

Optimal Design, Robustness, and Risk Aversion

M. E. J. Newman
Michelle Girvan
J. Dooyne Farmer

SFI WORKING PAPER: 2002-02-009

SFI Working Papers contain accounts of scientific work of the author(s) and do not necessarily represent the views of the Santa Fe Institute. We accept papers intended for publication in peer-reviewed journals or proceedings volumes, but not papers that have already appeared in print. Except for papers by our external faculty, papers must be based on work done at SFI, inspired by an invited visit to or collaboration at SFI, or funded by an SFI grant.

©NOTICE: This working paper is included by permission of the contributing author(s) as a means to ensure timely distribution of the scholarly and technical work on a non-commercial basis. Copyright and all rights therein are maintained by the author(s). It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may be reposted only with the explicit permission of the copyright holder.

www.santafe.edu



SANTA FE INSTITUTE

Optimal design, robustness, and risk aversion

M. E. J. Newman,¹ Michelle Girvan,^{1,2} and J. Doyne Farmer¹

¹*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501*

²*Department of Physics, Cornell University,*

Clark Hall, Ithaca, NY 14853-2501

(Dated: February 19, 2002)

Abstract

Highly optimized tolerance is a model of optimization in engineered systems, which gives rise to power-law distributions of failure events in such systems. The archetypal example is the highly optimized forest fire model. Here we give an analytic solution for this model which explains the origin of the power laws. We also generalize the model to incorporate risk aversion, which results in truncation of the tails of the power law so that the probability of disastrously large events is dramatically lowered, giving the system more robustness.

In a series of recent papers, Carlson and Doyle [1–3] have proposed a model for designed systems which they call “highly optimized tolerance” or HOT. The fundamental idea behind HOT is that systems designed for high performance naturally organize into highly structured, statistically unlikely states that are robust to perturbations they were designed to handle, yet fragile to rare perturbations and design flaws. As an example they consider an idealized model of forest fires [2]. In this model a forester is charged with finding the optimal distribution of the trees on a grid so as to maximize tree harvest in the face of occasional fires that burn complete connected clusters of trees and are started by sparks that arrive with a given spatial distribution. They find that optimizing the harvest, or yield, for the model gives rise to a segmented forest consisting of contiguous patches of trees separated by firebreaks, and that the resulting distribution of fire sizes usually follows a power law. While this type of configuration does typically achieve very good yields, the system is also fragile in the sense that perturbations to the firebreaks or changes in the spark distribution can lead to substantially sub-optimal performance. They argue that these are pervasive phenomena: high-performance engineering leads to systems that are robust to stresses for which they were designed but fragile to errors or unforeseen events.

In this paper we argue that simple yield maximization is problematic even if there are no errors in firebreaks or changes in the spark distribution. Because the power-law distributions generated by yield maximization have fat tails, disastrously large forest fires occur with non-negligible frequency—far greater frequency than one would expect from intuition based on normal distributions. This idea, that yield optimization can lead to ruinous outcomes, is not new. For the classic problem of gambler’s ruin, for example, it is well known that optimizing total return leads to ruin with probability one. By contrast, if one is willing to accept suboptimal returns it is possible to construct gambling strategies that are immune to ruin [4]. Applying similar ideas in the present context, we show that a risk-averse engineer who is willing to accept some loss in average system performance can effectively limit the large deviations in the event size distribution so that disasters are rare. We call this variation on the HOT theme “constrained optimization with limited deviations”, or COLD. By avoiding total ruin, a COLD design is more robust than a HOT one, even in a world of perfect error-free optimization.

To demonstrate the difference between HOT and COLD we first revisit the HOT forest

fire model. We give an analytic solution for the model which shows that the distribution of fire sizes does indeed follow a power law, the cumulative distribution having a exponent $-(1+1/d)$, where d is the dimensionality of the system. Using a fast percolation algorithm [5] we perform numerical simulations that confirm the value of this exponent. We then generalize our solution to include risk aversion in the design, thereby breaking the power law scaling and dramatically reducing the frequency of disastrously large events.

Following Refs. 1 and 2 then, we consider a forest divided into a large number of regions or patches, with firebreaks between them that prevent the spread of fire from one patch to another. Although the original forest fire model was based on a lattice, the model we consider is a continuum one, since this makes the mathematical treatment more tractable. For large system sizes, we expect the behavior of this continuum model to converge to that of the lattice model.

It is assumed that during the lifetime of the forest a single spark lands at a random position \mathbf{r} and starts a fire that burns the surrounding patch. Let us denote the area of this patch by $s(\mathbf{r})$. The forest is then harvested, giving a yield equal to the area of the remaining forest. In units where the total area of the forest is one, the yield is $1 - s(\mathbf{r}) - F$, where F is the cost in terms of yield of constructing the firebreaks.

Because dimensionality is an important property of HOT systems, we consider the model for general dimension d . If the cost of constructing firebreaks is a per unit length (or per unit surface area for $d > 2$), then the cost of the firebreak surrounding a patch m is $agd s_m^{(d-1)/d}$, where s_m is the value of $s(\mathbf{r})$ in patch m and g is a geometric factor of order 1 that depends on the geometry of the lattice and the shape of the patches. In the lattice version of the forest fire model, a is simply equal to the lattice parameter (i.e., the nearest-neighbor spacing), but in the continuum model we are at liberty to give a any value we feel to be appropriate. As we will see, as long as a is finite its value does not affect the shape of the distribution of fire sizes.

Because $s(\mathbf{r})$ is constant inside each patch, the integral of $1/s(\mathbf{r})$ over any patch is identically 1, and hence, summing over all patches, the total area occupied by firebreaks is

$$F = agd \sum_m s_m^{(d-1)/d} = agd \sum_m s_m^{(d-1)/d} \int_m \frac{d^d r}{s(\mathbf{r})} = agd \int s(\mathbf{r})^{-1/d} d^d r. \quad (1)$$

Letting the normalized probability distribution of sparks be $p(\mathbf{r})$, the mean yield is then

$$Y = 1 - \int p(\mathbf{r})s(\mathbf{r}) \, d^d r - agd \int s(\mathbf{r})^{-1/d} \, d^d r, \quad (2)$$

where the integrals run over the entire area of the forest.

To find the maximum yield with respect to the patch sizes $s(\mathbf{r})$, we set the functional derivative $\delta Y/\delta s(\mathbf{r}) = 0$, giving

$$s(\mathbf{r}) = \left[\frac{ag}{p(\mathbf{r})} \right]^{d/(d+1)}, \quad (3)$$

for all \mathbf{r} . The optimal yield is then given by substituting back into Eq. (2) to get

$$Y_{\text{opt}} = 1 - (d+1)(ag)^{d/(d+1)} \int p(\mathbf{r})^{1/(d+1)} \, d^d r, \quad (4)$$

and the optimal number of patches is

$$n = \int \frac{d^d r}{s(\mathbf{r})} = (ag)^{-d/(d+1)} \int p(\mathbf{r})^{d/(d+1)} \, d^d r. \quad (5)$$

For the lattice version of the model, a goes as L^{-1} , where L is the (linear) system size, and hence the number of patches should scale as $L^{d/(d+1)}$, i.e., as $L^{1/2}$ in one dimension or $L^{2/3}$ in two. Numerical experiments on a one-dimensional system confirm this, giving a measured exponent of 0.47 ± 0.03 .

Now we wish to calculate the distribution $\rho(s)$ of fire sizes that arises if we make this choice of patch sizes. We have

$$\rho(s) = p(\mathbf{r}) \frac{d^d r}{ds} = p(\mathbf{r}) \frac{d^d r}{dp} \frac{dp}{ds} = -ag \frac{d+1}{d} p(\mathbf{r}) \frac{d^d r}{dp} s^{-(2+1/d)}, \quad (6)$$

where $d^d r$ here represents the volume of the space between the contours p and $p+dp$ on the $p(\mathbf{r})$ surface. As we will show, the term $p(\mathbf{r}) \, d^d r/dp$ is constant or contributes logarithmic corrections for a wide selection of possible distributions $p(\mathbf{r})$, while the principal power-law behavior in the event size distribution comes from the factor $s^{-(2+1/d)}$ [8].

An alternative method for deriving Eq. (6) is to maximize the simple yield functional $Y = 1 - \int p(\mathbf{r})s(\mathbf{r}) \, d^d r$, subject to a constraint that fixes the volume F occupied by the firebreaks (Eq. (1)). This method, which is similar to the approach taken in Ref. 1, is equivalent to the method above, via a Lagrange transform, provided F is chosen so as to make the corresponding Lagrange multiplier equal to agd .

Consider then the case (which covers all the examples in Refs. 1 and 2) of a distribution of sparks with a single maximum at the origin, so that the volume $d^d r$ takes the form of an annulus enclosing the origin. If we denote by Ω_d a d -dimensional solid angle centered on the origin, then a volume element of the annulus is $r^{d-1} dr d\Omega_d$. In terms of the thickness dp of the annulus, we can write $dr = (r dp)/(\mathbf{r} \cdot \nabla p)$, and integrating our volume element over the contour $p = \text{constant}$ we find

$$\frac{d^d r}{dp} = \oint_p \frac{r^d d\Omega_d}{\mathbf{r} \cdot \nabla p}. \quad (7)$$

For example Carlson and Doyle [1] studied the case of a spark distribution in two dimensions having the form of the product of two Gaussians with different widths:

$$p(\mathbf{r}) = N \exp\left(-\left[\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2}\right]\right), \quad (8)$$

where N is a normalization constant. For this distribution the denominator of the integrand in Eq. (7) is

$$\begin{aligned} \mathbf{r} \cdot \nabla p &= -N \left[\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} \right] \exp\left(-\left[\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2}\right]\right) \\ &= 2p \log \frac{p}{N}, \end{aligned} \quad (9)$$

which is constant over our contour of constant p . The element of solid angle in two dimensions is simply the element of polar angle $d\theta$, and hence Eq. (7) simplifies in this case to

$$\frac{d^2 r}{dp} = \frac{1}{2p \log(p/N)} \oint_p r^2 d\theta = \frac{A(p)}{p \log(p/N)}, \quad (10)$$

where $A(p)$ is the area enclosed by the contour. This contour is a line of constant $(x/\sigma_x)^2 + (y/\sigma_y)^2$, i.e., an ellipse, which has major and minor axes $a = \sqrt{2\sigma_x^2 \log(N/p)}$ and $b = \sqrt{2\sigma_y^2 \log(N/p)}$. Thus the area enclosed by the contour is $A(p) = \pi ab = 2\pi\sigma_x\sigma_y \log(N/p)$. Combining Eqs. (6) and (10) we then find that the distribution of event sizes is

$$\rho(s) = 3\pi\sigma_x\sigma_y a g s^{-5/2}. \quad (11)$$

Thus, for the Gaussian case in two dimensions the model generates a perfect power-law with slope $-\frac{5}{2}$.

This argument is easily generalized to other spark distributions and other dimensions. We find that the HOT forest fire model generates a perfect power law with slope $-(2 + 1/d)$




FIG. 1: Cumulative distribution of fire sizes for simulations of the forest fire model in one and two dimensions, plotted on logarithmic scales. (a) One dimension with size 10 000 and an exponential spark distribution. (b) Two dimensions with size 128×128 and a Gaussian spark distribution. The results have been averaged over a number of different values of the parameters of the spark distributions to improve the statistics. The dotted lines show the expected slopes of -2 and $-\frac{3}{2}$.

for all dimensions when a spark distribution of the form $p(\mathbf{r}) = N \exp(-\sum_{i=1}^d [x_i/\sigma_i]^d)$ is used. As a test of this prediction, we show in Fig. 1 numerical results from direct simulations of the forest fire model in one and two dimensions with distributions of this type [9]. For better visualization and analysis the distributions pictured are cumulative, so the expected slope is $-(1 + 1/d)$, rather than $-(2 + 1/d)$. As the figure shows, the slopes of the observed distributions are in good agreement with this prediction.

We note in passing that the slope of $-(1 + 1/d)$ for the cumulative distribution of fire sizes seen in both our exact solution and our numerical results is different from the slope of approximately -1 found numerically by Carlson and Doyle [2] in two dimensions. The source of this discrepancy is unclear, although it may be that the simulations of Ref. 2 provided too few data points to make an accurate evaluation of the exponent possible. We

note also that the value $-\frac{3}{2}$ for the two-dimensional case is quite different from the slope of $-\frac{1}{2}$ measured for the cumulative size distribution of real forest fires [3, 6].

Other functions with exponential tails also generate power laws, but give logarithmic corrections as well. For instance, if $p(\mathbf{r}) = N \exp(-\sum_{i=1}^d [x_i/\sigma_i]^\gamma)$ with $\gamma \neq d$ then Eq. (10) still applies, but now $A(p) \sim [\log(N/p)]^{d/\gamma}$ and hence

$$p(\mathbf{r}) \frac{d^d r}{dp} \sim [\log(N/p)]^{d/\gamma-1}. \quad (12)$$

Thus the distribution of event sizes fundamentally still follows a power law with slope $-(2+1/d)$, but there is a logarithmic correction. Similar logarithmic corrections are noted in Ref. 1.

Spark distributions with power-law tails also (unsurprisingly) give power-law event size distributions, but in this case the exponent of the distribution is non-universal, varying with the exponent of the spark distribution. For example, if $p(\mathbf{r})$ takes the generalized Lorentzian form

$$p(\mathbf{r}) = \frac{N}{\sum_i (x_i/\sigma_i)^\nu + \Gamma^\nu}, \quad (13)$$

then we find

$$\frac{d^d r}{dp} \sim \frac{1}{p(\mathbf{r})} \left(\frac{N}{p} - \Gamma^\nu \right)^{d/\nu-1}, \quad (14)$$

which goes asymptotically as $p^{-d/\nu}$ in the power-law tail where $p(\mathbf{r})$ becomes small. Thus the tail of the distribution of event sizes goes as $\rho(s) \sim s^{-(2+1/d-d/\nu)}$.

We now turn to the COLD variant of the forest fire model, which incorporates risk aversion. In constructing this model we are guided by theories of risk aversion in economics, where the subjective benefit of outcomes is typically a nonlinear function of the loss s , which is captured by a utility function $u(s)$ [7]. Sensible utility functions are decreasing with increasing loss: $u' < 0$ [10]. Risk aversion also implies that $u'' < 0$, so that the negative utility of bad outcomes is weighted more strongly than the positive utility of good outcomes. One standard family of utility curves that achieves this is the one-parameter family

$$u(s) = \frac{(1-s)^\alpha}{\alpha}. \quad (15)$$

Note that, since we will be concerned only with maximizing utility, $u(s)$ is arbitrary to within both additive and multiplicative constants. For $\alpha = 1$, Eq. (15) gives $u = 1 - s$ and

maximizing utility is precisely equivalent to minimizing loss. For $\alpha < 1$, we have risk-averse utility functions and for $\alpha < 0$ we are infinitely averse to losing our entire investment.

Our goal now is to maximize the average utility functional

$$U = \int p(\mathbf{r})u(s(\mathbf{r})) d^d r, \quad (16)$$

subject to the constraint of fixed F , Eq. (1) (or equivalently maximize a combined utility functional similar to Eq. (2)). Carrying out the functional derivatives and using the utility function of Eq. (15), we find that the optimum U corresponds to

$$p(\mathbf{r})s(\mathbf{r})^{(d+1)/d}[1 - s(\mathbf{r})]^{\alpha-1} = \lambda, \quad (17)$$

where λ is a Lagrange multiplier whose value can be calculated from Eq. (1). The distribution of event sizes is given by Eq. (6) as before, and using Eq. (17) we find that the derivative dp/ds , which gives the principal variation in $\rho(s)$, is

$$\frac{dp}{ds} = \lambda \frac{(\alpha + 1/d)s - (1 + 1/d)}{(1 - s)^\alpha s^{2+1/d}}. \quad (18)$$

For $\alpha = 1$ our utility maximization is equivalent to simple yield maximization (HOT), so it is not surprising to observe that when we set $\alpha = 1$ in the above expression we recover our previous $s^{-(2+1/d)}$ power-law. For $\alpha < 1$, we have risk-averse utility functions (COLD), which give rise to event distributions following the $s^{-(2+1/d)}$ form for small event sizes, but having lower probability of large event sizes. When $\alpha < 0$, event probability tends to zero as $s \rightarrow 1$, as we would expect.

In Fig. 2 we compare the distribution of event sizes in HOT and COLD regimes for a variety of values of the risk-aversion parameter α . The figure shows that the COLD distribution approaches the HOT one as α approaches 1. For α large and negative the HOT power law is followed for only a small portion of the range of event sizes—about 20% in the case of $\alpha = -5$.

It is worth noting that while risk aversion truncates the power-law behavior in the event size distribution, the distribution of the *utilities* of events still follows a power law: we find that the tail of the distribution of utilities goes as $\rho(u) \sim u^{-\beta}$ with $\beta = (2\alpha - 1)/\alpha$. Note that this exponent is independent of the system dimension d [11].




FIG. 2: Event size distributions for HOT and COLD regimes in two dimensions. The dotted line is the distribution for the HOT regime ($\alpha = 1$) and the solid lines are for the COLD regime with (top to bottom) $\alpha = -1 \dots -5$. The distributions are not normalized—they cannot be since they diverge at the origin. In practice Eq. (1) provides a lower cutoff on s and makes the distribution normalizable. Inset: the same data on log-log scales.

While the introduction of the utility function has reduced the risk of large losses, the optimal utility solution does not normally coincide with the optimal yield solution, and hence we pay a cost for risk aversion in terms of yield. For the lattice forest fire model however we find that the cost paid is small [12]. For example, in a 128×128 two-dimensional system with a Gaussian spark distribution, we find numerically that the mean yield at the $\alpha = 1$ optimum (HOT) is 0.904, dropping to 0.900 for $\alpha = -3$ and 0.888 for $\alpha = -5$. It appears therefore that the introduction of risk aversion garners substantial benefits in terms of the reduction of large losses—and the complete elimination of 100% losses—while at the same time costing us only a few percent at most in terms of average system yield.

We conjecture that the suppression of power law tails in the COLD event size distribution will also make the system more robust against the other problems mentioned in the intro-

duction, namely errors in the design and changes in the spark distribution. The truncation of the power law means that the largest patches in the COLD solution are considerably smaller than those in the HOT solution. Thus, if a design flaw, such as a gap in one of the firebreaks, causes two patches to merge, the resulting combined patch is smaller too. Similarly, if the spark distribution is changed, the size of the resulting fires is smaller, and hence the effect on the average yield is not as catastrophic as in the HOT case. (A similar conjecture is also made in Ref. 1.)

In this paper we have examined in detail only the forest fire model, but similar principles should apply to other problems as well. We have shown that in order to produce power-law event size distribution, the HOT model requires the auxilliary assumption of risk-neutrality. If humans are risk-averse they will tend to prefer COLD designs, although this does not necessarily mean that HOT designs never occur. It might be, for instance, that blind evolutionary processes of the type found in natural systems would simply optimize yield, without risk aversion. On the other hand, COLD designs are more robust to rare events than HOT designs, and therefore might be selected for on long time-scales. Of course, in the real world, imperfect designs that fail to optimize either yield or utility are always a possibility too.

The authors thank Jean Carlson and John Doyle for many useful comments and suggestions. This work was supported in part by the National Science Foundation under grant number DMS-0109086 (MEJN and MG) and by McKinsey Corporation, Credit Suisse First Boston, Bob Maxfield, and Bill Miller (JDF).

-
- [1] J. M. Carlson and J. Doyle, Highly optimized tolerance: A mechanism for power laws in designed systems. *Phys. Rev. E* **60**, 1412–1427 (1999).
 - [2] J. M. Carlson and J. Doyle, Highly optimized tolerance: Robustness and design in complex systems. *Phys. Rev. Lett.* **84**, 2529–2532 (2000).
 - [3] J. Doyle and J. M. Carlson, Power laws, highly optimized tolerance, and generalized source coding. *Phys. Rev. Lett.* **84**, 5656–5659 (2000).
 - [4] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley, New York (1991).

- [5] M. E. J. Newman and R. M. Ziff, Efficient monte carlo algorithm and high-precision results for percolation. *Phys. Rev. Lett.* **85**, 4104–4107 (2000).
- [6] B. D. Malamud, G. Morein, and D. L. Turcotte, Forest fires: An example of self-organized critical behaviour. *Science* **281**, 1840–1842 (1998).
- [7] L. J. Savage, *The Foundations of Statistics*. Dover, New York, 2nd edition (1972).
- [8] As noted in Ref. 3, the exponent of this power law *decreases* with increasing dimension d , which is the opposite of the behavior seen in other models such as the self-organized critical forest fire model of B. Drossel and F. Schwabl, *Phys. Rev. Lett.* **69**, 1629 (1992).
- [9] The simulations were carried out using the greedy algorithm of Ref. 1, in which trees are placed on the lattice one by one, each in the position that maximizes the yield at that step. This is not a true global optimization, but appears to give results close to the global optimum and in good agreement with theory.
- [10] In traditional economics utility functions are usually defined to be increasing functions of their arguments; this could be arranged in the present case, if necessary, by substituting $s \rightarrow -s$.
- [11] A similar result can also be derived within the framework of the “PLR” model of Ref. 3, which is a general model of engineered systems subject to stresses that occur with specified probability. With this model a risk-averse utility function of the form (15) again produces a sharp cutoff in the event distribution, but gives a distribution of the utilities of events that has a dimension-independent power-law tail.
- [12] For other models or parameter regimes, such as the continuum version of the forest fire model with large a or the models of Ref. 3, the cost paid could certainly be larger.