



Procedure and guidelines for identifying significant spikes in suicide rates

US Army Medical Research Unit - Europe
Walter Reed Army Institute of Research
LTC Paul Bliese (paul.bliese@us.army.mil)
Dr. Amy Adler (amy.adler@us.army.mil)
Dr. Kathleen Wright (kathleen.wright@us.army.mil)
COL Charles Hoge (charles.hoge@us.army.mil)

The goal of this research report is to provide guidance for interpreting annual suicide numbers in the military to determine when numbers represent significant spikes. Using data collected between 1990 and 2000, Eaton, Messer, Wilson and Hoge (2006) estimated the average suicide rate per year in the US Army to be 12.36 Soldiers per 100,000. On a year-by-year basis, however, the observed rate per 100,000 varies within a predictable expected range. In this research report, we provide a statistical simulation based on the binomial distribution that can be used to determine whether an observed suicide rate for any given year is within normal expected variability or whether it represents a potentially significant increase. The simulation can be modified to provide specific values for a given population (e.g., 60,000 Soldiers in US Army, Europe or 130,000 Soldiers in Iraq) and used to focus resources to address this important mental health issue. Reference tables are provided for those without access to the software.

PURPOSE

Suicides are relatively rare events in the military and their absolute numbers vary from year to year. For example in a population of 50,000 Soldiers suicide rates could vary from 5 suicides to 9 suicides in any given year. The challenge in interpreting these numbers is to determine when such numbers represent normal fluctuations and when such numbers signal important changes. The goal of this research report is to provide guidance for interpreting annual suicide numbers to determine when numbers represent significant spikes.

BACKGROUND

In 2006, Eaton, Messer, Wilson and Hoge estimated that the annual Army suicide rate between 1990 and 2000 was 12.36 per 100,000. Eaton et al. (2006) used this estimate to determine the expected number of suicides in an Army population of 480,000. They estimated that

in any given year, the Army would experience 59 suicides a year. Perhaps more importantly, though, they also used the estimate to provide some guidelines as to when an increase might be considered statistically significant. They reported that values of 73 or higher would be strong evidence that a significant increase had occurred using a 95% confidence interval.

In this report we build on the work of Eaton et al. (2006) and show how statistical simulations can be used to identify significant spikes in suicide numbers in military populations of different sizes.

BINOMIAL SIMULATION

Binomial simulations can be illustrated using the same logic as the probability underlying a coin-flip. Coin flips can be represented as 1 (heads) or 0 (tails) with a .5 probability of heads following each flip. With 100,000


flips of the coin, a simulation would expect 50,000 heads, but the observed number might vary across each replication of 100,000 flips. The code using the open-source language R (R Development Core Team, 2006) available at <http://cran.r-project.org/> illustrates this process:

```
># Simulation 1
> TDAT<-rbinom(size=1,prob=.5,n=100000)
> sum(TDAT)
[1] 49850

># Simulation 2
> TDAT<-rbinom(size=1,prob=.5,n=100000)
> sum(TDAT)
[1] 50040
```

In this example, the first simulation came out with 150 heads fewer than would be expected based on the probability of .5 (50,000 – 49,850). The second simulation came up with 40 extra heads.

If the simulation process is run 10,000 times (Simulation 1, Simulation 2, Simulation 3....Simulation 10,000) the results from each of the 10,000 simulations will result in a distribution of values. This distribution can be used to estimate what is within normal expected range and what is outside of the normal expected range. The figure below shows the histogram for 10,000 runs of the simulation process.



Notice that the peak of this histogram is centered around 50,000 (in this example the average was 50,002.52). Importantly, however, the process also gives some indication about what values would be considered unexpected. In this example, only five percent of the simulated values were above 50,263 (approximate arrow location), so any coin tossed 100,000 times which produces 50,264 heads is not behaving like a normal coin. That is, it is more than 95% certain that this coin is somehow more prone to come up heads when flipped. The original two simulations (one with 150 heads too few and one with 40 heads too many) are well within the normal expected range.


BINOMIAL SIMULATION APPLIED TO SUICIDE RATES

The same process can be applied to suicide prediction rates by changing the probability to 12.36 per 100,000 instead of using a .5 probability as illustrated with the coin example. Two simulations representing populations of 100,000 are illustrated:

```
># Simulation 1
> TDAT<-rbinom(size=1,prob=12.36/100000,
n=100000)
> sum(TDAT)
[1] 5

># Simulation 2
> TDAT<-rbinom(size=1,prob=12.36/100000,
n=100000)
> sum(TDAT)
[1] 14
```

The first simulation estimated 5 suicides per 100,000 and the second estimated 14. If we run the simulation 10,000 times instead of two, we observe the following distribution of the number of suicides per 100,000 assuming a rate of 12.36 per 100,000. The histogram is basically equivalent to having 10,000 years of data on a population of 100,000 individuals.



The range of estimated numbers of suicides is from 1 to 27. The average is 12.41 (close to the target of 12.36). The 95% confidence interval is 19 and above. This means that with 100,000 Soldiers, 19 or more suicides would be outside the normal expected range ($p < .05$). Alternatively, given the importance of monitoring suicides, it may be useful to adopt a 90% confidence interval ($p < .10$). In this case, 18 or more suicides would be outside the expected range signaling a spike in rates.

ADAPTING THE SIMULATION

To make the simulation useful in the military, it needs to be able to make estimates for different sized populations. For instance, if the Commander, US Army Europe (USAREUR) wanted to know if suicide rates significantly increased in the previous year, the simulation would need to provide an estimate based on the USAREUR population.

The two functions in the appendix can be used to make predictions for different sized populations. To illustrate the use of the functions, we run the simulation using an Army size of 480,000 reported by Eaton et al., (2006). The simulation is run 10,000 times as was done in the previous examples.

The simulation indicates that the annual expected number of suicides for the US Army (population 480,000) is 59. The 95% confidence interval is 73 and

above. This means that with 480,000 Soldiers, 73 or more suicides would be outside the normal expected

```
> ARMY<-sim.suic(nrep=10000,
  rate=12.36/100000,
  popsize=480000)
> median(ARMY)
[1] 59
> quantile.suic(ARMY,c(.90,.95))
  quantile.values  confint.estimate
1             0.90                70
2             0.95                73
```

range ($p < .05$). The values of 59 and 73 match those reported in Eaton et al. (see their Table 3 on page 188). Again, it may be useful to adopt a 90% confidence interval ($p < .10$). In this case, 70 or more suicides would be outside the expected range signaling a spike in rates.

TABLE VALUES

The final table provides annual expected values and 90% and 95% confidence interval estimates for a range of population sizes from 10,000 to 150,000 using the simulation program:

Size of Population	Annual Expected Value	Value to be 90% Confident Increase is Significant	Value to be 95% Confident Increase is Significant
10,000	1	4	4
20,000	2	6	6
30,000	4	7	8
40,000	5	9	10
50,000	7	10	11
60,000	7	12	13
70,000	8	13	15
80,000	10	15	16
90,000	11	16	18
100,000	12	18	19
110,000	13	20	21
120,000	15	21	22
130,000	16	22	24
140,000	17	24	26
150,000	18	25	27

IMPLICATIONS

The simulation procedures and table of expected values provide estimates that can be useful to those who want

to distinguish between normal variations and unexpected variations in suicide rates. Unexpected variations warrant careful attention and planning strategies to trigger an appropriate community mental health response.

The illustrated procedures were based on historical Army suicide data, but can easily be adapted for other services by changing the “rate” option in the simulation. The procedures can also be adapted to detect significant decreases in suicide rates. For example, to be 95% confident an Army-wide suicide prevention program had significantly reduced the number of suicides, the annual suicide rate would need to drop to 48 or fewer. To be 90% confident a program had reduced suicides, the annual rate would need to drop to 51 (see box below).

```
> ARMY<-sim.suic(nrep=10000,
  rate=12.36/100000,
  popsize=480000)
> quantile.suic(ARMY,c(.05,.10))
  quantile.values  confint.estimate
1             0.05             48
2             0.10             51
```

REFERENCES

Eaton, K. M., Messer, S. C., Wilson, A. L. G. & Hoge, C. W. (2006). Strengthening the validity of population-based suicide rate comparisons: An illustration using U.S. military and civilian data. *Suicide and Life-Threatening Behavior*, 36, 182-191.

R Development Core Team (2006). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL <http://www.R-project.org>.

Material has been reviewed by the Walter Reed Army Institute of Research. There is no objection to its presentation and/or publication. The opinions or assertions contained herein are the private views of the authors, and are not to be construed as official, or as reflecting true views of the Department of the Army or the Department of Defense.

APPENDIX

```
sim.suic<-function(nrep, rate, popsize){
  OUT<-rep(NA,nrep)
  for(i in 1:nrep){
    OUT[i]<-sum(rbinom(size=1, prob=rate,n=popsize))
  }
  return(OUT)
}

quantile.suic<-function (x, confint, ...){
  out <- data.frame(quantile.values = confint, confint.estimate = rep(NA,
    length(confint)))
  cumpct <- cumsum(table(x)/length(x))
  lag1 <- c(NA, cumpct[1:length(cumpct) - 1])
  lag2 <- c(NA, lag1[1:length(lag1) - 1])
  TDAT <- data.frame(agree.val = as.numeric(names(cumpct)),
    cumpct, lag1, lag2)
  for (i in 1:length(confint)) {
    out[i, 2] <- TDAT[TDAT$cumpct > confint[i] & TDAT$lag1 >=
      confint[i] & TDAT$lag2 < confint[i], 1]
  }
  return(out)
}
```