IMPROVED RADAR CROSS-SECTION "TARGET TYPING" FOR SPACECRAFT

M.D. Hejduk*

Statistical histories of radar cross-section (RCS) data of satellites are used widely, both for the main application of tracking radar optimization and for tangential uses such as satellite characterization, size estimation, and observation correlation. The Swerling target typing models, however, have not evolved since 1954, when Peter Swerling's two chi-squared PDFs, based on the examination of aircraft data, were first proposed. In an earlier restricted paper only recently published, Swerling repudiated his previous position and advocated a broader set of types, which included additional chi-squared forms and the lognormal distribution. Other recent studies also support the use of lognormal distributions, suggesting that a systematic investigation of target typing for spacecraft is in order.

Hit-level RCS histories from the Eglin FPS-85 spacetrack radar on over 8000 objects were obtained and, after appropriate filtering, were examined with the empirical distribution function (EDF) goodness-of-fit technique to determine conformity to the classical two Swerling target type models, an expanded set of chi-squared models including Swerling's own recommendations and other promising variants, and the lognormal distribution. The results revealed that, contrary to the author's expectations, the traditional Swerling types fared reasonably well, adequately representing some 35% of the objects; and the addition of the expanded chi-squared types increased this number to 46%. The lognormal distribution can provide about the same marginal gain, accounting uniquely for about 10% of the objects tested; but its solo performance against the entire dataset reached only the 25% figure, indicating that it in no way is a good replacement candidate for the traditional Swerling types. Future research should thus focus on the some 42% of objects that cannot be represented by any one the chi-squared or lognormal distributions to determine whether a new canonical distribution type should be introduced.

INTRODUCTION AND HISTORY

There are a number of reasons to wish to know the probability density function (PDF) of the radar cross-section (RCS) of a space object. The best known and most common application is to determine the amount of power that a radar must expend in tracking the object in order to achieve a desired probability of detection. Such a consideration is especially import for phased-array radars, which can track many targets simultaneously; applying the correct amount of power allows all attempted targets to be acquired and maximizes the number of targets that can be simulta-

* Adjunct scientist, SRA International, 1795 Jet Wing Drive, #160, Colorado Springs, CO 80916.

neously kept under track. However, in addition to this more classical application, additional uses have also arisen. Size estimates of satellites from signature data (such as RCS) are becoming increasingly common, and to improve fidelity these estimates require consideration of the signature's PDF. Additionally, the signature PDF is being examined for satellite characterization purposes, to include the determination of satellite object type (*e.g.*, payload, rocket body, debris) and the inference of a satellite configuration change by observing a change in the PDF. An understanding of the canonical RCS PDF types—if indeed canonical types truly exist—is an important prerequisite to these applications.

In 1954 Peter Swerling completed a RAND Corporation report (which in 1960 was published as an IRE article) that, as part of outlining probability-of-detection calculation approaches for tracking radars, addressed the question of the expected RCS PDFs that these radars would encounter.[1,2] In the absence of significant observational data, Swerling deployed the following *a priori* reasoning to develop canonical RCS PDF forms:

1) A target type that is a collection of small scatterers of approximately equal areas will create a Rayleigh scattering pattern and thus produce a RCS PDF that follows a Rayleigh distribution, the familiar exponential pattern. Rather than describe it as an exponential decay, for reasons that will be clearer below Swerling chose to frame it as a chi-squared distribution with two degrees of freedom (which is the same as a gamma distribution with a shape parameter of 1). He further chose to make a distinction between slowly-changing cases in which there was RCS correlation pulse to pulse and quickly-changing cases in which correlation was observed only scan to scan. These two types—in which an exponential RCS PDF is expected and the additional difference is only in the expected correlation between pulses—constitute the Swerling I and II target fluctuation models.

2) A target that is a combination of large and small scatterers shows a combination of Rayleigh scattering and more systematic effects. For this case Swerling proposed a chi-squared distribution with four degrees of freedom: an evolution of the Rayleigh case but with the bounded "hump" appearance, not dissimilar to the typical PDF of a lognormal distribution. There was no specific mathematical reason to promote this particular distribution as the natural representative for this case other than the heuristic argument of an "evolved" chi-squared distribution. Swerling applied the same additional distinction between pulse-to-pulse versus scan-to-scan correlation, and this particular case became the Swerling III and IV models.

3) A target that shows no RCS dependence on target aspect, such as a calibration sphere, constitutes a non-fluctuating case and yields a symmetric PDF (sometimes called a Swerling 0 or Swerling V case). To remain consistent with the above paradigm, Swerling described this case as a chi-squared distribution with infinite degrees of freedom, which reduces to the standard normal distribution.

This collection of target types, which actually comprises only three different PDFs, quickly became canonical in radar theory and design for nearly all radar applications, despite its development in the 1950's when nearly all radar applications were for aircraft and ship tracking. To this author's knowledge, it has never been rigorously verified for spacecraft. But perhaps most injurious to its canonical status is Swerling's own "repudiation" of the adequacy of his two models and pursuit of additional target types in order to assemble a more complete portfolio. These particular developments have passed largely unobserved by the discipline, primarily due to their being documented in private reports and restricted conference proceedings before finally appearing in a trade journal in 1997.[3,4,5*] In these subsequent treatments Swerling recognized the failure of the two Swerling models to account for all of the targets encountered. Among his proposed remediations are the adoption of an additional canonical distribution of a chi-squared distribution with one degree of freedom (labeled the "Weinstock" distribution after the author of a 1956 unpublished dissertation that argued for it),[7] the outlining of a procedure for fits of object-by-object RCS empirical PDFs to tailored chi-squared distributions, and—significantly—consideration of the use of the lognormal distribution to model target types not well served by the chi-squared family. In this last case, he focuses on a procedure to distinguish when lognormal PDFs can be acceptably modeled by curves in the chi-squared family and when they cannot; but among his conclusions is the realization that there is a place for the lognormal distribution in fluctuating target typing. This contention has received some support from other studies, among them one of radar tracking data of mechanized infantry tanks[8] and one of a limited study of distilled radar tracking data on spacecraft.[9]

Two general realizations thus emerge from the above discussion. First, it may well be necessary to expand the canonical target type repertoire beyond the two main models given in Swerling's 1954 report; and second, in addition to merely a proliferation of additional chi-squared PDFs the lognormal distribution should also be considered as a type that may provide better matching to certain classes of targets. A specialized investigation is thus required two consider these two issues in the context of tracking data on spacecraft, a broad and considerable application of radar technology for which the particular issue of target typing has never been fully investigated. An additional consideration that would be helpful for such a study effort to address is whether any *a priori* knowledge about the target, such as its spacecraft object type (payload, rocket body, debris), can aid in the initial selection of this increasing set of canonical target types.

EVALUATION DATASET

The present investigation of the issues outlined in the preceding paragraph is made possible by the availability of an expansive dataset of individual radar hit data on spacecraft. The Eglin FPS-85 radar, located at Eglin AFB near Fort Walton Beach, FL, is a high-capacity phased-array radar that can track as many as 7000 space objects per day in a variety of orbits. Six months of individual radar hit data was made available for the present study. Hit data were winnowed by further requirements that the assigned satellite tag be properly assigned and that each hit participate in a metric observation formation process with a favorable resultant covariance. Object RCS hit histories were used only if more than 1000 hit data points were available for that object. The application of these additional criteria produced 8605 objects' worth of hit histories for the present study, with a total data volume of some 20 million radar hits. Throughout the study, the dataset was often divided by object type: active and inactive payloads (1828), rocket bodies (1097), and debris objects (5302).

TARGET TYPES DEFINED

Swerling chose to define target types with reference to the chi-squared distribution, but its more general form is the three-parameter gamma distribution, given below:

* A very helpful literature survey on the issue of radar target typing is provided in Reference 6.

$$f(x;\gamma,\beta,m)=\frac{1}{\beta^m\Gamma(m)}\left(x-\gamma\right)^{m-1}\exp\left(\frac{-(x-\gamma)}{\beta}\right) \qquad (1)$$

In this formulation m is the shape parameter, which governs the overall appearance of the PDF and thus the type of gamma distribution one is employing. If m is an integer, the doubling of m indicates the number of degrees of freedom of the distribution if one wishes to consider it to be a chi-squared distribution. β is the scale parameter, for which the maximum likelihood estimator (MLE) is m divided by the sample mean. γ is the location parameter for this distribution and is set to zero for simplicity of parameter estimation.[*] For a typical target type application, m would be chosen *a priori* or from the examination of the sample under consideration, β would be estimated using the MLE approach given above, and the properly-scaled RCS PDF would thus be available for subsequent calculations or applications.

Which "additional" gamma target types should be included in the present broader investigation of target typing? One way to bound the problem is to examine each object's hit history and estimate the m-parameter for a two-parameter gamma distribution. The ability to estimate the m-parameter does not guarantee that the dataset will actually be well-represented by a gamma distribution with that shape parameter; but it does mean that if it indeed is well-represented by a gamma distribution, it will be a gamma distribution with a shape parameter close to the one estimated.

At least the above would be true if the m-parameter could be reliably estimated, but one problem that plagues the use of the gamma distribution is difficulties in the efficient and reliable estimation of parameters. Moment estimators (MM) are the time-honored method of parameter estimation introduced by Karl Pearson, who included the gamma distribution as part of his Type III collection of frequency distributions. They are calculated by equating the distribution moments to estimators of the sample moments; but since sampling errors can often be magnified in the third sample moment (because of this moment's requiring the calculation of the cube of the difference between each data point and the sample mean), these large errors can be propagated into the estimated distribution parameters.[10] Maximum likelihood estimation (MLE) contains in principle many of the general virtues of maximum likelihood, but in the particular case of the gamma distribution it miscarries of necessity when the true m is less than 1 and will generally produce instabilities if the estimated m is simply near 1. Some authors recommend avoiding MLE unless the true m is believed to be greater than a much larger value, such as 2.5.[11] The Cohen-Whitten modified moment technique (MME), which modifies the moment estimators through the use of the first-order statistic, is also available, although like MLE it is subject to some restrictions in its application.[10] One advantage to the use of the two-parameter lognormal distribution, and thus yet another reason to determine whether it could serve as a wholesale replacement of the chi-squared family, is that parameter estimation is stable and greatly simplified: after a change of variable to re-express the dataset in logarithm space (dBsm), one simply computes the sample mean and variance for a normal distribution. If the RCS PDFs of many objects can be adequately

---

[*] The proper setting of the location parameter of the gamma distribution is a vexing problem that greatly complicates parameter estimation, hence the desire simply to render it as zero. This simplification becomes less realistic as the size of the object increases and advances the candidacy of the lognormal distribution in that, once the change of variable is made, parameter estimation of the mean and variance of the resulting (presumed) normal distribution is straightforward and unproblematic.

represented by the lognormal distribution, then this may be an overall superior paradigm. These two curve families can often produce very similar behavior, as shown in Figure 1 below.

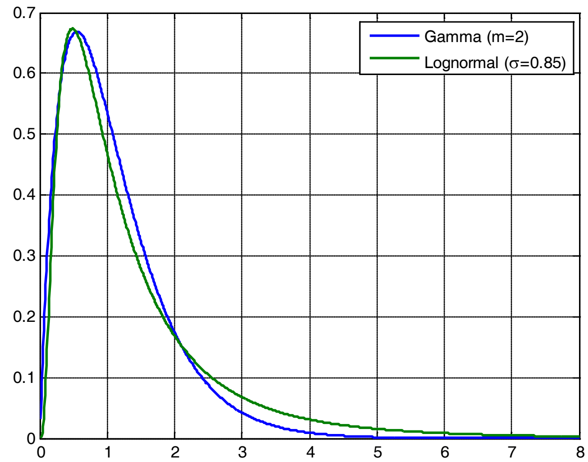

Figure 1. Example Comparison between Gamma and Lognormal Distributions

Figure 2 below gives a CDF plot of the estimated *m* for all 8600 objects in the present study, estimated by the three approaches described above. The propensity of the MM approach to produce very large *m*-estimates is obvious enough, and the disagreement between MLE and MME for the smaller values of *m* is apparent also. However, the two more reliable approaches (MLE and MME) both show that about 90% of the cases are adequately covered by shape parameters smaller than 3. To cover this range and to include both the canonical forms and the additional ones specifically called for by Swerling in his later work, a consideration of gamma PDFs in *m*-increments of 0.5, from a range of 0.5 to 2.5, should be adequate.

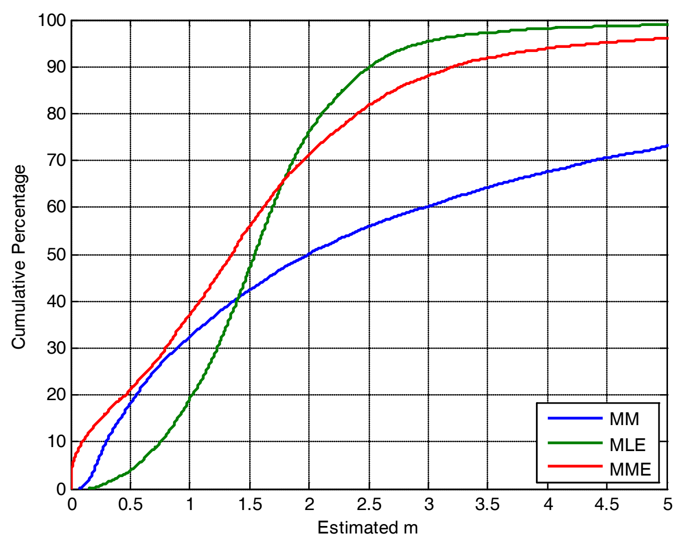

Figure 2. CDF of Estimates of Gamma-distribution m-Parameter

One wishes to include the lognormal distribution in the survey in a manner that can determine with some definiteness whether it should play a role in target typing for spacecraft. The most generous way to do this is to allow the analytical software to estimate the sample mean and variance of each data sample in attempting to fit a normal distribution (in log space) to each. This method of comparison in a way favors the lognormal distribution because it allows two parameters to be estimated, whereas the gamma distribution target typing allows only one to be estimated (because *m* is given *a priori* for each target type). However, if the lognormal distribution is compared to the entire battery of gamma distributions (shape parameters from 0.5 to 2.5), then the comparison is reasonably unbiased if it is believed that this distribution span truly covers most of the encountered target types. Table 1 below gives a summary of the different types of distributions thus to be considered in the present study, and Figure 3 (also below) shows representative plots of the six distributions under consideration. The two Swerling canonical types are shown with dashed lines.

Table 1. Distributions to be Investigated

Abbreviation	Type	Common Name	*m*-value
S0.5	gamma	Weinstock	0.5
S1	gamma	Swerling I, II	1
S1.5	gamma	none	1.5
S2	gamma	Swerling III, IV	2
S2.5	gamma	none	2.5
Log	lognormal	none	n/a

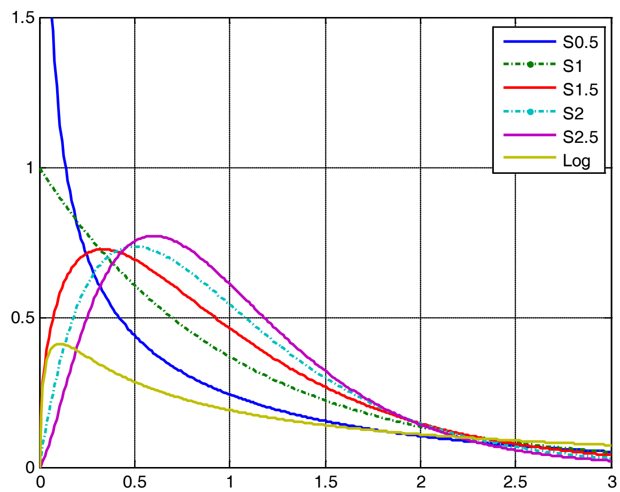

Figure 3. Representative CDFs of the Six Target Types under Investigation

GOODNESS-OF-FIT TESTING APPROACH

Most goodness-of-fit approaches are based on what is called a "weak-hypothesis" test in that one sets the null hypothesis to the desired outcome—that the sample conform to the hypothesized distribution—and rejects this only if goodness-of-fit testing assesses its likelihood to be below a threshold significance level. For example, suppose one wishes to establish whether a certain data set conforms to a Gaussian distribution. The procedure is to begin with the assumption that it in fact does conform to a Gaussian distribution. A goodness-of-fit test is then run to determine the likelihood that the dataset, with its particular statistical properties, could have been drawn from a population that conformed to a Gaussian distribution. If this likelihood is below a certain significance level (5% is a typical figure), then the null hypothesis is rejected, forcing the conclusion that the sample is not in fact Gaussian. The weak-hypothesis approach is often criticized as being too permissive; in the present case, a 6% chance that the sample could have been drawn from a Gaussian population would result in the retention of the null hypothesis and thus the conclusion that the sample is in fact Gaussian—not a conclusion of overwhelming strength. However, the goal of such testing often is not to establish definitively the parent population of a given sample but simply an adequate statistical model to describe the sample. This is in fact the purpose for the present application, and to this end weak-hypothesis testing is quite adequate. Furthermore, it is often used because there is really no other viable procedure for assessing goodness-of-fit in the general case.

Many specialty goodness-of-fit tests have been developed for particular distributions, especially the normal (Gaussian), exponential, and uniform. In the present case, it is necessary to test for conformity to both gamma and lognormal distributions; and as there are not (to this author's knowledge) any specialty tests for the gamma distribution, it is best to select a general technique that can be applied to both distributions, allowing for a more direct comparison of results. There are two mainstream techniques suitable for this: chi-squared and empirical distribution function (EDF) approaches. The Pearson chi-squared approach has a grand history in statistical inference; the earliest goodness-of-fit test, it has proven both reliable and versatile over the hundred years of its use. However, since the middle of this last century it has come under increased scrutiny and subjected to numerous analyses of its power; and these investigations have concluded that in nearly every case it is outperformed by specialty tests and alternative techniques, such as the EDF.[12] Its weakness is due primarily to the fact that continuous data must be divided up into "bins" in order to run the test, and this division masks additional information about the distribution that other techniques can exploit. If the data sample is already divided into discrete bins, then use of the chi-squared approach is often desirable; but if the data are continuous, a significant arbitrariness in introduced by the imposed binning—altering the resultant significance level by factors of 2 or 3 is not uncommon. Thus, for the case of continuous data, the literature strongly recommends the choice of a specialty test or EDF technique. Since the EDF approach can accommodate both the gamma and lognormal distributions, it is selected here as the preferred approach.

The EDF technique proceeds by computing a cumulative distribution function (CDF) empirically from the sample data and comparing this to the CDF of the hypothesized distribution. The two curves will never match exactly, but EDF techniques establish methods for evaluating the differences between these two curves and transforming these evaluations into significance levels, which give the likelihood that the sample is drawn from the hypothesized distribution. The graph below shows an empirical CDF (ECDF) constructed from 2000 hit RCS values for a randomlychosen satellite on the same graph as the CDF for a normal distribution with mean and variance estimated from the 2000-hit sample. The y-axis is cumulative probability.

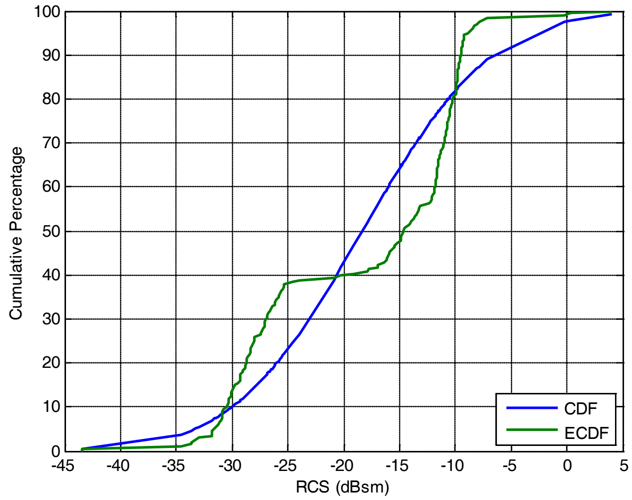

Figure 4. Comparison Plot of ECDF and CDF of Estimated (normal) Distribution

Two general types of EDF difference-quantification approaches exist: the *supremum* set, which examine the maximum horizontal and vertical difference between the two curves; and the *quadratic* set, which sum the squares of the differences over the entire curve set. While both are widely used, the quadratic set is generally considered the more powerful and was the one selected for the present analysis. There are three sub-types of quadratic analysis, all three of which are described by the general summation formula

$$Q = n\int_{-\infty}^{\infty}\left[F_n(x) - F(x)\right]^2 \psi(x)dx \tag{2}$$

and differ only in the weighting function ψ that is applied. The Cramér – von Mises statistic is the simplest:

$$\psi(x) = 1 \tag{3}$$

setting ψ simply to unity. The Anderson-Darling is the most complex, prescribing a function that weights data in the tails of the distribution more heavily than those nearer the center:

$$\psi(x) = \{F(x)[1 - F(x)]\}^{-1} \tag{4}$$

The Watson statistic is a modification of the Cramér – von Mises statistic to allow evaluation of points on a circle and is not relevant to the present investigation. The Anderson-Darling is generally considered the most powerful of the three, but it also the most heavily affected by outliers. After performing some sensitivity runs with RCS data, it was determined that the Anderson-Darling statistic was perhaps too exclusive for the present purpose and that a change to the Cramér – von Mises statistic was warranted.

The general procedure for testing a data sample using the EDF method is as follows. First, a hypothesized distribution is chosen; this choice is generally based on *a priori* conviction or graphical analysis. Next, parameters are estimated for the distribution. In the present situation, the five gamma distributions each has a fixed *m*-value, so the only parameter remaining to be estimated is the scale parameter β; for the lognormal distribution, both mean and variance (of the transformed dataset) are to be estimated. Once parameters are known, the CDF for the hypothesized distribution can be constructed. After this, the EDF for the sample, based only on the sample data, can also be constructed and compared to the CDF of the hypothesized distribution, with differences assessed using equation (2) and the expression for ψ appropriate to the test statistic being calculated (in the present case, which will use the Cramér – von Mises statistic, ψ will equal unity). Once the test statistic value is known, tables of significance points for this statistic are calculated to transform the statistic value into a probability that the sample was derived from the hypothesized distribution. This probability can either be compared to a significance value (*e.g.*, 5%) and a binary decision made of whether or not the sample conforms to the hypothesized distribution, or it can be examined absolutely to make relative decisions among competing distributions.[12]

The present case presents some special EDF data sampling issues. Most of the RCS data histories used in the present study are tens of thousands of hits in size; and it would seem appropriate, although perhaps not computationally efficient, to use these entire data histories as the sample to be analyzed. However, it has been documented that EDF techniques often miscarry when sample sizes become large, on the order of several hundred members; this is because even small variations between the two CDF curves, if evaluated at every one of a large number of data points, become large in the summation. Thus, some method was necessary to limit the size of each sample to be analyzed to about 100 points. At the same time, it is not desirable to limit the sample to a single track; nor would it be representative to choose 100 hits at random and allow this to serve as the single sample for an object. The emergent solution is to use a “bootstrap” resampling technique to choose multiple samples of an acceptable size and aggregate the results. For the present investigation, the approach to testing a particular dataset for conformity to a particular distribution was to test 1000 samples of 100 values drawn at random from the dataset. Parameters were estimated not repeatedly from the 100-point samples but once from the entire dataset. Results were evaluated at the 0.02, 0.05, and 0.10 significance levels (“p-values”). The 0.05 threshold is a good compromise between what might be thought an excessive permissiveness of the 0.02 value and the somewhat demanding 0.10 value; 0.05 is also more widely used as a goodness-of-fit threshold, which reinforces its choice here.

RESULTS

The first task is to assess the general adequacy of the two canonical Swerling target types (S1 and S2, corresponding to Swerling I, II and Swerling III, IV). Figure 5 below shows the percentage of objects in each class whose RCS history can be adequately described by the S1 type, the S2 type, or a combination of S1 and S2 (meaning that at least one of the two types, and possibly both, adequately described the RCS history). The number of compliant cases predictably drops as the p-value requirement is increase from 0.02 to 0.05 to finally 0.10, and the decrease is somewhat larger from 0.02 to 0.05 than from 0.05 to 0.10. The two main Swerling models, taken together, can account for about 34% of the objects, performing somewhat better for rocket bodies than for payloads and debris.

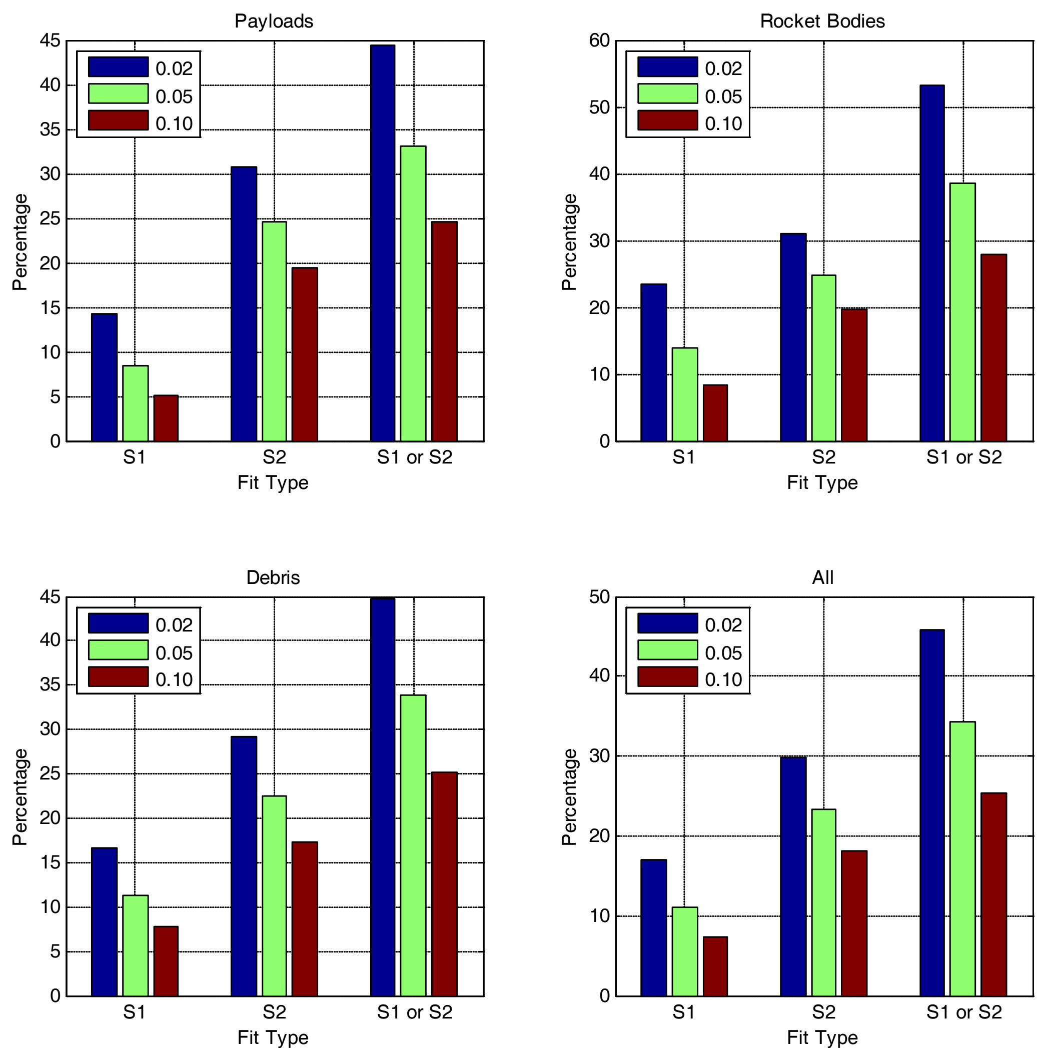

Figure 5. Performance of the Two Main Swerling Models

Figure 6 now shows the Swerling models extended in the matter proposed earlier (S0.5 to S2.5, with a 0.5 increment in shape parameter). The overall performance improves from 34% to 46%—notable but not overwhelming. The case that stands out as the most helpful is the addition of the S1.5 distribution to model rocket bodies, adding about ten percentage points of performance to its nearest competitor. The S0.5 case seems to add very little, but compromises were necessary in order to apply the EDF test to this distribution, and the suspicion is that the test is not entirely reliable for this case.* One may say in the main that an extended Swerling set of gamma distributions can account for about half of the objects in the test set.

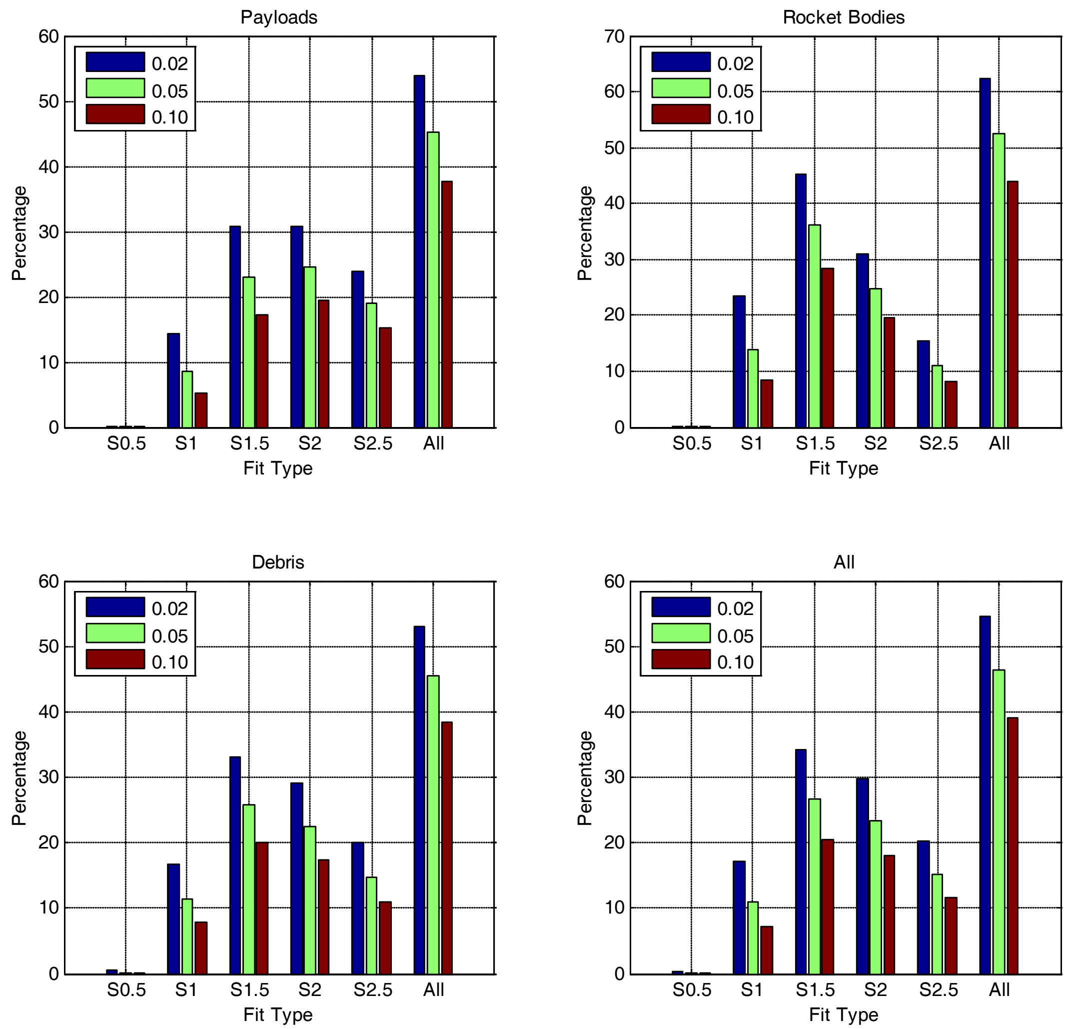

Figure 6. Performance of the Extended Swerling Models

* The tables of percentile points for assessing the EDF test results for the gamma distribution do not extend to situations in which the shape parameter is less than one. The present study attempted to extrapolate to this case using cubic splines, recognizing that this was at best an uncertain practice. The results appear to have confirmed this uncertainty, as one would have expected at least somewhat better performance for this type.

About half of the objects thus remain to be fitted to a target type, and the obvious next step is to see how the lognormal distribution will fare, comparing its results to the main Swerling types, the extended types, and the group of objects that, even after the introduction of the lognormal distribution, still cannot be assigned to a target type (the "none" column, conceptually an opposite of the "all" column in Figure 6). It should be remembered that for the lognormal case, the two governing parameters of this distribution (mean and variance) were estimated for each object, so the chances of finding an acceptable lognormal distribution in any given case are greater than for the applicability of a standard target type. Figure 7 below gives the results. The lognormal distribution compares favorably with the S1 and S2 models taken together for the payload and rocket body cases but notably underperforms in the debris case. Given that the number of unassigned objects comes in at a little over 40%, the addition of the lognormal distribution was thus able to improve the situation by only ten to fifteen additional percentage points—not nearly the improvement that was hoped.

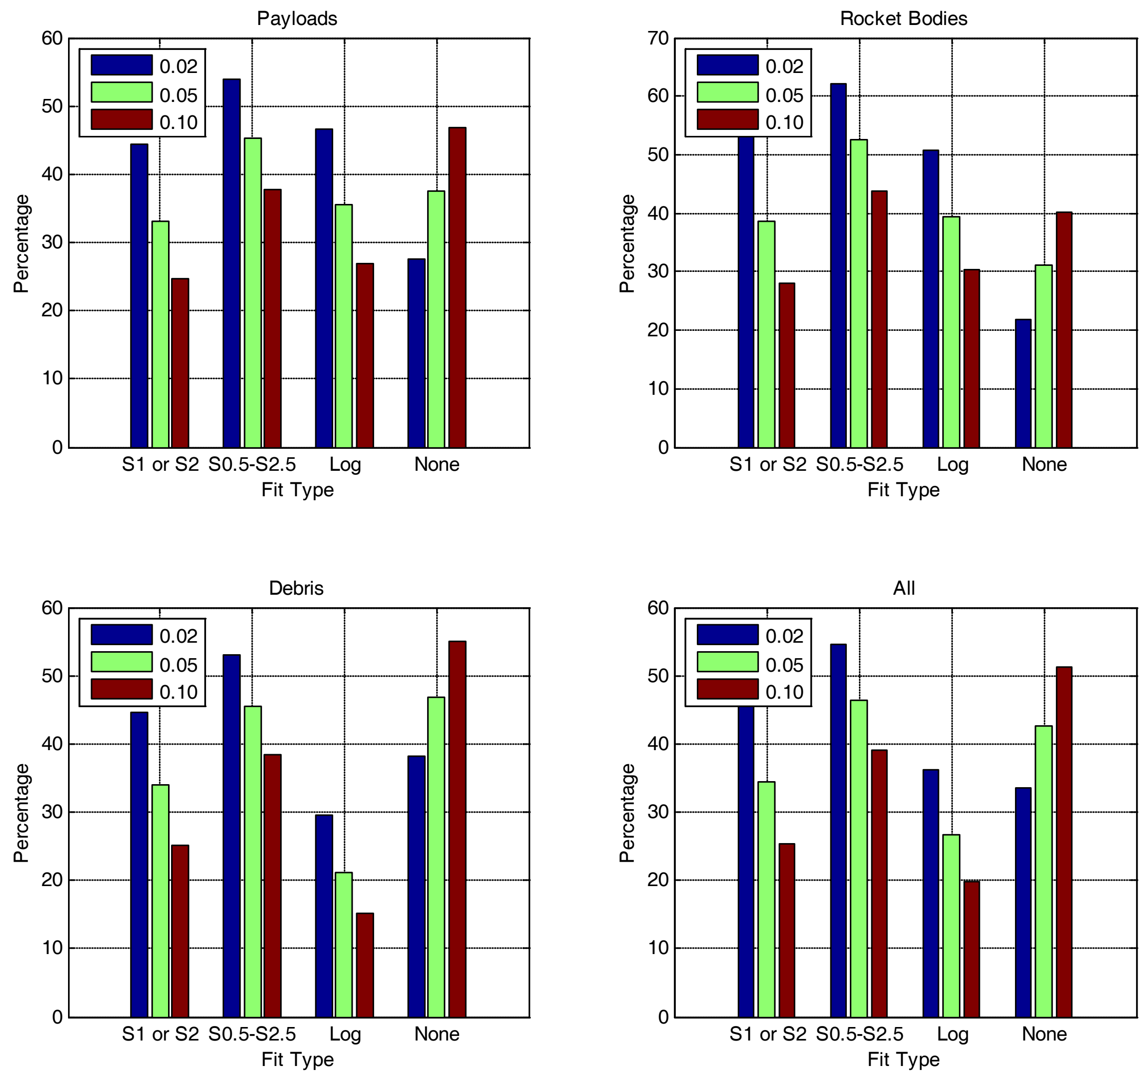

Figure 7: Results with Addition of Lognormal Distribution

Finally, one wishes to isolate the particular improvements that the introduction of the lognormal distribution actually brought. Figure 8 attempts to isolate any such improvements by separately representing cases in which the gamma-distribution models characterize the distribution but the lognormal does not, and vice versa. Since the bar-graph group labels may be somewhat confusing, Table 2 below gives a prose amplification.

Table 2. Explanation of Bar Cluster Definitions in Figure 8

Notation	Gloss
(S1 or S2) & ~Log	Either S1 or S2 (or both) passes, but Log does not
(S0.5 – S2.5) & ~Log	At least one of the distributions in the set S0.5 to S2.5 passes, but Log does not
Log & ~(S1 or S2)	Log passes, but neither S1 nor S2 does
Log & ~(S0.5 – S2.5)	Log passes, but none of the distributions in the S0.5 to S2.5 set does

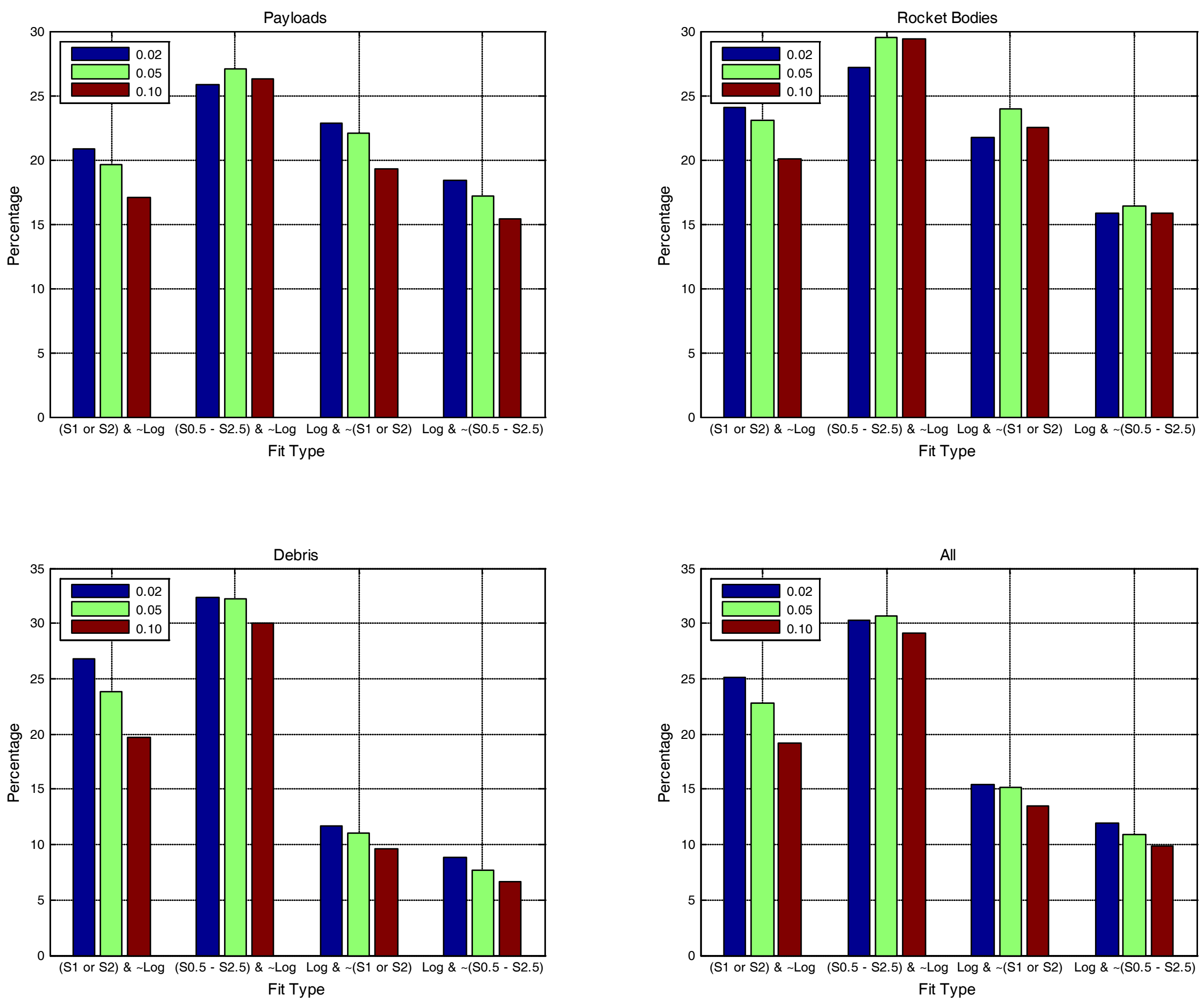

Figure 8. Chi-Squared and Lognormal Unique Contributions

One observes that about 23% of the objects can be modeled by either S1 or S2 that cannot be modeled by a lognormal distribution. About 15%, conversely, can be modeled by the lognormal case that cannot be accommodated by the S1-S2 combination. If one examines the expanded gamma distribution family to the full S0.5-S2.5, then about 30% of the objects can be handled that cannot be modeled at all by the lognormal distribution, and only 10% can be modeled only by lognormal. A similar phenomenon is observable across all three object types. For payloads and rocket bodies, the lognormal distribution is competitive with the S1-S2 combination but inferior to the S0.5-S2.5 combination; for debris, the lognormal distribution is substantially inferior to either the S1-S2 or the full S0.5-S2.5 combination. This last result is not entirely surprising, as debris objects are more irregularly shaped and are thus more likely to produce a Rayleigh or quasi-Rayleigh RCS distribution; but the difference is more marked than this author expected.

CONCLUSION AND FUTURE WORK

The research hypothesis, or perhaps more accurately the research hope, was that the lognormal distribution would be at the least competitive with, and perhaps superior to, the gamma distribution family for the purposes of spacecraft target typing; such a result would have allowed the substitution of a set of lognormal target types, facilitating parameter estimation and bringing many of the other favorable properties of the normal distribution. The present analysis clearly shows, however, that such a substitution is not warranted. It will be recalled that the lognormal distribution was given a certain advantage in this analysis by allowing both location and scale parameters (mean and variance) to be estimated from the sample; so if it merely achieves performance parity with the S1-S2 combination, it in fact underperforms because to make the lognormal distribution into actual target types, fixed variance values would need to be assigned. The underperformance is quite significant with debris objects, which may be the most important object class to consider because of its much greater object count and generally smaller object size; radars will struggle the most in tracking this object class, so the broad target typing paradigm should work to accommodate this class as its priority. It is true that about 10% of the objects are well modeled by the lognormal distribution and not well modeled at all by the gamma distributions, but one would need to consider whether it makes sense to introduce an entire new target typing distribution, and define specific target types within that distribution by establishing canonical variances, to serve a relatively small additional number of objects. Before taking so bold a step, one should examine the some 40% of the remaining objects that failed to fit either of these distributions (chi-squared or lognormal) to determine whether yet a third distribution type may account for a larger percentage, and perhaps a substantially larger one. In the absence of this additional work, the recommendation must be to leave the current Swerling target typing arrangement in place, inadequate though it may be for some applications.

If a minor repair is desired, then one could consider adding the S1.5 target type to the mix. Of the three non-canonical types considered in the S0.5-S2.5 set, this type appeared to make the greatest contribution. However, it must be remembered that testing for the Weinstock case (S0.5) was suspect, and Swerling thought that this additional type showed the most promise. It would seem ill-advised to make any changes until a more thorough investigation has been undertaken.

The first step in such an investigation, is to examine the four Pearson moments (especially the third and fourth moment) and associated histograms for the 40% of objects that do not conform to either the gamma or lognormal target types to see what can be learnt about their general behavior and whether they suggest the introduction of another standard distribution to model them. Such results will govern whether it is believed that the Swerling models and their natural extensions are about as far as one can go with *a priori* target types or whether there are other natural distributions to be found within the data that can profitably extend the target typing enterprise.

REFERENCES

[1] Swerling, P. "Probability of Detection for Fluctuating Targets." RAND Corporation Research Memorandum RM-1217, 17 MAR 1954, DTIC AD080638.

[2] Swerling, P. "Probability of Detection for Fluctuating Targets." *IRE Transactions on Information Theory Special Monograph Issue*, IT-6 (April 1960) : 269-308.

[3] Swerling, P. "Probability of Detection for Some Additional Fluctuating Target Cases." Aerospace Corporation Los Angeles Report TOR-669(9990), 14 MAR 1966, DTIC AD-488613.

[4] Swerling, P. "Recent Developments in Target Models for Radar Detection Analysis." *Proceedings of AGARD Avionics Technical Symposium*, Istanbul, Turkey (May 1970) : 7-1 – 7-21.

[5] Swerling, P. "Radar Probability of Detection for Some Additional Fluctuating Target Cases." *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, no. 2 (April 1997) : 698-709.

[6] Johnson, S. "Target Fluctuation Models for Radar System Design and Performance Analysis: An Overview of Three Papers." *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, no. 2 (April 1997) : 696-7.

[7] Weinstock, W. *Target Cross-Sections Models for Radar Systems Analysis*. Ph.D. dissertation, University of Pennsylvania, Philadelphia, 1956.

[8] Seybold, J. and Weeks, K. "Arithmetic versus Geometric Mean of Target Radar Cross Section." *Mircowave and Optical Technology Letters* vol. 11, no. 5 (5 APR 1996) : 265-270.

[9] Hejduk, M. and Lambert, J. "Statistical Distributions of Satellite RCS Histories." 2003 MIT/Lincoln Labortory Space Control Conference, Bedford, MA. April 2003.

[10] Cohen, A. and Whitten, B. *Parameter Estimation in Reliability and Life Span Models*. Statistics: Textbooks and Monographs, vol. 96. New York: Marcel Dekker, Inc., 1988.

[11] Johnson, N. and Kotz, S. *Continuous Univariate Distributions I*. Boston: Houghton Mifflin Company, 1970.

[12] D'Agostino, R. and Stephens, M. *Goodness-of-Fit Techniques*. Statistics: Textbooks and Monographs, Volume 68. New York: Marcel Dekker, Inc., 1986.