

Chapter 2

Shooting Without Feedback

*One or two sarcastic spirits
Pointed out to him, however,
That it might be much more useful
If he sometimes hit the target.*

*From “Hiawatha Designs an Experiment”, a statistical poem by
Maurice Kendall (1959)*

2.1 Introduction

In combat, man specializes in exerting lethal force at a distance. The mechanism for exerting this force has progressed from stones to spears to firearms to rockets, but the basic shooting problem for the marksman has always been to effectively combine accuracy and lethality at long range. This chapter is devoted to abstract models of the shooting process. Such models have several possible purposes. One purpose is simply to determine the probability of killing the target for use in a higher level combat model. Another is to influence the design of a weapon system – this purpose will be aided by the dependence of kill probability on fundamental parameters that quantify accuracy and lethality. A third purpose is to influence the shooting process itself, since there are tactical decisions to be made about where to aim and how often to shoot. We assume throughout that the marksman gets no feedback between shots, so questions about how to adjust the aim point or when to stop shooting must be reserved until Chapter 3.

The objective of the shooting process is to kill the target. The event “kill” may have several meanings such as “immobilize,” “damage,” “find,” or even “rescue,” but the only important thing is that the event should either happen to the target or not with each shot, with no residual effect if the target is not killed.

Each shot is aimed at a certain aim point that the marksman controls, but the shot usually hits a different impact point because of dispersion errors that are assumed to be independent from shot to shot. The marksman may have a faulty notion of where the target is located, with the difference between the marksman’s estimate and the actual target location being the target location error, one of several possible contributors to a bias error that affects all shots. We assume throughout that dispersion errors and bias errors are normally distributed in two dimensions. The world that we know has three dimensions, but most targets are known to lie on the surface of the Earth somewhere, which fixes the vertical dimension. Even problems that involve aerial or submarine targets have in most cases been reduced to two-dimensional problems by the invention of the proximity fuse.

While bivariate normality is not always the case in reality, it is usually close to the truth, except for outliers, and anyway the normal distribution should still be our default assumption because of the central limit theorem (see Appendix A).

The plan of this chapter is to first deal with problems involving only a single shot, and then generalize to multi-shot problems. Section 2.2 deals with problems where there is only one shot. Section 2.3 generalizes to include many shots, as long as they are all independent. Section 2.4 deals with multi-shot problems where bias errors are significant and thus the shots are not independent. In such situations pattern firing is useful.

It will be useful to have the *Chapter2.xls* Excel™ workbook available.

2.2 Single-Shot Kill Probability

In this section we present some basic firing models that describe what happens when a single shot is fired. The shot will kill the target if, in spite of various errors in firing and in locating the target, the shot's distance from the target is sufficiently small. Our object is to determine the probability of kill as a function of more basic parameters.

2.2.1 Damage Functions and Lethal Area

Definition: The *miss distance* r is the distance between the impact point of the shot and the location of the target.

Definition: The *damage function* $D(r)$ is the probability that the target is killed by the shot, as a function of the miss distance r . “Kill function” might be a more natural name, but our usage reflects common practice. Figure 1 shows three examples of damage functions.

The damage function has only one argument, so we are implicitly assuming a radially symmetric situation where damage to the target is invariant to the position of the impact point, except for the miss distance. Damage functions are in practice measured through some combination of theory and experiment. In this book, we will take them to be given.

The damage function is a conditional kill probability. The unconditional kill probability P_K is obtained by averaging over the miss distance. Let $f(x, y)$ be the bivariate density function of the position of the weapon's point of impact, relative to the target's location. Then, since $r = \sqrt{x^2 + y^2}$,

$$P_K = \iint D(\sqrt{x^2 + y^2}) f(x, y) dx dy, \quad (2.1)$$

where the lack of limits means that the integral is to be taken over the whole plane. Sections 2.2.2 through 2.2.4 deal with various special cases of (2.1).

If the position of the target were uniformly distributed within some large area A , then (2.1) would be (substituting $f(x, y) = 1/A$),

$$P_K = \frac{1}{A} \iint_A D(\sqrt{x^2 + y^2}) dx dy, \quad (2.2)$$

where the notation indicates that the integral is now taken only over the area A . However, since A is by assumption large, (2.2) is approximately the same as $P_K = a / A$, where

$$a = \iint D(\sqrt{x^2 + y^2}) dx dy \quad \text{or} \quad (2.3)$$

$$a = 2\pi \int_0^{\infty} r D(r) dr. \quad (2.4)$$

Formula (2.4) was obtained from (2.3) by introducing polar coordinates. The quantity “ a ” is the *lethal area* of the weapon and serves as a scalar measure of weapon effectiveness with respect to a certain target.

Although it is not logically necessary, the damage function $D(r)$ is typically a non-increasing function of its argument. As long as this is true, it is sometimes convenient to describe a damage function as follows: imagine that the weapon has a random “lethal radius” R associated with it, and that a target will be killed if and only if it lies within a distance R of the weapon’s point of impact. Recalling the meaning of $D(r)$, it must evidently be the case that

$$D(r) = P(R > r). \quad (2.5)$$

If $D(r)$ is differentiable, one can go further and discover the probability density function of R :

$$f_R(r) = -\frac{d}{dr} D(r). \quad (2.6)$$

The area covered by the weapon is πR^2 , so it should come as no surprise that $a = \pi E(R^2)$, where $E(\)$ denotes expectation. This provides another interpretation of the term “lethal area.”

2.2.2 Cookie-Cutter Damage Function

The conceptually simplest kind of weapon is one for which the lethal radius R is a constant, in which case the lethal area is $a = \pi R^2$. If the firing errors are circular

normal (by which we mean that the standard deviation of the error in all directions is the same number σ) and centered on the target, then the two-dimensional density function of the error is $\frac{1}{2\pi\sigma^2} \exp(-\frac{x^2 + y^2}{2\sigma^2})$, and (2.1) reduces after the introduction of polar coordinates to

$$P_K = \int_0^{2\pi} \int_0^R \frac{1}{2\pi\sigma^2} \exp(-\frac{r^2}{2\sigma^2}) r \, dr \, d\theta = 1 - \exp\left(-\frac{1}{2} R^2 / \sigma^2\right). \quad (2.7)$$

This is a very simple expression for kill probability. A large number of military analyses have been based on it. Since (2.7) is a formula for the probability that the miss distance does not exceed R , it is also a formula for the cumulative distribution function of the miss distance when firing errors are circular normal, a Rayleigh type of random variable.




Figure 1: The Diffuse Gaussian (alpha=1) damage function is contrasted with two more definite damage functions with higher values of alpha. All three have the same lethal area. Dimensionless range is the ratio of miss distance to the scale parameter b .

Unfortunately, most departures from the circular normal assumption about errors result in significantly more complicated expressions for P_K . If the circular normal error distribution is offset from the target by some distance h , for example,

then evaluation of (2.1) involves the integral of a Bessel function that must be done numerically. One of the early products of the RAND corporation was a table of such probabilities (Marcum, 1950). Sheet “OffsetQ” of *Chapter2.xls* employs the function $OffsetQ(R/\sigma, h/\sigma)$ to calculate the miss probability (the probability that the miss distance exceeds R) in this case. If the normal distribution is centered on the target, but not circular, then the function $EllipQ(R, \sigma_1, \sigma_2)$ can be used instead. VBA code for both functions can be found in module 1 of *Chapter2.xls*. The code is based on Gilliland (1962), which also covers the general case where the normal distribution is neither centered nor circular.

Example 1: If $R = 100$ m and $\sigma = 50$ m, the kill probability according to (2.7) is 0.865. If the aim point is offset from the center of the distribution by $h = 25$ m, then the miss probability is $OffsetQ(2, 0.5) = 0.169$, so the kill probability is only 0.831. The bigger the offset, the smaller the kill probability. See sheet “OffsetQ” in *Chapter2.xls* for more calculations of this kind.

Example 2: Suppose $R = 100$ m, but that the standard deviation of the error along the line between marksman and target (the downrange direction) is 90 m, while the standard deviation of the error in the direction perpendicular to that (the crossrange or deflection direction) is only 30 m. It is typical for downrange and crossrange errors to differ by about a factor of 3, as assumed here. The miss probability is $EllipQ(100, 90, 30) = EllipQ(100, 30, 90) = 0.293$, so $P_K = 0.707$. An approximation can be made by letting the error in both directions be whatever is required to preserve the product of $(30 \text{ m})(90 \text{ m}) = (2700 \text{ m}^2)$. That error is $\sigma = 51.96$ m. Employing (2.7), we find that $P_K = 0.843$ and see that it can be significantly optimistic to assume that errors are circular when they are not. Sheet “EllipQ” of *Chapter2.xls* generalizes this example.

Formula (2.7) is sometimes expressed in the form

$$P_K = 1 - (0.5)^{(R^2/CEP^2)}, \quad (2.8)$$

where CEP stands for “circular error probable.”

Definition: *Circular Error Probable* (CEP) is the radius of the smallest circle that contains the two-dimensional error with probability 0.5.

In (2.8), CEP references a circular normal firing error, but the definition of CEP applies to any kind of a two-dimensional error distribution. For a circular normal distribution, CEP is related to σ by $CEP = \sigma\sqrt{2\ln 2} = 1.1774\sigma$ (see Exercise 14).

Dealing with CEP instead of σ has the advantage of easing explanations to novices, since the idea of standard deviation is not involved. CEP can itself be confusing, however. The corresponding notion in one dimension would be an

interval that contains the error half the time, the half-length of which is sometimes called the linear error probable or LEP. A novice might expect LEP and CEP to be equal, since each corresponds to a region that contains half the shots, but they are not. A square with side $2 \times \text{LEP}$ will contain only $(0.5)^2$ of the shots, since both coordinates must fall within the relevant side of the square. A circle with radius LEP, since it falls within that square, will contain even less than 25% of the shots. CEP is thus considerably greater than LEP, and SEP (spherical error probable) is larger yet. In other words, the magnitude of the “error probable” must depend on the number of dimensions that one is interested in. It is simpler to say that the error has standard deviation σ along any line, regardless of the number of dimensions.

2.2.3 Diffuse Gaussian (DG) Damage Function

The DG damage function has the form $D(r) = \exp\left(-\frac{r^2}{2b^2}\right)$ for some scale factor

b . Unlike the cookie-cutter damage function, there is a positive probability of killing the target no matter what the miss distance. The lethal area of such a weapon is $2\pi b^2$. Figure 1 compares $D(r)$ for a DG damage function with two other damage functions that will be discussed in the next section. The weapon with the DG function is evidently “sloppier” than the others, in the sense that results at any given range are harder to predict. Whether this feature makes the DG assumption more realistic depends on the damage mechanism. Weapons that kill by fragmentation (e.g. artillery using “fragmentation” rounds) generally have a sloppier damage function than those that kill by overpressure (e.g., artillery using high explosive rounds). The cookie-cutter damage function is the least sloppy of them all because it suddenly falls from 1 to 0 as the miss distance increases.

The DG assumption combines very nicely with the assumption of normal errors to produce a simple, general expression for P_K . If the center of the error distribution is (μ_X, μ_Y) , and if the standard deviations of the X and Y errors are (σ_X, σ_Y) , then (2.1) can be evaluated analytically. In fact, (2.1) can be integrated in closed form even when the damage function is not radially symmetric, so we record here the result for

the asymmetric DG damage function: If $D(x, y) = \exp\left(-\frac{1}{2}\left(\left(\frac{x}{b_X}\right)^2 + \left(\frac{y}{b_Y}\right)^2\right)\right)$, then

$$P_K = \frac{b_X b_Y}{\sqrt{(b_X^2 + \sigma_X^2)(b_Y^2 + \sigma_Y^2)}} \exp\left(-\frac{1}{2}\left(\frac{\mu_X^2}{b_X^2 + \sigma_X^2} + \frac{\mu_Y^2}{b_Y^2 + \sigma_Y^2}\right)\right). \quad (2.9)$$

In the special case where $\mu_X = \mu_Y = 0$, $b_X = b_Y = b$, and $\sigma_X = \sigma_Y = \sigma$, (2.9) reduces to

$$P_K = \frac{b^2}{b^2 + \sigma^2}, \quad (2.10)$$

which is comparable to (2.7).

There is no cookie-cutter counterpart to (2.9); that is, there is no simple analytic expression for the cookie-cutter kill probability with the generality of (2.9). While it is true that the cookie-cutter damage function is conceptually simpler than the DG damage function, it is equally true that the DG function is analytically simpler than the cookie-cutter function. The simplicity of the DG assumption was first taken advantage of in 1941 by von Neumann (Taub, 1962), who used it in determining optimal bomb spacing in World War II.

The “DGGenrl” sheet of *Chapter2.xls* calculates P_K from (2.9), as well as performing a Monte Carlo simulation whose object is to estimate the same quantity. The agreement of the two answers is part of the process of verifying that each is correct.

2.2.4 Other Damage Functions

It was pointed out in Section 2.1 that any non-increasing damage function can be interpreted as the probability law for a random lethal radius R . The Diffuse Gaussian damage function, for example, has associated with it the density function

$f_R(r) = \frac{r}{b^2} \exp\left(-\frac{r^2}{2b^2}\right)$, which is a Rayleigh density function. It is perhaps more

natural to deal with the random variable R^2 , since R^2 is directly related to area covered. For the DG damage function, R^2 is an exponential random variable with mean $2b^2$.

It is possible, of course, to reverse the process: begin with some convenient density for R or R^2 and then discover the associated damage function by integration. One convenient class of damage functions (the Gamma class) can be obtained by assuming that $\frac{1}{2}R^2/b^2$ has the Gamma density $\frac{\alpha(\alpha x)^{\alpha-1}}{\Gamma(\alpha)} \exp(-\alpha x)$ for

some $\alpha > 0$, in which case the DG damage function is the special case $\alpha = 1$ and the cookie-cutter damage function is obtained in the limit as $\alpha \rightarrow \infty$. Every member of the class has the same lethal area $2\pi b^2$. The associated damage function is

$$D_\alpha(r) = 1 - \Gamma\left(\alpha, \frac{\alpha r^2}{2b^2}\right) \quad (2.11)$$

where $\Gamma(\alpha, x)$ is the incomplete Gamma function (in Excel™, $\Gamma(\alpha, x)$ is `GAMMADIST(x, α , 1, TRUE)` — see Figure 1 or sheet “Gamma” of *Chapter2.xls* for plots of $D_\alpha(r)$ versus $\frac{r}{\sqrt{2}b}$).

The Gamma class is convenient because it has both scaling (b) and shaping (α) parameters, and also because there is a simple formula for P_K when the firing error is circular normal with standard deviation σ and centered on the target:

$$P_K = 1 - \left(1 + \frac{b^2}{\alpha\sigma^2}\right)^{-\alpha}; \quad \alpha > 0. \quad (2.12)$$

Formula (2.10) is the special case where $\alpha = 1$, and (2.7) is the limiting case as $\alpha \rightarrow \infty$. Except in the case $\alpha = 1$, where (2.9) applies, the centered circular normal assumptions are essential for having a simple expression for P_K .

Another class of density functions for R^2 with both a shaping and a scaling parameter is the class of lognormal densities. There turns out to be little to recommend this class in terms of analytic convenience – there are no counterparts to (2.11) and (2.12), for example. Nonetheless, the class has been widely used to model the effects of nuclear weapons (DIA, 1974).

2.3 Multiple-Shot Kill Probability

In this section we consider multiple shots made in a salvo.

Definition: Regardless of the number of targets, a “salvo” is a group of shots all taken on the basis of the same information, with no information feedback between shots. We often imagine that the shots are all made at the same time, even though that is not necessarily the case.

2.3.1 Simultaneous Independent Shots

Suppose that a salvo of n independent shots is fired at a target, and let q_i be the probability that the i th shot fails to kill it. The numbers q_i may be obtained from one of the formulas in Section 2.2 or by some other method. Since all shots are by assumption independent, the probability that all n miss the target is the product of the miss probabilities, so

$$P_K = 1 - q_1 q_2 \dots q_n. \quad (2.13)$$

In the case where all the shots have the same miss probability q , this reduces to $P_K = 1 - q^n$, a formula that is used so often that it has a name “powering up.” If each of three independent shots has a kill probability of 0.3, one might naively expect the kill probability of the collection to be 0.9. The correct answer, however, is $1 - 0.7^3 = 0.657$. The miss probability should be powered up, rather than making the kill probability proportional to the number of shots.

Formula (2.13) takes on a particularly simple form if the shots are all of the cookie-cutter type, and the firing errors are circular normal centered on the target. Let R_i and σ_i be the lethal radius and error standard deviation of the i^{th} shot. Then

$q_i = \exp\left(-\frac{1}{2}R_i^2/\sigma_i^2\right)$ from (2.7), and therefore

$$P_K = 1 - \exp(-X/2), \text{ where} \quad (2.14)$$

$$X = R_1^2/\sigma_1^2 + \dots + R_n^2/\sigma_n^2.$$

The quantity X can be thought of as a measure of the effectiveness of an arsenal of weapons against a particular target. The target dependence of this effectiveness can be eliminated if lethal radius scales in a known manner with the energy yield Y of the weapons. If the kill mechanism is overpressure, for example, then $R_i = KY_i^{1/3}$ where K is a target-dependent constant, and therefore $X = K^2[Y_1^{2/3}/\sigma_1^2 + \dots + Y_n^{2/3}/\sigma_n^2]$. The quantity in [] is a target-independent measure of effectiveness for the group of weapons taken as a whole. It differs from “counter military potential” (CMP, see below) only in the scale factor required to convert standard deviation to circular error probable (CEP) for circular normal weapons (see Section 2.2.2).

The CMP of a group of n weapons is

$$\text{CMP} \equiv Y_1^{2/3}/\text{CEP}_1^2 + \dots + Y_n^{2/3}/\text{CEP}_n^2, \quad (2.15)$$

where yield is measured in equivalent kilotons of TNT, and CEP is measured in nautical miles. CMP is one of several measures that have been used to compare arsenals of nuclear weapons. This measure is very sensitive to accuracy – doubling all yields increases CMP by the factor $2^{2/3} = 1.6$, whereas halving all CEP’s increases CMP by the larger factor $2^2 = 4$. In the 1970s, this fact was sometimes used to make the point that the relatively small but accurate nuclear arsenal of the United States was actually more potent than the large but inaccurate arsenal of the Soviet Union. Tsipis (1974), for example, estimated that CMP was 22000 for the US and 4000 for the SU in 1974.

An alternative measure of effectiveness for an arsenal is “equivalent megatons” (EMT). The definition of EMT is

$$\text{EMT} \equiv Y_1^{2/3} + \dots + Y_n^{2/3}. \quad (2.16)$$

Since $Y_i^{2/3}$ is proportional to R_i^2 , EMT is essentially a target-independent measure of the total lethal area of the arsenal. The “EMTCMP” sheet of *Chapter2.xls* compares the 1978 ICBM arsenals of the USA and the USSR using both measures.

If a total of C units of CMP are applied to a target, then the kill probability is of course still a function of the hardness of the target. For nuclear weapons making

overpressure kills, with hardness h being measured in pounds per square inch (psi), an approximate formula (Matlin, 1972) valid for $30 \leq h \leq 1000$ psi is

$$P_K = 1 - \exp(-0.0583Ch^{-0.7}). \tag{2.17}$$

For example, a 1000 kiloton weapon with a CEP of 0.25 nautical miles will kill a 1000 psi target with probability $1 - \exp(-(0.0583)(400)(0.00794)) = 0.52$. Sixteen such weapons would be equally effective if the CEP were 1 nautical mile.

Figure 2 shows an application of the EMT idea to various historical ships, except that EMT is replaced by E8RPM. E8RPM is the equivalent rate of applying 8" rounds per minute, plotted against firing range for each of five ships. Since yield (Y) is proportional to weight, and weight to the cube of dimension, $Y^{2/3}$ is proportional to dimension squared. Thus one 16" round is roughly equivalent to four 8" rounds, etc. Reducing all firepower to a common scale permits each ship to have a single "weight of broadside" versus range curve, independent of the target. Aircraft carriers can be portrayed on the same scale by calculating the rate at which aircraft can deliver 500-pound bombs, each of which is equivalent to two 8" rounds. At a range of 20 nm, any of the battleships shown in Figure 2 could have quickly sunk the USS Enterprise. The battleship's problem in World War II was that an aircraft carrier could deliver enough firepower, even at long range, to put a battleship out of commission before any of its guns could come into play.




Figure 2: Equivalent 8" rounds per minute for five historical ships.

Figure 2 does not tell the whole story, of course, since no information about accuracy or armor is given. Figure 2 does make it clear that the USS Enterprise and the USS New Jersey were very different weapon systems, and that the passage of 35 years produced an aircraft carrier much more powerful than the Enterprise.

2.3.2 Salvos of Dependent Shots

The firing errors dealt with in the previous section were *dispersion errors*, whereas in this section we assume the additional presence of a *bias error*.

Definition: *Dispersion errors* are firing errors that are independent and identically distributed among multiple shots. Such errors are sometimes also called ballistic errors.

Definition: *Bias errors* are errors common to all shots. Such errors are sometimes also called systematic errors or aiming errors.

Bias error might be due to a misalignment between the aiming and launching systems, to an error in target location, or to any other effect that introduces an error component common to all shots. The result is frequently that the impact points relative to the target are tightly grouped (indicating small dispersion errors) but in the wrong place due to bias, as in Figure 3. One can think of the bias error as being the center of gravity of the group, and of the dispersion errors as being deviations from the center of gravity. We shall use the notation that (σ_U, σ_V) are the (horizontal and vertical, say) standard deviations of the bias error, whereas the independent dispersion error for each shot has standard deviations (σ_X, σ_Y) .




Figure 3: Illustrating a pattern with small dispersion and large bias. This marksman might be more effective if he had a larger (!) dispersion.

It is no longer possible to proceed by first finding the single-shot kill probability and then invoking (2.13) to obtain a simple expression for P_K , since the required independence assumption is falsified by the bias error. The following example should make this clear.

Example 3: Suppose there are two shots, and consider a one-dimensional problem where the bias error is X and the dispersion errors are Y_1 and Y_2 for shots 1 and 2. To keep the computations simple, assume that all errors are either -1 , 0 or $+1$, rather than being normally distributed, and that the target will be killed if and only if the sum of the errors is 0 for some shot. Thus, the probability of killing the target with two shots is $P_K(2) = P(X + Y_1 = 0 \text{ or } X + Y_2 = 0)$. Assume that X is equally likely to be any of the three possibilities, whereas Y_1 and Y_2 are independent random variables with probabilities $1/6$, $2/3$, and $1/6$ of being -1 , 0 , or $+1$. Using the theorem of total probability, the single-shot kill probability is $P_K(1) = P(X + Y_1 = 0) = (1/3)(1/6) + (1/3)(2/3) + (1/3)(1/6) = 1/3$. Each of the two shots will kill the target with probability $1/3$. If we were to power up the miss probability, the kill probability for two shots would be $P_K(2) = 1 - (1 - P_K(1))^2 = 5/9$. However, powering up is not justified in this case because the two shots are not independent – they share the common error X . The correct answer can be obtained by conditioning on the value of X ; that is, by

powering up for each of the three possible values of X , and then averaging: $P_K(2) = (1/3)(1 - (5/6)^2) + (1/3)(1 - (1/3)^2) + (1/3)(1 - (5/6)^2) = 1/2$. It is typical for the faulty computation to produce a kill probability that is too high.

In general, let (U, V) be the two-dimensional bias error, and let $f(u, v)$ be its density function. If there are n shots and $Q_i(u, v)$ is the miss probability of the i th shot, given that the bias error is (u, v) , the expression for P_K is

$$P_K = E(1 - \prod_{i=1}^n Q_i(U, V)) = \iint (1 - \prod_{i=1}^n Q_i(u, v)) f(u, v) du dv \quad (2.18)$$

We will have little direct application for (2.18), the main reason being the complicated nature of the functions $Q_i(u, v)$, which in general will depend on the aim points for each of the n shots. In fact, we will find no simple exact expressions for P_K in this section. The best that can be hoped for, other than solutions to specific problems that are important enough to justify the work involved in evaluating (2.18) for all aiming patterns of interest, is some rules of thumb that take the form of approximations. In deriving these approximations, it will be convenient to imagine that the only source of bias is an error in target location. The approximations are actually valid regardless of the source of bias or even if there are several sources (see Section 2.3.4).

Our first approximation to P_K is an upper bound obtained by making two unrealistic assumptions that are clearly favorable to the marksman. One assumption is that there are no dispersion errors, and the other assumption is that the marksman can exchange his weapons for any other weapon or weapons with the same total lethal area. The second assumption permits the marksman to avoid all of the problems associated with trying to fill up space with circles, since he can reshape the lethal area in any convenient manner. In the circularly symmetric case where $\sigma_U = \sigma_V = \sigma$, the marksman would prefer to have a single large cookie-cutter megaweapon that he would aim directly at the target or more precisely at the mean location of the target. If the total lethal area of n weapons is na , then the lethal radius of such a megaweapon would be $R = \sqrt{na/\pi}$, and the resulting kill probability would be (from (2.7)) $1 - \exp(-R^2/2\sigma^2) = 1 - \exp(-na/(2\pi\sigma^2))$. More generally, the best megaweapon for our privileged marksman is a cookie-cutter with the same elliptical shape as the iso-probability contours of the error distribution. The resulting bound is

$$P_K \leq 1 - \exp(-z), \text{ where } z = \frac{na}{2\pi\sigma_U\sigma_V}. \quad (2.19)$$

Formula (2.19) was obtained by essentially assuming away all the overlap that is caused by dispersion errors, circle-packing problems, and noncookie-cutter weapons. The expression $1 - \exp(-z)$ should therefore be expected to be an accurate approximation in circumstances where overlap is expected to be a minor problem. Seven circles, for example, pack rather nicely into one circle without very much overlap. Figure 4, which is taken from the “Patterns” sheet of *Chapter2.xls*, shows the upper bound lying well above two other approximations that are introduced below.




Figure 4: The two locally random approximations, SULR and SOLR, lie well below the upper bound on kill probability.

A different kind of approximation is based on the idea that overlap is inevitable, and that one should expect the amount of overlap to be whatever happens “at random.” More precisely, the total lethal area na is assumed to be *in effect* so much confetti, with the marksman being able to control the density of confetti on a large “strategic” scale, but not the small-scale tendency of the flakes to overlap one another. Now, if d square inches of confetti are scattered uniformly at random on a one-inch square or in other words if the coverage ratio is d , then the fraction of the square that is covered is $1 - \exp(-d)$, as long as the flakes are sufficiently small. This expression is most easily derived by imagining that there are k independently located flakes of confetti with total area d . The probability that each flake covers the target is d/k , the fraction of the square covered by the flake. By powering up, the probability that some flake covers the target is $1 - (1 - d/k)^k$. The

limit of this expression as k approaches infinity is $1 - \exp(-d)$. This is true even if the density of confetti (d) depends on location. The marksman's problem is then to distribute a fixed amount of confetti (na) over the plane in such a manner that the (unconditional) kill probability is maximized.

Assume that $\sigma_U = \sigma_V = \sigma$, and that the marksman scatters all na units of confetti uniformly over a circle with radius r in the hope that some flake covers the target. This is the strategically uniform, locally random (SULR) case. Within the circle, the coverage ratio is $d = na / (\pi r^2)$. The probability that the target is actually in the circle is (from (2.7)) $1 - \exp(-r^2 / 2\sigma^2)$, so the unconditional kill probability is

$$p(r) \equiv [1 - \exp(-r^2 / 2\sigma^2)] [1 - \exp(-na / \pi r^2)]. \quad (2.20)$$

The first factor in (2.20) is 0 if $r = 0$, while the second is 0 as $r \rightarrow \infty$, so there must be a maximizing value for r . The value turns out to be $r^* = \alpha(4z)^{1/4}$, where $z = na / (2\pi\sigma^2)$, as can be verified by showing that $(d/dr)p(r^*) = 0$. Upon substituting r^* into (2.20), one obtains the strategically uniform locally random (SULR) formula

$$P_K = p(r^*) = (1 - \exp(-\sqrt{z}))^2. \quad (2.21)$$

Formula (2.21) also holds when $\sigma_U \neq \sigma_V$, provided that $z = na / (2\pi\sigma_U\sigma_V)$ and that the confetti is scattered uniformly over an optimally sized ellipse. Figure 4 shows that the SULR formula (2.21) provides a much smaller estimate of P_K than does (2.19).

The final approximation is the same confetti approximation, except that the coverage ratio is allowed to be any function $d(x, y)$ of two spatial coordinates, subject of course to being nonnegative and to the constraint that the total amount of confetti used must be na . This includes the case where $d(x, y)$ is constant within some region and 0 outside it, so we should expect the current approximation to be larger than (2.21). This is the strategically optimal, locally random (SOLR) case. Formally, the optimization problem is

$$\begin{aligned} & \text{maximize } \iint f(x, y) [1 - \exp(-d(x, y))] dx dy \\ & \text{subject to } d(x, y) \geq 0 \text{ for all } x, y \\ & \text{and } \iint d(x, y) dx dy = na, \end{aligned} \quad (2.22)$$

where $f(x, y)$ is the bivariate normal density function with standard deviations (σ_U, σ_V) . The solution can be found in Morse and Kimball (1950, Chapter 5), together with a discussion of how the optimal coverage ratio function $d^*(x, y)$ can

be used as a guide in designing effective patterns. The optimal function $d^*(x, y)$ is

$$d^*(x, y) = \frac{1}{2} \left(\sqrt{8z} - \frac{x^2}{\sigma_U^2} - \frac{y^2}{\sigma_V^2} \right)^+, \quad (2.23)$$

where the $+$ indicates that $d^*(x, y)$ is to be 0 rather than negative, and where as usual $z = na / (2\pi\sigma_U\sigma_V)$. Note that the confetti should be most dense at the origin, with the density falling off gradually to 0 on the $(8z)^{1/4}$ -standard deviation ellipse. Outside of this ellipse there should be no confetti at all. The result of substituting $d^*(x, y)$ into the objective function is the SOLR formula

$$P_K = 1 - (1 + \sqrt{2z}) \exp(-\sqrt{2z}), \quad (2.24)$$

The SOLR formula is also shown in Figure 4. There is not much difference between (2.21) and (2.24). Once the total lethal area has been conceptually reduced to confetti, it turns out not to be crucial that its distribution be exactly (2.23).

Example 4: Suppose that there are four weapons with cookie-cutter damage functions, with $R = 7.5$, and that the error standard deviations are $\sigma_U = \sigma_V = 7.5$, $\sigma_X = \sigma_Y = 1$. By exhaustive trial and error computations, it can be determined that the exact best pattern is a square of side 11.7, and that the associated kill probability is 0.80. Since $z = 4\pi(7.5)^2 / (2\pi(7.5)^2) = 2$, the three approximations are 0.865, 0.594 (SOLR), and 0.573 (SULR). The upper bound is considerably closer to the truth than either of the confetti approximations. The confetti approximations can be made to look better by letting the weapons be diffuse Gaussian with the same lethal area, in which case the approximations do not change but exact computations reveal that the best P_K is only 0.69, achieved by aiming the four weapons in a square of side 10. If the dispersion error is in addition increased from 1 to 5, the approximations still do not change, but the best possible P_K decreases to 0.62.

Since neither σ_X , σ_Y , nor any feature of the damage function other than lethal area enters the computation of z , it is clear that one could find cases where the actual kill probability is even smaller than the confetti approximations. In fact, one has only to consider any problem where the shots are nearly independent, since $z = \infty$ when σ_U or σ_V is 0. In problems where the bias errors dominate the dispersion errors, however, the confetti approximations can usually be thought of as lower bounds on P_K .

There is one other bias-compensation technique that is worthy of mention. Instead of aiming the several shots in a pattern, one can simply aim them all at the target, but inaccurately. A shotgun blast is an example of this; the marksman aims at only one place, but the shotgun pellets form a random pattern around it. Sheet

“SULR_SOLR” of *Chapter2.xls* demonstrates this for the SULR and SOLR approximations by generating the shot locations randomly from the appropriate density, and then testing whether any of the shots actually kills the randomly located target. This is a Monte Carlo simulation (Appendix C). If the deliberate inaccuracy itself takes the form of a bivariate normal error for each shot, the technique goes by the name of “artificial dispersion” and is amenable to analysis because only the sum of artificial and real dispersions is important. Significant work on determining the right amount of artificial dispersion was performed during World War II in Russia (Kolmogorov, 1948). The overall effect should be similar to one of the confetti approximations.

Given all the above considerations, we offer the following procedure for obtaining an approximate P_K in the general case where both bias and dispersion are present:

- (a) If dispersion dominates bias, determine the “equivalent” dispersion standard deviations $\sigma'_X = \sqrt{\sigma_X^2 + \sigma_U^2}$ and $\sigma'_Y = \sqrt{\sigma_Y^2 + \sigma_V^2}$ for each shot, and then use (2.13) to obtain an approximate P_K .
- (b) If bias dominates dispersion, and if the “packing problem” can probably be solved without much overlap (nearly cookie-cutter weapons, dispersion small compared to lethal radius as well as bias, etc.), use (2.19). This is an upper bound.
- (c) If bias dominates dispersion, and if it is clear that the best pattern will involve substantial overlap, use one of the confetti approximations.

The above rules are not exhaustive, since there are certainly cases where neither type of error dominates the other, and in any case the resulting estimate of P_K is only an approximation. An accurate P_K can only be obtained by evaluating (by Monte Carlo simulation, for example – see Appendix C and Exercise 8) sufficiently many patterns to be sure of having discovered the best one.

2.3.3 The Diffuse Gaussian Special Case

The case of the DG damage function with normally distributed errors is exceptional in that P_K can be evaluated analytically for any given pattern of shots, even when the shots have different parameters. As a result, it is possible to evaluate and even optimize a given pattern using a spreadsheet and some VBA code. This is an important enough capability that we include the required mathematics in this section, but the reader may wish to skip it in favor of sheet “DGPattern” of *Chapter2.xls*, which employs the mathematics in the form of VBA code, or to refer to the expanded and slightly generalized treatment in Washburn (2003a).

Formula (2.9) is a product of two factors, one depending on X-parameters and the other on Y-parameters. Let

$$K(b, \mu, \sigma) = \frac{b}{\sqrt{b^2 + \sigma^2}} \exp\left(-\frac{1}{2} \frac{\mu^2}{b^2 + \sigma^2}\right). \quad (2.25)$$

Then (2.9) states that $P_K = K(b_X, \mu_X, \sigma_X) K(b_Y, \mu_Y, \sigma_Y)$. Now, let n shots be aimed at (x_i, y_i) , $i=1, \dots, n$. If (U, V) is the bias error common to all shots, then the i th shot is offset from the target by $(\mu_X, \mu_Y) = (x_i + U, y_i + V)$, and the kill probability for the i th shot is therefore $P_i(U, V) = K(b_{X_i}, x_i + U, \sigma_{X_i}) K(b_{Y_i}, y_i + V, \sigma_{Y_i})$.

Here we have changed the name of the product because we wish to reserve the symbol P_K for the kill probability of the whole collection of n shots. Given the bias error, the shots are independent and we can employ (2.13), so

$$1 - P_K = E\left(\prod_{i=1}^n \{1 - P_i(U, V)\}\right). \quad (2.26)$$

In (2.26), which is an application of the theorem of total probability, the expected value is with respect to the normal distribution of the common bias error (U, V) . Note that each shot can potentially have distinct parameters for dispersion and lethality, with b_{X_i} and σ_{X_i} being the lethality and dispersion of the i th shot in the X -direction, etc.

The first step in evaluating (2.26) is to expand the product into a sum of terms, each of which is associated with one of the 2^n subsets of the n shots. For the subset S , the term is

$$t_S \equiv E\left(\prod_{i \in S} P_i(U, V)\right). \quad (2.27)$$

The plan is to analytically evaluate each of these terms, and then sum up all 2^n of them, paying proper attention to signs. Our first observation is that, since U and V are independent errors, and since $P_i(U, V)$ can be factored into an X -part and a Y -part, so can t_S . In fact, let

$$t(\mathbf{b}, \boldsymbol{\mu}, \boldsymbol{\sigma}, s, S) \equiv E\left(\prod_{i \in S} K(b_i, \mu_i + W, \sigma_i)\right), \quad (2.28)$$

where \mathbf{b} , $\boldsymbol{\sigma}$, and $\boldsymbol{\mu}$ are n -vectors indexed by shot number, and the expectation is taken with respect to a normal distribution of W that has standard deviation s . Then $t_S = t(\mathbf{b}_X, \mathbf{x}, \boldsymbol{\sigma}_X, \sigma_U, S) t(\mathbf{b}_Y, \mathbf{x}, \boldsymbol{\sigma}_Y, \sigma_V, S)$. Here $\mathbf{b}_X = (b_{X_1}, \dots, b_{X_n})$, and similarly $\boldsymbol{\sigma}_X, \mathbf{x}, \mathbf{b}_Y, \boldsymbol{\sigma}_Y$, and \mathbf{y} are also n -vectors. The central problem is now to evaluate $t(\mathbf{b}, \boldsymbol{\mu}, \boldsymbol{\sigma}, s, S)$, since this will lead to the evaluation of P_K through (2.26).

Let $M_k \equiv \sum_{i \in S} \mu_i^k / (b_i^2 + \sigma_i^2)$; $k = 0, 1, 2$, let $C \equiv \prod_{i \in S} \left(\frac{b_i}{\sqrt{b_i^2 + \sigma_i^2}} \right)$, and recall the definition of $K()$ in (2.25). The product in (2.28) is

$$\prod_{i \in S} K(b_i, \mu_i + W, \sigma_i) = C \exp \left(-\frac{1}{2} (W^2 M_0 + 2WM_1 + M_2) \right) \quad (2.29)$$

The expectation of (2.29) with respect to W can be accomplished analytically because (2.29) involves a quadratic expression in W lying within an exponential. The result is

$$t(\mathbf{b}, \boldsymbol{\mu}, \boldsymbol{\sigma}, s, S) = \frac{C}{\sqrt{1 + s^2 M_0}} \exp \left(-\frac{1}{2} \left(M_2 - \frac{s^2 M_1^2}{1 + s^2 M_0} \right) \right). \quad (2.30)$$

Calculating P_K is now just a matter of going through the steps outlined earlier. One can even include a shot reliability by modifying the C -factor. An implementation of the method outlined above is the VBA function $PK()$ included with *Chapter2.xls*. The method is essentially that of Bressel (1971), who goes on to take the additional step of averaging over a rectangular target. The $PK()$ function uses double-precision arithmetic because the series being summed alternates in sign. See Grubbs (1968) for alternatives with better numerical stability.

Since P_K can be evaluated analytically as a function of the aim point vectors \mathbf{x} and \mathbf{y} , we can now consider the problem of finding the pattern that will maximize P_K . This is a nonlinear optimization problem. On sheet “DGPattn” of *Chapter2.xls*, Excel’s Solver is employed to accomplish this.

2.3.4 Area Targets and Multiple Error Sources

Section 2.3.2 is often applicable even when there are multiple sources of error. Suppose, for example, that

- (a) the location of a target relative to some known datum is E_1 .
- (b) all shots are to be fired from a platform whose location relative to the same datum is subject to an error E_2 .
- (c) each shot has an individual firing error E_3 due to trembling on the part of the marksman.
- (d) an additional firing error is introduced due to an unknown wind velocity E_4 .
- (e) E_1, E_2, E_3 , and E_4 are all independent, normal random variables with 0 mean and variances σ_i^2 , $i = 1, 2, 3, 4$.

It is necessary to classify each of the four errors as either “bias” or “dispersion.” E_1 and E_2 are clearly bias errors, since we assume that the positions of the target and the firing platform are the same for each shot. E_3 is clearly a dispersion error, since each shot has an independent dispersion that is different from all the rest. E_4 might be a bias error if the unpredictable part of wind velocity were constant in space over the length of time required to fire the shots (the predictable part is irrelevant, since the marksman could allow for it in aiming) or it might be a dispersion error if the wind were very gusty. Assume the latter. Then, making the natural assumption that the four error types are independent of each other, and noting that it is only the total bias error and the total dispersion error that affect the fate of the target, the equivalent bias and dispersion error variances are $\sigma_1^2 + \sigma_2^2$ and $\sigma_3^2 + \sigma_4^2$, respectively, and Section 2.3.2 can be applied to the equivalent errors. The principle being used is the theorem that the variance of a sum of independent random variables is the sum of the variances.

All of the above applies to targets whose only property is a location; i.e., to point targets. However, it is not difficult to handle area targets within this scheme, provided we are interested only in the average fraction of the target killed. In (2.18), if $f(u, v)$ is interpreted to be the amount of target value per unit area at point (u, v) , then P_K has the meaning “total amount of value killed, on the average.” Furthermore, if we normalize $f(u, v)$ so that the total target value is unity, then P_K means “average fraction of the target killed,” or, equivalently, the probability of killing a test element of the target, where the location of a test element is governed by the density function $f(u, v)$. In other words, the test element location (U, V) can be regarded as just another bias error, to be combined with any other bias errors as necessary. Any area target can be handled by converting the value density of the area target to an equivalent density function of a bias error, and then proceeding as if the target were a point target. This is especially easy to do, of course, if $f(u, v)$ happens to be bivariate normal.

Example 5: Suppose that (U, V) is circular normal with standard deviation 80 ft; that is, assume that the target’s value is spread out in a bell-shaped manner. Also assume that E_1 , E_2 , E_3 , and E_4 are all circular normal with standard deviations 10, 20, 30, and 40 ft, respectively. Assuming that the wind error is dispersion, the equivalent dispersion is $\sigma_X = \sigma_Y = \sqrt{30^2 + 40^2} = 50$ ft, and the equivalent bias is $\sigma_U = \sigma_V = \sqrt{10^2 + 20^2 + 80^2} = 83$ ft. One could now proceed as in Section 2.2, probably by ignoring the dispersion error and using the SOLR formula to estimate P_K , which is now interpreted as the maximum possible expected fraction of the target killed by an optimal pattern. If the pattern consists of 20 shots with lethal radius 10 ft each, we find that $z = 20(10)^2 / (2(83)^2) = 1/69$. The resulting P_K is very small, even according to the optimistic (2.19). Twenty shots of this size are simply not capable of doing much damage to a target as spread out as the one postulated.

The fact that area targets introduce an effective bias error is important in determining whether CMP or EMT is a better measure of effectiveness for an arsenal of weapons (see Section 2.3.1). Since (2.14) was derived under the assumption that the only firing error was dispersion, we can say that CMP is the proper measure if the targets are point targets and if the bias errors are negligible. If the effective bias (including the effects of target size) dominates the effective dispersion, however, then EMT is more appropriate. Thus (to conclude the comparison that was begun in Section 2.3.1), the United States nuclear arsenal in 1978 was more effective against well located, hard targets such as ICBM silos, but the Soviet Union arsenal was more effective against cities, which are well-located area targets, or against submarines, which are poorly located point targets. Dispersion is almost irrelevant for either of these latter target types, even though it is crucial for the former.

2.4 Multiple Shots, Multiple Targets, One Salvo

In this section we consider situations where there are multiple targets, all of which must be attacked “simultaneously” in the sense that no feedback about results is available between shots (the next chapter considers feedback). If several shots are available, the question of how they ought to be distributed over the targets arises.

2.4.1 Identical Shots, Identical Targets, Optimal Shooting

If all shots in a salvo are identical and all targets are identical, then any reasonable measure of effectiveness will be maximized when the shots are spread as evenly as possible over the targets. If the number of shots is an integer multiple of the number of targets, then the number of surviving targets is a binomial random variable, since all targets have the same survival probability. The situation is somewhat more complicated when the number of shots is not an integer multiple of the number of targets, since some targets are shot at one more time than others. Suppose there are b shots and n targets, with each shot having a miss probability q against its target. Let k be the integer part of b/n , so that some targets are shot at k times while others (r of them) are shot at $k+1$ times; r is the number of shots left over after every target is shot at k times, so $r = b - kn$. The total number of surviving targets is the sum of two binomial random variables, the number of survivors from the r targets shot at $k+1$ times and the number of survivors from the $n-r$ targets shot at k times. The distribution of the total number of survivors X is therefore the convolution of two binomial distributions. The probability mass function of X is given by

$$Surv(x; b, n, q) \equiv \sum_{j=\max(0, x-n+r)}^{\min(r, x)} Bin(j; r, q^{k+1}) Bin(x-j; n-r, q^k), \quad (2.31)$$

where $Bin(y; t, p)$ is the binomial probability of y successes out of t trials when the success probability is p . The average number of survivors is

$$E(X) = rq^{k+1} + (n-r)q^k. \quad (2.32)$$

The function $Surv()$ is included in module 1 of *Chapter2.xls*.

Example 6: Suppose that $n = 3$ targets are attacked by $b = 4$ weapons, each of which has an individual miss probability of $q = 0.6$. Then $k = 1$ and $r = 1$, so one target is attacked twice while the rest are attacked only once. The probability of three survivors is $q^4 = 0.1296$, since all four shots must miss if all targets are to survive (the sum in (2.30) contains only the term for $j = 1$). The rest of the probability distribution of X can best be obtained using the $Surv()$ function defined in (2.30). From (2.31), the average number of survivors is 1.56. See sheet “Surv” of *Chapter2.xls*.

2.4.2 Identical Shots, Diverse Targets, Optimal Shooting

Here we consider the case where the shots are all identical, but not the targets. This kind of situation is roughly the case in planning a nuclear first strike, so considerable work was done on the problem during the Cold War.

When targets are diverse, there are many possible measures of effectiveness. We will assume that each target has a given value, and that the marksman’s goal is to kill as much value as possible, on the average. Determining target values is in practice problematic, but nonetheless necessary if an optimal attack is to be planned – only scalar measures can be optimized, and assigning values is the most direct way to put all targets on the same scale. Let v_j be the value of target j , and let q_j be the probability that one shot at target j will fail to kill it (the kill probability is $1 - q_j$).

Since all shots are assumed to be independent, the probability that x_j shots will fail to kill target j is $q_j^{x_j}$. If there are n targets and b shots in total, this leads to the problem of minimizing the total average surviving value, subject to not using more shots than are available, which we name problem P1:

$$\text{Minimize } \sum_{j=1}^n v_j q_j^{x_j}, \quad (2.33)$$

subject to $\sum_{j=1}^n x_j \leq b$, and all variables x_j must be nonnegative integers.

Problem P1 has a nonlinear objective function, but nonetheless turns out to be an easy optimization. It can be solved by a “greedy” technique where shots are allocated sequentially, with each shot being allocated to the target for which the incremental target value killed is maximal. Imagine that the problem has already been solved, with \mathbf{x} being the vector of optimal allocations, and that one more shot becomes available. If the shot is allocated to target j , the decrease in surviving value will be $v_j q_j^{x_j} - v_j q_j^{x_j+1} = v_j q_j^{x_j} (1 - q_j)$. The greedy technique starts out with $\mathbf{x} = 0$, and then follows the principle “always assign the next shot to the target for which the decrease in surviving value is largest.” The procedure is efficient, and can easily find optimal allocations when targets and shots number in the thousands.

Example 7: Suppose there are three targets, all with value 1, and $\mathbf{q} = (0.1, 0.5, 0.9)$. Using the greedy method, the first 4 shots should be at targets 1, 2, 2, and 2, with the total surviving value after 4 shots being $0.1 + 0.125 + 0.9 = 1.125$. Note that most shots go to the target that is neither too easy nor too hard to kill. Target 1 is so easy to kill that a single shot will suffice, and target 3 is so hard to kill that shots against it are nearly wasted. The fifth shot, however, should be on either target 1 or target 3, since either shot will reduce the surviving value by 0.09, which exceeds the 0.0625 reduction available from target 2.

2.4.3 Diverse Shots, Diverse Targets, Optimal Shooting

In the 1991 Gulf War, the allies were confronted with the problem of attacking Iraq’s air defense system in a single, coordinated attack. An air defense system consists of multiple targets, and weapons of various kinds (aircraft, helicopters, cruise missiles) were available to attack it. The planners of that attack had to decide which weapons should attack which targets, and to do so without the prospect of getting much information feedback, since it was imperative that the attack happen quickly. This section deals with such situations, but one important aspect of that attack will be missing in this chapter. The actual attack anticipated losses to the attackers, with the losses depending very much on who attacked what. Such two-sided models will not be considered until Chapter 4.

Since the shots are diverse, the miss probability now depends on both the target and the weapon. Let q_{ij} be the probability that a weapon of type i will miss target j if assigned to attack it. All attacks are assumed to be independent, so powering up still applies. If x_{ij} weapons of type i are assigned to attack target j , the miss

probability is therefore $\prod_i q_{ij}^{x_{ij}}$. The object is now to solve problem P2 defined below:

$$\begin{aligned} &\text{Minimize } \sum_j v_j \prod_j q_{ij}^{x_{ij}}, \\ &\text{subject to } \sum_{j=1}^n x_{ij} \leq b_i \text{ for all weapons } i, \text{ and all variables nonnegative integers.} \end{aligned} \quad (2.34)$$

The greedy method of always making the next assignment to reduce the total surviving value as much as possible does not work on problem P2, as shown in the following example:

Example 8: Suppose there are 2 targets, each with value 1, and 2 weapons. Weapon A has miss probabilities of 0 and 0.1 on the 2 targets, while weapon B has miss probabilities of 0.1 and 0.5. The greedy method would first assign A to target 1, since A will certainly kill it, and would then assign B to target 2. The total surviving value would be $0 + 0.5 = 0.5$. The optimal allocation is just the reverse – if B is assigned to 1 and A is assigned to 2, the surviving value is only $0.1 + 0.1 = 0.2$. Weapon A is more effective than weapon B against all targets, and the greedy mistake is to lose sight of all other facts in assigning weapon A. A realistic version of this example might have cruise missiles playing the role of weapon A. Effective but expensive weapons such as cruise missiles should be used on high-value targets that are hard for other weapons to kill, so one needs to have a global perspective in assigning them.

Problem P2 is sometimes referred to as the weapon target assignment problem. It is known to be of a fundamentally difficult type, so no simple method like the greedy method will be able to solve it. Optimal methods can be expected to be lengthy on large problems, but a variety of efficient approximations and bounds have been investigated (Ahuja et al., 2003; Washburn, 1995a).

2.4.4 Identical Shots, Identical Targets, Random Shooting

Combat sometimes happens fast enough or communications are sufficiently difficult that deliberate attempts to jointly optimize fires are impossible. Here we investigate the consequences if several identical shots are randomly aimed at several identical targets, rather than optimally aimed as in Section 2.4.1. The assumption of randomness reflects the idea that coordination is impossible, so that a marksman is as likely to shoot at one target as another, independent of the other marksmen.

Thomas (1956), in the context of a situation where several bombers are suddenly ambushed by a group of interceptors, derives the probability distribution of

the number of surviving bombers. Using the notation of Section 2.4.1, let $RSurv(x; b, n, q)$ be the probability that x targets survive when b shots with individual miss probability q are randomly distributed over n targets. Thomas shows that the probability mass function of X , the number of survivors, is given by

$$RSurv(x; b, n, q) \equiv \binom{n}{x} \sum_{j=0}^{n-x} (-1)^j \binom{n-x}{j} \left(1 - \frac{(1-q)(x+j)}{n} \right)^b; \quad x = 0, \dots, n \quad (2.35)$$

The function $RSurv()$ is also provided in *Chapter2.xls*. See Exercise 13.

The expected value of X is comparatively easy to derive. Consider any particular target. That target will be killed by any particular shot with probability $(1-q)/n$, since the target is chosen by the shot with probability $1/n$. Since there are b shots, each of which has an independent chance of killing the target, we can power up to find the target's survival probability, which is $(1 - (1-q)/n)^b$. Now, let I_x indicate whether target x is killed ($I_x = 1$) or not ($I_x = 0$), so that $X = I_1 + \dots + I_n$. We know that $E(I_x) = (1 - (1-q)/n)^b$ for every x . Since expected values and sums commute, even when the summed random variables are dependent, as they are here, the expected number of survivors is

$$E(X) = n \left(1 - \frac{1-q}{n} \right)^b. \quad (2.36)$$

A comparison between random and uniform allocation of shots is of interest for the case where all shots and targets are otherwise identical. We expect that (2.32) will show a larger average number of survivors than (2.36), since surely there is no advantage to the marksman for aiming randomly. This turns out to be true. When $q = 0.6$ and $n = 100$ (or any large number), the average survival probability $E(X)/n$ is shown in Figure 5 for varying numbers of shots per target (k). The curves differ, but by surprisingly little, especially when k is either small or large. See sheet "OptRand" of *Chapter2.xls* to experiment with changes to n or q .




Figure 5: $E(X)/n$ for optimal and random shooting, $q = 0.6$.

2.5 Further Reading

A great deal of work on aiming and kill probability computation was inspired by World War II. Morse and Kimball (1950, Chapter 5) record some methods used in the United States during that war. Wartime Soviet work on artificial dispersion and other topics can be found in Kolmogorov (1948). Eckler and Burr (1972) is a good summary of subsequent work up to 1972, or see Przemieniecki (2000), which includes some basic computer programs for computing kill probability.

Although this chapter emphasizes firing problems in two dimensions, there are also one-dimensional problems of interest. Morse and Kimball (1950, Chapter 5) consider the problem of bombing a railroad track, and other linear structures such as roads or power lines may also prompt analysis. One advantage of one-dimensional problems is that the two-dimensional difficulty of packing the plane with circles disappears. With cookie-cutter weapons, in the absence of significant dispersion errors, the best pattern in one dimension is typically a “stick” that covers an interval with no gaps. Morse and Kimball observe this, and David and Kress (2005) generalize to situations where the bombs are heterogeneous and asymmetric in their effects.

The first few sections of this chapter treat firing errors as being one of two types. The first type is where all firing errors are mutually independent, and the second type adds the possibility that there may be a common component. Neither of these assumptions fit well for rapid-fire weapons, where the sequence of errors is better viewed as a stochastic process. Work on applications of such processes


was inspired by aerial combat in World War II in Great Britain, and subsequently reported in Cunningham and Hynd (1946). See also Fraser (1951, 1953).

There is much in common analytically between the problem of firing at a set of diverse targets and the problem of trying to detect a stationary target that is lost among a set of cells. Just as shots are allocated to targets with the object of killing them, units of search effort are allocated to cells with the object of detecting the lost target. Problem P1 in Section 2.4.1 serves both points of view and has benefited analytically from analysts in both camps. Perusal of the part of search theory literature devoted to stationary targets may prove fruitful to someone interested in firing theory. See Stone (1975), for example.

All military services include killing targets among their duties, and many weapons are shared among the services. It therefore makes sense to have a joint organization responsible for determining how to calculate kill probabilities. In the United States, this is the Joint Technical Coordinating Group for Munitions Effectiveness, or JCTG/ME. The principal products of the JTCG/ME are Joint Munitions Effectiveness Manuals (JMEMs), some of which are in the form of standardized computer programs. JMEMs employ some of the methods covered earlier in this chapter, but also other methods that are more computer intensive and include more detail. These manuals (there are hundreds of them) also include classified data for specific weapons and targets. At the unclassified level, a good way to find out more about JCTG/ME or JMEM is to go to the Federation of American Scientists (FAO, 2008) and search on either of those acronyms.

Exercises

- (1) Suppose $D(r) = 1 - r$ if $r < 1$; 0 if $r > 1$. What is the lethal area?
Ans. $a = \pi/3$
- (2) Plot $D(r)$ for the target illustrated below, assuming that the weapon must hit the shaded area and that the impact point is (r, θ) with θ uniformly random in $[0, 2\pi]$. Show that the lethal area is equal to the area of the target.
Ans. $D(r)$ is a step function, $a = 2.5\pi$



- (3) Show that (2.4) produces $\pi E(R^2)$ for lethal area, where $E(R^2)$ is computed using the density function given by (2.6). Hint: use integration by parts.
- (4) Derive (2.7).
- (5) When aiming errors are basically angular, the miss distances should increase with range. Suppose several independent shots are taken at a target, with $\sigma_i = 0.1r_i$, where r_i is the i^{th} range, and that the cookie-cutter lethal radius is 1. If the successive ranges are 10, 11, 12, ..., increasing by unity with each shot, compute P_K for the first shot, the first five shots as a group, and the first 10 shots as a group.
Ans. ($P_K(1) = 0.39$, $P_K(5) = 0.84$, $P_K(10) = 0.93$).
- (6) A weapon with a radially symmetric DG damage function is aimed at a terrorist camp that is located 100 m from a baby-milk factory. The weapon has different lethalties for camp and the factory. The parameter b (the common value of b_X and b_Y) is 50 m for the camp and 80 m for the factory. The circular normal firing error has standard deviation 50 m. What are the probabilities of killing the camp and the factory?
Ans. (0.50 for the camp, 0.41 for the factory)
- (7) An aircraft attempts to kill a tank as follows: It first drops a canister of "stickers" in the hope that one will hit the tank and activate. If a sticker acti-

vates, it can guide a projectile to the tank. The canister opens and scatters 1000 stickers, with the amount of scatter being under the control of the designer. The exposed area of the tank is 100 m^2 . The aircraft makes a two-dimensional error with standard deviations (100 m, 300 m) in dropping the canister. What is the probability that a sticker hits the tank, assuming a well-designed canister? If the tank is longer than it is wide, does the direction of the aircraft's approach matter?

Ans. (0.275, no)

- (8) Suppose you are given 16 detection devices, each of which is guaranteed to detect a target if and only if the relative distance is either less than 4 miles or between 30 and 33 miles (the “convergence zone” phenomenon in the ocean might be one explanation for such an assumption). The devices can be placed in any pattern whatever, and the object is to detect a target whose location relative to some known point is circular normal with standard deviation 30 miles in each direction. There are no dispersion errors.

(a) Estimate p_K .

(b) Make up a pattern and test it by writing a 5000-replication computer simulation.

Ans. The lethal area is $\pi(4^2 + 33^2 - 30^2) = 205\pi$ square miles, so $z = (16)(205)/1800 = 1.82$. Given that the shape of the lethal area makes considerable overlap inevitable even in the absence of dispersion, a confetti approximation is natural. The SOLR formula produces $P_K \approx 0.57$. This example has been the result of considerable experimentation, with the best pattern as of this writing having a detection probability of 0.64. The exact answer is unknown.

- (9) Suppose 10 cookie-cutter shots are available, with the lethal radius being 30 ft for each. Estimate P_K for the area target and errors considered in Section 2.3, assuming that

(a) the wind error is dispersion

(b) the wind error is bias

Ans. Using the SOLR formula in both cases, the expected fraction of the target killed with an optimal pattern would be approximately 0.316 in case a), or 0.275 in case (b).

- (10) Use sheet “DGPattern” of *Chapter2.xls* to verify the two DG claims made in Example 4. Hint: To make the weapons have the right lethal area, you must set the DG lethality parameter b so that $2\pi b^2 = \pi(7.5)^2$.

- (11) Consider problem P1 of Section 2.4.2, with $n = 3$, $\mathbf{v} = (1, 2, 3)$, and $\mathbf{q} = (0.3, 0.5, 0.8)$. Allocate the first four shots using the greedy algorithm.

Ans. The shots should be in the order 2, 1, 3, 2, and the average value surviving should be 3.2.

- (12) Write a computer program to solve problem P1 of Section 2.4.2, using the greedy algorithm. Use it to solve Example 7.
- (13) Insert a new worksheet into *Chapter2.xls* and use it to analyze a situation where there are $n = 3$ targets and $b = 4$ shots, each of which has a miss probability of $q = 0.6$.
- (a) Use the $Surv(x; b, n, q)$ function to compute the probability distribution of the number of survivors when b shots with miss probability q are distributed evenly over n targets. Use that distribution to compute the mean number of survivors, and verify that (2.31) gives the same result.
- (b) Use the $RSurv(x; b, n, q)$ function to compute the probability distribution of the number of survivors when b shots with miss probability q are distributed randomly over n targets. Use that distribution to compute the mean number of survivors and verify that (2.35) gives the same result.
- Ans. A partial answer is that the mean number of survivors should be 1.56 and 1.69 in the even and random cases, respectively. See sheet “Surv” of *Chapter2.xls* for a solution.
- (14) If formulas 2.7 and 2.8 are both true, show that CEP must be related to σ as claimed in Section 2.2.2.

- (15) Consider problem P2 of Section 2.4.3, with $\mathbf{v} = (3, 1)$ and $\mathbf{q} = \begin{bmatrix} 0.5 & 0.0 \\ 0.6 & 0.5 \\ 0.7 & 0.5 \end{bmatrix}$.

There is a single weapon of each of the three types, and you are to assign weapons to maximize the total value killed, or equivalently to minimize the total surviving value. The best that can be done with one weapon is to assign weapon 1 to target 1, since this will kill $3(0.5)$ units of target value. Complete the application of the greedy algorithm to this problem, find the optimal assignment, and compare the two.

Ans. The greedy algorithm results in weapons 1 and 3 assigned to target 1, with a total surviving value of $3(0.5)(0.7) + 1(0.5) = 1.55$. The optimal assignment has weapons 2 and 3 assigned to target 1, with a total surviving value of $3(0.6)(0.7) + 1(0) = 1.26$.

- (16) Page “Shotgun” of *Chapter2.xls* is a Monte Carlo simulation of a shotgun with 50 pellets trying to hit a target in the face of bias error. Relative to the bias, each pellet hits independently with a uniformly distributed angle and a dispersion distance X whose cumulative distribution function is $F(x) = 1 - \exp(-(x/b)^3)$. This is a Weibull distribution, which is often realistic and always easy to simulate because it is possible to solve the equation for x . Now suppose that the only control you have over the pattern is through the dispersion parameter b . This control might be exercised through a choke on a shotgun or through the dispersement altitude if the pellets are actually bomblets. Experiment with the sheet to find the best value of b when the bias

standard deviation is 50 and the lethal radius is 10 and compare the resulting kill probability with the SOLR formula. Use *SimSheet.xls* to replicate the experiment a large number of times.

Ans. With b set to 50, 60, 80, and 100, the kill probability after 3000 replications is approximately 0.376, 0.416, 0.393, and 0.328, respectively. The best setting for b is about 60, and the SOLR formula comes close to predicting the optimized kill probability.

- (17) The DG damage function (2.9) can be reduced to one dimension (the x -dimension, say) by setting $\mu_Y = \sigma_Y = 0$, so Section 2.3.3 and its accompanying sheet “DGPattn” of *Chapter2.xls* also apply to finding multiple-shot patterns in one dimension. Use that sheet to find the optimal placement of four identical bombs with unit reliability, $b_X = 30$ and $\sigma_X = 30$. Assume that the common error has standard deviations of $(\sigma_U, \sigma_V) = (100, 0)$; that is, there is no bias error in the vertical dimension. In using that sheet, the value of b_Y is irrelevant as long as it is not 0.

Ans. The optimized pattern is $(-100, -30, 30, 100)$, with a kill probability of 0.68. You will not discover that fact if you start at $(0, 0, 0, 0)$, since all derivatives are 0 there and Solver will therefore not move the aim points, so use some other starting point. If you increase the dispersion to $\sigma_X = 100$ and re-optimize, all four bombs should be aimed at the origin. This is typical when the dispersion error is large.



<http://www.springer.com/978-1-4419-0789-9>

Combat Modeling

Washburn, A.R.; Kress, M.

2009, VIII, 280 p. 54 illus., Hardcover

ISBN: 978-1-4419-0789-9