distribution, their product will be lognormally distributed. The self-regulatory behaviour of human motivation is often modelled as the product of influencing factors (for example, in the TMT), so a lognormal distribution of an ensemble is a direct consequence. On account of the above considerations, and because there is no minimum time limit imposed (Mitzemmacher, 2004), we chose the lognormal distribution over multi-power-law distributions, noting that even four-parameter multi-power-law models have been suggested in the literature (Reed & Jorgensen, 2004).

	μ^* (days)		σ^* (days)	1-σ^* interval	2-σ^* interval
Last	0.1086	(2h 36m)	0.1349	19h 19m–21m 5s	5d 23h–2m 50s
EA	0.0779	(1h 52m)	0.1334	14h 40m–14m 57s	4d 9m–1m 59s
EU	0.1051	(2h 31m)	0.1399	18h 1m–21m 9s	5d 8h–2m 57s
NA	0.1535	(3h 41m)	0.1369	1d 2h–30m 16s	8d 4h–4m 8s
Student project	0.0715	(1h 42m)	0.1215	14h 7m–12m 30s	4d 20h–1m 31s
Long projects	0.1023	(2h 27m)	0.1382	17h 45m–20m 20s	5d 8h–2m 48s
Short projects	0.1173	(2h 48m)	0.1333	21h 7m–22m 31s	6d 14h–3m 0s
Large collaborations	0.1002	(2h 24m)	0.1409	17h 4m–20m 20s	5d 1h–2m 51s
Small collaborations	0.1195	(2h 52m)	0.1291	22h 12m–22m 13s	7d 3h–2m 52s
Successful	0.1230	(2h 57m)	0.1381	21h 22m–24m 27s	6d 10h–3m 22s
Female	0.0974	(2h 20m)	0.1432	16h 19m–20m 5s	4d 17h–2m 52s
Male	0.1233	(2h 57m)	0.1343	22h 1m–23m 50s	6d 19h–3m 12s
First submission	0.4931	(11h 50m)	0.2073	2d 9h–2h 27m	11d 11h–30m 30s
NSF	0.8769	(21h 2m)	0.2013	4d 8h–4h 14m	21d 15h–51m 9s

Table 1. μ^* (median) and σ^* (multiplicative standard deviation) values for the various PI sub-groups. The one- and two-σ^* boundaries are very asymmetric owing to the log-time distribution and are given by $[\mu^*/\sigma^* .. \mu^* \sigma^*]$ and $[\mu^*/\sigma^{*2} .. \mu^* \sigma^{*2}]$, respectively. The 95 % confidence interval on μ^* is of the order of +/– 12 minutes.

For future research, it would be interesting to analyse the submission behaviour towards a fixed deadline for different amounts of work to be submitted and for longer and shorter time intervals in which the proposals can be submitted. This might allow prediction of the two free parameters of the DFF given that the behaviour of the PIs appears to be so universal and time-invariant. Practically, our analysis could also be useful for the design of future proposal submission systems, like the new system which is currently being developed for ESO.

As the time just before the proposal deadline is often associated with stress for ALMA PIs, they should be encouraged to make even more use of the possibility of submitting a draft version of the proposal well before the deadline.

Conclusions

We have analysed the behaviour of people submitting work before a deadline using metadata from the ALMA proposal submission process since 2011. We find that the proposal submission evolution is nearly perfectly identical over the years, which suggests it is linked to an intrinsic property of human nature. We also find that this evolution can be extremely well described by a simple lognormal distribution with the time-to-deadline as random variable; the deadline flurry formula.

With the large sample to hand we have also studied the very small deviations from the general submission behaviour for subgroups of PIs. While our result that male PIs submit slightly earlier than female PIs is opposite to the findings published in the literature, we cannot exclude that it is at least in part caused by different age demographics within these subgroups. Our results with respect to PI age and the success of a proposal are consistent with earlier findings.

Acknowledgements

We are indebted to the creators of the ALMA Observing Tool for having the foresight to record the precise timestamps of the proposal submissions. We thank Gautier Mathys, Tomasz Durakiewicz, Wolfgang Kerzendorf, Martin Zwaan, John Carpenter and Gaitee Hussain for very helpful comments. ALMA is a partnership of ESO (representing its Member States), NSF (USA) and NINS (Japan), together with NRC (Canada), MOST and ASIAA (Taiwan), and KASI (Republic of Korea), in cooperation with the Republic of Chile. The Joint ALMA Observatory is operated by ESO, AUI/NRAO and NAOJ.

References

Durakiewicz, T. 2016, Physics Today, 69, 11
Evelyn-White, H. G. 1936, *Hesiod, the Homeric Hymns, and Homerica*, 3rd rev. ed., (Cambridge, USA: Loeb Classical Library), lines 405–413
Kim, K. & Seo, E. 2015, Personality and Individual Differences, 82, 26
Limpert, E., Stahel, W. A. & Abbt, M. 2001, BioScience, 51, 341
Madap, C. 2016, International Journal of Education and Research, 4, No. 4
Mitzenmacher, M. 2004, Internet Mathematics, 1, 485
Reed, W. J. & Jorgensen, M. 2004, Commun. Stat.-Theory & Methods, 33, 1733
Rotenstein, A., Davis, H. Z. & Tatum, L. 2009, Journal of Accounting Education, 27, 223
Steel, P. & König, C. J. 2006, Academy of Management Review, 31, 889
Tice, D. M. & Baumeister, R. F. 1997, Psychological Science, 8
Tuckman, B. W. 1999, Educational and Psychological Measurement, 51, 473

Links

[1] Python name gender resolver: https://pypi.python.org/pypi/SexMachine