

German tank problem

In the [statistical theory of estimation](#), the **German tank problem** consists of estimating the maximum of a [discrete uniform distribution](#) from [sampling without replacement](#). In simple terms, suppose there exists an unknown number of items which are sequentially numbered from 1 to *N*. A random sample of these items is taken and their sequence numbers observed; the problem is to estimate *N* from these observed numbers.

The problem can be approached using either [frequentist inference](#) or [Bayesian inference](#), leading to different results. Estimating the population maximum based on a *single* sample yields divergent results, whereas estimation based on *multiple* samples is a practical estimation question whose answer is simple (especially in the frequentist setting) but not obvious (especially in the Bayesian setting).

The problem is named after its historical application by [Allied forces](#) in [World War II](#) to the estimation of the monthly rate of German tank production from very limited data. This exploited the manufacturing practice of assigning and attaching ascending sequences of serial numbers to tank components (chassis, gearbox, engine, wheels), with some of the tanks eventually being captured in battle by Allied forces.

Suppositions


The adversary is presumed to have manufactured a series of tanks marked with consecutive whole numbers, beginning with serial number 1. Additionally, regardless of a tank's date of manufacture, history of service, or the serial number it bears, the distribution over serial numbers becoming revealed to analysis is uniform, up to the point in time when the analysis is conducted.

Example

Assuming tanks are assigned sequential serial numbers starting with 1, suppose that four tanks are captured and that they have the serial numbers: 19, 40, 42 and 60.

A *frequentist* approach (using the [minimum-variance unbiased estimator](#)) predicts the total number of tanks produced will be:

A *Bayesian* approach (using a uniform [prior](#) over the integers in for any suitably large Ω) predicts that the [median](#) number of tanks produced will be very similar to the frequentist prediction:



During [World War II](#), production of German tanks such as the [Panther](#) was accurately estimated by Allied intelligence using statistical methods.

whereas the Bayesian mean predicts that the number of tanks produced would be:

Let N equal the total number of tanks predicted to have been produced, m equal the highest serial number observed and k equal the number of tanks captured.

The frequentist prediction is calculated as:

The Bayesian median is calculated as:


The Bayesian mean is calculated as:

These Bayesian quantities are derived from the Bayesian posterior distribution:

This probability mass function has a positive skewness, related to the fact that there are at least 60 tanks. Because of this skewness, the mean may not be the most meaningful estimate. The median in this example is 74.5, in close agreement with the frequentist formula. Using Stirling's approximation, the posterior may be approximated by an exponentially decaying function of n ,

which results in the following approximation for the median:

and the following approximations for the mean and standard deviation:




Estimated population size (N). The number of observations in the sample is k . The largest sample serial number is m . Frequentist analysis is shown with dotted lines. Bayesian analysis has solid lines with mean and shading to show range from minimum possible value to mean plus 1 standard deviation). The example shows if four tanks are observed and the highest serial number is "60", frequentist analysis predicts 74, whereas Bayesian analysis predicts a mean of 88.5 and standard deviation of 138.72 – 88.5 = 50.22, and a minimum of 60 tanks. In the SVG file (https://upload.wikimedia.org/wikipedia/commons/3/3e/German_tank_problem_graphs.svg), hover over a graph to highlight it.

Historical example of the problem

During the course of the Second World War, the Western Allies made sustained efforts to determine the extent of German production and approached this in two major ways: conventional intelligence gathering and statistical estimation. In many cases, statistical analysis substantially improved on conventional intelligence. In some cases, conventional intelligence was used in conjunction with statistical methods, as was the case in estimation of Panther tank production just prior to D-Day.

The allied command structure had thought the Panzer V (Panther) tanks seen in Italy, with their high velocity, long-barreled 75 mm/L70 guns, were unusual heavy tanks and would only be seen in northern France in small numbers, much the same way as the Tiger I was seen in Tunisia. The US Army was confident that the Sherman tank would continue to perform well, as it had versus the Panzer III and Panzer IV tanks in North Africa and Sicily.^[a] Shortly before D-Day, rumors indicated that large numbers of Panzer V tanks were being used.



Panther tanks are loaded for transport to frontline units, 1943.

To determine whether this was true, the Allies attempted to estimate the number of tanks being produced. To do this, they used the serial numbers on captured or destroyed tanks. The principal numbers used were gearbox numbers, as these fell in two unbroken sequences. Chassis and engine numbers were also used, though their use was more complicated. Various other components were used to cross-check the analysis. Similar analyses were done on wheels, which were observed to be sequentially numbered (i.e., 1, 2, 3, ..., N).^{[2][b][3][4]}

The analysis of tank wheels yielded an estimate for the number of wheel molds that were in use. A discussion with British road wheel makers then estimated the number of wheels that could be produced from this many molds, which yielded the number of tanks that were being produced each month. Analysis of wheels from two tanks (32 road wheels each, 64 road wheels total) yielded an estimate of 270 tanks produced in February 1944, substantially more than had previously been suspected.^[5]

German records after the war showed production for the month of February 1944 was 276.^{[6][c]} The statistical approach proved to be far more accurate than conventional intelligence methods, and the phrase "German tank problem" became accepted as a descriptor for this type of statistical analysis.

Estimating production was not the only use of this serial-number analysis. It was also used to understand German production more generally, including number of factories, relative importance of factories, length of supply chain (based on lag between production and use), changes in production, and use of resources such as rubber.

Specific data

According to conventional Allied intelligence estimates, the Germans were producing around 1,400 tanks a month between June 1940 and September 1942. Applying the formula below to the serial numbers of captured tanks, the number was calculated to be 246 a month. After the war, captured German production figures from the ministry of Albert Speer showed the actual number to be 245.^[3]

Estimates for some specific months are given as:^[7]

Month	Statistical estimate	Intelligence estimate	German records
June 1940	169	1,000	122
June 1941	244	1,550	271
August 1942	327	1,550	342

Similar analyses

Similar serial-number analysis was used for other military equipment during World War II, most successfully for the V-2 rocket.^[8]

Factory markings on Soviet military equipment were analyzed during the Korean War, and by German intelligence during World War II.^[9]

In the 1980s, some Americans were given access to the production line of Israel's Merkava tanks. The production numbers were classified, but the tanks had serial numbers, allowing estimation of production.^[10]

The formula has been used in non-military contexts, for example to estimate the number of Commodore 64 computers built, where the result (12.5 million) matches the low-end estimates.^[11]

Countermeasures

To confound serial-number analysis, serial numbers can be excluded, or usable auxiliary information reduced. Alternatively, serial numbers that resist cryptanalysis can be used, most effectively by randomly choosing numbers without replacement from a list that is much larger than the number of objects produced, or by producing random numbers and checking them against the list of already assigned numbers; collisions are likely to occur unless the number of digits possible is more than twice the number of digits in the number of objects produced; this is valid for serial numbers in any base ^[d] For this, a cryptographically secure pseudorandom number generator may be used. All these methods require a lookup table (or breaking the cypher) to back out from serial number to production order, which complicates use of serial numbers: a range of serial numbers cannot be recalled, for instance, but each must be looked up individually, or a list generated.

Alternatively, sequential serial numbers can be encrypted with a simple substitution cipher, which allows easy decoding, but is also easily broken by frequency analysis: even if starting from an arbitrary point, the plaintext has a pattern (namely, numbers are in sequence). One example is given in Ken Follett's novel Code to Zero, where the encryption of the Jupiter-C rocket serial numbers is given by:

H	U	N	T	S	V	I	L	E	X
1	2	3	4	5	6	7	8	9	0

The code word here is Huntsville (with repeated letters omitted) to get a 10-letter key.^[12] The rocket number 13 was therefore "HN", and the rocket number 24 was "UT".

V-2 rocket production was accurately estimated by statistical methods.

Frequentist analysis

Minimum-variance unbiased estimator

For point estimation (estimating a single value for the total, θ), the minimum-variance unbiased estimator (MVUE, or UMVU estimator) is given by:^[e]

where m is the largest serial number observed (sample maximum) and k is the number of tanks observed (sample size).^{[10][13]} Note that once a serial number has been observed, it is no longer in the pool and will not be observed again.

This has a variance^[10]

so the standard deviation is approximately N/k , the expected size of the gap between sorted observations in the sample.

The formula may be understood intuitively as the sample maximum plus the average gap between observations in the sample, the sample maximum being chosen as the initial estimator, due to being the maximum likelihood estimator,^[f] with the gap being added to compensate for the negative bias of the sample maximum as an estimator for the population maximum,^[g] and written as

This can be visualized by imagining that the observations in the sample are evenly spaced throughout the range, with additional observations just outside the range at 0 and $N + 1$. If starting with an initial gap between 0 and the lowest observation in the sample (the sample minimum), the average gap between consecutive observations in the sample is $N/(k+1)$, the $+1$ being because the observations themselves are not counted in computing the gap between observations.^[h] A derivation of the expected value and the variance of the sample maximum are shown in the page of the discrete uniform distribution.

This philosophy is formalized and generalized in the method of maximum spacing estimation; a similar heuristic is used for plotting position in a Q–Q plot, plotting sample points at $k / (n + 1)$, which is evenly on the uniform distribution, with a gap at the end.

Confidence intervals

Instead of, or in addition to, *point* estimation, *interval* estimation can be carried out, such as confidence intervals. These are easily computed, based on the observation that the probability that k observations in the sample will fall in an interval covering p of the range ($0 \leq p \leq 1$) is p^k (assuming in this section that draws are *with* replacement, to simplify computations; if draws are without replacement, this overstates the likelihood, and intervals will be overly conservative).

Thus the sampling distribution of the quantile of the sample maximum is the graph $x^{1/k}$ from 0 to 1: the p -th to q -th quantile of the sample maximum m are the interval $[p^{1/k}N, q^{1/k}N]$. Inverting this yields the corresponding confidence interval for the population maximum of $[m/q^{1/k}, m/p^{1/k}]$.

For example, taking the symmetric 95% interval $p = 2.5\%$ and $q = 97.5\%$ for $k = 5$ yields $0.025^{1/5} \approx 0.48$, $0.975^{1/5} \approx 0.995$, so the confidence interval is approximately $[1.005m, 2.08m]$. The lower bound is very close to m , thus more informative is the asymmetric confidence interval from $p = 5\%$ to 100% ; for $k = 5$ this yields $0.05^{1/5} \approx 0.55$ and the interval $[m, 1.82m]$.

More generally, the (downward biased) 95% confidence interval is $[m, m/0.05^{1/k}] = [m, m \cdot 20^{1/k}]$. For a range of k values, with the UMVU point estimator (plus 1 for legibility) for reference, this yields:

<i>k</i>	Point estimate	Confidence interval
1	$2m$	$[m, 20m]$
2	$1.5m$	$[m, 4.5m]$
5	$1.2m$	$[m, 1.82m]$
10	$1.1m$	$[m, 1.35m]$
20	$1.05m$	$[m, 1.16m]$

Immediate observations are:

- For small sample sizes, the confidence interval is very wide, reflecting great uncertainty in the estimate.
- The range shrinks rapidly, reflecting the exponentially decaying probability that *all* observations in the sample will be significantly below the maximum.
- The confidence interval exhibits positive skew, as N can never be below the sample maximum, but can potentially be arbitrarily high above it.

Note that m/k cannot be used naively (or rather $(m + m/k - 1)/k$) as an estimate of the standard error SE , as the standard error of an estimator is based on the *population* maximum (a parameter), and using an estimate to estimate the error in that very estimate is circular reasoning.

Bayesian analysis

The Bayesian approach to the German tank problem^[14] is to consider the posterior probability that the number of enemy tanks is N , when the number of observed tanks is K , and the maximum observed serial number is M . Note that this does not discard the relevant information contained in the samples, as we will see that the likelihood function can be constructed using only the number of samples and the max observed value.

The answer to this problem depends on the choice of prior for N . One can proceed using a proper prior over the positive integers, e.g., the Poisson or Negative Binomial distribution, where a closed formula for the posterior mean and posterior variance can be obtained.^[15] Below, we will instead adopt a bounded uniform prior.

For brevity, in what follows, $P(N|K,M)$ is written $P(N)$.

Conditional probability

The rule for conditional probability gives

Probability of M knowing N and K

The expression

is the conditional probability that the maximum serial number observed, m , is equal to k , when the number of enemy tanks, N , is known to be equal to M , and the number of enemy tanks observed, K , is known to be equal to k .

It is

where $\binom{M}{k}$ is a binomial coefficient and $[k \leq M]$ is an Iverson bracket.

The expression can be derived as follows: $P(m=k|K=k)$ answers the question: "What is the probability of a specific serial number k being the highest number observed in a sample of k tanks, given there are M tanks in total?"

One can think of the sample of size k to be the result of k individual draws without replacement. Assume k is observed on draw number k . The probability of this occurring is:

As can be seen from the right-hand side, this expression is independent of M and therefore the same for each k . As k can be drawn on k different draws, the probability of any specific k being the largest one observed is k times the above probability:

Seen differently, after k draws without replacement, one draw was for object k , and the other $k-1$ were for objects in the range 1 to $k-1$; there are $(k-1)!$ ways to achieve this. And there are $\frac{M!}{(M-k)!}$ ways to draw without replacement objects from a collection of M objects.

Probability of M knowing only K

The expression $P(M|M \leq k)$ is the probability that the maximum serial number is equal to k once tanks have been observed but before the serial numbers have actually been observed.

The expression $P(M|M \leq k)$ can be re-written in terms of the other quantities by marginalizing over all possible k .

Prior probability of N knowing only K

We assume that N is fixed in advance so that we do not have to consider any distribution over N . Thus, our prior can depend on K .

The expression

is the credibility that the total number of tanks, N , is equal to n when the number k tanks observed is known to be k , but before the serial numbers have been observed. Assume that it is some discrete uniform distribution

The upper limit Ω must be finite, because the function

is not a mass distribution function. Our result below will not depend on Ω .

Posterior probability of N knowing M and K

Provided that $k \geq 1$, so that the prior is consistent with the observed data:

As $k \rightarrow \infty$, the summation approaches $\frac{1}{k-1}$ (which is finite if $k \geq 2$). Thus, for suitably large Ω , we have

For $k \geq 1$ the mode of the distribution of the number of enemy tanks is m .

For $k \geq 2$, the credibility that the number of enemy tanks is *equal to* n , is

The credibility that the number of enemy tanks, N , is *greater than* n , is

Mean value and standard deviation

For $k \geq 3$, N has the finite mean value:

For $k \geq 4$, N has the finite standard deviation:

These formulas are derived below.

Summation formula

The following binomial coefficient identity is used below for simplifying series relating to the German Tank Problem.

This sum formula is somewhat analogous to the integral formula

These formulas apply for $k > 1$.

One tank

Observing one tank randomly out of a population of n tanks gives the serial number m with probability $1/n$ for $m \leq n$, and zero probability for $m > n$. Using Iverson bracket notation this is written

This is the conditional probability mass distribution function of m .

When considered a function of n for fixed m this is a likelihood function.

The maximum likelihood estimate for the total number of tanks is $N_0 = m$, clearly a biased estimate since the true number can be more than this, potentially many more, but cannot be fewer.

The marginal likelihood (i.e. marginalized over all models) is infinite, being a tail of the harmonic series.

but

where H_k is the harmonic number.

The credibility mass distribution function depends on the prior limit Ω :

The mean value of is

Two tanks

If two tanks rather than one are observed, then the probability that the larger of the observed two serial numbers is equal to m , is

When considered a function of n for fixed m this is a likelihood function

The total likelihood is

and the credibility mass distribution function is

The median satisfies

so

and so the median is

but the mean value of is infinite

Many tanks

Credibility mass distribution function

The conditional probability that the largest of k observations taken from the serial numbers $\{1,...,n\}$, is equal to m , is

The likelihood function of n is the same expression

The total likelihood is finite for $k \geq 2$:

The credibility mass distribution function is

The complementary cumulative distribution function is the credibility that $N > x$

The cumulative distribution function is the credibility that $N \leq x$

Order of magnitude

The order of magnitude of the number of enemy tanks is

Statistical uncertainty

The statistical uncertainty is the standard deviation , satisfying the equation

So

and

The variance-to-mean ratio is simply

See also

- Mark and recapture, other method of estimating population size
- Maximum spacing estimation, which generalizes the intuition of "assume uniformly distributed"
- Copernican principle and Lindy effect, analogous predictions of lifetime assuming just one observation in the sample (current age).
 - The Doomsday argument, application to estimate expected survival time of the human race.
- Generalized extreme value distribution, possible limit distributions of sample maximum (opposite question).
- Maximum likelihood
- Bias of an estimator
- Likelihood function

Further reading

- Goodman, L. A. (1954). "Some Practical Techniques in Serial Number Analysis". *Journal of the American Statistical Association*. **49** (265). American Statistical Association: 97–112. doi:10.2307/2281038 (<https://doi.org/10.2307%2F2281038>). JSTOR 2281038 (<https://www.jstor.org/stable/2281038>).

External links

- Grime, James (31 July 2024). "The Clever Way to Count Tanks - Numberphile" (<https://www.youtube.com/watch?v=WLCwMRJBhul>) (video). *YouTube*. Brady Haran. Retrieved 18 October 2024.

Notes

- a. An Armored Ground Forces policy statement of November 1943 concluded: "The recommendation of a limited proportion of tanks carrying a 90 mm gun is not concurred in for the following reasons: The M4 tank has been hailed widely as the best tank of the battlefield today. ... There appears to be no fear on the part of our forces of the German Mark VI (Tiger) tank. There can be no basis for the T26 tank other than the conception of a tank-vs.-tank duel – which is believed to be unsound and unnecessary."^[1]
- b. The lower bound was unknown, but to simplify the discussion, this detail is generally omitted, taking the lower bound as known to be 1.
- c. Ruggles & Brodie is largely a practical analysis and summary, not a mathematical one – the estimation problem is only mentioned in footnote 3 on page 82, where they estimate the maximum as "sample maximum + average gap".
- d. As discussed in birthday attack, one can expect a collision after $1.25\sqrt{H}$ numbers, if choosing from H possible outputs. This square root corresponds to half the digits. For example, in any base, the square root of a number with 100 digits is approximately a number with 50 digits.
- e. In a continuous distribution, there is no -1 term.
- f. Given a particular set of observations, this set is most likely to occur if the population maximum is the sample maximum, not a higher value (it cannot be lower).
- g. The sample maximum is never more than the population maximum, but can be less, hence it is a biased estimator: it will tend to *underestimate* the population maximum.
- h. For example, the gap between 2 and 7 is $(7 - 2) - 1 = 4$, consisting of 3, 4, 5, and 6.

References

1. AGF policy statement. Chief of staff AGF. November 1943. MHI
2. Ruggles & Brodie 1947, pp. 73–74.
3. "Gavyn Davies does the maths – How a statistical formula won the war" (<https://www.theguardian.com/world/2006/jul/20/secondworldwar.tvandradio>). *The Guardian*. 20 July 2006. Retrieved 6 July 2014.
4. Matthews, Robert (23 May 1998), "Data sleuths go to war, sidebar in feature "Hidden truths" " (<https://web.archive.org/web/20010418025817/http://www.newscientist.com/ns/980523/features.html#data>), *New Scientist*, archived from the original (<https://www.newscientist.com/article/mg15821355.000-hidden-truths.html>) on 18 April 2001
5. Bob Carruthers (1 March 2012). *Panther V in Combat* (<https://books.google.com/books?id=99JRxKz4Da4C&pg=PT94>). Coda Books. pp. 94–. ISBN 978-1-908538-15-4.
6. Ruggles & Brodie 1947, pp. 82–83.
7. Ruggles & Brodie 1947, p. 89.
8. Ruggles & Brodie 1947, pp. 90–91.
9. Volz 2008.
10. Johnson 1994.
11. "How many Commodore 64 computers were really sold?" (<https://web.archive.org/web/20160306232450/http://www.pagetable.com/?p=547>). *pagetable.com*. 1 February 2011. Archived from the original (<http://www.pagetable.com/?p=547>) on 6 March 2016. Retrieved 6 July 2014.
12. "Rockets and Missiles" (<https://www.spaceline.org/rocketsum/jupiter-c.html>). *www.spaceline.org*.
13. Joyce, Smart. "German Tank Problem" (<https://web.archive.org/web/20120424231135/http://www.lhs.logan.k12.ut.us/~jsmart/tank.htm>). Logan High School. Archived from the original (<http://www.lhs.logan.k12.ut.us/~jsmart/tank.htm>) on 24 April 2012. Retrieved 8 July 2014.
14. Simon, Cory (2023). "A Bayesian Treatment of the German Tank Problem" (<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11147940>). *The Mathematical Intelligencer*. **46** (2): 117–127. arXiv:2301.00046 (<https://arxiv.org/abs/2301.00046>). doi:10.1007/s00283-023-10274-6 (<https://doi.org/10.1007/s00283-023-10274-6>). PMC 11147940 (<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11147940>). PMID 38841650 (<https://pubmed.ncbi.nlm.nih.gov/38841650>).
15. Höhle, M.; Held, L. (2006). "Bayesian Estimation of the Size of a Population" (https://epub.ub.uni-muenchen.de/2094/1/paper_499.pdf) (PDF). *Technical Report SFB 386, No. 399, Department of Statistics, Ludwig-Maximilians-Universität München*. Retrieved 17 April 2016.

Works cited

- Johnson, R. W. (Summer 1994). "Estimating the Size of a Population" (<https://web.archive.org/web/20140223104835/http://www.rsscse-edu.org.uk/tsj/wp-content/uploads/2011/03/johnson.pdf>) (PDF). *Teaching Statistics*. **16** (2): 50–52. doi:10.1111/j.1467-9639.1994.tb00688.x (<https://doi.org/10.1111/j.1467-9639.1994.tb00688.x>). Archived from the original (<http://www.rsscse-edu.org.uk/tsj/wp-content/uploads/2011/03/johnson.pdf>) (PDF) on 23 February 2014.
- Ruggles, R.; Brodie, H. (1947). "An Empirical Approach to Economic Intelligence in World War II". *Journal of the American Statistical Association*. **42** (237): 72. doi:10.1080/01621459.1947.10501915 (<https://doi.org/10.1080/01621459.1947.10501915>). JSTOR 2280189 (<https://www.jstor.org/stable/2280189>).
- Volz, A. G. (July 2008). "A Soviet Estimate of German Tank Production". *The Journal of Slavic Military Studies*. **21** (3): 588–590. doi:10.1080/13518040802313902 (<https://doi.org/10.1080/13518040802313902>). S2CID 144483708 (<https://api.semanticscholar.org/CorpusID:144483708>).

Retrieved from "https://en.wikipedia.org/w/index.php?title=German_tank_problem&oldid=1358905339"